From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2025

Essays in Information Economics: Attention, Emotions, and Social Image

Vaeth, Martin

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp016d570104f

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Essays in Information Economics: Attention, Emotions, and Social Image

Martin Vaeth

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Economics
Adviser: Roland Bénabou, Alessandro Lizzeri

May 2025

© Copyright by Martin Vaeth, 2025.

All rights reserved.

Abstract

The first chapter models electoral competition between two parties when voters can rationally learn about their political positions through flexible information acquisition. Rational voter learning generates polarized and aligned political preferences, even when voters' true positions are unimodally distributed and independent across policy issues. When parties strategically select their positions, voter and party polarization mutually reinforce each other, and both rise as information costs decline. Because voters learn exclusively about the axis of party disagreement, party positions respond to only one dimension of aggregate shocks to voter preferences.

The second chapter explains regret as an optimal self-control mechanism to motivate attention and thereby improve decision-making. The model endogenizes the optimal emotions as incentives for an agent who overweights the cost of attention, for example due to temptation or present bias. If, ex post, the realized state is observable, the model provides a foundation for regret theory, including disproportionate aversion to large regrets. Advancing regret theory, the model explains why regret is stronger than rejoicing and why it is stronger in simpler decision problems.

The third chapter studies the optimal design of public recognition schemes to incentivize agents who care about their social image—the public's belief about their private type, such as ability or prosociality. We allow public recognition schemes to take the form of any signal structure, employing an information design approach. When agents are risk neutral over image, we show optimality of monotone partitional recognition schemes. If agents are risk averse over image, it may be optimal to maintain full privacy and not screen the agent. This highlights a difference between public recognition schemes and monetary incentive schemes, where screening remains optimal even under risk aversion.

Acknowledgements

Many people have contributed to this dissertation. I am especially indebted to Roland Bénabou, who supported me from the beginning and guided me throughout the PhD. He encouraged me to stay true to my ideas and gave me the confidence to pursue them.

I am deeply grateful to Alessandro Lizzeri, whose advice was invaluable before and during the job market. He taught me a lot about how to think about research and how to communicate it.

Fedor Sandomirskiy's curiosity and excitement for theory have been a source of inspiration. Thinking through problems with him has shaped how I approach mathematical challenges.

I have spent countless hours in meetings with Pietro Ortoleva, Wolfgang Pesendorfer, Faruk Gul, Xiaosheng Mu, Matias Iaryczower, Sylvain Chassang, and Leeat Yariv. These conversations have shaped my research, and I am grateful to all of them.

I have greatly benefited from discussions with Mira Frick, Germán Gieczewski, Tom Griffiths, Ryota Iijima, Gleason Judd, John Londregan, Nolan McCarty, Nick Netzer, Christopher Sims, Can Urgun, and Juan-Pablo Xandri.

I gratefully acknowledge helpful feedback from Alexander Bloedel, Simon Bornschier, Steven Callander, Mark Dean, Laura Doval, Armin Falk, Jiadong Gu, En Hua Hu, Andreas Kleiner, Botond Kőszegi, Stephan Lauermann, Elliot Lipnowski, Filip Matejka, Benny Moldovanu, Jacopo Perego, Doron Ravid, Ali Shourideh, Aidan Smith, Guido Tabellini, Mark Whitmeyer, Florian Zimmermann, seminar participants in Berkeley, Bonn, CERGE-EI, Copenhagen, Princeton, UCL, and Warwick, and participants at the Calgary/SFU/UBC Theory Workshop, the European Winter Meeting of the Econometric Society 2023, the Stony Brook Game Theory Conference 2024, and the BFI Rational Inattention Conference 2025.

I have also benefited from discussions with fellow Princeton students, including Narek Alexanian, Francesco Fabbri, Dan McGee, Paola Moscariello, Joseph Ruggiero, Alejandro Sabal, Hassan Sayed, and João Thereze.

I would also like to thank Laura Hedden, the most amazing administrator, who was of great help throughout the PhD.

Financial support was provided by The William S. Dietrich II Economic Theory Center.

Finally, I would like to thank Jess, my family, and my friends for making this journey enjoyable.

Contents

Abstract 3

Acknowledgements 4

List of Tables 8

List of Figures 9

1 Rational Voter Learning, Issue Alignment, and Polarization 10

1.1 Introduction 10

1.2 Model 19

1.2.1 Discussion 25

1.3 Voter Learning 28

1.3.1 Issue Alignment 28

1.3.2 Polarized Ideology 33

1.3.3 Comparative Statics 36

1.3.4 Evidence 39

1.4 Electoral Competition 44

1.4.1 Platform Polarization 46

1.4.2 Aggregate Uncertainty 49

1.5 Alternative Timing 52

1.5.1 Model 52

1.5.2 Voter Learning 55

1.5.3 Equilibrium 58

1.6 Horizontally Differentiated Goods 61
1.7 Conclusion 67
1.8 Appendix: Main Proofs 68
1.9 Appendix: Alternative Timing 87
1.10 Appendix: IO Extension 105
1.11 Appendix: Other Results and Proofs 112

2 Attention and Regret **136**
2.1 Introduction 136
2.2 Model 144
2.3 Results 149
2.3.1 Two Actions 152
2.3.2 More than Two Actions 162
2.4 Interpretation as Contracting with an Expert 164
2.4.1 Hiring an Expert with Conflicting Interests 165
2.5 Conclusion 169
2.6 Appendix 170
2.7 Online Appendix 200

3 The Optimal Design of Public Recognition Schemes **212**
3.1 Introduction 212
3.2 Baseline Model 216
3.3 Risk Neutrality over Image 219
3.3.1 Implications 223
3.4 Risk Aversion over Image 226
3.5 Appendix 233

Bibliography **240**

List of Tables

2.1 Exemplary payoff-matrix 150

List of Figures

1.1 Exemplary policy platforms x_a and x_b of parties a and b, respectively, under two policy issues. Voter learning induces revealed ideal points on the diagonal line. 12

1.2 Reflection argument underlying the proof of Theorem 1. The dark clouds visualize the distribution over posterior means. 29

1.3 Symmetric equilibrium with party platforms (x_a, x_b) and party ideal points (x_a^*, x_b^*) 47

1.4 Party platforms as a function of info cost parameter κ 59

1.5 Consumer welfare increases as information becomes more costly. 66

1.6 The value function $v(\theta)$ 88

1.7 Value function and feasible posterior means under threshold signal t' 90

1.8 The value function $v(\hat{\theta}, \alpha)$ as a function of $\hat{\theta}$ for two values of α. The higher alpha corresponds to the steeper value function, which offers a higher return to information. 127

3.1 Example of an extreme $P(\theta)$ under $\theta \sim U([0,1])$. 223

3.2 Pareto frontiers under incentive provision through image and a non-informational good. 231

Chapter 1

Rational Voter Learning, Issue Alignment, and Polarization

1.1 Introduction

Voter positions in the United States display two puzzling features. First, one can predict a voter's position on most policy issues remarkably well by knowing just their location on a one-dimensional left-right axis. As a consequence, positions on different issues are strongly aligned, which is surprising given the wide variety of issues, such as taxation, immigration, and the environment. Second, evidence suggests that voter positions are increasingly polarized, in the sense of being clustered around two poles on the left-right axis. In most settings, we expect the distribution of characteristics to have a unimodal distribution by considerations such as the central limit theorem.[1]

This paper provides a joint explanation of issue alignment and polarization based on rational voter learning. Previous research and public discourse attribute issue alignment

[1]Section 1.3.4 examines the evidence for issue alignment and ideological polarization in detail. While the evidence for issue alignment is strong, the evidence on bimodality is more speculative; there is ongoing debate about this issue (Fowler et al., 2022). Understanding the sources of potential polarization is important because of its relevance for political tensions, as emphasized by the literature on polarization (Esteban and Ray, 1994): a polarized distribution can lead to the formation of two homogeneous groups with few individuals bridging the divide, increasing the risk of conflict.

and polarization to voter biases like confirmation bias, herding due to echo chambers, or partisan news media. By contrast, we show issue alignment and polarized ideology emerge naturally from rational (i.e., no biases) and individual (i.e., no herding) voter learning, driven by voters' own information choices (i.e., no media effects). Central to the mechanism is that voters learn about their political position through flexible information acquisition in order to decide between two parties. Voter learning can involve understanding the effects of policies—for example, learning about the effects of tariffs to inform their position on trade policy. We assume that such learning is costly, whether in time, effort, or money. This paper shows that the resulting cost-minimization motive structures rational voter learning in a way that generates issue alignment and polarization.

Why study rational voter learning even though voters may be influenced by various biases? An explanation based on rational learning can help disentangle behaviors driven by biases from those arising from rational ideology formation. This distinction is important for assessing the functioning of democratic elections (Achen and Bartels, 2017). If biases dominate, candidates may cater to these biases rather than advancing policies that address societal needs. Conversely, if voter positions reflect rational learning, we may be more optimistic that elections produce policies aligned with voters' interests. Whether this optimism is warranted is explored in the second part of this paper, which studies the effects of voter learning on party positions.

We analyze rational voter learning in an otherwise standard political-economy model. To study issue alignment, we incorporate a multidimensional policy space, where each dimension represents a different policy issue. Voters face two parties, each adopting a policy platform in this space. These policy platforms are exogenous in the first part of our paper. Following the literature, a voter's utility decreases quadratically in the distance between a policy and her *ideal point*, which reflects her political position. In our model, voters initially lack knowledge of their ideal point but can form political opinions through learning. We assume that voters start with a homogeneous prior that conforms to the true distribution of ideal points, which

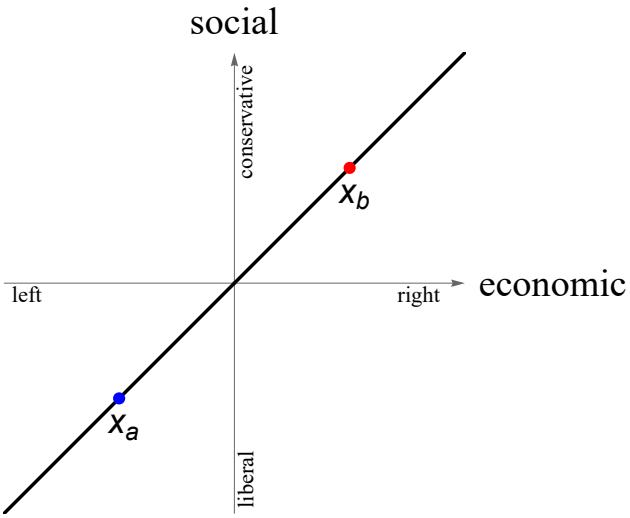


Figure 1.1: Exemplary policy platforms x_a and x_b of parties a and b, respectively, under two policy issues. Voter learning induces revealed ideal points on the diagonal line.

in the simplest case is normal and independent across issues.[2] Voters then update their beliefs by acquiring information flexibly, following the rational-inattention framework (Sims, 2003). This framework allows voters to choose both how much and what type of information to acquire, subject to a cost increasing in information. After acquiring information, and given quadratic utility, a voter selects the party whose policy bundle is closer to her expected ideal point. We refer to this expected ideal point as her *revealed ideal point*, distinguishing revealed ideology—the distribution of revealed ideal points induced by learning—from the distribution of true ideal points.

Our first result is that rational voter learning generates issue alignment and polarization: revealed ideal points align along a left-right axis and cluster around two poles, even if true ideal points are independent across issues and centrally clustered. At the core of this mechanism is that voters trade off making an informed choice and minimizing the cost of information. By acquiring only the information necessary to determine which party position is closer, voters align their preferences along a single axis and polarize—a process we explain in more detail next.

[2]More generally, the paper allows for any elliptical distribution with arbitrary correlation across issues. Section 1.3.1 discusses the case of heterogeneous priors.

To understand the high-level intuition for why cost-effective learning generates issue alignment, consider a two-dimensional policy space with an economic and a social issue. Suppose the two parties propose policy platforms, x_a and x_b, as shown in Figure 1.1: one platform is more economically left and socially liberal than the other. A unique line passes through these platforms, the direction of which we refer to as the *axis of disagreement*. This axis of disagreement is the direction along which parties differentiate, combining the two issues in proportion to the extent of disagreement on each. Crucially, voters learn only where their ideal point lies along this axis, as this determines which platform is closer. Their position along directions orthogonal to the axis is irrelevant, as it does not affect the *relative* proximity of the platforms. Consequently, voters align preferences across issues, learning only whether they lean economically left and socially liberal or economically right and socially conservative. Revealed ideal points lie on a line, which can be interpreted as an endogenous left-right axis. This left-right axis reflects the direction of party disagreement and passes through voters' prior expectation of their ideal points. This result holds for any number of policy issues, assuming a reflection-invariant information cost and an elliptical prior.

The mechanism by which cost-effective learning generates ideological polarization is as follows. Voters only need to determine which party they are closer to, not by how much. To do so, they optimally acquire a *binary* signal about their ideal points. This signal sorts voters into two groups: those who believe they are closer to one party and those who believe they are closer to the other. Even centrist voters categorize themselves based on this binary information, causing them to perceive their preferences as more extreme than they truly are. We also examine how robust this mechanism is to shocks about the desirability of the two parties that occur after learning. For example, voters may receive new information about candidates' competence closer to the election. These "valence shocks" make it valuable for voters to learn *how much* they prefer one party over the other, which previously held no value. We show that the distribution of revealed ideal points remains confined to the endogenous left-right axis. For small valence shocks, this distribution is bimodal.

The second part of this paper examines whether elections produce policies that serve voters' interests when voter positions result from rational learning. To address this, we endogenize party platforms and explore their interaction with voter learning. As standard, parties strategically position themselves in the policy space, balancing electoral prospects against their own policy preferences. In equilibrium, voters learn optimally in anticipation of party platforms, and parties choose their platforms based on the revealed ideology shaped by this learning. We show two results about party platforms: more voter learning can hurt voters by increasing party polarization, and parties respond to only one dimension of aggregate shocks to voter preferences.

First, ideological polarization of voters and platform polarization are mutually reinforcing, and perhaps counterintuitively, less costly information increases both. When party platforms are more polarized—that is, farther apart—voters face higher stakes in the election and are motivated to learn more about their ideal points. More information leads to a more polarized distribution of voter ideal points, creating more extreme voters who are less responsive to party platforms. As a result, parties can move their platforms closer to their ideal policies without losing as many votes, reinforcing platform polarization. Cheaper information amplifies this cycle: it encourages voters to learn more about their preferences, leading to more voter and platform polarization. Paradoxically, better access to information about political preferences may harm voters, as the equilibrium platforms end up farther from the welfare-maximizing policy. This mechanism may help explain the increasing platform polarization in the US over recent decades (McCarty et al., 2016), during advances in information technology, such as the internet, which has made information more accessible.

Second, due to rational voter learning, policy platforms respond solely to a single dimension of aggregate shocks to voter preferences. Aggregate shocks to voter preferences affect what the optimal policy should be across multiple dimensions; ideally, we would want voters to learn about these shocks so that parties can adjust their platforms accordingly. However, voters' optimal learning reduces politics to a single dimension—the axis of party

disagreement—and neglects all other dimensions. As a result, party platforms respond solely to one dimension of aggregate shocks. This inefficiency is particularly problematic because it is not apparent from revealed voter ideology: aggregate shocks manifest as one-dimensional changes along the axis of disagreement, and party platforms adjust in response to these observed changes. Meanwhile, the other dimensions of aggregate shocks remain latent—they do not influence revealed voter ideology and thus go unnoticed in empirical data.

We also study an alternative timing where parties move before voters learn, allowing parties to strategically affect voter learning through their choice of policy platform. This extension reveals two novel strategic forces: a moderation force and a differentiation force. First, parties may moderate their policy platforms to skew voter learning in their favor. Second, the more extreme party may find it optimal to differentiate from the moderate party to trigger more voter learning, thereby reducing the skew toward their opponent. As information costs decline, the moderation force weakens, leading to increased platform polarization, as in our main timing. The second force implies that parties may adopt positions more extreme than their own ideal policies to strategically affect voter learning.

To illustrate the broader applicability of our results, we adapt our model to a market setting where consumers learn about their preferences for horizontally differentiated products. In this context, firms not only choose product attributes—similar to how parties select policy platforms—but also set prices to maximize profits. We show that cheaper information leads consumers to become better informed, prompting firms to increase product differentiation. While this could benefit consumers by better matching them to products, product differentiation harms consumers because firms exploit it to raise prices. As a result, despite the reduction in information costs, consumer welfare decreases overall.

Our results call for caution when interpreting empirical findings about voter ideology being influenced by party elites as a sign of voter irrationality. Political scientists have long argued that political elites have a large influence on the ideology of voters (Campbell et al., 1960; Zaller, 1992; Lenz, 2012). Such findings have traditionally been interpreted as evi-

dence of voter irrationality. This interpretation is shared by Achen and Bartels (2017), who argue such voter behavior presents a serious threat to democracy. If parties can shape the ideology of voters instead of merely responding to it, it is unclear whether elections produce governments responsive to the preferences of voters. We hope to contribute to this debate by showing that some forms of party influence on voter ideology are consistent with voter rationality and do not preclude that policy is responsive to voters' true preferences. In our model, both issue alignment and polarization of voter ideal points depend on party platforms. First, the alignment of voter ideal points across policy issues is determined by the relative positions of parties. As illustrated in Figure 1.1, because party a is more left and liberal than party b, in the resulting voter ideology a left economic position aligns with a liberal social position. Second, more polarized party platforms result in more polarized voters (Proposition 1). Although in both cases voters seem to simply follow party positions, our model shows such effects result from rational voter learning. Moreover, in our equilibrium, party platforms do respond to voters' true ideal points, namely to the center of their distribution (Theorem 3). However, parties do not respond to aggregate shocks to voter preferences in more than one dimension (Theorem 4). On a higher level, our model illustrates how revealed preferences may differ systematically from true preferences and how they may do so in a context-dependent way.

Our model has policy implications for addressing issue alignment and polarization. Traditional approaches—such as improving political knowledge (Carpini and Keeter, 1996) or breaking up echo chambers (Sunstein, 2018)—may be ineffective if issue alignment and polarization stem from rational voter learning. Instead, more fundamental changes to voters' choice sets would be necessary to incentivize multidimensional and non-binary learning. Multidimensional learning could be incentivized through electoral reforms that increase the number of parties, such as transitioning from plurality elections to proportional representation (as suggested by Corollary 1).[3] Non-binary learning could be incentivized through

[3]Under $k > 2$ parties, a weaker form of issue alignment is predicted by our model: revealed voter ideology is at most $(k-1)$-dimensional.

voting or participation mechanisms that elicit the intensity of voter preferences (e.g., Casella, 2005).

The paper is organized as follows. The rest of this section discusses related literature. Section 1.2 introduces the model. Section 1.3 analyzes voter learning and discusses the related evidence. Section 1.4 studies electoral competition when voter ideology results from optimal learning. Section 1.5 studies the model under an alternative timing, in which parties move before voters learn. Section 1.6 adapts our model to an industrial organization setting with horizontally differentiated products. Section 1.7 concludes.

Related Literature This paper contributes to the growing literature on rationally inattentive voters. We show rational inattention explains properties of voter ideology by studying *flexible* learning about *ideal points*. Matějka and Tabellini (2021) study electoral competition where voters are inattentive to *party platforms*. They show more attentive voters influence platforms more strongly, as they respond more to them. By contrast, we show attention to *ideal points* reduces responsiveness to platforms, which increases polarization when information becomes cheaper. Matějka and Tabellini (2021) find that multidimensional policies are inefficient because voters pay excessive attention to divisive issues. In our model, inefficiency arises because voters focus on a *single* dimension, the axis of party disagreement. Yuksel (2022) analyzes voters learning under partitional signals and finds that more specialized learning increases polarization. We allow for flexible learning and show that cheaper information increases polarization through a different mechanism. Li and Hu (2023) study attention to implemented policies in an electoral-accountability setting. They show that the welfare effects of increased attention and mass polarization are ambiguous. Hu et al. (2023) analyze learning about valence through an attention-maximizing intermediary. They show it generates policy divergence even with office-motivated candidates. By contrast, we study divergence as arising from ideologically motivated candidates and study how it interacts with voter learning.

The literature has proposed several explanations for issue alignment and polarized ideology. Converse (1964) introduced the concept of issue alignment, which he termed ideological constraint, attributing it to logical, psychological, and social sources. Enke et al. (2023) propose that moral universalism, the degree of altruism toward strangers versus in-group members, explains correlations in policy views across domains. Spector (2000) derives one-dimensional ideology from cheap talk between two groups with differing priors, while DeMarzo et al. (2003) attribute it to networks and a persuasion bias. Bayesian persuasion can also generate low-dimensional types (Rayo and Segal, 2010; Malamud and Schrimpf, 2022).[4] Ideological polarization has been linked to cognitive limitations, such as correlation neglect (Levy and Razin, 2015; Ortoleva and Snowberg, 2015) and bounded rationality (Eguia and Hu, 2022). Other papers have also proposed rational inattention as a mechanism for belief polarization, albeit through ex-ante heterogeneity in preferences (Novák et al., 2024) or path-dependent sequential information acquisition (Nimark and Sundaresan, 2019). Callander and Carbajal (2022) explain dynamic polarization through voters adjusting their ideal points toward the party they voted for to rationalize their choice. Our model complements these works by providing a unified explanation for both issue alignment and polarization.

The literature has provided many mechanisms for platform divergence, breaking the median voter result by Downs (1957). We do not propose a new mechanism but show how the new ingredient of our model—endogenous ideology formation through voter learning—interacts with perhaps the most prominent mechanism for platform divergence: ideologically motivated parties and probabilistic voting. Ideological parties still converge to the median voter unless the electoral outcome is uncertain (Wittman, 1983; Hansson and Stuart, 1984; Calvert, 1985). The two common ways to introduce electoral uncertainty are through uncertainty about the ideological position of the median voter (Roemer, 1994) and through valence shocks (Hinich, 1977; Lindbeck and Weibull, 1987). Our model falls into the latter

[4]More broadly related is McMurray (2023), who shows that pivotality considerations in multidimensional common-value elections drive party platforms to bundle logically related issues, reducing platforms to a persistent one-dimensional axis.

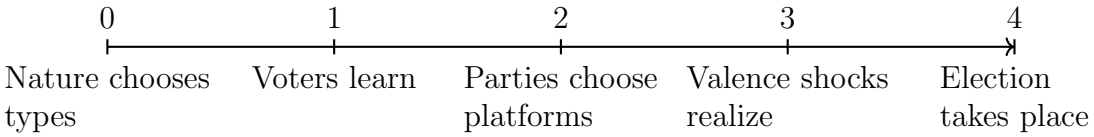
category, which, according to Duggan (2017), has seen little, if any, formal analysis under ideologically motivated parties. We introduce a version of this model that is tractable even in a multidimensional policy space. The model generalizes the mean-voter theorem by Hinich (1977) and allows comparative statics with regards to platform polarization.

Finally, we contribute to the burgeoning theoretical literature on rationally inattentive consumers in industrial organization. As we illustrate in section 1.6, our model can be adapted to a setting of consumer learning in the face of horizontally differentiated products. Of particular relevance are Albrecht and Whitmeyer (2023) and Biglaiser et al. (2024), who study a duopoly with consumers learning about their preferences. Albrecht and Whitmeyer (2023) show consumers learn only about the relative value of products and that, in contrast to Ravid et al. (2022), an ex-post efficient equilibrium exists as the information cost converges to zero. Biglaiser et al. (2024) study comparative statics of the unique symmetric equilibrium and show an application to platform design. In both papers, as is standard, consumers learn directly about their valuations of products, whereas we assume they learn about their ideal products in an attribute space. The additional structure on preferences facilitates an analysis of endogenous product *attributes*, whereas the literature typically focuses on the effect of attention on firm's *pricing* decisions. An exception to this is Cunha et al. (2022), who study a spatial setting where consumers pay attention to product attributes and prices, whereas in our model consumers learn about their preferences. As a result, our model makes predictions on the distribution of consumer preferences, which are endogenously one-dimensional allowing us to tractably study a multidimensional attribute space.

1.2 Model

We employ a standard probabilistic voting model with valence shocks (Hinich, 1977; Lindbeck and Weibull, 1987) and add to it an earlier stage in which voters learn about their ideal points, anticipating the election. We discuss our assumptions at the end of this section.

Game The policy space is $\mathbb{R}^n$ with $n \in \mathbb{N}$, where each dimension corresponds to a policy issue. Voter ideal points as well as platforms live in this space. There is a continuum of voters $i \in [0,1]$ and two parties, a and b. The timing is as follows.



(0) Nature draws voters' ideal points $\theta_i \in \mathbb{R}^n$ independently from an elliptical distribution $\mu \in \Delta(\mathbb{R}^n)$, with mean normalized to 0 and an arbitrary positive definite covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$.[5]

An example is the normal distribution $\mathcal{N}(0, \Sigma)$, but elliptical distributions can also have bounded support.

(1) Each voter $i \in [0,1]$ chooses how to learn about their ideal point $\theta_i \in \mathbb{R}^n$.

(2) Parties a and b observe the realized distribution of voter preferences and commit to policy platforms $x_a \in \mathbb{R}^n$ and $x_b \in \mathbb{R}^n$, respectively.

(3) Each voter observes party platforms (x_a, x_b) and the realization of the valence shocks $\nu_i \in \mathbb{R}$, which can be interpreted as a signal about the competence differential of party candidates.

(4) Voters elect their preferred parties. The party with the majority wins, and each party wins with probability one half if a tie occurs.

Because each voter is infinitesimal and there is no aggregate uncertainty about voter ideal points, this timing is equivalent to one where stages one and two occur simultaneously. If, alternatively, parties moved before voters learned, the results on voter learning (Theorem 1 and 2) would remain unchanged. The qualitative predictions on voter ideology do not

[5]For any topological space X, we denote by $\Delta(X)$ the set of Borel probability measures on X. A measure $\mu \in \Delta(\mathbb{R}^n)$ with mean 0 and covariance matrix Σ is *elliptical* if its characteristic function Φ takes the form $\Phi(\theta) = \psi(\theta^\top \Sigma \theta)$ with $\psi \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. If μ admits a density f, it must be of the form $f(\theta) = g(\theta^\top \Sigma^{-1} \theta)$ with $g \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. This symmetry assumption states that the isodensity curves are ellipses. Independence of ideal points is not necessary for our results on voter ideology in section 1.3.

depend on whether voters learn given observed or given anticipated platforms. Our timing rules out that parties choose platforms in order to affect voter learning. In section 1.5, we show the comparative statics on polarization (Theorem 3) extend to the reversed timing.

Voters The utility U_i of voter $i \in [0,1]$ has three components, which we expand on below. Utility depends on the implemented policy x and the voter's ideal point θ_i via the policy utility $u(x, \theta_i)$, on the net valence shock ν_i for candidate b, and on the information cost $c(\tau_i)$ of signal structure τ_i scaled by κ,

$$U_i(x, \tau_i) = u(x, \theta_i) + \nu_i \cdot \mathbb{1}_{\{\text{vote } b\}} - \kappa \cdot c(\tau_i). \tag{1.1}$$

Following the workhorse model of spatial voting, we assume voters' policy utility is quadratic in the difference between their ideal point and the policy; that is,

$$u(x, \theta) = -(x - \theta)^\top A (x - \theta), \tag{1.2}$$

where $A \in \mathbb{R}^{n \times n}$ is an arbitrary symmetric, positive definite matrix. Although this assumption is restrictive, most of the evidence on voter ideal points assumes quadratic utility with a homogeneous A and shows it can explain voters' survey responses well (see section 1.3.4).[6] Further, quadratic utility allows us to speak about "revealed ideal points" of voters who have a non-degenerate belief π about their true ideal point θ, as the following remark shows.

Remark 1. *A voter with belief π over her ideal point θ votes as if she had the known ideal point $\mathbb{E}_\pi[\theta]$.*

By a bias-variance decomposition of the expected utility from policy x under belief π over θ,

$$\mathbb{E}_\pi[u(x, \theta)] = u(x, \mathbb{E}_\pi[\theta]) - \mathbb{E}_\pi\big[(\theta - \mathbb{E}_\pi[\theta])^\top A (\theta - \mathbb{E}_\pi[\theta])\big]. \tag{1.3}$$

[6] Moreover, in recent work, Bachmann et al. (2024) find that out of a range of commonly used algorithms, the one based on quadratic utility performs best at predicting survey responses.

The latter term does not depend on x, so when a voter compares two platforms (or survey response items), they choose the one that is closer to their posterior mean. Hence, a voter with belief π acts like a voter with a known ideal point of $\mathbb{E}_\pi[\theta]$. Accordingly, we call the posterior mean of a voter's belief π over θ her *revealed ideal point*. This ideal point is the one that is estimated from survey responses, which is important when we interpret empirical findings about voter ideology. We refer to the distribution $\rho \in \Delta(\mathbb{R}^n)$ over posterior means induced by learning as the *revealed ideology* in the population, to distinguish it from the true distribution over ideal points.

The valence shock ν_i is to be interpreted as the valence difference between parties b and a. It has the same distribution for all voters $i \in [0, 1]$. The distribution of ν has a finite first absolute moment and admits a continuous density f_ν that is symmetric around 0, and strictly quasiconcave. The symmetry of the valence shocks means that no party has a valence advantage, which simplifies our analysis of electoral competition. Strict quasiconcavity together with symmetry implies the density of the valence shock is maximal at 0. We show later that this assumption implies more extreme voters are less sensitive to party platforms. Because we assume parties care only about their expected vote shares, we do not need to specify the joint distribution of valence shocks.

Learning Voters share a homogeneous prior μ, conforming to the true distribution, before learning.[7] Each voter can acquire *any* signal structure (Blackwell experiment) about her ideal point at a cost proportional to mutual information, as in the rational-inattention literature (Sims, 2003; see also the survey Maćkowiak et al., 2023). The information cost captures that learning takes time and effort. To define mutual information, recall that a signal structure specifies a conditional distribution over signal realizations given any ideal point. Upon a signal realization, the agent forms a posterior via Bayesian updating. Thus, a signal structure induces a distribution over posteriors. Bayesian updating implies this distribution averages to the prior, also called Bayes consistency. In fact, following the posterior approach (Kamenica

[7] In section 1.3.1, we show Theorem 1 can be extended to allow for ex-ante heterogeneity in beliefs.

and Gentzkow, 2011; Caplin and Dean, 2013), we can represent signal structures as Bayes-consistent distributions $\tau \in \Delta(\Delta(\mathbb{R}^n))$ over posteriors $\pi \in \Delta(\mathbb{R}^n)$.[8] The mutual-information cost can then be defined as the expected Kullback-Leibler divergence[9] of posterior π from prior μ,

$$c(\tau) = \mathbb{E}_\tau[D_{\mathrm{KL}}(\pi||\mu)]. \tag{1.4}$$

Intuitively, the Kullback-Leibler divergence defines a "distance" on beliefs, and mutual information measures how much the acquired information moves the voter's belief away, on average, from her prior according to this "distance." We assume different voters' signal realizations are independent. The cost parameter κ in (1.1), which we vary for comparative statics, translates mutual information into utils.

For voters to acquire costly information despite never being pivotal among the continuum of voters, we assume voters engage in *expressive voting*, as is standard in the literature on rationally inattentive voters (Matějka and Tabellini, 2021; Hu et al., 2023; Li and Hu, 2023).[10] That is, voters genuinely care about voting for the correct candidate given their true preferences, for which they are willing to incur an information cost. The reason may be that voters derive a psychological benefit from doing so or they consider it their civic duty (see also Feddersen and Sandroni, 2006).

Formally, voter i first chooses information τ_i and, after the observation of valence ν_i and platforms (x_a, x_b), votes for $x \in \{x_a, x_b\}$ to maximize $U_i(x, \tau_i)$. Dropping indices, the voter's choice of information, that is, distribution $\tau \in \Delta(\Delta(\mathbb{R}^n))$ over posteriors $\pi \in \Delta(\mathbb{R}^n)$ that is

[8] That Bayesian updating imposes only Bayes-consistency of τ holds for general Polish state spaces, which includes $\mathbb{R}^n$, as a consequence of the disintegration theorem, as shown by Lipnowski and Ravid (2023), Appendix C.2.

[9] The Kullback-Leibler divergence of π from μ is defined as

$$D_{\mathrm{KL}}(\pi||\mu) = \begin{cases} \int_{\mathbb{R}^n} \log\left(\frac{d\pi}{d\mu}\right) d\pi & \text{if } \pi \ll \mu \\ \infty & \text{else,} \end{cases}$$

where $\frac{d\pi}{d\mu}$ is the Radon-Nikodym derivative and $\pi \ll \mu$ means π is absolutely continuous with respect to μ.

[10] Martinelli (2006) studies information acquisition in large elections assuming the pivotal-voter model. In large electorates, all voters are nearly uninformed.

Bayes-consistent (BC), must solve the following problem:

$$\sup_{\tau \in \Delta(\Delta(\mathbb{R}^n))} \int \left(\mathbb{E}_\nu \Big[\max \Big\{ \mathbb{E}_\pi[u(x_a, \theta)], \mathbb{E}_\pi[u(x_b, \theta)] + \nu \Big\} \Big] - \kappa D_{\mathrm{KL}}(\pi || \mu) \right) d\tau \tag{P}$$

$$\text{s.t.} \int \pi d\tau = \mu. \tag{BC}$$

The integrand of (P), which we call the *value function*, has the following interpretation. Given a posterior π, the voter anticipates that for each realization of the valence shock ν, they will choose the maximum out of the expected policy utility of party a, $\mathbb{E}_\pi[u(x_a, \theta)]$, and the expected policy and valence utility of party b, $\mathbb{E}_\pi[u(x_b, \theta)] + \nu$. Further, they incur a cost proportional to the Kullback-Leibler divergence $D_{\mathrm{KL}}(\pi||\mu)$. In Appendix 1.11, we establish that an optimal distribution τ over posteriors exists, despite the infinite and non-compact state space.

The distribution τ over posteriors induces a distribution ρ over posterior means, which are well-defined by existence of the prior mean. Because voters are ex-ante homogeneous, we assume all voters acquire the same information τ.[11] Then, given voters' ideal points and signal realizations are uncorrelated, the population distribution of revealed ideal points equals ρ (Uhlig, 1996).[12]

Parties Two parties, a and b, choose platforms, x_a and x_b, respectively, to maximize a weighted sum of their expected vote share and their ideological utility.[13] Their utilities, U_a and U_b, as a function of platforms, x_a and x_b, and the population distribution of revealed

[11] We expect this assumption to be without loss. Even if multiple optimal τ's existed and different voters acquired different ones, the resulting population distribution of revealed ideal points should be equivalent to one where each voters chooses the population-mixture of τ's. Such a choice of information τ is also optimal because the posterior-separable information cost implies indifference to mixing between optima.

[12] The distribution ρ is necessarily a mean-preserving contraction of the prior, which has finite second moments, so ρ has finite second moments. Thus, the law of large numbers by Uhlig (1996) applies if we interpret the population distribution ρ as a Pettis integral.

[13] That only two relevant parties exist is typically understood as a consequence of plurality voting systems (Duverger, 1954). For multiple parties, see also Corollary 1 and the discussion after Theorem 2.

ideal points $\rho \in \mathbb{R}^n$, are

$$U_a(x_a, x_b, \rho) = m \int_{\mathbb{R}^n} F_\nu\big(u(x_a, \theta) - u(x_b, \theta)\big) d\rho(\theta) + u(x_a, x_a^*) \tag{1.5}$$

$$U_b(x_a, x_b, \rho) = m \left(1 - \int_{\mathbb{R}^n} F_\nu\big(u(x_a, \theta) - u(x_b, \theta)\big) d\rho(\theta)\right) + u(x_b, x_b^*) \tag{1.6}$$

where $m > 0$ is the weight on vote share and x_j^* is the known ideal point of party $j \in \{a, b\}$. The probability of voting for a given revealed ideal point θ is the probability that the valence shock ν does not exceed $u(x_a, \theta) - u(x_b, \theta)$, that is, $F_\nu\big(u(x_a, \theta) - u(x_b, \theta)\big)$. The expected vote share is simply this vote probability integrated over all voters. We assume the parties have different ideal points, $x_a^* \neq x_b^*$, which guarantees platform divergence in equilibrium. Otherwise, voters would have no incentive to learn, resulting in a trivial equilibrium.

Equilibrium We study pure-strategy perfect Bayesian equilibria. In the last period, voters vote for their preferred platform given their revealed ideal point and the realized valence shock. Before that, parties simultaneously choose platforms x_a and x_b given the distribution of revealed ideology ρ induced by voter learning. Voters, in turn, learn optimally anticipating platforms x_a and x_b. Equilibria can be characterized by a triple (ρ, x_a, x_b), where ρ is the distribution over posterior means induced by a solution τ to (P) given (x_a, x_b), x_a maximizes (1.5) given (ρ, x_b), and x_b maximizes (1.6) given (ρ, x_a). Intuitively, revealed voter ideology ρ results from optimal voter learning given the anticipated platforms (x_a, x_b), which optimally respond to each other and to revealed voter ideology.

1.2.1 Discussion

Learning about Ideal Points Voter learning about ideal points can be interpreted as (i) introspecting on how to *value* the consequences of policies, (ii) learning about the private *consequences* of policies (recall ideal points are private), or (iii) a combination of both. As an example, consider a voter's position on income taxation. To determine her optimal tax policy, the voter may want to introspect on her values for equity versus efficiency, as well

as learn about what tax bracket she is in and what other economic consequences the policy has. We remain agnostic as to which interpretation should be adopted.

We assume voters can acquire costly information about their ideal points, but they observe party platforms and a valence signal for free. This approach allows us to make clear what mechanisms result from endogenous voter learning about ideal points, as opposed to learning about platforms or valence (for the latter, see Matějka and Tabellini, 2021; Hu et al., 2023). However, as we briefly illustrate in Appendix 1.11, our results on voter ideology also hold when voters do not know platforms but observe a common signal about platforms, based on which they choose how to learn about ideal points. What is more importantly ruled out by our assumption is that voters learn *jointly* about platforms and ideal points, which is an interesting avenue for future research.

Flexible Information Acquisition The rational-inattention approach preserves tractability while allowing complete flexibility in what kind of information voters can acquire. The flexibility assumption ensures the optimal signal structure is determined endogenously and not through exogenous restrictions. In particular, we are not imposing that signals about different policy issues need to be independent. That is, voters can, for example, acquire a signal that informs them about whether they are left or right, when aggregating their positions on multiple policy issues. We show such signals are, in fact, optimal.

The substantive meaning of this assumption depends on which of the above-mentioned interpretations of voter learning we adopt. When we interpret voter learning as learning about values, one can think of the voter imagining two policies that differ on multiple issues and introspecting on their relative desirability, similar to drift-diffusion models, widely used in psychology and neuroscience. By not imposing any restrictions on information, our approach stays true to the original motivation of rational inattention as modeling the brain as an efficient information processor subject to only information capacity constraints (Sims, 2003). On the other hand, when we interpret voter learning as learning about private con-

sequences of policies, another way for such learning to be aggregated across dimensions is through information intermediaries, as in Hu et al. (2023). Voters may learn about private policy consequences from sufficiently personalized media outlets, such as news feeds or newspapers catering to specific demographics. Such media outlets may aggregate information about different policy issues to a one-dimensional signal, as other models of media assume (Duggan and Martinelli, 2011; Yuksel, 2022; Perego and Yuksel, 2022).

Mutual Information Cost Although we assume the standard mutual-information cost, our results hold more generally. We use only posterior separability, Blackwell monotonicity, (reflection)-invariance, and continuity properties of the information cost. A posterior separable cost (Caplin et al., 2022) is linear under mixing between distributions over posteriors, which we use in our proof of Theorem 1. Posterior separability has foundations from information theory (Sims, 2003), sequential sampling (Morris and Strack, 2019; Bloedel and Zhong, 2020; Hébert and Woodford, 2023), and constant marginal cost of experimentation (Pomatto et al., 2023). Blackwell monotonicity means less information, in the sense of a garbling, is less costly. This property implies agents will not acquire information that does not affect their behavior, because ignoring such information would be cheaper. Reflection-invariance of the Kullback-Leibler divergence is used for Theorem 1. In Appendix 1.11, we discuss this property further and show it is satisfied by certain versions of distance-based information costs that the literature has recently proposed. Our results assuming a normal distribution (Proposition 1 and Theorem 3), use the stronger property of invariance stemming from information geometry (Amari, 2016; Caplin et al., 2022). Finally, we use lower semicontinuity of the Kullback-Leibler divergence to establish existence and continuity results.

Party Objective Our party objective makes two notable assumptions. First, as is common to the probabilistic voting literature, we assume parties care about their expected vote share instead of the probability of winning (for examples, see the references in Duggan, 2017). The expected vote share is in general a less complex object and, under some conditions,

equivalent to the probability of winning.[14] Second, we model ideological motivation through an additively separable objective. The more common approach, following Wittman (1973), assumes parties care about the implemented policy. Our objective captures in a simpler way a party (or party candidate) that cares both about votes and about not deviating too far from the parties' ideology. The advantage of our party objective is that it provides greater tractability—see our discussion under related literature—while capturing the main trade-off between vote share and ideology.

1.3 Voter Learning

We characterize optimal voter learning given equilibrium party platforms, x_a and x_b, assuming $x_a \neq x_b$ (otherwise, voters will learn nothing). In section 1.4 on electoral competition, we show parties indeed choose distinct platforms in equilibrium if the party ideal points are distinct.

This section can be seen independent of the political-economy application, and results apply analogously to an industrial organization setting with horizontally differentiated goods in a product attribute space. In that setting, valence shocks can be seen as uncertainty about prices, for example.

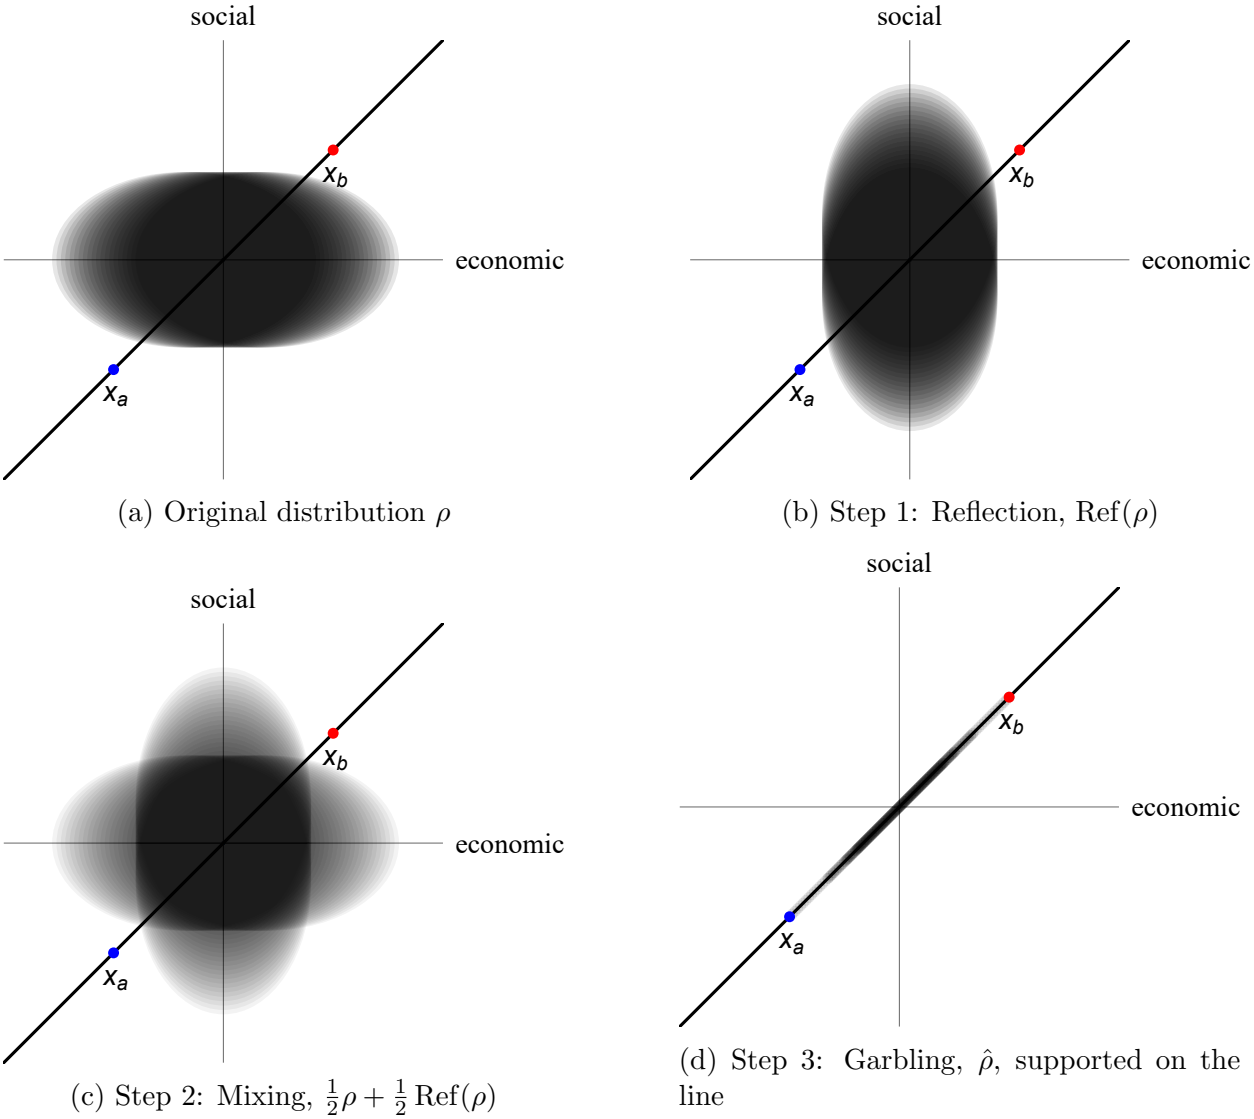
1.3.1 Issue Alignment

Our first result shows the revealed ideal points of voters (their posterior means) are on a line. Thus, even though the true distribution of ideal points is multidimensional, the revealed ideology in the population is one-dimensional. By implication, the revealed ideal points are perfectly aligned across policy issues, which holds even if the true ideal points

[14] Patty (2002) and Patty (2005) provide conditions for equivalence between maximizing probability of winning and expected vote share under office-motivated candidates. Yuksel (2022) gives a condition under which probability of winning equals the expected vote share under ideologically motivated candidates. More generally, one could assume parties care non-linearly about their expected vote share. We expect Theorem 3 to be robust to this extension.

are independent across dimensions. We show in section 1.3.4 the data on voter ideal points indicates that ideal points are on a line (Proposition 2). All proofs are relegated to the Appendix.

Theorem 1 (Issue Alignment)**.** *The distribution of revealed ideal points ρ has support inside the line through the prior mean with direction $\Sigma A(x_b - x_a)$.*



(a) Original distribution ρ

(b) Step 1: Reflection, $\text{Ref}(\rho)$

(c) Step 2: Mixing, $\frac{1}{2}\rho + \frac{1}{2}\,\text{Ref}(\rho)$

(d) Step 3: Garbling, $\hat{\rho}$, supported on the line

Figure 1.2: Reflection argument underlying the proof of Theorem 1. The dark clouds visualize the distribution over posterior means.

On a high level, the intuition of this result is that only one dimension of the ideal point is relevant for voting. We outline the logic of the proof of Theorem 1 more carefully for the

special case $A = \Sigma = I_n$, namely, that the matrix A associated with the policy utility and the prior covariance matrix Σ are equal to the identity matrix. In this case, the line of voter ideal points is parallel to the platform difference $x_b - x_a$, as in Figure 1.1. The first part of the proof shows that under quadratic utility, the instrumental value of information depends only on the projection of the posterior mean on the platform difference $x_b - x_a$.[15] For the second part, suppose by way of contradiction that voters acquired some information such that the induced distribution ρ over posterior means was not supported on the diagonal line in Figure 1.2, (a). The proof constructs through a reflection argument in three steps a distribution over posteriors that has the same instrumental value but that is cheaper. For the first step, the voter is indifferent between the original information and acquiring the "reflected" distribution over posterior means in Figure 1.2, (b). This distribution is also Bayes-consistent due to the spherical prior. It induces the same projection of the posterior mean on the platform difference (or equivalently, on the diagonal line) and hence has the same instrumental value.[16] And the Kullback-Leibler divergence is invariant under coordinate transformations, so reflections preserve the cost of information. For the second step, because the voter's information cost is posterior separable, she is indifferent to mixing between equivaluable distributions over posteriors and hence to acquiring the mixed distribution in Figure 1.2, (c), instead. For the third and last step, the voter prefers to acquire the distribution in Figure 1.2, (d), which presents a mean-preserving contraction of the mixed distribution and hence a garbling of the information. Thus, this distribution is cheaper to acquire while having the same instrumental value, because it has the same projection of posterior means on the platform difference. By symmetry of the mixed distribution constructed by the second step,

[15]This statement holds also if voters anticipate valence shocks, because the utility difference between platforms, which depends on said projection only, is still a sufficient statistic for voting. It also holds, when instead of valence shocks, voting-cost shocks are present, provided they induce what is called abstention due to indifference (e.g. Ledyard, 1984) and not abstention due to alienation (Smithies, 1941). In the former case, the utility difference between parties is again a sufficient statistic for behavior, because agents vote if the utility difference exceeds the voting cost. Under abstention due to alienation, voters care not only about the relative but also about the absolute utility from parties, so they would be motivated to learn how far they are from the closer party platform.

[16]For general Σ and A, we construct a reflection that preserves both the elliptical prior and the payoff-relevant projection on the platform difference $x_b - x_a$ with respect to A.

the mean-preserving contraction in the third step results in a distribution supported on the line through the prior mean with direction $x_b - x_a$.

Theorem 1 is related to but distinct from two other results in the rational-inattention literature that can explain a reduction of dimensionality: learning about the partition of payoff-equivalent states only and the so-called water-filling algorithm.

Under rational inattention with an entropy-based cost, agents learn only about the partition of payoff-equivalent states (Sims, 2003; Caplin et al., 2022), which implies agents neglect payoff-irrelevant dimensions. This result does not necessarily hold when the information cost depends on the distance between states, as in recent contributions to the literature (Hébert and Woodford, 2021; Pomatto et al., 2023). A concern about a result based on this mechanism is that it may require that voters are able to differentiate well between arbitrarily close states. By contrast, our proof builds on reflection-invariance of the information cost and holds for some plausible distance-based information costs as well.[17] Furthermore, we show the induced distribution over posterior means is supported on a certain line, which makes predicted survey behavior indistinguishable from that under a one-dimensional policy space (Proposition 2), as used in much of formal political economy.

Theorem 1 is also reminiscent of the water-filling algorithm, which applies in linear-quadratic Gaussian tracking problems, that is, decision problems where agents choose a continuous action $x \in \mathbb{R}^n$ to track the state $\theta \in \mathbb{R}^n$ under a quadratic loss, $u(x, \theta) = -(x-\theta)^\top A(x-\theta)$, and a normal prior (Kőszegi and Matějka, 2020). According to the water-filling algorithm, attention is allocated to a subset of dimensions according to a particular order of priority, which is determined by how payoff-relevant these dimensions are. Further, the agent pays attention to more dimensions when the attention cost is lowered. By contrast, in our case, agents learn about at most one dimension, regardless of the information-cost

[17] We show in Corollary 6 in Appendix 1.11 that Theorem 1 holds for appropriate versions of the posterior-variance, neighborhood-based (Hébert and Woodford, 2021), and log-likelihood-ratio (Pomatto et al., 2023) costs, which are all distance-based. The important condition is that the information cost is preserved under reflections that preserve the prior. Thus, our result does not allow for comparative statics under changing the information-cost distance and the prior separately.

parameter, because this dimension is sufficient for decision-making purposes. The reason is that, in contrast to tracking problems, in our model, the agent can choose only from a discrete subset of the vector space.

Robustness and Extensions Theorem 1 is robust to several generalizations. The proof works for any distribution of valence ν, under correlated ideal points across voters, and under a heterogeneous information cost parameter κ in the population if κ is independent of ideal points (otherwise, voters could infer something about their ideal points from observing κ). Although our stark result relies on ex-ante homogeneity of voters and the existence of only two parties, appropriate extensions hold when we drop these assumptions.

First, the analysis can be extended to heterogeneous priors. One way to model heterogeneous priors is to suppose voters start with some homogeneous prior but obtain exogenous signals before learning, which creates heterogeneous interim beliefs. In particular, suppose voters have a normal prior and obtain normal signals $S_i = \theta_i + \varepsilon_i$, where the noise term ε_i has a common normal distribution and is independent of θ_i. Then, the interim beliefs after observing the exogenous signals are normal, that is, elliptical, with a covariance matrix that is common to all voters. Thus, Theorem 1 applies and all voters learn in the same direction. This implies learning does not change the marginal of the distribution of revealed ideal points on the hyperplane *orthogonal* to the direction $\Sigma A(x_b - x_a)$. By contrast, the marginal on the line increases in the mean-preserving spread order through learning. Thus, learning simply "stretches out" the distribution of revealed ideal points in the direction $\Sigma A(x_b - x_a)$ and does not increase the mean-squared error of predicting voter ideal points through their projection on said line.

Second, the result can be generalized to more than two parties. First, under plurality rule, in a Duvergerian equilibrium where voters decide between the two front-runners, our mechanism would still apply.[18] Second, under electoral rules featuring proportional representation instead of plurality rule, expressive voting may remain a good assumption. Maintaining

[18] Although we assume expressive voting, strategic voting does not undermine our result under private values as long as voters perceive a positive probability of being pivotal.

expressive voting and assuming $k > 2$ parties, we show by an analogous reflection argument that voters' ideal points lie on an at most $(k-1)$-dimensional hyperplane. If k platforms were to lie on a line, however, voter ideology would still be one-dimensional.[19] If the platforms lie in general position, the model predicts that the dimensionality of voter ideology is increasing in the number of parties. Unfortunately, we are not aware of any systematic evidence regarding this prediction.

Corollary 1. *If voters face k party platforms, the distribution of revealed ideal points ρ has support inside a subspace of dimension at most $k-1$.*

This result aligns with evidence from Western European countries, which have more than two parties and where a two-dimensional political space is needed to capture party and voter positions (Kriesi et al., 2006; Kriesi et al., 2008; Bornschier, 2010a; Bornschier, 2010b; however, see Van Der Brug and Van Spanje, 2009).

1.3.2 Polarized Ideology

Our second main result shows that without uncertainty about valence, the optimal signal structure induces a binary distribution of revealed ideology. This result holds even if the true distribution of ideal points is continuous and unimodal. With uncertainty about valence, this result does not necessarily hold, but we show for "small" valence shocks, the distribution of revealed ideal points is "almost" binary.

Theorem 2 (Polarization)**.** *Absent valence shocks, voters' revealed ideal points are supported on at most two points. If the distribution of the valence shock converges in mean to zero, any selection of optimal distributions over revealed ideal points converges weakly to a binary distribution.*

[19] The dimensionality of voter ideology may also be strictly smaller than the number of parties minus 1 if voters neglect some party (e.g., because it is too far away from their prior mean) in their learning problem. In the terminology of Caplin et al. (2019), this happens if a party is not in the voter's consideration set. It follows that the dimensionality of voter ideology is less than or equal to the size of their consideration set minus 1.

Without valence shocks, voters only want to learn what party they are closer to. In other words, the voter faces a binary decision problem after learning. In rational-inattention problems with k actions, an optimal signal structure that induces at most k distinct posteriors is known to exist.[20] The reason is that if the agent acquired more posteriors—and thus signals—than actions, they could garble the signal structure based on the action recommendation. This garbling would maintain the instrumental value of information while saving on the information cost, because the garbling leads to a Blackwell-dominated signal structure.

To our knowledge, the rational-inattention literature has not emphasized the implication of this result for polarization. Under rational inattention, the necessity to take an action makes agents learn about their preferences in a way that divides them into discrete groups, one for each action—or, in our case, one for each party. In fact, this mechanism holds under flexible information acquisition for any strictly Blackwell-monotone information cost, that is, for any cost that makes a strict garbling of the signal structure strictly cheaper.

When valence shocks realize after learning, they effectively enlarge the choice set and break the mechanism for binary learning. The choice set is larger because the voter can now decide for each realization of valence ν who to vote for. Or, equivalently, voters now care about learning *how much* they prefer one party to the other. Such learning informs them for what size of the valence shock they should start voting for party b. This results in a continuous rational-inattention problem, which generally do not have closed-form solutions (Jung et al., 2019).

However, Theorem 2 shows a continuity result for valence shocks close to degenerate. If the valence shock converges in mean to zero, the distribution over revealed ideal points converges to a binary distribution. This result implies that for any two open neighborhoods of the two points of the binary distribution, as valence converges to zero, the mass of these two neighborhoods converges to 1. That is, for small-enough valence, almost all revealed

[20]For the Shannon-entropy cost, the result has been observed by Sims (2003). In Bayesian persuasion, it has been observed by Kamenica and Gentzkow (2011).

ideal points will be very close to one of the two points, so we can talk essentially of a bimodal distribution.

To show the continuity result, we prove a more general continuity result for information design problems in Appendix 1.11 (Proposition 4). Proposition 4 applies to any information design problem with state space $\mathbb{R}^n$ and an upper semicontinuous value function that is bounded from above. It establishes that the solution is upper hemicontinuous in the topology of weak convergence under uniform convergence of the value function, which may be useful beyond our application. The proof of Proposition 4 is complicated by the fact that, unlike existing result (Caplin et al., 2022; Dworczak and Kolotilin, 2023), we do not restrict ourselves to a finite or compact state space. Moreover, we cannot assume a continuous value function, because with an infinite state space, the Kullback-Leibler divergence is only lower semicontinuous rather than continuous. We show that using a generalization of Berge's maximum theorem due to Tian and Zhou (1992), we can nonetheless obtain our result. Our result shows a sense in which it is not true that "similar decision problems may lead to sharply different [behavior]" (Jung et al., 2019), which is reassuring for the theory of rational inattention.

Comparing our result with the well-studied tracking problems in the rational-inattention literature, introduced above, is again instructive. Under quadratic loss and normal prior, the agent is known to optimally acquire a normal signal, resulting in a normal distribution over posterior means that, of course, cannot be bimodal. Although the presence of continuous valence shocks makes voters' choice set effectively continuous, the utility is not quadratic in the distance of action and state, which allows for a *bimodal* distribution over posterior means. Relatedly, Jung et al. (2019) show in tracking problems, when the utility depends on the distance between the action and the state but not in a quadratic way, agents will often choose from a discrete set of actions only. In our case, the utility is not a function of the distance between action and state, so their result does not apply. Instead, the bimodality is driven by the existence of two underlying options, as explained above.

Two related results study belief polarization over a common state. Nimark and Sundaresan (2019) show that the beliefs of a population of rationally inattentive agents can become polarized over time, as agents information acquisition is path-dependent. Eguia and Hu (2022) show beliefs can become polarized if agents are boundedly rational in the sense of a finite memory and have heterogeneous preferences. Our result that revealed voter ideal points are binary without valence also holds if ideal points are common. We expect that the continuity result also generalizes, but this part requires additional work. By comparison to above papers, we show a polarized distribution can result without ex-ante heterogeneity and dynamics or bounded rationality.

1.3.3 Comparative Statics

Broadly speaking, polarization of a distribution is understood as capturing how bimodal and how spread out the distribution is (Esteban and Ray, 2012). We have shown above flexible information acquisition predicts bimodal ideology when valence shocks are small. Here, we show how a smaller cost of information or more distant party platforms can polarize voters, in the sense of leading to a more spread out distribution of revealed ideal points. We use this comparative statics result for our third main result, Theorem 3.

For the comparative statics result, we assume a normal prior and restrict voter learning to *normal* signals while maintaining that the information cost is proportional to mutual information. Formally, we define a normal signal as a random vector S such that (S, θ) is jointly normal. We make this simplification because comparative statics under flexible information acquisition are notoriously difficult due to the high dimensionality of the signal choice.[21] By contrast, under the restriction to normal signals, and because voters learn only

[21] Yoder (2022), who provides a comparative statics result under a small state space, notes that the value and cost of information need not be quasisupermodular in τ, so one cannot apply the comparative statics by Milgrom and Shannon (1994). See also the discussion in Curello and Sinander (2024) on costly information acquisition, which shows that even under posterior-mean separable information costs, increasing comparative statics hold only under very special conditions.

in a one-dimensional way, their candidate signals are completely Blackwell-ordered, which we exploit for the proof.

For this comparative statics result, we also assume the party platforms are equally distant from the voter's prior mean under the distance relevant to voter preferences, $x_a^\top A x_a = x_b^\top A x_b$, an assumption we revisit in section 1.4.1. We formalize this by supposing that party platforms (x_a, x_b) are a scaled version of (x, y), $(x_a, x_b) = \alpha(x, y)$, with $x^\top A x = y^\top A y$. We vary the scalar α, which we call the degree of platform polarization.

Proposition 1 (Comparative Statics)**.** *Restrict the prior μ and feasible voter signals to be normal and let $(x_a, x_b) = (\alpha x, \alpha y)$ with $\alpha \in \mathbb{R}_{\geq 0}$. The variance of the optimal distribution of revealed ideal points strictly increases in the strong set order when the information cost parameter κ decreases and when the degree of platform polarization α increases.*

Because the optimal signal structure may not be unique, the comparative statics result is expressed in terms of the standard strong set order. The intuition is as follows.

First, smaller κ or larger α encourage voters to acquire more information. As information becomes cheaper, voters learn more by supermodularity of their objective in the parameter κ and the cost of information $c(\tau)$, using the fact that the candidate signal structures are completely Blackwell ordered. If party platforms were very close to each other, it would not matter much for voters who to vote for, so they would learn little about their ideal points. As party platforms are more polarized, voters face larger stakes in the election and acquire more informative signals.

Second, more information leads to a distribution of voter ideal points with higher variance. While one might expect that more information leads to more agreement, here it leads to more disagreement simply because voters learn about their idiosyncratic ideal points. A more informative signal leads to a mean-preserving spread of the distribution of posterior means. Because voters learn about their *independent* ideal points, this translates to a mean-preserving spread of revealed ideology. We show, after the proof of Proposition 1 in Appendix 1.11, that this conclusion is robust to some correlation between ideal points

through a common component. The result is robust as long as the variance of the common component is smaller than the variance of the idiosyncratic component.

The prediction that more information leads to greater polarization is consistent with evidence. Palfrey and Poole (1987) develop an index of voter information and find more informed voters tend to be more extreme. Abramowitz and Saunders (2008) and Abramowitz (2010) find that more educated and engaged voters are more ideologically extreme. Lauderdale (2013) provides causal evidence that increasing information leads to ideological polarization. We discuss the evidence on the comparative statics regarding platform polarization in section 1.3.4.

A sizeable literature studies how information can lead beliefs about a *common* state to diverge. The beliefs of agents with heterogeneous priors can diverge when observing a common signal, due to ambiguity aversion (Baliga et al., 2013) or uncertainty about the signal structure (Acemoglu et al., 2016). Novák et al. (2024) also studies rationally inattentive agents, but with a common prior and heterogeneous preferences for the status quo in a binary decision problem. They show beliefs may diverge in expectation, conditional on the true state of the world, as agents acquire different signal structures. In contrast to these papers, our agents learn about *idiosyncratic states*, namely, their independent ideal points. However, as mentioned above, our monotone comparative statics would also hold in the presence of a common component of ideal points, provided the variance of the common component is smaller than the variance of the idiosyncratic component. Our focus on idiosyncratic ideal points is motivated by our application. While above papers aim at explaining persistent disagreement about facts, we focus on political positions, which are naturally heterogeneous due to conflicting interests. We therefore take seriously that voters need to learn about idiosyncratic factors affecting their political positions. This provides a simple and natural explanation of how information leads to increasing spread of voter ideal points.

1.3.4 Evidence

We relate our results to the existing evidence on voter ideology.

Issue Alignment

Recent evidence shows that the ideology of US voters is approximately one-dimensional (Jessee, 2009; Jessee, 2012; Tausanovitch and Warshaw, 2012; Shor and Rogowski, 2018; Fowler et al., 2022; Hare et al., 2023). These studies use voter surveys, such as the American National Election Studies, to estimate voter ideal points, similar to the ideal point estimation of legislators from roll-call data (Poole and Rosenthal, 1985). Specifically, these studies estimate a one-dimensional spatial model with quadratic utility to predict the binary responses $y_{ij} \in \{0, 1\}$ of each individual i to each question j (e.g. should the minimum wage be raised).[22] According to the model, the likelihood is $\Pr(y_{ij} = 1) = \Phi(u(x_{j1}, \theta_i) - u(x_{j2}, \theta_i))$, where Φ is the logistic or normal cumulative distribution function and utility is quadratic, $u(x, \theta) = -(x - \theta)^2$. The to-be-estimated parameters are the ideal points $\theta_i \in \mathbb{R}$ of each individual i and the positions $x_{j1}, x_{j2} \in \mathbb{R}$ of the policies corresponding to two responses of each question j (e.g. a minimum wage raise and the status quo). This is a standard logit or probit discrete choice model, where a voter responds more likely with the policy closer to their ideal point. These studies find that such a one-dimensional model explains voter responses well (typically about 80% of binary responses) and that adding more dimensions only marginally increases the explanatory power of the model. They conclude that ideology is well described by a one-dimensional ideological spectrum.[23]

[22] To be even more precise, these models estimate one-dimensional item-response theory models, which are known to be equivalent to one-dimensional spatial models with quadratic utility (e.g. Ladha, 1991). The only exception is Hare et al. (2023), who use a different methodology but also conclude that voter ideology is approximately one-dimensional.

[23] This finding stands in contrast to the older literature, starting with Converse (1964), which studies correlation between voter responses to different policy questions instead of estimating ideal points. These papers typically find low correlation and conclude that ideology not well represented by a one-dimensional spectrum, or that there is little *constraint* on voter ideology in the terminology of Converse. Later research has found that this conclusion is partly driven by response mistakes such as arising from inattentiveness of respondents, which reduce correlation (Ansolabehere et al., 2008). Further, the literature seems to have overlooked another reason for why the correlation between responses is a poor measure of to what extent

Upon closer examination, it is not clear whether the prediction of Theorem 1 aligns with the evidence that survey responses are well explained by a one-dimensional spatial model. First, would it not be sufficient for voter ideal points to be on some one-dimensional *curve* for a one-dimensional spatial model to explain voter's survey responses? In that case, Theorem 1 would be proving too much. Second, does equivalence to a one-dimensional spatial model require not just that the ideal points but also the *policies* are in a one-dimensional space? Theorem 1 predicts one-dimensional ideal points within a multidimensional policy space, while in one-dimensional spatial models *both* the ideal points as well as the policies live in a one-dimensional space. If the answer to the second question is affirmative, then Theorem 1 would be proving too little to explain the evidence.

In the following, we show neither is the case and Theorem 1 proves the property of voter ideology identified by the evidence, namely the property that ensures that voters' survey responses can be explained by a one-dimensional spatial model.

First, we need additional definitions. A *multidimensional spatial model* with quadratic utility is defined identically to the one-dimensional spatial model described above, except for the parameters $\{\theta_i, x_{j1}, x_{j2}\}$ being elements of $\mathbb{R}^n$ and $u(x, \theta) = -(x - \theta)^\top (x - \theta)$ being the multidimensional analogue of quadratic utility.[24] It turns out that the property of a multidimensional spatial model identified by the evidence is that respondents' *ideal points are on a line when projected onto the space spanned by the survey questions.* Formally, this property states that there exist $\lambda_i \in \mathbb{R}$, $\Delta\theta \in \mathbb{R}^n$ and $\theta_i^\perp \in \mathbb{R}^n$, such that

$$\forall i \colon \theta_i = \theta_1 + \lambda_i \Delta\theta + \theta_i^\perp,$$

$$\forall i, j \colon (x_{j1} - x_{j2})^\top \theta_i^\perp = 0.$$

voter ideology is one-dimensional. Even if voters respond to questions according to a one-dimensional spatial model, they may not consistently give left or right responses but respond with whichever response option is closer to their ideal point. In the one-dimensional spatial model, the response to question j depends on whether the voter's ideal point θ is below or above the question midpoint $\frac{1}{2}(x_{j1} + x_{j2})$. If different questions have different midpoints, the voter would be expected to choose the left or right response depending on the question.

[24] The result remains the same if we assume a general quadratic form $u(x, \theta) = (x - \theta)^\top A (x - \theta)$ because we can switch to an orthonormal basis of A.

That is, each ideal point θ_i is on the line through θ_1 with direction $\Delta\theta$, modulo a component $\theta_i^{\perp}$ that is orthogonal to the policy-differences $x_{j1} - x_{j2}$ for each question j. We say that a multidimensional spatial model with (n-dimensional) parameters $\{\theta_i, x_{j1}, x_{j2}\}$ is *observationally equivalent* to a one-dimensional spatial with (one-dimensional) parameters $\{\hat{\theta}_i, \hat{x}_{j1}, \hat{x}_{j2}\}$ if they predict the same likelihood $\Pr(y_{ij} = 1)$ over survey responses for all i and j.

Proposition 2. *Under quadratic utility, a multidimensional spatial model is observationally equivalent to some one-dimensional spatial model if and only if the multidimensional ideal points are on a line when projected onto the space spanned by the survey questions.*

Proposition 2 shows ideal points being on a line (when projected onto the space spanned by the survey questions) is the property of ideal points that makes survey behavior explainable by a one-dimensional spatial model. The parenthesized caveat holds because, naturally, survey responses are not affected by policy dimensions that are orthogonal to all survey questions. Because voter surveys try to cover most relevant policy dimensions, we take this caveat to be of limited importance. Since the above-mentioned papers show voters' survey responses are well-explained by a one-dimensional spatial model (and not much better by higher-dimensional models), we conclude they confirm the prediction of Theorem 1.

The high-level intuition for Proposition 2 is as follows. For the "only if"-direction, suppose voter ideal points were not on a line but, say, on a U-shaped curve. Then, the extreme voters on both sides of the U may prefer some policy to a policy at the bottom of the U, that is preferred by the centrist voters. Such non-monotonic behavior is ruled out by one-dimensional ideological spectrum. For the "if"-direction, for any survey question, one can find suitable projections of its two policies onto the voter line that do not change predicted behavior and make the model one-dimensional.

Polarized Ideology

Whether voters are ideologically polarized, that is, have a bimodal distribution, has lead to an academic debate between Abramowitz and Saunders (e.g. Abramowitz and Saunders,

2008) on the affirmative side, and Fiorina, Abrams, Pope, and Levendusky (e.g., Fiorina and Abrams, 2008) on the other (see Lelkes, 2016 for a critical overview of the debate). Unfortunately, neither side of the debate estimates voter ideal points but only uses "raw" survey evidence, so it is not clear how to interpret their findings. For example, much of the evidence against ideological polarization stems from evidence on ideological self-placements on 7-point scales (e.g. Fiorina and Abrams, 2008). These validity of ideological self-placements has been criticized for several reasons but the literature has, to our knowledge, overlooked a more fundamental problem. A 7-point scale is a categorical, *ordinal* scale. To assess the bimodality of the distribution of voter ideology, a *cardinal* scale is needed. The reason is that one can always monotonically transform the scale to make a distribution bimodal or unimodal. An ideological 7-point scale is only meaningful if one assumes the 7 categories correspond to intervals of the same size on the appropriate cardinal scale of ideology, for which the authors provide no evidence. This underlines the importance of estimating ideal points from survey responses to multiple questions, as this results on ideal points on a meaningful cardinal scale. On the other hand, the evidence *for* polarization relies on measures of individual-level correlation between left and right responses to different policy questions (Abramowitz and Saunders, 2008). However, without estimation of ideal points it is not clear whether their findings relate to issue alignment or polarization of ideology.

Several newer papers do, however, estimate voter ideal points from survey responses. Bafumi and Herron (2010) find a bimodal distribution of voter ideal points, while most papers find a unimodal distribution (e.g., Jessee, 2012; Hill and Tausanovitch, 2015; Dun and Jessee, 2020). However, there are reasons to believe that current estimation algorithms underestimate ideological polarization of voters. Survey respondents who do not pay much attention to the survey are artificially placed in the middle of the distribution, because this best explains their random responses (McCarty, 2019, 204). Indeed, when Fowler et al. (2022) screen for inattentive respondents (and for respondents that are not well-represented by a one-dimensional ideal point), they find a more bimodal distributions of ideal points in

most survey years. Moreover, Abramowitz (2010) finds the distribution of actual or engaged voters is more polarized. We conclude that the matter is not settled yet.

Party Influences on Voter Ideology

Political scientists have long argued that mass opinion is heavily influenced by the elite political discourse (Zaller, 1992; Lenz, 2012), yet the underlying mechanisms remain debated (Leeper and Slothuus, 2014). Our model provides a mechanism through which both issue alignment and polarization of voters is affected by parties. Perhaps surprisingly, this mechanism is consistent with voter rationality.

Specific to issue alignment, Malka et al. (2019) write "political scientists generally agree that [issue alignment] among politically attentive citizens results from such citizens following elite political cues." This idea is also motivated by findings such as that the meaning of left and right changes over time and space, in congruence with party positions (Inglehart and Klingemann, 1976). For example, whether protectionism is associated with the left or right in the US has evolved over time (McCarty, 2011). Theorem 1 predicts the orientation of the ideological spectrum, that is, $\Sigma A(x_b - x_a)$, is determined by party platforms, x_a and x_b. As argued in the introduction, this orientation determines what issues go together.

Implicit in the understanding by Malka et al. (2019) is that the issue alignment among voters is consistent with relative party platforms. That is, for example, if one party is more left on economics and more liberal on social issues than the other party, then voters who are more left are also more liberal. More precisely, we say *issue alignment is consistent with party platforms* if the sign of the k-th component of relative party platforms, $x_b - x_a$, equals the sign of the k-th component of the orientation of the ideological spectrum, $\Sigma A(x_b - x_a)$, for all $k = 1, ..., n$. Then, for any two dimensions the alignment of relative party platforms conforms to the issue alignment of voters. While this is not a necessary prediction of our model, it holds in important special cases. If $x_b - x_a$ is an eigenvector of ΣA, then the ideological spectrum is exactly parallel to the platform difference $x_b - x_a$, so issue alignment

is consistent with party platforms. In section 1.6, we give a microfoundation for $x_b - x_a$ being an eigenvector of ΣA if party objectives are driven by valence competition. This alignment may also occur in a richer model in which party ideology arises from the ideology of voters who are party members, resulting in party platforms that are on the line of voter ideal points. If the covariance matrix Σ of true ideal points and the matrix A associated with the policy utility are both diagonal, the issue alignment is also consistent with party platforms.[25] Diagonality holds if voter positions on different policy issues are uncorrelated and there are no preference interdependencies between issues. Broadly speaking, as long as such correlations and interdependencies are not strongly enough misaligned with relative party positions, issue alignment should be expected to be consistent with party platforms.

Regarding polarization, Proposition 1 shows how platform polarization can lead to polarization of voters. This is not because voters blindly follow party positions but instead as a consequence of rational learning. Again, the political environment can affect revealed ideology even when true ideology remains unchanged. This is consistent with the finding of Bischof and Wagner (2019) that voters ideal points diverge immediately after new radical parties enter parliament.

1.4 Electoral Competition

We are interested in welfare properties of the equilibrium platforms, from the viewpoint of voters' policy utility. For voters, it is crucial how much parties polarize their platforms, moving away from the policy that maximizes voter's aggregate policy utility. Before we turn to this question in section 1.4.1, we highlight some important forces at place in electoral competition, holding voter preferences fixed.

Recall that parties choose their platforms, x_a and x_b, in a Nash equilibrium of the electoral-competition game given the distribution ρ of voters' revealed ideal points. It is

[25] If $\Sigma = \text{diag}(\Sigma_{11}, ..., \Sigma_{nn})$ and $A = \text{diag}(A_{11}, ..., A_{nn})$, with $\Sigma_{11}, ..., \Sigma_{nn}, A_{11}, ..., A_{nn} > 0$ by positive definiteness, the k-th component of $\Sigma A(x_b - x_a)$ is simply $\Sigma_{kk}A_{kk}(x_{b,k} - x_{a,k})$, which has the same sign as $x_{b,k} - x_{a,k}$ for $k = 1, ..., n$.

typically hard to obtain characterizations of equilibrium platforms when parties are motivated both by vote share and ideology. However, our party objective, which is linear in vote share and ideological utility, allows such a characterization. The following lemma shows (1) party platforms are a weighted mean of voter and party ideal points, and (2) the weight on a voter is decreasing in the "extremeness" of the voter. We use these properties subsequently and relate them to platform polarization.

Lemma 1. *In any equilibrium, party platforms are a weighted average of voter and party ideal points,*

$$x_a = \frac{m \int w(\theta)\theta d\rho(\theta) + x_a^*}{m \int w(\theta) d\rho(\theta) + 1}, \tag{1.7}$$

$$x_b = \frac{m \int w(\theta)\theta d\rho(\theta) + x_b^*}{m \int w(\theta) d\rho(\theta) + 1}. \tag{1.8}$$

where

$$w(\theta) = f_\nu\big(u(x_a, \theta) - u(x_b, \theta)\big).$$

The result generalizes the mean-voter theorem by Hinich (1977), which assumes purely office-motivated candidates. The mean-voter theorem states that under quadratic voter utility and probabilistic voting, party platforms converge at the mean of voter ideal points. Because parties are office- and ideologically-motivated in our model, their platforms are affected by voter ideal points as well as the party's own ideal point in an intuitive way.

In particular, by symmetry of f_ν, the weight $w(\theta)$ on a voter with ideal point θ depends only the size of the utility difference between party platforms, $|u(x_a, \theta) - u(x_b, \theta)|$. Voters that have a larger utility difference can be seen as more "extreme" or "ideologically entrenched" relative to the party platforms. By strict quasi-concavity of f_ν, a voter with a larger utility difference has a smaller weight $w(\theta)$. Intuitively, more extreme voters are less sensitive to platform changes (the probability that they change their vote due to a small platform change is small), so they have less influence on equilibrium platforms. While this

observation is not new (Persson and Tabellini, 2002, 57), most models of probabilistic voting rule this effect out by focusing on a uniform distribution of valence shocks to improve tractability. We use this observation later to show voter polarization amplifies platform polarization: If voters are more extreme on average, parties moderate less and choose policies closer to their own ideal points. This mechanism is consistent with the finding by McCarty et al. (2019) that more ideologically heterogeneous districts have more extreme legislators.[26]

Lemma 1 only speaks to *necessary* conditions of equilibrium platforms, as derived from first-order conditions. Hence, additional work is necessary to show that equilibrium candidates that satisfy the first-order conditions constitute actual equilibria. For example, they constitute equilibria if the party objectives are quasi-concave, in which case the first-order conditions are sufficient for optimality. Along these lines, in Appendix 1.11, we give a condition that ensures that our equilibrium candidates in the context of Theorem 3 are equilibria. We also show this condition is satisfied when the weight m on vote share is small enough or when the valence shock ν is large enough, echoing observations by Lindbeck and Weibull (1987) and Enelow and Hinich (1989).

1.4.1 Platform Polarization

To illustrate the mechanisms underlying platform polarization, we simplify to a symmetric setup. A symmetric setup is obtained when the party ideal points x_a^* and x_b^* are equally far from the origin according to the distance relevant to voter preferences, $x_a^{*T} A x_a^* = x_b^{*T} A x_b^*$. We show below under this assumption all equilibria are symmetric in the following sense. We say (ρ, x_a, x_b) is a *symmetric* equilibrium if

$$(x_a, x_b) = (\alpha x_a^*, \alpha x_b^*)$$

[26]McCarty et al. (2019) interpret this finding through the Calvert-Wittman model, in which greater uncertainty about the location of the median voter leads to greater platform polarization. They theoretically connection voter polarization to uncertainty about the median voter through the informativeness of a public poll. Our model provides a more direct mechanism.

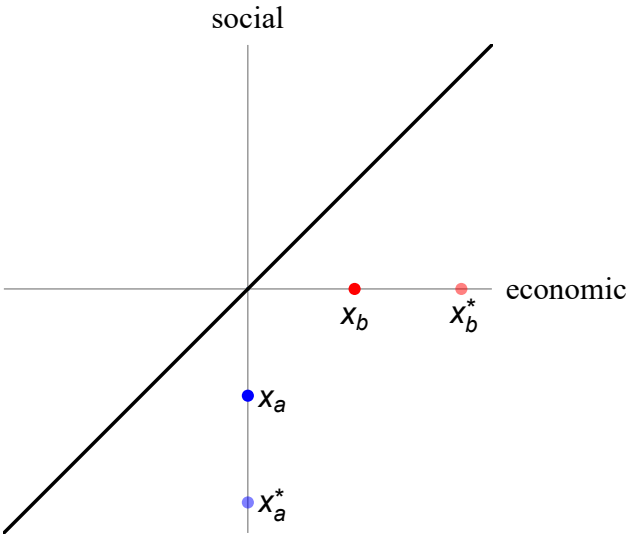


Figure 1.3: Symmetric equilibrium with party platforms (x_a, x_b) and party ideal points (x_a^*, x_b^*)

with $\alpha \in (0, 1)$, which we call the *degree of platform polarization*. Theorem 3 below identifies comparative statics for the degree of platform polarization α as the information cost parameter κ changes. Figure 1.3 visualizes an example of a symmetric equilibrium. Note that a higher degree of platform polarization not only increases the distance between party platforms but also makes party platforms move further away from the ideological spectrum of voter ideal points.

Bringing together endogenous voter ideology and endogenous party platforms, we show the following result. We restrict again to normal distributions to make use of the comparative statics result Proposition 1. Because voter and platform polarization are mutually reinforcing, there may be multiple equilibria, which we order by their degrees of platform polarization α. Therefore, as usual, our comparative statics are expressed in terms of the smallest and largest equilibrium.

Theorem 3. *Restrict the prior μ and feasible voter signals to be normal. There exists an equilibrium and every equilibrium is symmetric. Cheaper information increases polarization: The smallest and largest equilibrium degree of platform polarization α weakly increase as κ decreases.*

Theorem 3 combines our earlier results on voter ideology and on party platforms. To prove the theorem, we show voter polarization and platform polarization are mutually reinforcing: if voters are more extreme, their voting is less sensitive to party platforms, allowing parties to polarize more (Lemma 1). If platforms are more polarized, then voters face larger stakes in the election, inducing them to learn better and become more extreme (Proposition 1). One can think of cheaper information to start this self-reinforcing process by allowing voters to learn at a lower cost (Proposition 1). On a formal level, we establish existence of pure-strategy equilibria and the comparative statics result, through monotonicity arguments similar to those in supermodular games (despite our game not being supermodular).

The theorem implies better availability of information makes the equilibrium policy worse for voters. Platform polarization hurts voters in our model because the utilitarian optimum for voters is the policy coinciding with the mean ideal point, which is at the origin. A higher degree of platform polarization, α, implies that any implemented policy (x_a or x_b) moves further away from the origin, decreasing voter's aggregate policy utility. While cheaper information allows voters to learn more accurately about their ideal points, this makes voters less responsive to party platforms, leading to greater platform polarization. Voters do not internalize this information externality of their learning strategy on party platforms because each voter is infinitesimal.

Theorem 3 underscores the different implications of learning about preferences versus learning about equilibrium actions of other agents. Matějka and Tabellini (2021) show more informed voters are *more* responsive to party platforms, when voters learn about party platforms knowing their ideal points. This would suggest decreasing platform polarization in equilibrium as information becomes more accessible. By contrast, in our model, better informed voters are more extreme and therefore *less* responsive to party platforms. Furthermore, we show in section 1.5 this difference is not due to our timing assumption. In a symmetric equilibrium, the vote share is less responsive to the platform choice under cheaper

information, also if parties publicly commit to their platforms *before* voters learn about their preferences.

The theorem demonstrates one mechanism that may have contributed to increasing party polarization in the US, as observed in the past decades (McCarty et al., 2016). Information can become cheaper due to advances in information technology, such as the internet. Theorem 3 shows better availability of information can lead to more platform polarization. The underlying mechanism operates through increasing polarization of voters. While the empirical evidence on increasing polarization of US voters is somewhat mixed, it suggests that voter polarization may have increased more recently. Hill and Tausanovitch (2015) find that the variance of US voter ideal points is generally stable from 1956 to 2012, but their point estimates for variance increase after the year 2000.[27] The Pew Reseach Center (2014) find similar spread of voter position in 1994 and 2004 but a significant increase in 2014. Thus, one may take Theorem 3 to suggest a contributor to platform polarization in the post-2000 era.

1.4.2 Aggregate Uncertainty

Until now, we have assumed that there is no aggregate uncertainty about voter preferences. Therefore, the optimal policy was always at the mean of voter ideal points, that is, the origin. In general, however, the optimal policy may depend on such aggregate uncertainty. This opens up the new question whether under endogenous voter learning, elections aggregate preferences, in the sense of making policy responsive to aggregate uncertainty. We show a novel failure of information aggregation: because voter learning is one-dimensional, policy responds to only one dimension of aggregate uncertainty.

We model aggregate uncertainty about voters' ideal points through an aggregate state ω, which enters voters ideal points as a common component. The ideal point of voter i is $\theta_i = \omega + \delta_i$, where the idiosyncratic components $\{\delta_i\}$ are identically distributed and $(\omega, \{\delta_i\})$

[27] Moreover, in subsequent unpublished work, they find more evidence of increasing polarization in the post-2012 data, but the conclusion depends on the measure and statistical model (personal communication).

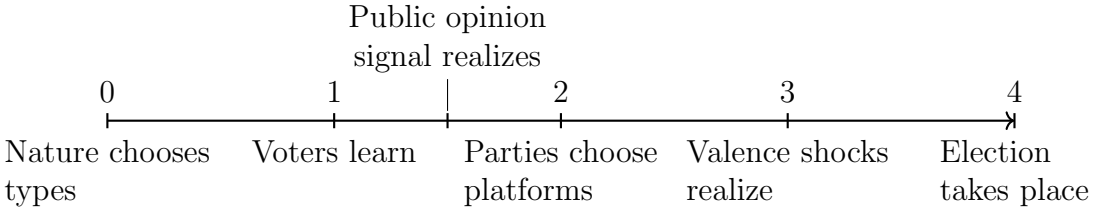
are jointly independent. For our result in this section, we do not need to impose elliptical distributions for ω and δ_i. However, to simplify the proof of equilibrium existence and because it comes with little loss of economic substance, we assume the support Ω of ω and the support D of δ_i are finite. We allow voters i to acquire any signal structure about (ω, δ_i) and maintain the assumption that all voters acquire the same signal structure. Formally, each voter $i \in [0,1]$ chooses a signal structure, that is, a stochastic kernel $\sigma_i : \Omega \times D \to \Delta(\mathcal{S})$ that maps each state (ω, δ_i) into a distribution over signals of a sufficiently rich signal space $\mathcal{S}$.[28] We understand the common component ω as a way to model aggregate uncertainty about voter preferences and assume that it does not affect party ideal points.

For parties to be able to respond to realized voter preferences, they must obtain information about voter preferences. If parties had *private* information about the realized distribution of revealed voter ideology, then platforms could convey information about the common component ω to voters. This would introduce a signalling motive into electoral competition (see Martinelli, 2001). However, since this signaling motive is not the focus of our analysis, we assume instead that both parties and voters observe a *public* signal $s \in S$ about voter preferences. This public signal could represent a poll or, more generally, any channel through which information about public opinion is disseminated.

Formally, the public signal is a stochastic kernel $\sigma_p : \Delta(\mathcal{S}) \to \Delta(S)$ that maps the realized distribution over voters' private signals into a distribution over public signals. As the public signal depends only on the distribution of voters' private signals, the realization of the public signal cannot be affected by a single infinitesimal voter. To show equilibrium existence, we assume the public signal space S is finite and the probability of any public signal $s \in S$, $\sigma_p(s|\pi)$, is continuous under weak convergence of the distribution $\pi \in \Delta(\mathcal{S})$ over private signals. This weak assumption is satisfied if the public signal contains some amount of noise.

[28] While up till now, we have modelled signal structures as distributions over posteriors, to prove equilibrium existence in the context of this section, it is more useful to model signal structures as Blackwell experiments.

Our extended game thus contains an additional stage between voter learning and platform choice where the public opinion signal realizes. As a result, parties can condition their platforms (x_a, x_b) on the realization of the public signal $s \in S$ and voters can condition their voting behavior both on the chosen platforms (x_a, x_b) and on the realization of the public signal s.[29]



The main result of this section shows that while we introduce a channel through which party platforms can respond to the aggregate state ω, they respond to only one dimension of ω, namely the projection of ω on the ideological difference between parties.

Theorem 4. *There is an equilibrium in which party platforms are affected by the aggregate state ω only through its A-projection on $x_b^* - x_a^*$, $(x_b^* - x_a^*)^\top A\omega$. That is, the distribution over equilibrium policy is the same under aggregate states ω and ω' if $(x_b^* - x_a^*)^\top A\omega = (x_b^* - x_a^*)^\top A\omega'$.*

The intuition is as follows. Parties only respond to components of the aggregate state ω that voters learn about. Consequently, if voters learn only about the A-projection of ω on the ideological difference $x_b^* - x_a^*$ between parties, then party platforms are only affected by this component. While Theorem 1 suggests that voters optimally learn only about one component of ω, the public opinion signal introduces two complications for showing this. First, at the time of learning, voters do not know yet what party platforms will be, giving them potentially an incentive to learn about multiple dimensions of their ideal points. However, independent of the public opinion signal, the platform difference $x_b - x_a$ is *parallel* to the ideological difference $x_b^* - x_a^*$ of parties, by Lemma 1. Therefore, voters have no incentive to learn about components of their ideal points orthogonal to $x_b^* - x_a^*$ for the purpose of voting. Second, the

[29] See the Appendix for a formal definition of strategies.

information about ω obtained through the public opinion signal could be complementary to private learning about orthogonal components of ω. However, we show that there exists an equilibrium in which no voter learns about such orthogonal components of ω, which rules out such complementarities. It is an open question whether there are equilibria in which party platforms respond to more than one component of ω.

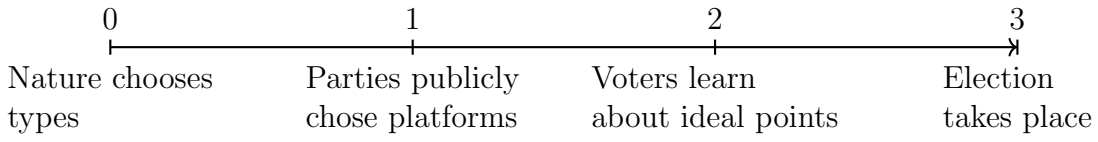
Theorem 4 presents a severe inefficiency of preference aggregation due to endogenous voter learning. Because of independent idiosyncratic components δ_i, the average ideal point equals the common component ω. Because of quadratic preferences this makes ω the policy that maximizes unweighted aggregate voter welfare. However, voters learn only about the dimension of the common state ω along which parties disagree and, as a consequence, equilibrium policies respond only to this *one dimension* of ω, even if the policy space, and thus ω, is high-dimensional.

The result is particularly relevant because the failure of preference aggregation and resulting welfare loss might not show in the data and therefore go unnoticed. Judging on the basis of *revealed* ideology, it seems that parties do respond to voter preferences. However, the two-party system prevents voter learning about their preferences in more than one dimension, so there is a large scope for *unrevealed* voter ideology that policy does not respond to.

1.5 Alternative Timing

1.5.1 Model

Game:



We simplify to a one-dimensional policy space.[30] After nature has chosen the types of voters, parties simultaneously commit to policy platforms $x_a, x_b \in \mathbb{R}$, observed by voters. Next, voters learn about their own ideal points. Finally, the election takes place.

Voters: Voter preferences are as in the baseline model, except for the absence of a valence shock. The utility U_i of voter $i \in [0,1]$ has two components, which we expand on below. Utility depends on the implemented policy x and the voter's ideal point $\theta_i \in \mathbb{R}$ via the policy utility $u(x, \theta_i)$ and on the information cost $c(\tau_i)$ of signal structure τ_i scaled by κ,

$$U_i(x, \tau_i) = u(x, \theta_i) - \kappa \cdot c(\tau_i). \tag{1.9}$$

where

$$u(x, \theta) = -(x - \theta)^2.$$

Voters ideal points are independent and the prior is symmetric around 0. Furthermore, we assume in this section that the prior is log-concave. The information cost is the expected reduction of variance, from the prior μ to the posterior π,

$$c(\tau) = \mathbb{E}_\tau[\mathrm{Var}_\mu(\theta) - \mathrm{Var}_\pi(\theta)].$$

This information cost is used in recent contributions due to its tractability (e.g., Ravid et al., 2022; Thereze, 2022b) and belongs to the class of posterior-mean separable information-costs axiomatized in Mensch and Malik (2024). The information cost is tractable because it is equivalent to the variance of the posterior means by the law of iterated variance,

$$\mathbb{E}_\tau[\mathrm{Var}_\mu(\theta) - \mathrm{Var}_\pi(\theta)] = \mathrm{Var}_\tau(\mathbb{E}_\pi[\theta]).$$

[30] Results generalize to a multidimensional policy space (coming soon).

Together with the quadratic policy utility, this allows to express the voter's objective, up to constants, as a simple function of the distribution over posterior means ρ, induced by τ:

$$\mathbb{E}_\rho\left[\max\left\{(x_b - x_a)\left(\theta - \frac{x_a + x_b}{2}\right), (x_a - x_b)\left(\theta - \frac{x_a + x_b}{2}\right)\right\} - \kappa\theta^2\right] \tag{1.10}$$

By Strassen's theorem, a distribution $\rho \in \Delta(\mathbb{R})$ over posterior means is induced by some signal structure if and only if ρ is a mean-preserving contraction of the prior μ, $\rho \leq_{\text{MPS}} \mu$. So, voters choose $\rho \in \Delta(\mathbb{R})$ to maximize (1.10) subject to $\rho \leq_{\text{MPS}} \mu$. Due to the simple form of this objective, consisting of a piecewise linear and a quadratic function, we can obtain a closed-form solution for the agent's information acquisition problem. While the qualitative results on voter learning hold for a more general class of information costs, this cost function proves particularly tractable to solve for endogenous party positions.

Parties We assume that parties are policy-motivated (Calvert, 1985; Wittman, 1973). That is, the payoffs U_a and U_b of parties a and b, respectively, are

$$\begin{aligned} U_a(x_a, x_b) &= P_a(x_a, x_b)u(x_a, x_a^*) + (1 - P_a(x_a, x_b))u(x_b, x_a^*) \\ U_b(x_a, x_b) &= P_a(x_a, x_b)u(x_a, x_b^*) + (1 - P_a(x_a, x_b))u(x_b, x_b^*), \end{aligned}$$

where P_a is the probability that party a's policy x_a gets implemented. We make the common assumption that this implementation probability is the vote share (Wittman, 1983; Wittman, 1990; Callander and Carbajal, 2022; Yuksel, 2022). A continuous mapping from the vote share to the implementation probability can result from additional noise voters, such as random partisans (e.g., Feddersen and Pesendorfer, 1996), or non-majoritarian institutions in parliament. Computations suggest that our results generalize to S-shaped mappings from vote share to implementation probability instead of linear mappings, however, an S-shaped mapping may lead to equilibrium multiplicity.

Equilibrium We focus on subgame perfect pure-strategy equilibria, in which both parties receive positive expected vote shares.[31]

1.5.2 Voter Learning

We first characterize the optimal voter learning strategies and resulting expected vote shares given party platforms $x_a, x_b \in \mathbb{R}$.

As discussed above, the voter learning problem can be expressed as a function of the distribution ρ over posterior means, which can be generated by a Bayes-consistent distribution τ over posteriors if and only if $\rho \leq_{\text{MPS}} \mu$. The following definition will be useful for the following.

Definition 1 (Feasibility)**.** *We say a binary set $\{\theta_a, \theta_b\}$ is **feasible** if there exists a distribution ρ over $\{\theta_a, \theta_b\}$, $\rho \in \Delta(\{\theta_a, \theta_b\})$, that is a mean-preserving contraction of the prior μ.*

Note that if ρ is a mean-preserving contraction of μ, then in particular ρ has expectation 0, $\mathbb{E}_{\theta\sim\rho}[\theta] = \mathbb{E}_{\theta\sim\mu}[\theta] = 0$. Thus, if $0 \notin [\theta_a, \theta_b]$, then $\{\theta_a, \theta_b\}$ is not feasible. On the other hand, if $0 \in [\theta_a, \theta_b]$, then there is a unique distribution ρ over $\{\theta_a, \theta_b\}$ with expectation zero. Thus, feasibility boils down to whether this distribution ρ is a mean-preserving contraction of the prior μ.

Without loss, assume $x_a \leq x_b$ and define

$$\theta_a := \frac{x_a + x_b}{2} - \frac{x_b - x_a}{2\kappa}$$
$$\theta_b := \frac{x_a + x_b}{2} + \frac{x_b - x_a}{2\kappa}.$$

Proposition 3. *A solution $\tau \in \Delta(\mathbb{R})$ to the voters' information acquisition problem exists and is unique. There are three cases:*

[31] There exists a disconnected class of equilibria, where voters acquire no information and all vote for one party. These equilibria can be easily ruled out, for example by introducing an arbitrarily small office benefit.

(1) If $\{\theta_a, \theta_b\}$ is feasible, voters acquire the two posterior means θ_a and θ_b.

(2) If $\theta_a < 0 < \theta_b$ and $\{\theta_a, \theta_b\}$ is not feasible, voters acquire a threshold signal structure $s(\theta) = \mathbb{1}(\theta > t)$, where $t \in \mathbb{R} \cup \{-\infty, \infty\}$ is more extreme than the party midpoint, that is, we have $0 < \frac{x_a+x_b}{2} < t$, $0 = \frac{x_a+x_b}{2} = t$, or $t < \frac{x_a+x_b}{2} < 0$. As κ decreases, the threshold t moves weakly closer to the party midpoint $\frac{x_a+x_b}{2}$.

(3) If $0 \notin (\theta_a, \theta_b)$, voters acquire no information.

When voters acquire information, then one posterior mean θ_a is always closer to x_a and the other, θ_b, is closer to x_b. The expected vote shares of parties are simply the probabilities of posterior means θ_a and θ_b. These probabilities are uniquely pinned down by $\Pr_\rho(\theta_a)\theta_a + \Pr_\rho(\theta_b)\theta_b = 0$.

The following result summarizes the implications for parties' expected vote shares and is crucial for the analysis of endogenous party positions.

Corollary 2. *There are three cases.*

(1) If $\{\theta_a, \theta_b\}$ is feasible, the expected vote share P_a of party a is

$$P_a(x_a, x_b) = \frac{1}{2} + \frac{\kappa}{2}\frac{x_b + x_a}{x_b - x_a}.$$

(2) If $\theta_a < 0 < \theta_b$ and $\{\theta_a, \theta_b\}$ is not feasible, the expected vote share P_a of party a is $P_a = F_\mu(t)$, where t is as in Proposition 3. If party a is more extreme, that is, $|x_a| > |x_b|$, then $P_a(x_a, x_b) < \frac{1}{2} + \frac{\kappa}{2}\frac{x_b+x_a}{x_b-x_a}$.

(3) If $0 \notin (\theta_a, \theta_b)$, the party whose position is closer to 0 obtains all votes.

Interestingly, the expected vote share in the first case takes the same form as in the moderate utility model (Halff, 1976; He and Natenzon, 2024). The following qualitative property of expected vote shares shows that costly voter learning leads to a "bias" toward the moderate party.

Corollary 3. *Under $|x_a| < |x_b|$, party a obtains a higher expected vote share than the number of voters that are closer to x_a than to x_b, that is, $P_a > F_\mu(\frac{x_a+x_b}{2})$. The number of votes of party a decreases as κ decreases or as party polarization $|x_b - x_a|$ increases, holding $\frac{x_a+x_b}{2}$ constant.*

The mechanism behind this result is that more asymmetric signals are less informative and hence cheaper to acquire. Therefore, voters acquire "biased" signals to reduce their information cost. In more detail, under asymmetric party positions, voters acquire asymmetric signals: the signal favoring the extreme party is obtained less frequently and is more informative about voter preferences. When designing the distribution over signals in a given state, two factors influence the cost of information. First, the probabilities of signals shift. To economize on information costs, voters select the more informative "extreme signal" less frequently, inducing a bias toward the moderate party. Second, the content of signals shift. Under a log-concave prior, this reinforces the bias: the posterior variance of the moderate signal is less sensitive than that of the extreme signal, as the posterior mean for the moderate party is more central.

Corollary 3 implies that endogenous information creates two new forces in electoral competition: a **moderation motive** and a **differentiation motive**. The bias toward the moderate party creates an additional motive for parties to moderate. This motive is stronger, the more costly is information. Further, the extreme party has an incentive to differentiate, because party differentiation decreases the bias toward the moderate party. On the other hand, the moderate party has an incentive to move closer to the extreme party to increase the bias toward itself.

1.5.3 Equilibrium

Symmetric Setup

We first characterize equilibrium under a symmetric setup, that is, when parties are equally ideologically extreme from the viewpoint of voters, $|x_a^*| = |x_b^*|$. Define $\underline{\kappa}$ via

$$\frac{1-\underline{\kappa}}{\underline{\kappa}} x_a^* = \mathbb{E}[\theta|\theta > 0].$$

Theorem 5. *Let $x_a^* < 0 < x_b^* = -x_a^*$. The following is an equilibrium.*

		Party platforms	Revealed voter ideology
$\kappa \geq 1$:	Downsian equilibrium	$(x_a, x_b) = (0, 0)$	$\rho = \delta(0)$
$\kappa \in (\underline{\kappa}, 1)$:	Polarizing equilibrium	$(x_a, x_b) = (1-\kappa)(x_a^*, x_b^*)$	$\rho = \frac{1}{2}\delta\left(\frac{x_a}{\kappa}\right) + \frac{1}{2}\delta\left(\frac{x_b}{\kappa}\right)$

Voters vote for the party closer to their revealed ideal point. Equilibrium is unique for $\kappa \neq \frac{1}{2}$ if $|x_a^*| \leq \mathbb{E}[\theta|\theta > 0]$.[32]

There are two types of equilibria.

In the Downsian equilibrium, reminiscent of the median-voter theorem (Downs, 1957), parties fully converge on voter's prior expectation of the optimal policy. Voters acquire no information and split their votes between the two parties.

In the polarizing equilibrium, parties diverge and voters acquire information, which polarizes the population into two groups. Remarkably, the model allows for simple closed-form solution of the equilibrium platforms and voter positions.

In the Calvert-Wittman model with policy-motivated parties and uncertainty about the election winner, parties diverge (Calvert, 1985; Roemer, 1994). Therefore, it is perhaps surprising that there exists a Downsian equilibrium in which parties converge. This is driven by

[32]We conjecture that this condition, which we only use to show uniqueness for $\kappa > 1/2$, can be dropped. Further, we conjecture that the additional equilibria under the knife-edge case $\kappa = 1/2$ cannot be approximated by an arbitrarily small office benefit.

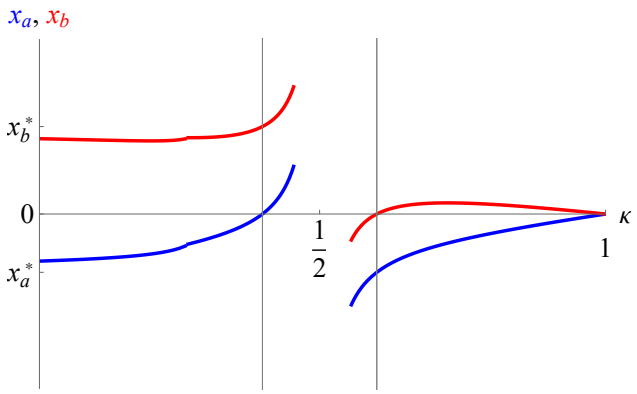


Figure 1.4: Party platforms as a function of info cost parameter κ

the novel moderation motive. The intuition for the comparative statics has two parts. First, party polarization increases as the vote share becomes less responsive to party platforms. Second, the vote share becomes less responsive as information costs decline.

First, party polarization increases as the vote share becomes less responsive to party platforms. Parties trade off votes (and probability of winning) and choose a policy closer to their own ideology. To gain votes parties move their policies closer to voters prior expectation and to gain ideological utility, they move the policy platforms closer to their own ideal points. The optimal policy equates the loss in votes with the gain in ideological utility from polarizing further. When the loss in votes becomes smaller, they polarize more.

Second, the vote share becomes less responsive as information costs decline, as is apparent from the first case of Corollary 2. What is driving this mechanism is that voter learning is biased towards the ex-ante better option, that is, the party that is closer to their prior expectation of the best policy (Corollary 3). This bias makes the vote share more responsive, thereby driving more moderation. However, as information costs decline, this motive recedes (Corollary 3).

Asymmetric Setup

In the following we assume without loss that party b is more ideologically extreme.

Theorem 6. *Suppose $x_a^* < 0 < x_b^*$ and $|x_a^*| \leq |x_b^*|$. The positions*

$$x_a(\kappa) = \frac{1-\kappa}{1-2\kappa}\Big((1-\kappa)x_a^* + \kappa x_b^*\Big),$$
$$x_b(\kappa) = \frac{1-\kappa}{1-2\kappa}\Big((1-\kappa)x_b^* + \kappa x_a^*\Big)$$

are an equilibrium if

$$\kappa \in \mathcal{K} = \left\{\kappa \in \left(0, \frac{x_b^* - x_a^*}{3x_b^* - x_a^*}\right) \cup \left(\frac{x_b^* - x_a^*}{x_b^* - 3x_a^*}, 1\right) \;\middle|\; \max\{|x_a(\kappa)|, |x_b(\kappa)|\} < \kappa\mathbb{E}[\theta|\theta > 0]\right\}.$$

This is the unique equilibrium if $|x_b^* - x_a^*| < \kappa\mathbb{E}[\theta|\theta > 0]$.[33]

If $\kappa \geq 1$, the unique equilibrium is $(x_a, x_b) = (0, 0)$.

Figure 1.4 plots the equilibrium positions (x_a, x_b) as a function of κ. For $\kappa \in \left[\frac{x_b^* - x_a^*}{3x_b^* - x_a^*}, \frac{x_b^* - x_a^*}{x_b^* - 3x_a^*}\right]$, there is no equilibrium in pure strategies where both parties obtain positive expected vote shares. We comment on this below. For small κ, we also plot numerically solved equilibrium positions.

A notable feature of the equilibrium is that in a certain range, both parties are on the same side of zero and one party is more extreme than its ideal point. This is contrasts with the model with exogenous voter positions, in which party positions are necessarily between their ideal points (Roemer, 1997).

This effect occurs because of the differentiation motive, which creates chase-and-evade incentives, similar to the literature on valence advantage (Groseclose, 2001; Aragones and Palfrey, 2002; Bernhardt et al., 2020). By the logic of the differentiation motive, in an asymmetric equilibrium, the extreme party has an incentive to moderate and the moderate party has an incentive to move closer to the extreme party. When κ is close enough to $\frac{1}{2}$, these forces become so strong that the extreme party moves beyond its own ideal policy to differentiate from the moderate party. The moderate party, on the other hand, moves close to the extreme party beyond the center of the voter distribution, zero. If κ is very close to

[33] We conjecture, again, that this condition is superfluous.

$\frac{1}{2}$, these forces become so strong that a pure-strategy equilibrium fails to exist as is known from the literature on valence advantage.

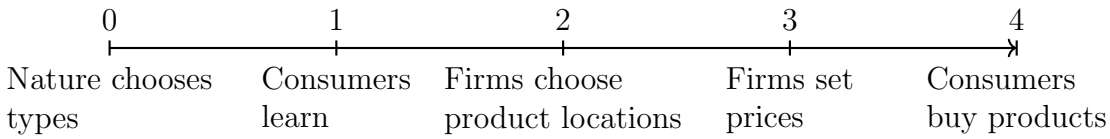
1.6 Horizontally Differentiated Goods

In this section, we adapt our model to a market context. We examine firms that produce horizontally differentiated products in some attribute space. A key difference between markets and politics is that consumption is private in markets, whereas policy is public: consumers can purchase the product that best fits their preferences, but while voters can support their preferred policy platform, ultimately only one policy is implemented for everyone. This distinction implies a potential benefit for product differentiation in markets—allowing consumers to select their optimal match—that is not present in politics. However, we show that even with this advantage, a lower cost of information still harms consumers. The reason is twofold: firms differentiate their products excessively from a social welfare standpoint, and unlike political parties, they also set prices. As consumers become more polarized due to lower information costs, firms can exploit their increased market power by raising prices, which further decreases consumer welfare. Methodologically, we show that endogenous consumer learning allows us to solve a multidimensional Hotelling model by reducing it to one dimension.

Our adaptation of the model can also be interpreted within a political-economy framework that includes valence competition. In this context, political parties not only choose policy platforms but also compete based on valence attributes such as competence of their candidates (see Ashworth and Bueno de Mesquita, 2009). Similar to how firms set prices in our industrial organization adaptation, parties may invest in valence to attract voters.

Adapting our model to a market context is largely a matter of reinterpretation. We reinterpret voters as consumers with unit demand, the policy space $\mathbb{R}^n$ as a product attribute space, and parties as firms. However, we need to make two important modifications to the

model. First, we replace the exogenous valence shocks with prices, which are chosen by firms. Second, firms maximize profits rather than a combination of vote share and ideological utility, which our parties maximized. Apart from these two adjustments—the substitution of valence shocks with endogenous pricing and the change in the firms' objectives—the models remain essentially the same, except for the previously mentioned welfare difference between private consumption and public policy. We also maintain the timing structure, adhering to the standard sequence in Hotelling models, where firms first choose product locations in attribute space and then set prices.



Consumers The utility U_i of consumer $i \in [0,1]$ from consuming one unit of product x at price p given their preference θ and information cost $c(\tau)$ is

$$U_i(x, \theta, p, \tau) = u(x, \theta) - p - \kappa c(\tau).$$

We interpret $u(x, \theta)$ as the utility of consuming the good with attribute x. Consumers purchase whichever product gives them the higher expected utility.[34] We call the expected ideal point of a consumer their *revealed preference*.

Firms There are two firms, labelled a and b, in the market. Firms simultaneously choose their respective product locations, $x_a \in \mathbb{R}^n$ and $x_b \in \mathbb{R}^n$, and afterwards simultaneously choose their respective prices, p_a and p_b. Both firms have identical constant marginal costs, which we normalize to zero, so prices should be interpreted as markups. Firms maximize profits, that is market share times price. Given the distribution ρ of revealed preferences,

[34] In other words, we assume that the utility from not purchasing any product is sufficiently low that all consumers prefer to buy one of the available products in equilibrium.

the utility U_a of firm a is

$$U_a(x_a, x_b, p_a, p_b, \rho) = p_a \cdot \mathbb{E}_{\theta\sim\rho}\big[\mathbb{1}\big(u(x_a, \theta) - u(x_b, \theta) \geq p_a - p_b\big)\big],$$

and analogously for firm b.

We assume that consumers preferences are drawn from a normal distribution $\mathcal{N}(0, \Sigma)$ and consumers are restricted to normal signal structures. This is done, so the resulting distribution of revealed preferences is necessarily log-concave (it is normal), by which there exists a pure-strategy equilibrium of the price subgame for all product locations (Caplin and Nalebuff, 1991). We study the case of no aggregate uncertainty, assuming consumer ideal points are independently distributed, and maintain the focus on pure-strategy equilibria.

The following key lemma shows two important differences between the political and the market context. First, product locations are on the line of revealed preferences, whereas policy platforms are not necessarily on the line of revealed ideal points. Second, the direction of product differentiation is an eigenvector of ΣA, whereas the direction of platform differentiation is determined by the ideological difference between parties. Furthermore, the lemma allows us to reduce the model to one dimension and apply results from one-dimensional Hotelling models (Anderson et al., 1997) for our subsequent equilibrium characterization.

Lemma 2. *Consumers' revealed preferences and product locations are supported on the same line, the direction of which is an eigenvector of* ΣA.

The intuition is as follows. Analogous to the political-economy context, consumers' best response to product locations is to learn such that revealed preferences are on a line with the direction $\Sigma A(x_b - x_a)$. On the other hand, we show firms' best response is to locate their products on the line of consumer preferences: any other location is dominated by its projection on the line. Combining both, product differentiation $x_b - x_a$ must be parallel to $\Sigma A(x_b - x_a)$. That is, firms differentiate their products in a direction that is an eigenvector of ΣA.

Thus, there is a potential multiplicity of equilibria in this model. In fact, the following theorem shows that there is an equilibrium for any eigenvector of ΣA, provided the information cost parameter is small enough. To state the theorem, let $(v_1, ..., v_n)$ denote a basis of A-normalized eigenvectors of ΣA, that is, there exists $\lambda_i \in \mathbb{R}$: $\Sigma A v_i = \lambda_i v_i$ and $v_i^\top A v_i = 1$ for all $i \in \{1, ..., n\}$.[35] Let ρ denote the distribution of revealed consumer preferences in equilibrium.

Theorem 7. *The set of equilibria is fully characterized as follows:*

- *There is an equilibrium without learning, product differentiation, or markups:*

$$\sigma_\rho = 0, \quad x_a = x_b = 0, \quad p_a = p_b = 0.$$

- *For all $i \in \{1, ..., n\}$, if*

$$\kappa < \frac{3}{2} v_i^\top A \Sigma A v_i,$$

there is an equilibrium with positive learning, product differentiation, and markups, given by

$$\rho = \mathcal{N}\left(0, \sigma_\rho^2 v_i v_i^\top\right), \quad \sigma_\rho^2 = v_i^\top A \Sigma A v_i - \frac{2}{3}\kappa, \quad -x_a = x_b = \frac{3}{4}\sqrt{2\pi}\sigma_\rho v_i, \quad p_a = p_b = 3\pi\sigma_\rho^2.$$

The intuition for the equilibrium without differentiation is as follows. In standard Hotelling models, firms differentiate their products to soften price competition (see, for example, d'Aspremont et al., 1979). However, this mechanism relies on consumer heterogeneity. If consumers do not acquire information and their revealed preferences are all located at zero, firms do not differentiate their products and charge no markups. Anticipating identical product locations, consumers choose not to learn, confirming this as an equilibrium.

[35] Such a basis is given by the basis in which A and Σ^{-1} are simultaneously diagonalized as quadratic forms.

In the equilibrium with consumer learning, consumers acquire information that disperses their revealed preferences ($\sigma_\rho > 0$). Firms respond by differentiating their products, which allows them to charge positive markups and earn profits. As the information cost κ decreases, consumers learn more, increasing σ_ρ. This greater consumer differentiation leads to higher product differentiation, further softening price competition and resulting in higher prices. Notably, prices increase quadratically in the dispersion of consumer preferences σ_ρ because *both* more differentiated products and more dispersed consumer preferences reduce competitive pressures.

A priori, it is unclear whether lower information costs benefit or harm consumers. While prices increase as κ decreases, product differentiation may benefit consumers by allowing a better match to products, and there is a direct positive effect from the reduced information cost. To address this question, we focus on the firm-optimal equilibrium, where products differentiate along the eigenvector v_i that maximizes $v_i^\top A\Sigma A v_i$. However, the comparative statics would be the same if we considered the equilibrium associated with another eigenvector. We assess the aggregate utilitarian welfare of consumers, or, equivalently, their ex-ante utility. Define

$$\beta := 3e^{-(2+\frac{3}{4})\pi} \approx 0.04,$$
$$\overline{\kappa} := \frac{3}{2} v_i^\top A\Sigma A v_i.$$

Corollary 4 (Welfare Comparison)**.** *As the information cost parameter κ increases, consumer welfare*

- *decreases strictly for $\kappa \in [0, \beta\overline{\kappa})$,*
- *increases strictly for $\kappa \in (\beta\overline{\kappa}, \overline{\kappa})$,*
- *remains constant for $\kappa \geq \overline{\kappa}$.*

Consumer welfare is maximized in the last case, in which there is no product differentiation.

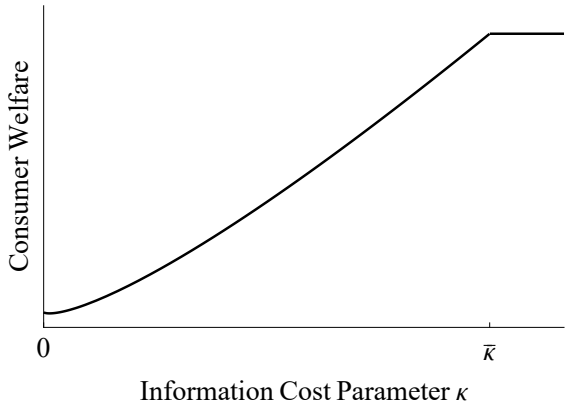


Figure 1.5: Consumer welfare increases as information becomes more costly.

Firm profits are decreasing in κ*.*

Figure 1.5 illustrates how consumer welfare varies with κ. Welfare decreases briefly when κ is small (from 0 to approximately 0.04 $\overline{\kappa}$) and then increases until κ reaches $\overline{\kappa}$.

Our result shows that, except for a negligibly small region, cheaper information harms consumers. Although product differentiation can theoretically benefit consumers, firms differentiate excessively from a social standpoint (Anderson et al., 1997), and increasingly so for lower information costs. Moreover, prices increase *quadratically* with σ_ρ because both more differentiated products and more dispersed consumer preferences reduce the intensity of competition. These negative effects outweigh the direct benefits consumers receive from lowering information costs. As a result, overall consumer welfare decreases when information becomes cheaper.

Our findings contribute to the literature on the welfare effects of information in markets. It is known that information can harm consumers by leading to higher prices (Moscarini and Ottaviani, 2001; Choi et al., 2018; Armstrong and Zhou, 2022; Albrecht and Whitmeyer, 2023; Biglaiser et al., 2024). Our analysis additionally incorporates endogenous product characteristics, which exacerbates the effect of consumer information on prices. New to this paper is, to the best of our knowledge, the comparative static result taking into account welfare effects of the cost of information, when both product characteristics and prices are endogenous.

1.7 Conclusion

In this paper, we dispense with the assumption that voters perfectly know their political preferences. Instead, voters can flexibly learn about their ideal points at a cost and do so for the purpose of expressing their political opinion in elections. Voters' choice set shapes—through learning about ideal points—the revealed ideology in the population. Because voters' choice set is constrained to the two party platforms in policy space, voters are not motivated to learn in a multidimensional or continuous way about their ideal points. As a result, revealed ideology displays issue alignment, and polarizes in the sense of approaching a binary distribution as valence becomes less uncertain. Voter learning predicts that polarization of voters and parties are mutually reinforcing and increase as information becomes cheaper. Finally, because voters only learn about the axis of party disagreement, policy is not responsive to dimensions of voter preferences that are orthogonal to this axis.

This paper opens several avenues for future research. One important question is what happens when voters learn jointly about their ideal points and party platforms. Such joint learning could explain the correlation between voters' perceptions of candidate positions and their own ideologies (Hare et al., 2015). Furthermore, while we considered voter learning occurring either before or after platform choice, in reality, both processes may coexist, especially since elections are held repeatedly. Some voters have acquired information during past elections, while other voters form opinions after observing current party campaigns. Additionally, examining repeated elections could shed more light on the dynamics of polarization.

1.8 Appendix: Main Proofs

Throughout we use the notation $\langle x, y\rangle_A := x^\top A y$ and $\langle x, y\rangle := x^\top y$ for $x, y \in \mathbb{R}^n$, $A \in \mathbb{R}^{n\times n}$.

Proof of Theorem 1

Proof. The proof consists of two parts. In the first part, we show that only the A-projection[36] of the posterior mean on the platform difference $x_b - x_a$ is payoff relevant. In the second part of the proof, we show via a reflection argument that a voter acquires a distribution over posteriors such that the distribution over posteriors means has support on the line through the origin and $\Sigma A(x_b - x_a)$.

Part I The instrumental utility of τ, that is the objective of (P) neglecting the information cost, can be rewritten as follows:

$$\begin{aligned}
&\mathbb{E}_\tau\Big[\mathbb{E}_\nu\Big[\max\Big\{\mathbb{E}_\pi\big[-\langle\theta - x_a, \theta - x_a\rangle_A\big], \mathbb{E}_\pi\big[-\langle\theta - x_b, \theta - x_b\rangle_A\big] + \nu\Big\}\Big]\Big]\\
=&\mathbb{E}_\tau\Big[\mathbb{E}_\nu\Big[\max\Big\{-\langle x_a, x_a\rangle_A + 2\langle x_a, \mathbb{E}_\pi[\theta]\rangle_A, -\langle x_b, x_b\rangle_A + 2\langle x_b, \mathbb{E}_\pi[\theta]\rangle_A + \nu\Big\}\Big] - \mathbb{E}_\pi\big[\langle\theta,\theta\rangle_A\big]\Big]\\
=&\mathbb{E}_\tau\Big[\mathbb{E}_\nu\Big[\max\Big\{-\langle x_a, x_a\rangle_A + 2\langle x_a, \mathbb{E}_\pi[\theta]\rangle_A, -\langle x_b, x_b\rangle_A + 2\langle x_b, \mathbb{E}_\pi[\theta]\rangle_A + \nu\Big\}\Big]\Big] - \mathbb{E}_\mu\big[\langle\theta,\theta\rangle_A\big]\\
=&\mathbb{E}_\tau\Big[\mathbb{E}_\nu\Big[\max\Big\{-\langle x_a, x_a\rangle_A + \langle x_a - x_b, \mathbb{E}_\pi[\theta]\rangle_A, -\langle x_b, x_b\rangle_A + \langle x_b - x_a, \mathbb{E}_\pi[\theta]\rangle_A + \nu\Big\}\Big]\Big] + C_1\\
=&\mathbb{E}_\tau\Big[\mathbb{E}_\nu\Big[\max\Big\{\Big\langle x_a - x_b, \mathbb{E}_\pi[\theta] - \frac{x_a + x_b}{2}\Big\rangle_A, \Big\langle x_b - x_a, \mathbb{E}_\pi[\theta] - \frac{x_a + x_b}{2}\Big\rangle_A + \nu\Big\}\Big]\Big] + C_1 + C_2
\end{aligned} \tag{1.11}$$

where

$$\begin{aligned}
C_1 &= \mathbb{E}_\tau[\langle x_a + x_b, \mathbb{E}_\pi[\theta]\rangle_A] - \mathbb{E}_\mu[\langle\theta,\theta\rangle_A] = \mathbb{E}_\mu[\langle x_a + x_b, \theta\rangle_A - \langle\theta,\theta\rangle_A]\\
C_2 &= -\frac{1}{2}(\langle x_a, x_a\rangle_A + \langle x_b, x_b\rangle_A)
\end{aligned}$$

[36]For any symmetric, positive definite matrix $B \in \mathbb{R}^{n\times n}$ and vector $v \in \mathbb{R}^n$, we refer to $\frac{\langle v,\theta\rangle_B}{\langle v,v\rangle_B} v = \frac{v^\top B\theta}{v^\top B v} v$ as the B-projection of θ on v.

are constants. In the second line, we used that the expectation is a linear operator and that the inner product is linear. In the third line, we used the law of iterated expectations to show only the first moment of the posterior is payoff-relevant.

Because $\mathbb{E}_\nu[\max\{x, y+\nu\}]$ is a function of x, y, and the distribution of ν only, (1.11) shows that the instrumental value of information depends only on the distribution of the A-projection of the posterior mean $\mathbb{E}_\pi[\theta]$ on the platform difference $x_b - x_a$.

Part II Suppose, for the sake of contradiction, that the voter acquires a distribution τ over posteriors that induces a distribution ρ of posterior means, which does not have has support on the line through the prior mean (origin) and $\Sigma A(x_b - x_a)$. We construct—through reflection, mixing, and garbling—a strictly cheaper but instrumentally as valuable distribution $\hat{\tau}$ over posteriors that induces a distribution $\hat{\rho}$ over posterior means supported on the line through the origin and $\Sigma A(x_b - x_a)$. The three steps of our construction of $\hat{\tau}$ are visualized in Figure 1.2.

Step 1 - Reflection: We define a Bayes-consistent distribution $\mathrm{Ref}(\tau)$ of the "reflected" posteriors that has the same information cost as well as instrumental value as τ.

To prove the result for general A and Σ, we make use of the Σ^{-1}-reflection Ref across the line through the origin and $\Delta\hat{x} := \Sigma A(x_b - x_a)$, defined as

$$\mathrm{Ref}(\theta) = 2\frac{\langle \Delta\hat{x}, \theta\rangle_{\Sigma^{-1}}}{\langle \Delta\hat{x}, \Delta\hat{x}\rangle_{\Sigma^{-1}}}\Delta\hat{x} - \theta.$$

This is a well-defined reflection because the symmetric, positive definite matrix Σ has a symmetric, positive definite inverse Σ^{-1}. In the simple case when Σ and A are equal to the identity matrix, Ref is simply the standard reflection across the line through the origin and $x_b - x_a$. In the general case, this reflection is useful for two reasons.

First, the Σ^{-1}-projection on $\Delta\hat{x}$ preserves the instrumental value. That is because the projection is equivalent to the A-projection on the platform difference $x_b - x_a$ by

$$\langle x_b - x_a, \theta\rangle_A = (x_b - x_a)^\top(\Sigma^{-1}\Sigma)A(x_b - x_a) = (\Sigma A(x_b - x_a))^\top\Sigma^{-1}\theta = \langle \Delta\hat{x}, \theta\rangle_{\Sigma^{-1}}. \quad (1.12)$$

That the reflection Ref preserves the Σ^{-1}-projection on $\Delta\hat{x}$ thus implies preserving the payoff-relevant A-projection on the platform difference $x_b - x_a$.

Second, Ref preserves the prior μ. Note that Ref is a linear function, which we can describe as multiplication by a matrix Q, $\text{Ref}(\theta) = Q\theta$. By Ref being a reflection, we have $Q = Q^{-1}$. Ref is a reflection with respect to inner product Σ^{-1}, so Ref preserves the distance induced by Σ^{-1}. Hence, we have $Q^\top \Sigma^{-1} Q = \Sigma^{-1}$, inverting which delivers $Q\Sigma Q^\top = \Sigma$. The characteristic function of $Q\theta$, as a random vector, satisfies for all $t \in \mathbb{R}^n$, $\Phi_{Q\theta}(t) = \Phi_\theta(Q^\top t) = \psi(t^\top Q\Sigma Q^\top t) = \psi(t^\top \Sigma t) = \Phi_\theta(t)$. Thus, $\text{Ref}(\theta) = Q\theta$ and θ have the same distribution.

The reflection Ref of the state space induces an according reflection on posteriors and on distributions over posteriors through the pushforward and iterated pushforward, which we both denote by Ref as well.[37] Intuitively, we are simply relabeling the states.

The distribution $\text{Ref}(\tau)$ is Bayes-consistent, since for all Borel sets $A \in \mathbb{R}^n$,

$$\int \pi(A) d(\text{Ref}(\tau)) = \int \text{Ref}(\pi)(A) d\tau = \text{Ref}(\mu)(A) = \mu(A),$$

where the last equality holds by the prior μ being invariant under Ref.

As argued above, Ref preserves the A-projection on $x_b - x_a$, and by linearity of Ref it maintains the distribution of the A-projection of the posterior *mean* on the platform difference. Thus, the instrumental value of τ is preserved under Ref. The information cost is also preserved under Ref since

$$\int D(\pi||\mu) d(\text{Ref}(\tau)) = \int D(\text{Ref}(\pi)||\mu) d\pi = \int D(\text{Ref}(\pi)||\,\text{Ref}(\mu)) d\tau = \int D(\pi||\mu) d\tau,$$

where the last equality holds because the Kullback-Leibler divergence is invariant under coordinate transformations. Thus, the voter is indifferent between τ and $\text{Ref}(\tau)$. Finally,

[37] The pushforward $\text{Ref}_*\colon \Delta(\mathbb{R}^n) \to \Delta(\mathbb{R}^n)$ is formally defined via $\text{Ref}_*(\pi)(A) = \pi(\text{Ref}^{-1}(A))$ for Borel sets $A \subseteq \mathbb{R}^n$. The iterated pushforward on $\Delta(\Delta(\mathbb{R}^n))$ is simply $(\text{Ref}_*)_*$. For ease of reading, we write Ref for both Ref_* and $(\text{Ref}_*)_*$.

the distribution over posteriors means induced by $\text{Ref}(\tau)$ is simply $\text{Ref}(\rho)$, the reflection of the distribution over posterior means induced by τ.

Step 2- Mixing: It follows immediately that the mixture $\frac{1}{2}\tau + \frac{1}{2}\text{Ref}(\tau)$ is also Bayes-consistent and has the same instrumental value. It also has the same information cost by posterior separability of the information cost, which implies that the cost is linear under mixing. By posterior separability, the information cost can be written as an expectation with respect to the distribution τ over posteriors, which is linear in the distribution τ.

Step 3 - Garbling: Finally, we take a certain mean-preserving contraction of $\frac{1}{2}\tau + \frac{1}{2}\text{Ref}(\tau)$ to reach $\hat{\tau}$ (which corresponds to a garbling of the corresponding signal structure), which is also Bayes-consistent and has the same instrumental value, *but has a lower information cost*. We use the mean preserving contraction that contracts all posteriors whose means have the same Σ^{-1}-projection on the line through $\Sigma A(x_b - x_a)$. Any mean-preserving contraction is Bayes-consistent. The contraction $\hat{\tau}$ preserves the instrumental value of information since it preserves the distribution of the A-projection of the posterior mean on $x_b - x_a$. Crucially, the distribution over posterior means $\hat{\rho}$ induced by $\hat{\tau}$ has support on the the through the origin and $\Sigma A(x_b - x_a)$. The reason is that $\frac{1}{2}\tau + \frac{1}{2}\text{Ref}(\tau)$ is constructed to be symmetric around this line with respect to the Σ^{-1}-projection. Finally, a mean-preserving contraction lowers the information cost by convexity of the Kullback-Leibler divergence $D_{KL}(\pi||\mu)$ in its first argument. In fact, the Kullback-Leibler divergence is strictly convex for π that are absolutely continuous with respect to μ by strict convexity of $x \log x$ and $D_{\text{KL}}(\pi||\mu) = \int \log\left(\frac{d\pi}{d\mu}\right)\frac{d\pi}{d\mu}d\mu$. Without loss, the posteriors π induced by τ are almost surely absolutely continuous with respect to prior μ, otherwise τ has infinite cost and is clearly suboptimal. Thus, $\hat{\tau}$ has a strictly lower information cost than the original distribution τ if τ did not have posterior means already on the line, in which case the mean-preserving contraction is strict. □

In Appendix 1.11, we discuss more general information costs, such as certain distance-based costs, for which this proof works.

Proof of Theorem 2

Proof. We show, for completeness, that *any* optimal signal structure has no more signals than actions under degenerate valence shock, $\nu \sim \delta_0$. The proof uses only convexity properties of the Kullback-Leibler divergence D_{KL}.

The voter acquires a distribution τ over posteriors π, such that the posterior is almost surely absolutely continuous with respect to the prior μ. Otherwise, the voter obtains a negative infinite payoff and could do better by acquiring no information. Suppose, for the sake of contradiction, that the voter acquires a signal structure such that the induced distribution τ over posteriors has a support with more than two posteriors. Then, we can strictly improve the voter's utility by garbling the signal structure based on the action recommendation (after resolving indifferences between parties a and b for a, say). More precisely, we partition the space of posteriors $\Delta(\mathbb{R}^n)$ into the subset Δ_a on which voting for a is weakly preferred and another subset Δ_b on with voting for b is strictly preferred. The garbling corresponds to a mean-preserving contraction of τ, namely contracting all posteriors in Δ_a and in Δ_b, respectively, inducing a binary distribution over posteriors. The voter's utility from a distribution τ over posteriors π is the expected value of the value function, which consists of the instrumental value and the Kullback-Leibler divergence,

$$\int \left(\max\left\{\mathbb{E}_\pi[u(x_a,\theta)], \mathbb{E}_\pi[u(x_b,\theta)]\right\} - \kappa D_{\mathrm{KL}}(\pi||\mu)\right) d\tau.$$

On each Δ_a and Δ_b, the instrumental value (the max term) is linear in the posterior. The divergence D_{KL} is strictly convex in posterior π for π absolutely continuous with respect to μ, as argued in step 3 of the proof of Theorem 1. Thus, the value function is strictly concave on the support of τ in Δ_a and in Δ_b, each. By Jensen's inequality, our mean-preserving contraction on Δ_a and Δ_b weakly improves the voter's utility. Because the mean-preserving contraction is strict on at least one of Δ_a and Δ_b, the voter's utility improves strictly.

The second part of Theorem 2 follows immediately from Proposition 4 and Lemma 17 in the Appendix 1.11. □

Proof of Lemma 1

Proof. First, we show the best response x_a of party a to x_b and ρ necessarily satisfies the first-order condition of party a's objective. The argument is analogous for party b. If party a chooses $x_a = x_a^*$, it obtains utility greater or equal to zero because x_a delivers zero ideological utility, $u(x_a, x_a^*) = 0$, and a non-negative vote share. If x_a is outside the ellipse described by $u(x_a, x_a^*) \geq -m$, then the resulting utility is negative, because the utility from the vote share can be at most m. Thus, all x_a outside this compact ellipse are suboptimal and by differentiability of the vote share in x_a, shown below, the maximum is obtained on this ellipse and necessarily satisfies the first-order condition, which we analyze next.

Taking the gradient ∇ with respect to x_a of the objective of party a, we obtain the necessary first-order condition of the optimal platform x_a given ρ and x_b:

$$
\begin{aligned}
&\nabla\left(m\int F_\nu\big(u(x_a,\theta)-u(x_b,\theta)\big)d\rho(\theta)+u(x_a,x_a^*)\right)=0\\
&\Leftrightarrow m\int \underbrace{f_\nu\big(u(x_a,\theta)-u(x_b,\theta)\big)}_{=:w(\theta)}\nabla u(x_a,\theta)d\rho(\theta)+\nabla u(x_a,x_a^*)=0\\
&\Leftrightarrow \int mw(\theta)2A(x_a-\theta)d\rho(\theta)+2A(x_a-x_a^*)=0\\
&\Leftrightarrow 2A\left(m\int w(\theta)(x_a-\theta)d\rho(\theta)+(x_a-x_a^*)\right)=0\\
&\Leftrightarrow x_a=\frac{m\int w(\theta)\theta d\rho(\theta)+x_a^*}{m\int w(\theta)d\rho(\theta)+1}
\end{aligned}
$$

The integral of the expected vote share is well-defined. We can exchange integration and differentiation because the partial derivative of the integrand exists and is bounded in absolute value by an integrable function in θ. The latter holds because ν has finite first absolute moment and $\nabla u(x_a, \theta)$ is linear in θ. The last equivalence uses that A is symmetric and positive definite, so its kernel is $\{0\}$.

The result for platform x_b is analogous. Together, this implies

$$x_b - x_a = \frac{x_b^* - x_a^*}{m \int_{\mathbb{R}^n} w(\theta) d\rho(\theta) + 1}, \tag{1.13}$$

so $x_b - x_a$ is parallel to $x_b^* - x_a^*$. □

Proof of Theorem 3

Proof. By Lemma 18 in Appendix 1.11, under the restriction to normal signal structures, it still holds that voters' revealed ideal points are on the line through the origin and $\Sigma A(x_b - x_a)$. The first-order conditions that characterize the equilibrium platforms (Lemma 1) are unaffected by the component of voter ideal points orthogonal (with respect to A) to $x_b - x_a$. That is because the ideal point θ enters the first-order conditions only via the utility difference $u(x_a, \theta) - u(x_b, \theta 2la x_a - x_b, \theta - \frac{x_a + x_b}{2}\rangle_A$. This utility difference is unaffected by the component of θ orthogonal (with respect to A) to $x_b - x_a$. Thus, while in the following proof we assume that the line of voter ideal points is parallel to $x_b - x_a$, all steps generalize to a line of voter ideal points that is slanted with respect to $x_b - x_a$. Furthermore, in any equilibrium, $x_b - x_a$ is parallel to $x_b^* - x_a^*$ by (1.13). To simplify exposition we change into an orthonormal basis of A in which $x_b^* - x_a^*$, and hence $x_b - x_a$, is parallel to the first basis vector. Such a basis exists by the Gram-Schmidt algorithm.

We show all equilibria are symmetric. By Lemma 1, equilibrium platforms can be written as a weighted average of voter ideal points and an *aggregate voter ideal point* $\overline{\theta}$,

$$\overline{\theta} := \frac{\int w(\theta)\theta d\rho(\theta)}{\int w(\theta) d\rho(\theta)},$$

by

$$x_j = \frac{m \int w(\theta)\theta d\rho(\theta) + x_j^*}{m \int w(\theta) d\rho(\theta) + 1} = \frac{m\overline{\theta} + \frac{1}{\int w(\theta) d\rho(\theta)} x_j^*}{m + \frac{1}{\int w(\theta) d\rho(\theta)}} \tag{1.14}$$

for $j = a, b$. We show in any equilibrium, the aggregate voter ideal point $\overline{\theta}$ is zero, which implies a symmetric equilibrium

$$(x_a, x_b) = \frac{\frac{1}{\int w(\theta) d\rho(\theta)}}{m + \frac{1}{\int w(\theta) d\rho(\theta)}} (x_a^*, x_b^*).$$

Suppose, for the sake of contradiction, that the aggregate voter ideal point $\overline{\theta}$ was not zero. By the paragraph above, the revealed voter ideal points are on a line parallel to the first basis vector. Thus, the aggregate voter ideal point $\overline{\theta}$ is on this line. Suppose without loss that its first component is positive, $\overline{\theta}_1 > 0$. We show this implies that the platform midpoint $\overline{x} := \frac{x_a + x_b}{2}$ must be to the left of the aggregate voter ideal point $\overline{\theta}$, which in turn must be to the left of the platform midpoint, creating a contradiction. By (1.14) and $x_{a,1}^* = -x_{b,1}^*$, the first component of the platform midpoint $\overline{x}_1$ is positive by

$$\overline{x}_1 = \frac{m\overline{\theta}_1 + \frac{1}{\int w(\theta) d\rho(\theta)} (x_{a,1}^* + x_{b,1}^*)}{m + \frac{1}{\int w(\theta) d\rho(\theta)}} = \frac{m}{m + \frac{1}{\int w(\theta) d\rho(\theta)}} \overline{\theta}_1 \in (0, \overline{\theta}_1). \tag{1.15}$$

Under normal signals, the distribution over posterior means is symmetric around 0 and quasi-concave. Thus, a positive party midpoint, $\overline{x}_1 > 0$, implies that the weighted mass of ideal points to the left of $\overline{x}$ is greater than to the right, so the aggregate voter ideal point must be to the left of the party midpoint. Formally, writing vectors in row-notation, $\forall y > 0,\ w((\overline{x}_1 - y, 0, ..., 0)) = w((\overline{x}_1 + y, 0, ..., 0))$ but the density of revealed voter ideal points is greater at $(\overline{x}_1 - y, 0, ..., 0)$ than at $(\overline{x}_1 + y, 0, ..., 0)$. Thus, $\overline{\theta}_1 = \int w(\theta) \theta d\rho(\theta) < \overline{x}_1$ in contradiction to (1.15).

Symmetric party platforms

$$(x_a, x_b) = \alpha(x_a^*, x_b^*) \tag{1.16}$$

satisfy the first-order conditions of optimality if

$$\alpha = \frac{1}{m \int f_\nu(\alpha \langle x_b^* - x_a^*, \theta \rangle) d\rho(\theta) + 1}. \tag{1.17}$$

We call α the degree of platform polarization. We establish equilibrium existence by showing that there exists a degree of platform polarization $\alpha \in (0,1)$ that satisfies (1.17) where ρ is the induced distribution over revealed ideal points when the party platforms satisfy (1.16). We do so by constructing an equilibrium correspondence whose fixed point exists by monotonicity properties. Conceptually, this proof of pure-strategy equilibrium existence and the subsequent comparative statics result are similar to those in supermodular games.

We construct a correspondence G from $[0,1]$ to $(0,1)$ as a concatenation of two correspondences, g_1 and g_2. Let g_1 map $\alpha \in [0,1]$ to set of distributions ρ over posterior means induced by some optimal learning strategy given the degree of platform polarization α. We can restrict attention to one-dimensional normal signals, which can be parametrized and ordered by the variance σ_ρ^2 of ρ. Because this variance is bounded by the prior variance in that dimension, standard arguments deliver that g_1 is nonempty compact-valued. Let g_2 map some distribution ρ of revealed ideal points to the set of equilibrium α that satisfy the first-order condition (1.17). The function g_2 is nonempty-valued with values in $(0,1)$ because both the left-hand side and right-hand side of (1.17) are continuous and at $\alpha = 0$ the right-hand side is larger while at $\alpha = 1$, the left-hand side is larger. By the intermediate value theorem, a solution α exists. Define the correspondence $G = g_2 \circ g_1$,

$$\begin{aligned} G\colon \quad [0,1] \quad &\longrightarrow \quad 2^{[0,1]} \\ \alpha \overset{g_1}{\longmapsto} \{\rho\} &\overset{g_2}{\longmapsto} \{\alpha\} \end{aligned}$$

Lemma 3. *Platform polarization increases voter polarization, that is* $\min g_1$ *and* $\max g_1$ *are strictly increasing.*

This lemma follows from Proposition 1.

Lemma 4. *Voter polarization increases platform polarization, that is,* $\min g_2$ *and* $\max g_2$ *are strictly increasing in the variance* σ_ρ^2 *of the symmetric, normal distribution* ρ *of voter ideal points.*

Proof. The smallest and largest α that solve (1.17) exist because g_2 is nonempty by the above and because the left- and right-hand side of (1.17) are continuous, so the preimage of $\{0\}$ under the continuous difference between the left- and right-hand side is closed.

A higher variance σ_ρ^2 of ρ implies that voters are strictly further away from 0 (the projection of the party midpoint) in first-order stochastic dominance. This implies that the term $\int f_\nu(\alpha\langle x_b^* - x_a^*, \theta\rangle)d\rho(\theta)$ strictly decreases by strict quasi-concavity of f_ν. Thus, the right-hand side of (1.17) strictly increases pointwise, which implies a greater smallest and largest α that solves (1.17). □

Together, this implies that the minimum and maximum of $G : [0,1] \to [0,1]$,

$$(\min G)(\alpha) := \min\{G(\alpha)\} = \min g_2(\min g_1(\alpha))$$
$$(\max G)(\alpha) := \max\{G(\alpha)\} = \max g_2(\max g_1(\alpha))$$

are strictly increasing. Fixed points of $\min G$ and $\max G$ correspond to equilibrium degrees of platform polarization. These exist by Tarski's fixed point theorem, viewing $[0,1]$ with the usual ordering as a complete lattice using monotonicity of $\min G$ and $\max G$.

Finally, we show the comparative statics result.

Lemma 5. *A smaller κ implies that the smallest and largest fixed point of G weakly increase.*

Proof. A smallest and largest fixed point of G exist by Tarski's fixed point theorem.

The functions $\min G$ and $\max G$ weakly increase pointwise as κ decreases. This implies that the smallest and largest intersection of G with the identity function on $[0,1]$ increase strictly, leading to a higher smallest and largest equilibrium α. To prove this, note first that g_2 is unchanged. The functions $\min g_1$ and $\max g_1$ weakly increase pointwise by Proposition 1. Together with $\min g_2$ and $\max g_2$ being non-decreasing, this implies $\min G$ and $\max G$ are weakly higher pointwise. □

This concludes the proof of Theorem 3. □

Proof of Theorem 4 Before we prove Theorem 4, we formally describe the strategies and payoffs of players. We also pave the way for our proof by introducing a way to represent the extensive-form strategies of our two parties and continuum of voters as a static game between only four players.

Players and Strategies: Parties $j \in \{a, b\}$ choose their platforms conditional on the realized public opinion signal $s \in S$. Formally, a strategy of party j is a function $x_j : S \to \mathbb{R}^n$.

After learning, voters choose who to vote for conditional on the public signal s and the realized party platforms $(x_a(s), x_b(s))$. Formally, voters choose vote choice functions $v : S \times \mathbb{R}^n \times \mathbb{R}^n \to \{a, b\}$ that map $(s, x_a(s), x_b(s))$ into a choice among parties. Because the posterior mean is a sufficient statistic for optimal voting behavior (Remark 1), we can code subgame-perfect vote choice functions in the following strategically equivalent reduced-form way: voters choose a posterior mean $p \in \mathbb{R}^n$ conditional on each public signal s, incorporating that, subsequently, they choose optimally between $x_a(s)$ and $x_b(s)$ given posterior mean p and their valence shock. We denote a generic reduced-form strategy by $p_S \in (\mathbb{R}^n)^S$. As usual in rational inattention, it is without loss to identify the signal space with the set of actions, $\mathcal{S} = (\mathbb{R}^n)^S$. Thus, voter i's extensive-form strategy is reduced to a signal structure (stochastic kernel) $\sigma_i : \Omega \times D \to \Delta(\mathbb{R}^n)^S$.

Because voters are ex-ante homogeneous and we assume that all voters acquire the same signal structure, we model our continuum of voters through a *representative voter*, who chooses a signal structure $\sigma : \Omega \times D \to (\mathbb{R}^n)^S$.

As mentioned in the main text, a single, infinitesimal voter cannot affect the realized signal distribution. Modeling voters as a representative voter, we have to ensure that, when we define payoffs, the representative voter's signal structure does not affect the public signal. Therefore, we introduce a fourth, *fictional player* who also chooses a signal structure $\sigma_f : \Omega \times D \to \Delta((\mathbb{R}^n)^S)$. We will define the payoff of the representative voter such that in equilibrium, it mimics the signal structure of the representative voter, ensuring consistency of the public opinion signal with voters' learning strategies.

In particular, the public signal is obtained as follows from σ_f. By the continuum of voters and the compact support of p_S, which implies finite second moments, we can apply a law of large numbers for a continuum of random variables (Uhlig, 1996) if we interpret realized population distributions as Pettis integrals. In particular, if voters acquire the signal structure σ_f, then, conditional on the aggregate state ω, the realized distribution of signals is deterministically $\sum_\delta \mu(\omega, \delta)\sigma_f(\cdot|\omega, \delta)$. Thus, the probability of a certain public opinion signal s conditional on aggregate state ω and learning strategy σ_f is

$$\sigma_p(s|\omega, \sigma_f) := \sigma_p\left(s \middle| \sum_\delta \mu(\omega, \delta)\sigma_f\big(\,\cdot\,|\omega, \delta\big)\right).$$

From this definition, it follows that the probability $\sigma_p(s|\omega, \sigma_f)$ is continuous under pointwise weak convergence of σ_f.

To sum up, we have represented our game as a static game between four players (two parties, a representative voter, and a fictional player), who choose strategies

$$(x_a, x_b, \sigma, \sigma_f) \in (\mathbb{R}^n)^S \times (\mathbb{R}^n)^S \times (\Delta((\mathbb{R}^n)^S))^{\Omega\times D} \times (\Delta((\mathbb{R}^n)^S))^{\Omega\times D}.$$

Payoffs: The payoff U_a of party a is

$$U_a(x_a, x_b, \sigma, \sigma_f) :=$$
$$\sum_{\omega,\delta,s} \mu(\omega, \delta)\sigma_p(s|\omega, \sigma_f)\left(m \int F_\nu\Big(u\big(x_a(s), p_S(s)\big) - u\big(x_b(s), p_S(s)\big)\Big)d\sigma(p_S|\omega, \delta) + u(x_a(s), x_a^*)\right)$$

The payoff U_b for party b is defined analogously.

The utility U_v of the representative voter is

$$U_v(x_a, x_b, \sigma, \sigma_f) := \sum_{\omega,\delta,s} \mu(\omega,\delta)\sigma_p(s|\omega,\sigma_f) \int v\Big(p_S(s), \omega+\delta, x_a(s), x_b(s)\Big) d\sigma\big(p_S(s)|\omega,\delta\big) \\ -c(\sigma) \tag{1.18}$$

where

$$\begin{aligned} v(p,\theta,x,y) &:= \int_{\mathbb{R}} \begin{cases} u(x,\theta) & \text{if } u(x,p) \geq u(y,p)+\nu \\ u(y,\theta)+\nu & \text{if } u(x,p) < u(y,p)+\nu \end{cases} dF(\nu) \\ &= F_\nu\big(u(x,p)-u(y,p)\big)u(x,\theta) + \big(1-F_\nu(u(x,p)-u(y,p))\big)u(y,\theta) \\ &\quad + \int_{u(x,p)-u(y,p)}^{\infty} \nu dF(\nu), \\ c(\sigma) &:= D_{\mathrm{KL}}(P^{(\omega,\delta),p_S} || P^{(\omega,\delta)} \otimes P^{p_S}). \end{aligned}$$

The utility $v(p,\theta,x,y)$ captures that the voter votes optimally between platforms x and y given reported posterior mean p. The information cost $c(\sigma)$ is mutual information, which is the Kullback-Leibler divergence of the joint distribution $P^{(\omega,\delta),p_S}$ of state (ω,δ) and voter signal p_S from the product distribution $P^{(\omega,\delta)} \otimes P^{p_S} = \mu \otimes P^{p_S}$.

The fictional fourth player has payoff

$$U_f(x_a, x_b, \sigma, \sigma_f) := \begin{cases} 1 & \text{if } \sigma_f = \sigma, \\ 0 & \text{else.} \end{cases}$$

Thus, in equilibrium $\sigma_f = \sigma$.

Proof. First, we prove equilibrium existence through a fixed point theorem. Second, we prove existence of an equilibrium in which the desired statement of Theorem 4 holds.

Fixed Point Theorem: We apply the Kakutani-Fan-Glicksberg fixed point theorem (Aliprantis and Border, 2006, Theorem 17.55) to show existence of a pure-strategy equilibrium. It states that a correspondence Φ with closed graph and nonempty convex values on a nonempty compact convex subset K of a locally convex Hausdorff space has a fixed point.

Below, we define K as a nonempty compact convex subset of the strategy space

$$(\mathbb{R}^n)^S \times (\mathbb{R}^n)^S \times (\Delta((\mathbb{R}^n)^S))^{\Omega \times D} \times (\Delta((\mathbb{R}^n)^S))^{\Omega \times D},$$

by ruling out certain dominated strategies. Our strategy spaces are metrizable and hence Hausdorff. The weak topology is induced by a family of seminorms (the integral with respect to continuous bounded functions) and hence locally convex.

We construct Φ as the best-response correspondence. We show below that the best-response correspondences are upper hemicontinuous and nonempty compact-valued. Thus, by the closed graph theorem, the graph of Φ is closed. Finally, we show that the best-response correspondences are convex-valued, through showing that the payoffs are concave.

Compact and Convex Strategy Spaces: While strategy spaces are not compact, we can restrict attention to compact convex spaces of undominated strategies.

As shown in Appendix 1.11, party j would never choose a platform outside a certain compact ellipse around their ideal points, $\mathcal{E}_j$. By Tychonoff's theorme, the strategy space $\mathcal{E}_j^S$ is compact (under the topology of pointwise convergence) and because $\mathcal{E}_j$ is compact. It is convex because $\mathcal{E}_j$ is convex.

For a voter it is never optimal to report a posterior mean that lies outside of the convex hull $\text{conv}\Theta$ of the support Θ because the posterior mean given any belief must lie inside this convex hull. Because Θ is finite, $\text{conv}\Theta$ is compact, and therefore $(\text{conv}\Theta)^S$ is compact. Because $(\text{conv}\Theta)^S$ is also metrizable, $\Delta\left((\text{conv}\Theta)^S\right)$ is compact under the topology of weak convergence. The set $(\Delta((\text{conv}\Theta)^S))^{\Omega \times D}$ is compact under the pointwise topology of weak convergence.

The fictitional fourth player's best response is never outside the compact space of strategies $(\Delta((\text{conv}\Theta)^S))^{\Omega\times D}$ of the representative voter.

Formally, we can restrict attention to the compact and convex strategy space

$$K := \mathcal{E}_a^S \times \mathcal{E}_b^S \times (\Delta((\text{conv}\Theta)^S))^{\Omega\times D} \times (\Delta((\text{conv}\Theta)^S))^{\Omega\times D}.$$

Upper Hemicontinuous Best Response Correspondence: We show that the voter objective U_v is upper semicontinuous in $(x_a, x_b, \sigma, \sigma_f)$ and continuous in (x_a, x_b, σ_f). Together with the feasibility set of σ being nonempty compact-valued and constant, it follows from the generalization of Berge's maximum theorem due to Tian and Zhou (1992), analogous to our proof of Proposition 4, that the best-response correspondence is nonempty compact-valued and upper hemicontinuous.

As a first step, the instrumental value of information is jointly continuous in players' strategies. The function v in (1.18) is uniformly continuous in (x_a, x_b) due to continuous differentiability over a compact domain, and continuous in p by continuity of f_ν. Moreover, v is bounded by the compact domain of (p, θ, x_a, x_b) and by ν having finite absolute first moment. Thus, by the Portmanteau theorem, the integral in (1.18) is continuous under weak convergence of $\sigma(p_S(s)|\omega, \delta)$. By uniform convergence of the integrand in $(x_a(s), x_b(s))$, the integral is jointly continuous in $(\sigma, x_a(s), x_b(s))$ (see (1.32)). By continuity of $\sigma_p(s|\omega, \sigma_f)$ in σ_f and because the sum over (ω, δ, s) in (1.18) is finite, the instrumental value is jointly continuous in players' strategies.

As a second step, the information cost, which depends only on σ, is lower semicontinuous, making the voter objective jointly upper semicontinuous. By E. Posner (1975), the Kullback-Leibler divergence $D_{\text{KL}}(P||Q)$ is jointly lower semicontinuous under weak convergence of P and Q. The joint distribution $P^{(\omega,\delta),p_S}$ is just a finite average of the conditional distributions of p_S conditional on (ω, δ), that is, $\sigma(\cdot|\omega, \delta)$, so the joint distribution weakly converges as σ weakly converges pointwise. Similarly, the distribution P^{p_S} converges weakly

as σ does. The product measure $P^{(\omega,\delta)} \otimes P^{p_S}$ converges weakly if P^{p_S} does, which can be verified via the Portmanteau theorem by testing expectations $\mathbb{E}_{(\omega,\delta),p_S}[f((\omega,\delta),p_S)]$ under $((\omega,\delta),p_S) \sim P^{(\omega,\delta)} \otimes P^{p_s}$ of continuous bounded functions f. Such expectations converge because they are weighted averages of expectations that converge by the Portmanteau theorem, $\mathbb{E}_{(\omega,\delta),p_S}[f((\omega,\delta),p_S)] = \sum_\theta P^{(\omega,\delta)}(\omega,\delta)\mathbb{E}_{p_S}[f((\omega,\delta),p_S)]$.

Combining the first and second step, the voter objective is upper semicontinuous in $(x_a, x_b, \sigma, \sigma_f)$ and continuous in (x_a, x_b, σ_f).

The party objective U_j is jointly continuous in players' strategies by an analogous argument to the voter's instrumental value of information being jointly continuous. Thus, by Berge's maximum theorem, the best-response correspondence of parties is upper hemicontinuous and nonempty compact-valued and therefore has a closed graph.

The fictional player's utility is not continuous but their best response $\sigma_f = \sigma$ is nevertheless continuous in (x_a, x_b, σ).

Convex-valued Best Response Correspondence: The set of best responses σ to (σ_f, x_a, x_b) is convex, because U_v is concave in σ. The instrumental value of information is linear in σ and the Kullback-Leibler divergence $c(\sigma)$ is convex in the conditional distribution σ: both the joint and the product measure are linear in the conditional distribution σ and the Kullback-Leibler divergence is convex.

For the party objective to be concave in x_a, we require again that m is small enough or the valence shock is large enough, see Appendix 1.11. By compactness of convΘ, the set of distributions of posterior means supported on convΘ is compact. Then, by an argument as in Lemma 20, for m small enough or valence ν large enough, the party objective is strictly concave and the best response correspondence is single- and therefore convex-valued.

This concludes our proof of equilibrium existence. Next, we show that there exists an equilibrium where party platforms respond to ω only through $(x_b^* - x_a^*)^\top A\omega$.

Existence of an equilibrium with the desired properties: We show that there exists an equilibrium in which voters' acquired signal structures do not distinguish between aggregate

states ω that have the same A-projection on $x_b^* - x_a^*$, $\langle x_b^* - x_a^*, \omega \rangle_A$. Because party platforms respond to voter preferences (Lemma 1), this implies that party platforms do not distinguish between such states. Formally, there exists an equilibrium $(x_a, x_b, \sigma, \sigma_f)$ such that σ satisfies the measurability condition

$$\forall \omega, \omega', \delta, p_S \colon \langle x_b^* - x_a^*, \omega \rangle_A = \langle x_b^* - x_a^*, \omega' \rangle_A \to \sigma(p_s|\omega, \delta) = \sigma(p_s|\omega', \delta). \tag{1.19}$$

To show this, we restrict σ and σ_f to signal structures that satisfy the measurability condition (1.19) and x_a and x_b to functions such that the platform difference is necessarily parallel to the ideological difference of parties,

$$\forall s \colon x_a(s) - x_b(s) || x_a^* - x_b^*. \tag{1.20}$$

Formally, let $K' \subset K$ be the subset of strategies $(x_a, x_b, \sigma, \sigma_f) \in K$ that satisfy both the measurability condition (1.19) and the parallelity condition (1.20). The space K' is nonempty convex subset of K'. Because K' is a closed subspace of the compact space K, K' is compact. Then, we show that the best-response correspondence restricted to the subspace K' maps into K'. By the Kakutani-Fan-Glicksberg fixed-point theorem, there is an equilibrium where (1.19) holds.

First, parties' best responses satisfies the parallelity condition (1.20) by Lemma 1. Because parties care about the expected vote share, Lemma 1 also holds for any *belief* over the distribution of voter preferences that parties share.

Second, the best response of σ_f satisfies the measurability condition (1.19) if σ does, because the best response is simply $\sigma_f = \sigma$.

Third and finally, we prove that any best response σ to (x_a, x_b, σ_f) satisfies (1.19) if σ_f satisfies (1.19) and (x_a, x_b) satisfy (1.20). To prove this, we show that aggregate states with the same projection on the ideological difference of parties are payoff equivalent. Because

of the invariance property of mutual information, voters optimal learning strategy does not distinguish between payoff equivalent states.

Let ω and ω' be such that $\langle x_b^* - x_a^*, \omega\rangle_A = \langle x_b^* - x_a^*, \omega'\rangle_A$. If σ_f satisfies (1.19), then the public signal σ_p satisfies the measurability condition $\sigma_p(s|\omega, \sigma_f) = \sigma_p(s|\omega', \sigma_f)$ by

$$\begin{aligned}\sigma_p(s|\omega,\sigma_f) &= \sigma_p\left(s\Big|\sum_\delta \mu(\delta)\sigma_f\big(\,\cdot\,|\omega,\delta\big)\right)\\ &= \sigma_p\left(s\Big|\sum_\delta \mu(\delta)\sigma_f\big(\,\cdot\,|\omega',\delta\big)\right) = \sigma_p(s|\omega',\sigma_f).\end{aligned} \tag{1.21}$$

By $\langle x_b^* - x_a^*, \omega\rangle_A = \langle x_b^* - x_a^*, \omega'\rangle_A$, we have, for any δ, $\langle x_b^* - x_a^*, \theta\rangle_A = \langle x_b^* - x_a^*, \theta'\rangle_A$ where $\theta = \omega + \delta$ and $\theta' = \omega' + \delta$. By the parallelity condition (1.20), this implies for all $s \in S$

$$\begin{aligned}u(x_a(s),\theta) - u(x_b(s),\theta) &= \Big\langle x_a(s) - x_b(s), \theta - \frac{x_a(s)+x_b(s)}{2}\Big\rangle_A\\ &= \Big\langle x_a(s) - x_b(s), \theta' - \frac{x_a(s)+x_b(s)}{2}\Big\rangle_A = u(x_a(s),\theta') - u(x_b(s),\theta').\end{aligned}$$

Therefore,

$$\begin{aligned}v(p,\theta,x_a(s),x_b(s)) &= F_\nu\big(u(x_a(s),p) - u(x_b(s),p)\big)\underbrace{\big(u(x_a(s),\theta) - u(x_b(s),\theta)\big)}_{=\,u(x_a(s),\theta') - u(x_b(s),\theta')} + u(x_b(s),\theta)\\ &\quad + \int_{u(x_a(s),p)-u(x_b(s),p)}^{\infty} \nu dF(\nu)\\ &= v(p,\theta',x_a(s),x_b(s)) + u(x_b(s),\theta) - u(x_b(s),\theta').\end{aligned} \tag{1.22}$$

The voter's utility, neglecting the information cost, under action p_S and state (ω, δ), is

$$\sum_s \sigma_p(s|\omega,\sigma_f)v\big(p_S(s),\omega+\delta,x_a(s),x_b(s)\big).$$

Equations (1.21) and (1.22) imply that this voter's utility is the same under state (ω', δ)—up to a constant that does not interact with the action and is therefore immaterial:

$$\begin{aligned}
&\sum_s \sigma_p(s|\omega, \sigma_f) v\big(p_S(s), \omega + \delta, x_a(s), x_b(s)\big) \\
=&\sum_s \sigma_p(s|\omega, \sigma_f) \Big(v\big(p_S(s), \omega' + \delta, x_a(s), x_b(s)\big) - u(x_b(s), \theta') + u(x_b(s), \theta)\Big) \\
=&\sum_s \sigma_p(s|\omega', \sigma_f) \Big(v\big(p_S(s), \omega' + \delta, x_a(s), x_b(s)\big)\Big) + \sum_s \sigma_p(s|\omega', \sigma_f) \Big(u(x_b(s), \theta) - u(x_b(s), \theta')\Big)
\end{aligned}$$

This shows that states (ω, δ) and (ω', δ) are payoff-equivalent. By information monotonicity of mutual information (Amari, 2016; Caplin et al., 2022), any optimal signal structure σ does not distinguish between these states in the sense of (1.19). □

1.9 Appendix: Alternative Timing

Proof of Proposition 3

Proof. Using (1.11), the voter's maximization problem can be written, up to a constant, as a function of the acquired distribution ρ over posterior means. By Strassen's theorem, ρ is feasible if it is a mean-preserving contraction of the prior μ. The voter solves

$$\max_{\rho \in \Delta(\mathbb{R}^n)} \mathbb{E}_{\theta \sim \rho}\left[\max\left\{ (x_b - x_a)\left(\theta - \frac{x_a + x_b}{2}\right), -(x_b - x_a)\left(\theta - \frac{x_a + x_b}{2}\right) \right\} - \kappa\theta^2 \right]$$
$$\text{s.t. } \rho \leq_{\text{MPS}} \mu$$

We first solve the relaxed problem where we relax the constraint $\rho \leq_{\text{MPS}} \mu$ to $\mathbb{E}_{\theta \sim \rho}[\theta] = 0$. Through the coordinate change

$$\hat{\theta} = \theta - \frac{x_a + x_b}{2},$$

the voter's problem becomes

$$\max_{\rho \in \Delta(\mathbb{R}^n)} \mathbb{E}_{\hat{\theta} \sim \hat{\rho}}\Big[\overbrace{\max\left\{ (x_b - x_a)\hat{\theta}, -(x_b - x_a)\hat{\theta} \right\} - \kappa\hat{\theta}^2}^{v(\hat{\theta}) :=} \Big]$$
$$\text{s.t. } \mathbb{E}_{\hat{\theta} \sim \hat{\rho}}[\theta] = -\frac{x_a + x_b}{2}$$

Because the value function v is symmetric around 0 (see Figure 1.6), the solution is obtained simply by finding its maxima. The gradient of the value function is zero if $2\kappa\hat{\theta} = x_b - x_a$ or

$$\hat{\theta} = \frac{x_b - x_a}{2\kappa}.$$

Define

$$\theta_a := \frac{x_a + x_b}{2} - \frac{x_b - x_a}{2\kappa},$$
$$\theta_b := \frac{x_a + x_b}{2} + \frac{x_b - x_a}{2\kappa}.$$

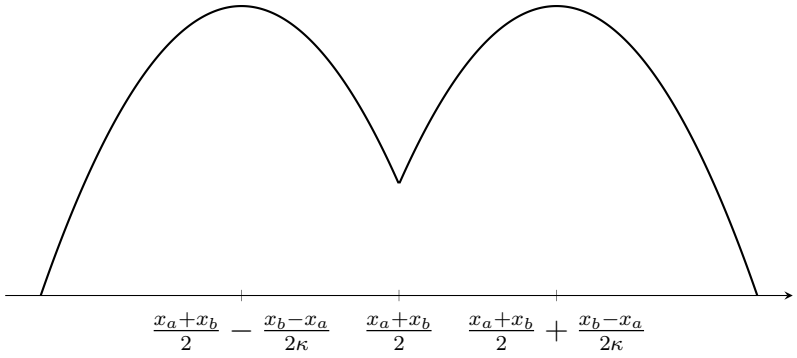


Figure 1.6: The value function $v(\theta)$

If the origin lies between θ_a and θ_b, then the voter acquires the posterior means θ_a and θ_b.

This is the solution if $\{\theta_a, \theta_b\}$ is feasible. If $\{\theta_a, \theta_b\}$ is not feasible, then a threshold signal $s(\theta) = \mathbb{1}(\theta > t)$ is optimal by the following. The distribution ρ over $\{\theta_a, \theta_b\}$ is the only solution to the voter learning problem which satisfies the first-order condition of optimality of the relaxed problem (see the Lagrangian Lemma in Caplin et al., 2022). Thus, the solution to the non-relaxed problem must have a binding mean-preserving contraction constraint. The mean-preserving contraction constraint $\rho \leq_{\text{MPS}} \mu$ can, by Strassen's theorem, be written as

$$\forall t \in \mathbb{R}: \quad \int_{-\infty}^{t} F_\rho(\theta)d\theta \leq \int_{-\infty}^{t} F_\mu(\theta)d\theta.$$

If ρ is binary and the constraint is binding at some t, then ρ is induced by the threshold signal with threshold t. Also, any threshold signal with threshold t induces a distribution ρ with binding mean-preserving contraction constraint. Note that we allow the threshold signal to have threshold $t = \pm\infty$, in which case the voter acquires no information and unconditionally votes for the party closer to the origin.

If a threshold signal is optimal, the threshold t is more extreme than $\frac{x_a+x_b}{2}$ by the following argument. It is easy to see that the *instrumental value of information* $b(t)$ is increasing as t

moves closer to $\frac{x_a+x_b}{2}$:

$$\begin{aligned} b(t) &:= \int_{-\infty}^{t} -(x_a-\theta)^2 f(\theta)d\theta + \int_{t}^{\infty} -(x_b-\theta)^2 f(\theta)d\theta \\ \Rightarrow b'(t) &= \left(-(x_a-t)^2 + (x_b-t)^2\right) f(t) \end{aligned}$$

Note that $b'(\frac{x_a+x_b}{2}) = 0$.

The *information cost* decreases as t becomes more extreme if the prior is log-concave. The information cost is

$$\begin{aligned} c(t) :=& \operatorname{Var}[\theta] - \left(F(t)\operatorname{Var}[\theta|\theta \le t] + (1-F(t))\operatorname{Var}[\theta|\theta > t]\right) \\ =& \operatorname{Var}[\theta] - \left(F(t)\min_E\left\{\frac{\int_{-\infty}^{t}(\theta-E)^2 dF(\theta)}{F(t)}\right\} + (1-F(t))\min_E\left\{\frac{\int_{t}^{\infty}(\theta-E)^2 dF(\theta)}{1-F(t)}\right\}\right) \\ =& \operatorname{Var}[\theta] - \left(\min_E\left\{\int_{-\infty}^{t}(\theta-E)^2 f(\theta)d\theta\right\} + \min_E\left\{\int_{t}^{\infty}(\theta-E)^2 f(\theta)d\theta\right\}\right) \end{aligned}$$

Using the envelope theorem, we obtain the derivative

$$\begin{aligned} c'(t) =& f(t)(t-\mathbb{E}[\theta|\theta>t])^2 - f(t)(t-\mathbb{E}[\theta|\theta\le t])^2 \\ =& f(t)\left(\mathbb{E}[\theta-t|\theta>t]^2 - \mathbb{E}[t-\theta|\theta\le t]^2\right). \end{aligned}$$

By Theorem 1.C.52 in Shaked and Shanthikumar (2007), the conditional expectation $\mathbb{E}[\theta - t|\theta > t]$ is decreasing in t for logconcave density of θ. By symmetry of f around 0, $\mathbb{E}[t-\theta|\theta \le t] = \mathbb{E}[\theta-(-t)|\theta > -t]$. Together, this delivers for $t > 0$,

$$0 < \mathbb{E}[\theta-t|\theta>t] < \mathbb{E}[\theta|\theta>0] < \mathbb{E}[\theta-(-t)|\theta>(-t)] = \mathbb{E}[t-\theta|\theta\le t].$$

Thus, $c'(t)$ is negative if $t > 0$.

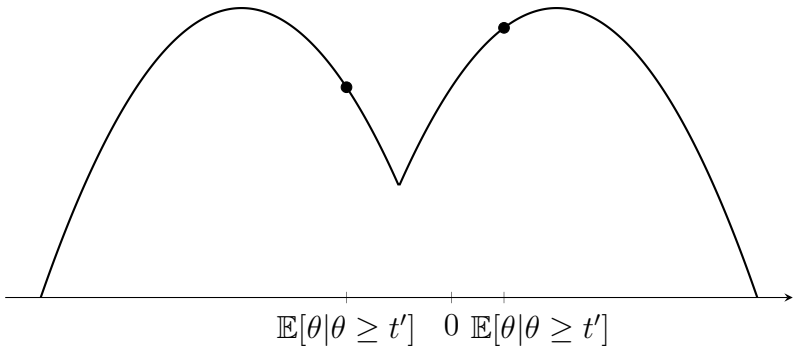


Figure 1.7: Value function and feasible posterior means under threshold signal t'

Note that for $\frac{x_a+x_b}{2} > 0$, $b'(t) > 0$ and $c'(t) < 0$ for all $t < \frac{x_a+x_b}{2}$, so $t = \frac{x_a+x_b}{2}$ dominates any $t < \frac{x_a+x_b}{2}$. Moreover, $b'(\frac{x_a+x_b}{2}) = 0$ and $c'(\frac{x_a+x_b}{2}) < 0$, so $t > \frac{x_a+x_b}{2}$ is optimal. Recall that the optimal distribution over posteriors is unique and thus the optimal t is unique.

Restrict t to $[\frac{x_a+x_b}{2}, \infty]$ ($t = \infty$ refers to always voting for party a). The utility $b(t) - \kappa c(t)$ is supermodular in (κ, t) by $c'(t) < 0$. Thus, the optimal t is weakly increasing as a function of κ. □

Proof of Corollary 2

Proof. The case (3) is immediate.

Case (1): By the law of iterated expectations, we have

$$P_a\left(\frac{x_a+x_b}{2} - \frac{x_b-x_a}{2\kappa}\right) + (1-P_a)\left(\frac{x_a+x_b}{2} + \frac{x_b-x_a}{2\kappa}\right) = \mathbb{E}[\theta] = 0 \Rightarrow P_a = \frac{1}{2} + \frac{\kappa}{2}\frac{x_a+x_b}{x_b-x_a}.$$

Case (2): That $P_a = F_\mu(t)$ is immediate from case (2) of Proposition 1. The second part of the statement follows from the following lemma.

Lemma 6. *Let x_a, x_b be such that $\{\theta_a, \theta_b\}$ is not feasible, $|x_a| > |x_b|$, and $x_a < x_b$. The expected vote share P_a of party a satisfies $P_a < \frac{1}{2} + \frac{\kappa}{2}\frac{x_a+x_b}{x_b-x_a}$.*

Proof. First, t is such that the induced posterior means $\mathbb{E}[\theta|\theta < t]$ and $\mathbb{E}[\theta|\theta > t]$ are in between the maxima $\{\theta_a, \theta_b\}$ of the value function,

$$\theta_a \;\leq\; \mathbb{E}[\theta|\theta < t] \;<\; \mathbb{E}[\theta|\theta > t] \;\leq\; \theta_b, \tag{1.23}$$

or the voter acquires no information. The reason is the following. If the posterior means were outside the maxima of the value function, then the maxima would be a garbling of the threshold signal, so the maxima would be feasible. Both posterior means need to be on different sides of $\frac{x_a+x_b}{2}$, otherwise the agent takes the same action regardless of the information they acquired and no information would be better. Finally, it cannot be the case that

$$\mathbb{E}[\theta|\theta < t] \;<\; \theta_a \;\leq\; \frac{x_a + x_b}{2} \;\leq\; \mathbb{E}[\theta|\theta > t] \;\leq\; \theta_b,$$

because increasing the threshold t marginally would increase both posterior means $\mathbb{E}[\theta|\theta < t]$ and $\mathbb{E}[\theta|\theta > t]$, which would result in a higher value because the derivative of the value function is positive at $\mathbb{E}[\theta|\theta < t]$ and non-negative at $\mathbb{E}[\theta|\theta > t]$, see also Figure 1.7. Analogously, the posterior mean $\mathbb{E}[\theta|\theta > t]$ cannot be to the right of θ_b. Thus, the posterior means must satisfy the ordering (1.23) or $t = \pm\infty$.

Second, consider the threshold t' that delivers the same vote shares as $\hat{P}_a = \frac{1}{2} + \frac{\kappa}{2}\frac{x_a+x_b}{x_b-x_a}$, $t' = F_\mu^{-1}(P_a)$. By $|x_a| > |x_b|$, $\hat{P}_a < 1/2$, so by symmetry of μ, $t' < 0$. By the law of iterated expectations

$$0 = \hat{P}_a\mathbb{E}[\theta|\theta < t'] + (1 - \hat{P}_a)\mathbb{E}[\theta|\theta \geq t'] = \hat{P}_a\theta_a + (1 - \hat{P}_a)\theta_b.$$

Hence, there exists an $\alpha > 0$ such that the induced posterior means satisfy

$$(\mathbb{E}[\theta|\theta < t'], \mathbb{E}[\theta|\theta \geq t']) = (\alpha\theta_a, \alpha\theta_b).$$

By $\{\theta_a, \theta_b\}$ not being feasible, $\alpha < 1$. We show that the optimal t satisfies $t < t'$, that is, the binding mean-preserving contraction constraint benefits the moderate party. For that we use the following Lemma.

Lemma 7 (Jewitt's Lemma)**.** $\Delta(t) := \mathbb{E}[\theta|\theta > t] - \mathbb{E}[\theta|\theta < t]$ *is weakly decreasing in* t *for* $t < 0$ *and weakly increasing in* t *for* $t > 0$*.*

Proof. Jewitt (2004) shows that $\Delta(t)$ is quasiconvex if the density of θ is quasiconcave. Because θ has a symmetric density around 0, $\Delta(t)$ is weakly decreasing for $t < 0$ and weakly increasing for $t > 0$. □

We prove $t < t'$ by showing that increasing t beyond t' lowers the voter's utility. Intuitively, increasing t has two effects on the voter's utility: it shifts the posterior means and it shifts their probabilities. We show that both effects have a negative effect on the voter's utility (consulting Figure 1.7 may help visualize the following steps). First, note that we have for $t > t'$,

$$\frac{-v'(\mathbb{E}[\theta|\theta < t])}{v'(\mathbb{E}[\theta|\theta \geq t]} \geq \frac{-v'(\mathbb{E}[\theta|\theta < t'])}{v'(\mathbb{E}[\theta|\theta \geq t']} = \frac{1 - F(t')}{F(t')}.$$

A $t > 0$ is clearly suboptimal because $-t$ would deliver an equal cost of information but higher higher instrumental value of information. Under $t < 0$, Jewitt's Lemma implies

$$\frac{d\Delta(t)}{dt} < 0 \Rightarrow \frac{d\mathbb{E}[\theta|\theta < t]}{dt} > \frac{d\mathbb{E}[\theta|\theta \geq t]}{dt}.$$

Now, for any t with $t' \leq t \leq 0$, as long as the ordering requirement (1.23) holds,

$$\begin{aligned}
&\frac{d}{dt}\Big(F(t)v(\mathbb{E}[\theta|\theta < t]) + (1 - F(t))v(\mathbb{E}[\theta|\theta \geq t])\Big) \\
=&F(t)v'(\mathbb{E}[\theta|\theta < t])\frac{d\mathbb{E}[\theta|\theta < t]}{dt} + (1 - F(t))v'(\mathbb{E}[\theta|\theta \geq t])\frac{d\mathbb{E}[\theta|\theta \geq t]}{dt} \\
&\qquad + f(t)\big(\underbrace{v(\mathbb{E}[\theta|\theta < t]) - v(\mathbb{E}[\theta|\theta \geq t])}_{<0}\big) \\
<&\Big(F(t)v'(\mathbb{E}[\theta|\theta < t]) + (1 - F(t))v'(\mathbb{E}[\theta|\theta \geq t])\Big)\frac{d\mathbb{E}[\theta|\theta \geq t]}{dt} \\
&\qquad + F(t)\underbrace{v'(\mathbb{E}[\theta|\theta < t])}_{<0}\underbrace{\left(\frac{d\mathbb{E}[\theta|\theta < t]}{dt} - \frac{d\mathbb{E}[\theta|\theta \geq t]}{dt}\right)}_{>0} \\
<&\Big(\underbrace{F(t)v'(\mathbb{E}[\theta|\theta < t]) + (1 - F(t))v'(\mathbb{E}[\theta|\theta \geq t])}_{\leq 0}\Big)\underbrace{\frac{d\mathbb{E}[\theta|\theta \geq t]}{dt}}_{>0} \leq 0
\end{aligned}$$

Therefore, $t < t'$ is optimal. □

This concludes the proof of Corollary 5. □

Proof of Corollary 3

Proof. We go through the three cases of Corollary 2 by increasing size of κ. In each case, the expected vote share is biased toward the moderate party and this bias increases in κ. Increasing party polarization $|x_b - x_a|$ while holding $\frac{x_a+x_b}{2}$ constant, has the same effect on the voter's objective as decreasing κ (and, at the same time, multiplying the objective by an irrelevant constant).

By Corollary 2, if (θ_a, θ_b) are not feasible, then the optimal signal is a threshold signal with a threshold more extreme than the party midpoint, thus there is a bias toward the moderate party. The threshold becomes more extreme as κ increases, so the bias increases.

Let κ be such that $\{\theta_a, \theta_b\} = \{\frac{x_a+x_b}{2} - \frac{\kappa}{2}\frac{x_a+x_b}{x_b-x_a}, \frac{x_a+x_b}{2} + \frac{\kappa}{2}\frac{x_a+x_b}{x_b-x_a}\}$, is just feasible. Then, these posterior means are induced by a threshold signal, which by the above, is biased toward the moderate party. For a higher κ, the posterior means $\{\theta_a, \theta_b\}$ move symmetrically closer to the party midpoint $\frac{x_a+x_b}{2}$. This increases the weight on the more likely posterior mean, so it increases the bias.

If κ is large enough, then voters acquire no information and just vote for the ex-ante more attractive party, which is the moderate party. □

Lemmas The following lemma describes properties of any equilibrium and partially characterize parties' best response. We use these lemmas in the proofs of Theorem 5 and Theorem 6 below.

Lemma 8. *If $x_a^* < x_b^*$, then $x_a \leq x_b$ holds in any equilibrium with positive expected vote shares. If additionally $x_a^* < 0 < x_b^*$, then $x_a \leq x_b^*$ and $x_b \geq x_a^*$.*

Proof. As argued above, because the parties are policy motivated, a party never finds it optimal to choose a policy x_j is further away from x_j^* than the other parties' policy x_{-j} and obtain a positive expected vote share. This implies $x_a \leq x_b$. To show this, distinguish

between two cases. First, if $x_a \leq x_b^*$, then $x_b \geq x_a$. Second, if $x_a > x_b^*$, then $x_b \leq x_a$. If $x_b < x_a$, then x_a is further away from x_a^* than x_b, which is suboptimal.

To prove the second statement note that in the first case, $x_a \leq x_b^*$. In the second case, $x_b = x_a$ cannot be a best response of party b if $0 < x_b^*$, because $x_b = x_a - \varepsilon$ for ε small enough would gain all votes and deliver a better policy. Thus, $x_a > x_b^*$ is impossible if additionally $0 < x_b^*$.

Analogously $x_b \geq x_a^*$ if additionally $x_a^* < 0$. □

Lemma 9. *Let (x_a, x_b) be an equilibrium with positive expected vote shares for both parties and $x_a^* < 0 < x_b^*$. Then, $x_a \geq 2x_a^*$ and $x_b \leq 2x_b^*$.*

Proof. As shown by Lemma 8, in equilibrium $x_b \geq x_a^*$. Because x_a is closer to x_a^* in an equilibrium with positive vote shares than is x_b, this implies that if $x_b < 0$, $x_a \geq 2x_a^*$.

If $0 < x_b$ and $x_a < 2x_a^*$ is more extreme than x_b, that is $|x_a| \geq |x_b^*|$, then $P_a(x_a, x_b) \leq \frac{1}{2}$. Then, $x_a = 0$ delivers a greater expected vote share for party a (namely a vote share greater than 1/2) and a greater policy utility when winning, so x_a is suboptimal.

If $x_b > 0$ and $x_a < x_a^*$ and feasibility is not binding at (x_a, x_b), then $P_a(x_a, x_b) = \hat{P}_a(x_a, x_b)$ is increasing in x_a (this can be obtained by taking the derivative of $\hat{P}_a(x_a, x_b)$ in x_a). Thus, increasing x_a would both increase the vote share and the policy utility when wining, so x_a is suboptimal. So, in equilibrium, $x_a \geq x_a^* > 2x_a^*$.

Consider, finally, that case that $x_b > 0$, feasibility is binding at $x_a < x_a^*$, and x_a is less extreme than x_b, that is, $|x_a| < |x_b|$. If x_a is less extreme than x_b by $|x_a^*| \geq x_b$, then either party a obtains expected votes share 1 (which we have ruled out) or a threshold signal is optimal with $t > \frac{x_a + x_b}{2} > x_a$ (that is, the signal is biased toward the moderate party a) by Proposition 3. One can easily show that the utility $b(t) - \kappa c(t)$ is supermodular in x_a and t, when t is restricted to $t > x_a$, so t increases in x_a. Again, both the expected vote share and the policy utility when winning increase when raising x_a, so $x_a < x_a^*$ is suboptimal.

The proof of $x_b \leq x_b^*$ is analogous. □

Next, we provide a lemma that partially characterizes parties best responses. We need a few definitions. Define

$$\theta_a(x_a, x_b) := \frac{x_a + x_b}{2} - \frac{x_b - x_a}{2\kappa},$$
$$\theta_b(x_a, x_b) := \frac{x_a + x_b}{2} + \frac{x_b - x_a}{2\kappa}.$$

Definition 2. *We say* ***feasibility is not binding at*** $(\mathbf{x_a}, \mathbf{x_b})$ *if the posterior means* $\{\theta_a(x_a, x_b), \theta_b(x_a, x_b)\}$ *are feasible. Otherwise, we say feasibility is binding at* (x_a, x_b)*.*

Further, define

$$\hat{P}_a(x_a, x_b) := \frac{1}{2} + \frac{\kappa}{2}\frac{x_b + x_a}{x_b - x_a},$$
$$\hat{U}_a(x_a, x_b) := \hat{P}_a(x_a, x_b)u(x_a, x_a^*) \ + \ \big(1 - \hat{P}_a(x_a, x_b)\big)u(x_b, x_a^*).$$

We call $\hat{P}_a$ and $\hat{U}_a$ the *pseudo vote share* and *pseudo utility*, respectively, because they are equal to the expected vote share and expected utility of party a only if feasibility is not binding at (x_a, x_b), by Corollary 2.

Lemma 10 (Best Response)**.** *Let* $x_a^* < x_b^*$*,* $\kappa < 1$*, and* $x_b \in \mathbb{R}$*. Define*

$$\hat{x}_a = x_a^* + \frac{\kappa}{1 - \kappa}x_b.$$

Necessity: *If* x_a *is the best response of party* a *to* x_b *and feasibility is not binding at* (x_a, x_b)*, then* $x_a = \hat{x}_a$*.*

Sufficiency: *If feasibility is not binding at* $(\hat{x}_a, x_b)$ *and* $|\hat{x}_a| \geq |x_b|$*, then* $\hat{x}_a$ *is the unique best response of party* a *to* x_b*.*

Analogous statements hold when exchanging a *and* b *(maintaining* $x_a^* < x_b^*$*).*

Proof. **Necessity:** If feasibility is not binding at (x_a, x_b), then party a's objective is

$$\begin{aligned}\hat{U}_a(x_a, x_b) &= \hat{P}_a(x_a, x_b)u(x_a, x_a^*) \ + \ \big(1 - \hat{P}_a(x_a, x_b)\big)u(x_b, x_a^*) \\ &= \big(u(x_a, x_a^*) - u(x_b, x_a^*)\big)\hat{P}_a(x_a, x_b) + u(x_b, x_a^*) \qquad (1.24) \\ &= 2\left(x_a^* - \frac{x_a + x_b}{2}\right)(x_a - x_b)\left(\frac{1}{2} + \frac{\kappa}{2}\frac{x_a + x_b}{x_b - x_a}\right) + u(x_b, x_a^*) \\ &= \left(\frac{x_a + x_b}{2} - x_a^*\right)\big(x_b(1+\kappa) - x_a(1-\kappa)\big) + u(x_b, x_a^*). \qquad (1.25)\end{aligned}$$

The last term is a constant and can be ignored. Since the domain of x_a is unbounded and the objective any x_a further from x_a^* than x_b is dominated, the best response x_a to x_b must satisfy the first-order condition $\frac{dU_a(x_a,x_b)}{dx_a} = 0$. Simple derivation shows that the unique solution to the first-order condition is

$$\hat{x}_a = x_a^* + \frac{\kappa}{1-\kappa}x_b.$$

Sufficiency: To show that $\hat{x}_a$ is the best response if feasibility is not binding at $(\hat{x}_a, x_b)$ and $|\hat{x}_a| \geq |x_b|$, we rule out step-by-step that any other x_a could deliver a higher utility $U_a(x_a, x_b)$ than $\hat{x}_a$.

First, we rule out x_a such that feasibility is not binding at (x_a, x_b). If feasibility is not binding at (x_a, x_b), then party a's objective is given by $\hat{U}_a(x_a, x_b)$, which is negative quadratic in x_a and maximized at $\hat{x}_a$. Thus, $\hat{x}_a$ is the best response among all such x_a.

Inserting $\hat{x}_a$ into (1.25), we obtain that the resulting utility is

$$U_a(\hat{x}_a, x_b) = \frac{1-\kappa}{2}\Big(\frac{x_b}{1-\kappa} - x_a^*\Big)^2 + u(x_b, x_a^*) > u(x_b, x_a^*), \qquad (1.26)$$

which we use below.

Second, we rule out x_a that are at least as far away from x_a^* than is x_b. Such x_a deliver utility $U_a(x_a, x_b) \leq u(x_b, x_a^*) < U_a(\hat{x}_a, x_b)$, using (1.26). Therefore, x_a delivers a lower utility than $\hat{x}_a$. Intuitively, because party a cares about the implemented policy, it never makes sense to propose a policy that is worse than its opponent's policy.

Third, we rule out x_a that deliver an expected vote share $P_a(x_a, x_b)$ of 0 or 1. If $P_a(x_a, x_b) = 0$, then the resulting utility is $U_a(x_a, x_b) = u(x_b, x_a^*) < U_a(\hat{x}_a, x_b)$, using (1.26). Therefore, x_a delivers a lower utility than $\hat{x}_a$. Consider $P_a(x_a, x_b) = 1$, which can be the case if 0 is not between $\theta_a(x_a, x_b)$ and $\theta_b(x_a, x_b)$ or if $\{\theta_a(x_a, x_b), \theta_b(x_a, x_b)\}$ is not feasible, by Corollary 2. We show that $1 < \hat{P}_a(x_a, x_b)$. Assume, for the sake of contradiction, that x_a delivers a higher utility than $\hat{x}_a$. x_a can only deliver an expected vote share of 1 if $|x_a| < |x_b|$. By assumption we have $|x_b| < |\hat{x}_a|$. Because $\hat{x}_a$ and x_a deliver greater utility U_a than $u(x_b, x_a^*)$, $\hat{x}_a$ and x_a are closer to x_a^* than is x_b, by (1.26). This leaves only the case that $\hat{x}_a$ and x_b are on different sides of the origin and x_a is between $\hat{x}_a$ and x_b. This implies that if x_a delivers an expected vote share of 1, then $\hat{P}_a > 1$. The reason is that if 0 is between θ_a and θ_b, then θ_a and θ_b are feasible. Finally, if $P_a(x_a, x_b) = 1 < \hat{P}_a(x_a, x_b)$, then party a's utility $U_a(x_a, x_b)$ is strictly less than $\hat{U}_a(x_a, x_b)$, because, by the above, x_a is strictly closer to x_a^* than x_b. Further $\hat{U}_a(x_a, x_b)$ is strictly less than $\max_{x'_a} \hat{U}_a(x'_a, x_b) = \hat{U}_a(\hat{x}_a, x_b) = U_a(\hat{x}_a, x_b)$. Therefore, x_a delivers a lower utility than $\hat{x}_a$.

Finally, consider the remaining possibility that feasibility is binding at (x_a, x_b), x_a is closer to x_a^* than is x_b, and both parties obtain positive expected vote shares. We distinguish between three subcases, depending on which of the three policies $\{x_a, \hat{x}_a, x_b\}$ is between the other two.

In the first subcase, x_a is between $\hat{x}_a$ and x_b. One can easily verify that $\theta_a(x_a, x_b)$ and $\theta_b(x_a, x_b)$ are between $\theta_a(\hat{x}_a, x_b)$ and $\theta_b(\hat{x}_a, x_b)$. Because $\{\theta_a(\hat{x}_a, x_b), \theta_b(\hat{x}_a, x_b)\}$ is feasible, $\{(\theta_a(x_a, x_b), \theta_b(x_a, x_b)\}$ is feasible (if 0 is between $\theta_a(x_a, x_b)$ and $\theta_b(x_a, x_b)$) or one party obtains all the votes (otherwise). We have already ruled out both of possibilities above.

In the second subcase, x_b is between $\hat{x}_a$ and x_a. Because $\hat{x}_a$ delivers higher utility than $x_a = x_b$, $\hat{x}_a$ is closer to x_a^* than is x_b. Thus, this subcase implies that x_a is further away from x_a^* than is x_b, which we have ruled out above.

In the third and final subcase, $\hat{x}_a$ is between x_a and x_b. By $|\hat{x}_a| \geq |x_b|$, this implies that $|x_a| \geq |\hat{x}_a| \geq |x_b|$. Thus, if feasibility is binding at (x_a, x_b), then, by Corollary 2,

$P_a(x_a, x_b) < \hat{P}_a(x_a, x_b)$, so $U_a(x_a, x_b)$ is strictly less than $\hat{U}_a(x_a, x_b)$, by x_a being closer to x_a^* than is x_b. Further, $\hat{U}_a(x_a, x_b)$ is strictly less than $\hat{U}_a(\hat{x}_a, x_b) = U_a(\hat{x}_a, x_b)$. Therefore, x_a delivers a lower utility than $\hat{x}_a$.

By symmetry, analogous statements hold when exchanging a and b. □

Proof of Theorem 5

Proof. First, we verify that the candidate equilibria are in fact equilibria. Second, we show uniqueness.

Solving for a solution to the equations $x_a = x_a^* + \frac{\kappa}{1-\kappa} x_b$ and $x_b = x_b^* + \frac{\kappa}{1-\kappa} x_a$, we obtain

$$(x_a, x_b) = (1 - \kappa)(x_a^*, x_b^*) \tag{1.27}$$

if $\kappa \in (0, 1)$ and $\kappa \neq 1/2$. By $\kappa > \underline{\kappa}$, the induced distribution over posterior means is feasible, so feasibility is not binding at (x_a, x_b). By $|x_a| = |x_b|$ and the sufficiency part of Lemma 10, this constitutes mutual best responses and hence an equilibrium.

Uniqueness By the necessity part of Lemma 10 and the solution to (1.27) being unique, there cannot be another equilibrium (x_a, x_b) where feasibility is not binding if $\kappa \neq 1/2$. Also, we have ruled out (by our refinement) equilibria where one party obtains zero votes. Thus, *it remains to show that there are no equilibria* (x_a, x_b) *where feasibility is binding and both parties obtain positive expected vote shares.* Denote the equilibrium positions by $(\tilde{x}_a, \tilde{x}_b) = (1 - k)(x_a^*, x_b^*)$.

First, we show that it cannot be the case in equilibrium that both parties propose policies further away from the origin than $\tilde{x}_a$ and $\tilde{x}_b$, respectively, that is, $x_a < \tilde{x}_a < \tilde{x}_b < x_b$. Suppose without loss that x_b is weakly closer to 0 than is x_a. Then, by feasibility being binding and Corollary 2, $P_a(x_a, x_b) \leq \hat{P}_a(x_a, x_b)$. Because in any equilibrium with positive expected vote shares, x_a must be weakly closer to x_a^* than is x_b, this implies that $U_a(x_a, x_b) \leq \hat{U}_a(x_a, x_b)$. By $x_a \leq -x_b < \tilde{x}_a$, $-x_b$ is closer to $\hat{x}_a = x_a^* + \frac{\kappa}{1-\kappa} x_b > \tilde{x}_a$ than is x_a. So, $\hat{U}_a(x_a, x_b) < \hat{U}_a(-x_b, x_b)$.

Finally, $\hat{U}_a(-x_b, x_b) = U_a(-x_b, x_b)$ because $P_a(-x_b, x_b) = \hat{P}_a(-x_b, x_b) = 1/2$. Together, $U_a(x_a, x_b) < U_a(-x_b, x_b)$ and (x_a, x_b) is not an equilibrium.

Second, we show that there cannot be an equilibrium where $x_a > \tilde{x}_a$ or $x_b < \tilde{x}_b$. To show this, we distinguish between $\kappa > 1/2$ and $\kappa < 1/2$, which affects whether best response of a party, according to Lemma 10, is responsive to other party's position less or more than 1-to-1 by $\frac{\kappa}{1-\kappa} > 1$ and $\frac{\kappa}{1-\kappa} < 1$, respectively.

Case 1: $1/2 < \kappa < 1$.

For the sake of contradiction, suppose $x_b < \tilde{x}_b$. The proof for $x_a > \tilde{x}_a$ is analogous.

If $(1-\kappa)x_a^* \leq x_b \leq \tilde{x}_b$, then $\hat{x}_a = x_a^* + \frac{\kappa}{1-\kappa}x_b$ is the best response of party a by the following. First, it can be shown that under these conditions, $\hat{P}_a(\hat{x}_a, x_b) > 0$, so $\theta_a(\hat{x}_a, x_b) \leq 0 \leq \theta_b(\hat{x}_a, x_b)$. Further, the distance $|x_b - \hat{x}_a|$ increases as x_b decreases by $\frac{\kappa}{1-\kappa} > 1$. The maximal distance between x_b and $\hat{x}_a$ is thereby reached at $x_b = (1-\kappa)x_a^*$, in which case $\hat{x}_a = (1+\kappa)x_a^*$ and $|x_b - \hat{x}_a| = 2\kappa x_a^*$. By

$$2x_a^* \leq \Delta \Rightarrow 2\kappa x_a^* \leq \kappa\Delta,$$

and Lemma 12, feasibility is not binding at $(\hat{x}_a, x_b)$. Further, by $|\hat{x}_a| \geq |x_b|$ and Lemma 10, $\hat{x}_a$ is the best response of party a to x_b. On the other hand, if $x_b < (1-\kappa)x_a^*$, then the best response x_a must be closer to x_a^* than x_b is, so $|x_b - x_a| < 2\kappa x_a^*$ as well. By the above, feasibility is not binding at (x_a, x_b) again. Thus, feasibility is not binding in the equilibrium, which we have ruled out above.

Case 2: $0 < \kappa < 1/2$.

Suppose party a chose a position x_a in the direction of the origin from $\tilde{x}_a$, $\tilde{x}_a < x_a$. Define $\Delta := x_a - \tilde{x}_a$ and $\hat{x}_b := x_b^* + \frac{\kappa}{1-\kappa}x_a$. Then, by $\frac{\kappa}{1-\kappa} < 1$, $\hat{x}_b - \tilde{x}_b < \Delta$. By $x_a < x_b^*$ (Lemma 8), we have $x_a \leq \hat{x}_b$, so together, $|\hat{x}_b - x_a| < |\tilde{x}_b - \tilde{x}_a|$. By Jewitt's Lemma, feasibility is not binding at $(x_a, \hat{x}_b)$ or party a obtains all the votes. In the former case, $\hat{x}_b$ is the best response of party b by $|\hat{x}_b| > |x_a|$ and the sufficiency part of Lemma 10, so feasibility would

not be binding at the equilibrium. In the latter case, the best response of party b would give party b no votes. We have ruled out both possibilities above. Analogously, we can rule out $x_b < \tilde{x}_b$ in equilibrium.

Lemma 11. *If $\kappa \geq 1$, then the unique equilibrium is $(x_a, x_b) = (0, 0)$.*

Proof. Under $\kappa \geq 1$, the implemented policy is zero if one party chooses policy 0. This follows from Corollary 2 and the fact that $\theta_a(x_a, x_b)$ and $\theta_b(x_a, x_b)$ are weakly between x_a and x_b for $\kappa \geq 1$. Therefore, $(x_a, x_b) = (0, 0)$ is an equilibrium.

By Lemma 8, $x_a \leq x_b$ in any equilibrium.

Suppose (x_a, x_b) is an equilibrium. The average implemented policy is weakly positive or weakly negative. Suppose without loss that it is weakly positive. Then party a can obtain policy 0 for certain by choosing $x_a = 0$. This weakly improves the utility of party a because the policy outcome is weakly better in expectation and has no variance. Thus, if the policy outcome under (x_a, x_b) is not 0 for certain, then this is a profitable deviation. Finally, there can be no equilibrium in which only one party chooses policy 0, because then this party could move its policy x_j closer to x_j^* by ε small enough, while still obtaining expected vote share 1, and improve its utility. □

This concludes the proof. □

Proof of Theorem 6

Proof. First, we verify that the above constitutes an equilibrium. Then, we show uniqueness.

Equilibrium verification If $\kappa \geq 1$, then the unique equilibrium is $(x_a, x_b) = (0, 0)$ by Lemma 11.

Suppose $\kappa < 1$. Using the mutual best responses

$$x_a = x_a^* + \frac{\kappa}{1-\kappa} x_a$$
$$x_b = x_b^* + \frac{\kappa}{1-\kappa} x_b,$$

one can solve for the unique solution

$$\tilde{x}_a = \frac{1-\kappa}{1-2\kappa}\Big((1-\kappa)x_a^* + \kappa x_b^*\Big),$$
$$\tilde{x}_b = \frac{1-\kappa}{1-2\kappa}\Big((1-\kappa)x_b^* + \kappa x_a^*\Big).$$

if $\kappa \neq 1/2$.

Part 1: We verify first that feasibility is not binding at $(\tilde{x}_a, \tilde{x}_b)$.

Under $\kappa > 1/2$, $\tilde{x}_a$ is more extreme than $\tilde{x}_b$ by $|x_b^*| > |x_a^*|$. The posterior means $\theta_a(\tilde{x}_a, \tilde{x}_b)$ and $\theta_b(\tilde{x}_a, \tilde{x}_b)$ are on opposite sides of the origin if

$$0 < \theta_b(\tilde{x}_a, \tilde{x}_b) = \frac{\tilde{x}_a + \tilde{x}_b}{2} + \frac{\tilde{x}_b - \tilde{x}_a}{2\kappa} \quad \Rightarrow \quad \kappa > \frac{x_b^* - x_a^*}{x_b^* - 3x_a^*} > \frac{1}{2}.$$

Under $\kappa < 1/2$, $\tilde{x}_b$ is more extreme than $\tilde{x}_a$ by $|x_b^*| > |x_a^*|$. The posterior means $\theta_a(\tilde{x}_a, \tilde{x}_b)$ and $\theta_b(\tilde{x}_a, \tilde{x}_b)$ are on opposite sides of the origin if

$$0 > \theta_a(\tilde{x}_a, \tilde{x}_b) = \frac{\tilde{x}_a + \tilde{x}_b}{2} - \frac{\tilde{x}_b - \tilde{x}_a}{2\kappa} \quad \Rightarrow \quad \kappa < \frac{x_b^* - x_a^*}{3x_b^* - x_a^*} < \frac{1}{2}.$$

It remains to show that the prior is wide enough, so $\{\theta_a, \theta_b\}$ are feasible.

Define

$$\Delta := \mathbb{E}[\theta|\theta > 0] - \mathbb{E}[\theta|\theta < 0].$$

Lemma 12. *If $|x_b - x_a| \leq \kappa\Delta$ and $\theta_a(x_a, x_b)$ and $\theta_b(x_a, x_b)$ are on different sides of the origin, then feasibility is not binding at (x_a, x_b).*

Proof. By Jewitt's Lemma, if θ_a and θ_b are on different sides of the origin and $|\theta_b - \theta_a| < \Delta$, then $\{\theta_a, \theta_b\}$ is feasible.

$$|\theta_b - \theta_a| = \frac{|x_b - x_a|}{\kappa} < \Delta \Leftrightarrow |x_b - x_a| < \Delta.$$

□

We have

$$\tilde{x}_b - \tilde{x}_a = (1-\kappa)(x_b^* - x_a^*).$$

Thus, as long as

$$(1-\kappa)(x_b^* - x_a^*) < \kappa\Delta \Leftrightarrow \kappa > \kappa^* = \frac{\frac{x_b^* - x_a^*}{\Delta}}{1 + \frac{x_b^* - x_a^*}{\Delta}},$$

feasibility is not binding at $(\tilde{x}_a, \tilde{x}_b)$ if $\theta_a(\tilde{x}_a, \tilde{x}_b) \leq 0 \leq \theta_b(\tilde{x}_a, \tilde{x}_b)$.

Part 2: It remains to show that $\tilde{x}_a$ and $\tilde{x}_b$ are mutual best responses when $\kappa \in (\kappa^*, \frac{x_b^* - x_a^*}{3x_b^* - x_a^*}) \cup (\frac{x_b^* - x_a^*}{x_b^* - 3x_a^*}, 1)$.

For mutual best responses, by Lemma 10, the party with the more extreme policy platform (party a if $\kappa > 1/2$ and party b if $\kappa < 1/2$) is playing a best response. To show that the party with the less extreme policy platform, say a, is playing a best response, it is sufficient that feasibility is not binding at $(-x_b, x_b)$, by the following lemma.

Lemma 13. *Let $\hat{x}_a = x_a^* + \frac{\kappa}{1-\kappa} x_b$. If $\hat{x}_a < x_b$ and feasibility is not binding at $(\hat{x}_a, x_b)$ nor $(-x_b, x_b)$, then $\hat{x}_a$ is a best response to x_b.*

Proof. We know that $\hat{x}_a$ is a best response among the x_a such that feasibility is not binding at (x_a, x_b). We rule out step-by-step other x_a.

Any $x_a \geq x_b$ is suboptimal because it is further from x_a^* than x_b and $U_a(\hat{x}_a, x_b) > u(x_b, x_a^*)$ by (1.26).

Any x_a with $\hat{x}_a < x_a < x_b$ is suboptimal by the following. The posterior means $\theta_a(x_a, x_b)$ and $\theta_b(x_a, x_b)$ are between $\theta_a(\hat{x}_a, x_b)$ and $\theta_b(\hat{x}_a, x_b)$. So, either feasibility is not binding, which we have ruled out above, or party a obtains an expected vote share of 0 or 1. An expected vote share of 0 delivers utility $u(x_b, x_a^*)$, which is smaller than $U_a(\hat{x}_a, x_b)$ and therefore suboptimal. If the expected vote share is 1, then $\hat{P}_a(x_a, x_b) \geq 1$ (by the posterior means (θ_a, θ_b) being feasible if on different sides of 0), so $U_a(x_a, x_b) \leq \hat{U}_a(x_a, x_b) < \hat{U}_a(\hat{x}_a, x_b) = U_a(x_a, x_b)$, so x_a is suboptimal.

Any x_a with $x_a < \hat{x}_a$ and $|x_a| < |x_b|$ has non-binding feasibility at (x_a, x_b), and is ruled out above, by the following. If $x_b < 0$, then $x_a < \hat{x}_a < x_b < 0$ implies that $|x_a| > |x_b|$, a contradiction. If, otherwise, $x_b > 0$, then, because feasibility is not binding at $(-x_b, x_b)$ and $(\hat{x}_a, x_b)$ has non-binding feasibility, feasibility must be non-binding at (x_a, x_b).

The remaining x_a satisfy $x_a < \hat{x}_a$, feasibility is binding at (x_a, x_b), and $|x_a| > |x_b|$. Then, by Corollary 2, either $P_a(x_a, x_b) = 0$, which we have ruled out above, or $P_a(x_a, x_b) < \hat{P}_a(x_a, x_b)$. If x_a is further away from x_a^* than x_b is, then x_a is clearly suboptimal. So, x_a is closer to x_a^* than is x_b, so $U_a(x_a, x_b) < \hat{U}_a(x_a, x_b) < \hat{U}_a(\hat{x}_a, x_b) = U_a(\hat{x}_a, x_b)$, so x_a is suboptimal. □

This implies the following lemma.

Lemma 14. *Let $\hat{x}_a = x_a^* + \frac{\kappa}{1-\kappa} x_b$. If (x_b, κ) and $(\hat{x}_a, \kappa)$ lie in the cone*

$$\mathscr{C} = \Big\{(x, \kappa) \;\Big|\; |x| \leq \kappa \mathbb{E}[\theta | \theta > 0]\Big\},$$

and $\theta_a(\hat{x}_a, x_b) \leq 0 \leq \theta_b(\hat{x}_a, x_b)$, then $\hat{x}_a$ is the best response to x_b. An analogous statement holds when replacing a and b.

Proof. Feasibility is not binding at $(-x_b, x_b)$ because $(-x_b, \kappa)$ lies in the cone $\mathscr{C}$ if (x_b, κ) does and $\theta_a(-x_b, x_b) \leq 0 \leq \theta_b(-x_b, x_b)$. If $(\hat{x}_a, \kappa)$ also lies in the cone, then $|x_b - x_a| \leq \kappa\Delta$. If $\theta_a(\hat{x}_a, x_b) \leq 0 \leq \theta_b(\hat{x}_a, x_b)$, then feasibility is not binding at $(\hat{x}_a, x_b)$ by Lemma 12. Then, by Lemma 13, $\hat{x}_a$ is a best response to x_b. □

As long as

$$\max\big\{|x_a|, |x_b|\big\} \leq \kappa \mathbb{E}[\theta | \theta > 0],$$

both (x_a, κ) and (x_b, κ) lie in the cone $\mathscr{C}$. If, additionally,

$$\kappa \in \left(0, \frac{x_b^* - x_a^*}{3x_b^* - x_a^*}\right) \cup \left(\frac{x_b^* - x_a^*}{x_b^* - 3x_a^*}, 1\right),$$

then $\theta_a(\hat{x}_a, x_b) \leq 0 \leq \theta_b(\hat{x}_a, x_b)$, so by Lemma 14, x_a and x_b are mutual best responses.

Uniqueness Define $\tilde{x}_a, \tilde{x}_b$ as the equilibrium positions.

By Lemma 9 and Lemma 8, we have $2x_a^* < x_a < x_b < 2x_b^*$. By assumption, we have $2|x_b^* - x_a^*| \leq \kappa\Delta$, so

$$|x_b - x_a| \leq 2|x_b^* - x_a^*| \leq \kappa\Delta,$$

so by Lemma 12, either feasibility is not binding at (x_a, x_b) or one party obtains all the votes. We have ruled out both cases. □

1.10 Appendix: IO Extension

Proof of Lemma 2

Proof. In a pure-strategy equilibrium, consumers anticipate product locations, which induces a one-dimensional distribution of revealed preferences (Theorem 1).

Suppose without loss that revealed preferences are supported on the line spanned by

$$e_1 = \begin{pmatrix} 1 \\ 0 \\ ... \\ 0 \end{pmatrix}.$$

Solving for Equilibrium Profits: First, we show an equivalence between multidimensional product locations $x_a, x_b \in \mathbb{R}^n$ and one-dimensional *pseudo-locations* $\hat{x}_a, \hat{x}_b \in \mathbb{R}$ in the sense that the the utility difference is the same,

$$\begin{aligned} u(x_b, \theta) - u(x_b, \theta) &= 2\langle \theta - \frac{x_a + x_b}{2}, x_b - x_a \rangle_A \\ &= 2\langle x_b - x_a, e_1 \rangle_A \theta_1 + (\langle x_a, x_a \rangle_A - \langle x_b, x_b \rangle_A) \\ &= 2\left(\theta_1 - \frac{\langle x_b, x_b \rangle_A - \langle x_a, x_a \rangle_A}{2\langle x_b - x_a, e_1 \rangle_A}\right) \langle x_b - x_a, e_1 \rangle_A \\ &= 2\left(\theta_1 - \frac{\hat{x}_a + \hat{x}_b}{2}\right)(\hat{x}_b - \hat{x}_a) = \hat{u}(\hat{x}_b, \theta_1) - \hat{u}(\hat{x}_a, \theta_1) \end{aligned}$$

where $\hat{x}_a, \hat{x}_b \in \mathbb{R}$ are uniquely determined by

$$\begin{aligned} \hat{x}_b - \hat{x}_a &= \langle x_b - x_a, e_1 \rangle_A \\ \frac{\hat{x}_a + \hat{x}_b}{2} &= \frac{\langle x_b, x_b \rangle_A - \langle x_a, x_a \rangle_A}{\langle x_b - x_a, e_1 \rangle_A}. \end{aligned}$$

Because the utility difference is a sufficient statistic for subsequent behavior of consumers, this implies behavioral equivalence of consumers. This behavioral equivalence allows us to

use results from Anderson et al. (1997) to determine the equilibrium of the pricing subgame and characterize equilibrium profits as a function of product locations x_a, x_b.

Assume without loss that x_a is to the left of x_b, that is, $\langle x_a, e_1\rangle_A < \langle x_b, e_1\rangle_A$. (Under equality, prices are zero, which is not an equilibrium.) Following Anderson et al. (1997), we define $\xi \in \mathbb{R}$ as the indifferent type given $\hat{x}_a, \hat{x}_b \in \mathbb{R}$ (and the resulting equilibrium prices), which is defined implicitly via

$$\xi = \frac{\hat{x}_a + \hat{x}_b}{2} + \frac{1 - 2F(\xi)}{f(\xi)}. \tag{1.28}$$

As Anderson et al. (1997) note, for log-concave F, the term $\frac{1-2F(\xi)}{f(\xi)}$ is non-increasing in ξ and there is a unique solution to (1.28). Moreover, it follows that the solution is strictly increasing in $\frac{\hat{x}_a+\hat{x}_b}{2}$. They show that the equilibrium price is $p_a = 2(\hat{x}_b - \hat{x}_a)F(\xi)/f(\xi)$ and therefore equilibrium utility (profits) U_a for firm a are

$$U_a = 2c(\hat{x}_b - \hat{x}_a)\frac{F(\xi)}{f(\xi)}F(\xi).$$

Projecting x_a: Suppose $x_a \in \mathbb{R}^n$ was not on the line of revealed consumer preferences. We show that given any $x_b \in \mathbb{R}^n$, product location x_a is dominated by its projection $\tilde{x}_a = \frac{\langle e_1, x_a\rangle_A}{\langle e_1, e_1\rangle_A} e_1$ on the line of revealed preferences. Recall that the one-dimensional pseudo-locations $\hat{x}_a, \hat{x}_b \in \mathbb{R}$, which we use to characterize equilibrium profits, are a function of the multidimensional real locations x_a, x_b. To show that equilibrium profits for firm a are higher under $\tilde{x}_a$ than under x_a, we show first note that the difference of the corresponding pseudo-locations is the same,

$$\hat{x}_b(\tilde{x}_a, x_b) - \hat{x}_a(\tilde{x}_a, x_b) = \langle x_b - \tilde{x}_a, e_1\rangle_A = \langle x_b - x_a, e_1\rangle_A = \hat{x}_b(x_a, x_b) - \hat{x}_a(x_a, x_b),$$

because $\tilde{x}_a$ and x_a have the same inner product with e_1 by construction of $\tilde{x}_a$. This implies that the implied midpoints of pseudo-locations are ordered,

$$\begin{aligned}\frac{\hat{x}_a(\tilde{x}_a, x_b) + \hat{x}_b(\tilde{x}_a, x_b)}{2} &= \frac{\langle x_b, x_b\rangle_A - \langle \tilde{x}_a, \tilde{x}_a\rangle_A}{\langle x_b - \tilde{x}_a, e_1\rangle_A} = \frac{\langle x_b, x_b\rangle_A - \langle \tilde{x}_a, \tilde{x}_a\rangle_A}{\langle x_b - x_a, e_1\rangle_A} \\ &> \frac{\langle x_b, x_b\rangle_A - \langle x_a, x_a\rangle_A}{\langle x_b - x_a, e_1\rangle_A} = \frac{\hat{x}_a(x_a, x_b) + \hat{x}_b(x_a, x_b)}{2},\end{aligned}$$

because $\langle x_a, x_a\rangle_A > \langle \tilde{x}_a, \tilde{x}_a\rangle_A$ and $\langle x_b - x_a, e_1\rangle_A > 0$.

We noted above that ξ is strictly increasing in $\frac{\hat{x}_a + \hat{x}_b}{2}$. Thus, the last equation implies that ξ is larger under $\tilde{x}_a$ and x_b than under x_a and x_b.

This implies that profits U_a are higher under $\tilde{x}_a$ than under x_a. The differentiation of pseudo-locations $\hat{x}_b - \hat{x}_a$ is the same under $\tilde{x}_a$ as under x_a. But the indifferent type ξ is higher. By $F(\xi)/f(\xi)$ and $F(\xi)$ being strictly increasing, equilibrium profits are strictly higher $\tilde{x}_a$ than under x_a.

We now know that, in equilibrium, firms choose product locations on the line of consumer's revealed preferences. We also know from Theorem 1 that the line of consumer preferences has direction $\Sigma A(x_b - x_a)$. If product locations x_a and x_b are on said line, then $x_b - x_a$ is parallel to $\Sigma A(x_b - x_a)$. This is the case if and only if $x_b - x_a$ is an eigenvector of ΣA. □

Proof of Theorem 7

Proof. Fix a v_i, $i \in \{1, \ldots, n\}$. We look for equilibria where product locations (x_a, x_b) are on the line spanned by v_i and consumer learning is to be about the A-projection

$$\frac{\langle v_i, \theta\rangle_A}{\langle v_i, v_i\rangle_A} \cdot v_i = \langle v_i, \theta\rangle_A \cdot v_i$$

of their ideal points on v_i. By Lemma 2 all equilibria are of this form.

To find such equilibria, we can reduce the model to one dimension by projecting the product space on v_i. The utility of consumers for product location $\hat{x} \cdot v_i$ given ideal point

$\hat{\theta} \cdot v_i$ is

$$-(\hat{x} \cdot v_i - \hat{\theta} \cdot v_i)^\top A(\hat{x} \cdot v_i - \hat{\theta} \cdot v_i) = -(\hat{x} - \hat{\theta})^2.$$

The distribution of ideal points projected on v_i is $\mathcal{N}(0, \sigma_\mu^2)$ with

$$\sigma_\mu^2 = v_i^\top A \Sigma A v_i.$$

Normal signal structures induce normal distributions $\rho = \mathcal{N}(0, \sigma_\rho^2)$ over posteriors means with posterior variance σ_π^2, where $\sigma_\mu^2 = \sigma_\rho^2 + \sigma_\pi^2$ by the law of total variance.

Anderson et al. (1997) analyze a one-dimensional Hotelling model under quadratic consumer preferences $u(x, \theta) = -(x - \theta)^2$, $x, \theta \in \mathbb{R}$, and an exogenous log-concave distribution of consumer preferences. Using their Corollary 1, we know that given the distribution over revealed preferences $\rho = \mathcal{N}(0, \sigma_\rho^2)$ with density f_ρ, the unique equilibrium product attributes x and prices p are characterized by

$$x := -x_a = x_b = \frac{3}{4f_\rho(0)} = \frac{3}{4}\sqrt{2\pi}\sigma_\rho, \tag{1.29}$$

$$p := p_a = p_b = c\frac{3}{2f_\rho(0)^2} = 3\pi\sigma_\rho^2. \tag{1.30}$$

To solve for the optimal standard deviation of consumer preferences σ_ρ given x_a and x_b, we write the consumer's instrumental value of information of the normal distribution τ over normal posteriors π as as function of σ_ρ:

$$\begin{aligned}
\mathbb{E}_\tau\Big[\mathbb{E}_\pi\Big[\max_{j\in\{a,b\}}\{-(\theta - x_j)^2\} - p\Big]\Big] &= \mathbb{E}_\tau\big[-(|\mathbb{E}_\pi[\theta]| - x)^2 - \sigma_\pi^2 - p\big] \\
&= \mathbb{E}_\tau\big[-\mathbb{E}_\pi[\theta]^2 + 2x\mathbb{E}_\pi[|\theta|] - x^2 - \sigma_\pi^2 - p\big] \\
&= -\sigma_\rho^2 + \sqrt{\frac{8}{\pi}}x\sigma_\rho - x^2 - \sigma_\pi^2 - p \\
&= \sqrt{\frac{8}{\pi}}x\sigma_\rho - x^2 - p - \sigma_\mu^2
\end{aligned}$$

The information cost is

$$\kappa \left(\log(2\pi\sigma_\mu^2) - \log(2\pi\sigma_\pi^2) \right) = \kappa \left(\log(\sigma_\mu^2) - \log(\sigma_\pi^2) \right).$$

Neglecting constants, the consumer chooses σ_ρ to solve

$$\begin{aligned} \max_{\sigma_\rho} \ & \sqrt{\frac{8}{\pi}} x \sigma_\rho + \kappa \log(\sigma_\mu^2 - \sigma_\rho^2) \\ \text{s.t. } & 0 \leq \sigma_\rho \leq \sigma_\mu. \end{aligned}$$

The second derivative of the objective

$$\frac{-2(\sigma_\mu^2 - \sigma_\rho^2) - 4\sigma_\rho^2}{(\sigma_\mu^2 - \sigma_\rho^2)^2} = \frac{-2\sigma_\pi^2 - 4\sigma_\rho^2}{(\sigma_\mu^2 - \sigma_\rho^2)^2} < 0,$$

is negative, so the first-order condition is sufficient for optimality. Supposing the inequality constraints are not binding, we get the first-order condition

$$\sqrt{\frac{8}{\pi}} \frac{x}{\kappa} = \frac{2\sigma_\rho}{\sigma_\mu^2 - \sigma_\rho^2}. \tag{1.31}$$

Defining $d := \sqrt{8/\pi} x/\kappa$, the equation has a unique positive solution

$$\sigma_\rho = -\frac{1}{d} + \sqrt{\frac{1}{d}^2 + \sigma_\mu^2}$$

that lies between 0 and σ_μ, so the inequality constraints are satisfied. Thus, the first-order condition (1.31) characterizes the unique optimum of consumer learning.

Inserting equilibrium product locations (1.29) into the first-order condition (1.31) of consumer learning, we obtain

$$\frac{3}{\kappa}\sigma_\rho = \frac{2\sigma_\rho}{\sigma_\mu^2 - \sigma_\rho^2} \quad \Longrightarrow \quad \sigma_\rho = 0 \quad \vee \quad \sigma_\mu^2 - \sigma_\rho^2 = \frac{2}{3}\kappa.$$

Thus, under $2\kappa/3 \geq \sigma_\mu^2 = v_i^\top A\Sigma A v_i$, there is a unique equilibrium without learning, product differentiation, and markups,

$$\sigma_\rho = 0, \qquad x_a = x_b = 0, \qquad p_a = p_b = 0.$$

Else, there is an additional equilibrium with learning, product differentiation, and markups,

$$\sigma_\rho^2 = \sigma_\mu^2 - \frac{2}{3}\kappa, \qquad -x_a = x_b = \frac{3}{4}\sqrt{2\pi}\sigma_\rho, \qquad p_a = p_b = 3\pi\sigma_\rho^2.$$

This concludes the proof. □

Proof of Corollary 4

Proof. By the proof of Theorem 7, consumer utility (whether ex ante or ex post) is

$$U_i = -\sigma_\mu^2 + \sqrt{\frac{8}{\pi}}x\sigma_\rho - x^2 - p - \kappa\left(\log\left(\sigma_\mu^2\right) - \log\left(\sigma_\mu^2 - \sigma_\rho^2\right)\right).$$

When $\frac{2}{3}\kappa > \sigma_\mu^2$, then a small change in κ does not change the no-learning equilibrium and consumer utility remains $-\sigma_\mu^2$.

When $\frac{2}{3}\kappa \leq \sigma_\mu^2$, a change in κ has a direct effect on consumer welfare through the information cost and an indirect effect through σ_ρ, x, and p. By the envelope theorem, the indirect effect through σ_ρ is zero. Using

$$\begin{aligned} x &= \frac{3}{4}\sqrt{2\pi}\sigma_\rho = \frac{3}{4}\sqrt{2\pi}\sqrt{\sigma_\mu^2 - \frac{2}{3}\kappa}, \\ p &= 3\pi\sigma_\rho^2 = 3\pi\left(\sigma_\mu^2 - \frac{2}{3}\kappa\right), \end{aligned}$$

the total derivative of consumer utility with respect to κ is

$$\begin{aligned}\frac{d}{d\kappa}U_i &= \frac{dx}{d\kappa}\frac{d}{dx}\left(-x^2\right) + \frac{dp}{d\kappa}\frac{d}{dp}(-p) + \frac{\partial}{\partial\kappa}\left(-\kappa(\log\left(\sigma_\mu^2\right) - \log\left(\sigma_\mu^2 - \sigma_\mu^2\right)\right) \\ &= \frac{dx}{d\kappa}(-2x) + \frac{dp}{d\kappa}(-1) - \left(\log\left(\sigma_\mu^2\right) - \log\left(\frac{2}{3}\kappa\right)\right) \\ &= -\frac{\sqrt{2\pi}}{8\sigma_\rho}\left(-\frac{3}{2}\sqrt{2\pi}\sigma_\rho\right) + \pi - \left(\log\left(\sigma_\mu^2\right) - \log\left(\frac{2}{3}\kappa\right)\right) \\ &= \left(1+\frac{3}{8}\right)\pi - \left(\log\left(\sigma_\mu^2\right) - \log\left(\frac{2}{3}\kappa\right)\right)\end{aligned}$$

Thus, increasing the cost parameter κ has constant positive marginal effect of $(1+3/8)\pi$ on consumer utility through lowering product differentiation and prices, and a negative effect through its direct effect on the cost of information. The latter effect is becoming less strong as κ increases. Thus, voter utility is quasi-convex in κ.

Setting the total derivative to zero, we obtain that minimal consumer utility is achieved at

$$\left(1+\frac{3}{8}\right)\pi = \log\left(\sigma_\mu^2\right) - \log\left(\frac{\kappa}{3}\right) \Rightarrow e^{\left(1+\frac{3}{8}\right)\pi} = \frac{3\sigma_\mu^2}{\kappa} \Rightarrow \kappa = 3e^{-\left(1+\frac{3}{8}\right)\pi}\sigma_\mu^2 \approx 0.04\sigma_\mu^2.$$

Thus, consumer utility is maximal under $\kappa \geq \frac{2}{3}\sigma_\mu^2$ or under $\kappa = 0$. Under $\kappa = \frac{2}{3}\sigma_\mu^2$, consumer utility is $-\sigma_\mu^2$. Under $\kappa = 0$, we have $\sigma_\rho = \sigma_\mu$ and consumer utility is

$$\begin{aligned}U_i &= -\sigma_\mu^2 + \sqrt{\frac{8}{\pi}}x\sigma_\rho - x^2 - p \\ &= -\sigma_\mu^2 + \sqrt{\frac{8}{\pi}}\frac{3}{4}\sqrt{2\pi}\sigma_\mu^2 - \frac{9}{8}\pi\sigma_\mu^2 - 3\pi\sigma_\mu^2 \\ &= \sigma_\mu^2\left(-1+3-\frac{9}{8}\pi - 3\pi\right) < -\sigma_\mu^2.\end{aligned}$$

Thus, consumer welfare is maximal under $\kappa \geq \frac{2}{3}\sigma_\mu^2$. □

1.11 Appendix: Other Results and Proofs

Throughout, we use the notation $\langle x, y\rangle_A := x^\top Ay$.

Imperfect Observation of Party Platforms Our results on voter learning are, under some assumptions, robust to voters observing a *noisy* signal about party platforms before learning about ideal points (under a timing where parties move before voters learn).

Suppose that platforms are stochastic and independent of each other and of voters' ideal points. This stochasticity may stem from parties making random errors when choosing their platforms or from parties having stochastic and private ideal points, similar to Matějka and Tabellini (2021). Further, assume that, after a common signal about platforms, voters acquire a signal about their ideal points. Because both signals preserve the independence of platforms and ideal points, the expected policy utility from voting for unknown platform x under unknown ideal point θ can be written as

$$\mathbb{E}[u(x,\theta)] = u(\mathbb{E}[x],\mathbb{E}[\theta]) - \mathbb{E}\big[(x-\mathbb{E}[x])^\top A(x-\mathbb{E}[x])\big] - \mathbb{E}\big[(\theta-\mathbb{E}[\theta])^\top A(\theta-\mathbb{E}[\theta])\big].$$

Up to the constant $\mathbb{E}\big[(x-\mathbb{E}[x])^\top A(x-\mathbb{E}[x])\big]$, the agent's utility is as under known platforms, (1.3), except for replacing the platform x with the expected platform $\mathbb{E}[x]$. Thus, our results from section 1.3 remain to hold after replacing party platforms with their expectation.

If voters observe heterogeneous signals about platforms, this creates heterogeneity in their learning strategies, such as the direction in which they learn (cf. Theorem 1). However, as long as their signals about platforms are similar enough, our results should carry over approximately. As a consequence, an extension of our model to heterogeneous signal may explain the empirical finding that the ideal points of politically better informed citizens are better described by a low-dimensional model (Converse, 1964; Hare, 2022), since better informed voters should have more homogeneous information about platforms.

Existence and Continuity We show the set of maximizers of the voter learning problem are nonempty and upper hemicontinuous in the appropriate topology. We cannot show this

using Berge's maximum theorem because, for an infinite state space, the Kullback-Leibler divergence, which is part of the objective, is only lower semicontinuous and not continuous. Berge's maximum theorem requires a continuous (and not just upper semicontinuous) objective to show upper hemicontinuity of the argmax correspondence. Although our objective is only upper semicontinuous in the choice variable, it is continuous in the parameter (the value function). Using this observation, we can apply the generalization of Berge's maximum theorem by Tian and Zhou (1992) to obtain our result.

To apply the result by Tian and Zhou (1992), we first note that the set of Bayes-consistent distributions

$$X := \{\tau \in \Delta(\Delta(\mathbb{R}^n))|\mathbb{E}_\tau[\pi] = \mu\}$$

is compact with respect to the weak topology by Kartik et al. (2022), Lemma 4. They show this result for any sigma-compact Polish state space by applying Prokhorov's theorem twice. For the interested reader, we include a shorter proof for the state space $\mathbb{R}^n$ by using a compactification argument.

Lemma 15. *The set of Bayes-consistent distributions X is compact with respect to the weak topology.*

Proof. Denote by $\mathbb{R}^n \cup \{\infty\}$ the one-point compactification of $\mathbb{R}^n$, which is an embedding. The space $\mathbb{R}^n \cup \{\infty\}$ is homeomorphic to the unit n-sphere, so it is a (compact) Polish space. By Aliprantis and Border (2006), Theorem 15.14, the pushforward of the embedding induces an embedding between Polish spaces $\Delta(\mathbb{R}^n) \hookrightarrow \Delta(\mathbb{R}^n \cup \{\infty\})$. By iteration of this argument, this induces an embedding $\Delta(\Delta(\mathbb{R}^n)) \hookrightarrow \Delta(\Delta(\mathbb{R}^n \cup \{\infty\}))$. Under this embedding, the image of X is the set $\hat{X} := \{\tau \in \Delta(\Delta(\mathbb{R}^n \cup \{\infty\}))| \int \pi d\tau = \mu\}$. The space $\Delta(\Delta(\mathbb{R}^n \cup \{\infty\}))$ is compact because $\mathbb{R}^n \cup \{\infty\}$ is a compact Polish space. The set $\hat{X}$ is the preimage of a singleton $\{\mu\}$ under a continuous function $\tau \mapsto \int \pi d\tau$, so it is closed. A closed subset of a compact space is compact, so $\hat{X}$ is compact. The set X is the preimage of $\hat{X}$ under an embedding, so X is compact. □

Next, we prove a general maximum theorem for information-design problems on non-compact state spaces and for upper semicontinuous and bounded-from-above value functions. As above, define X as the set of Bayes-consistent distributions $\tau \in \Delta(\Delta(\mathbb{R}^n))$ over posteriors $\pi \in \Delta(\mathbb{R}^n)$ endowed with the topology induced by weak convergence. Define Y as the set of upper semicontinuous and bounded-from-above value functions v from $\Delta(\mathbb{R}^n)$ to $\mathbb{R} \cup \{-\infty\}$, endowed with the topology induced by uniform convergence. Here, upper semicontinuity is defined with respect to the topology of weak convergence on $\Delta(\mathbb{R}^n)$. Define $f\colon X \times Y \to \mathbb{R}$ as the expected value

$$f(\tau, v) = \int v(\pi) d\tau(\pi).$$

Proposition 4. *The argmax correspondence of the information design problem,*

$$M\colon Y \rightrightarrows X$$

$$M(v) := \arg\max_{\tau \in X} f(\tau, v),$$

is nonempty compact-valued and upper hemicontinuous.

Proof. For the proof, we use Theorem 1 in Tian and Zhou (1992). It shows that in a maximization problem, if (1) the objective is upper semicontinuous and feasible path transfer lower semicontinuous and (2) the feasibility correspondence is nonempty compact-valued, closed and upper hemicontinuous, then the maximum correspondence is nonempty compact-valued and upper hemicontinuous.

First, we show upper semicontinuity of f in (τ, v). Suppose τ_n converges weakly to τ and v_n converges uniformly to v. We abbreviate $\int v(\pi) d\tau(\pi)$ by $\int v d\tau$ and show that

$\lim_{n\to\infty} f(\tau_n, v_n) - f(\tau, v) \leq 0$:

$$\begin{aligned}\lim_{n\to\infty} \int v_n d\tau_n - \int v d\tau &= \lim_{n\to\infty} \left(\int v_n d\tau_n - \int v d\tau_n \right) + \left(\int v d\tau_n - \int v d\tau \right) \\ &= \lim_{n\to\infty} \int (v_n - v) d\tau_n + \left(\int v d\tau_n - \int v d\tau \right) \\ &\leq \lim_{n\to\infty} \int |v_n - v| \, d\tau_n + \lim_{n\to\infty} \left(\int v d\tau_n - \int v d\tau \right) \leq 0 \qquad (1.32)\end{aligned}$$

In the last line, the first limit is zero because v_n converges uniformly to v and τ_n is a probability measure. By Villani (2009), Lemma 4.3, v being upper semicontinuous and bounded from above implies that $\tau \mapsto \int v d\tau$ is upper semicontinuous, so the second limit is less or equal to zero.

Second, we show feasible path transfer lower semicontinuity, which is introduced by Tian and Zhou (1992). The objective f is feasible path transfer lower semicontinuous in y with respect to feasibility correspondence F if for each $(x, y) \in X \times Y$ with $x \in F(y)$, there exists some neighborhood $\mathcal{N}(y)$ of y such that $\forall y' \in \mathcal{N}(y), \exists x' \in F(y')$ satisfying

$$f(x, y) \leq \liminf_{y' \to y} f(x', y').$$

Because in our case, $\forall y \in Y \colon F(y) = X$, we can choose $x' = x$ for all y'. Then, for any sequence $y_n \to y$

$$\lim_{n\to\infty} f(x, y_n) = \int y_n dx = \int y dx = f(x, y)$$

because y_n converges uniformly to y and x is a probability measure. So, $f(x, y) \leq \liminf_{y' \to y} f(x', y')$. Key is that while our objective may be discontinuous in the choice variable, it is continuous in the parameter (the value function).

Finally, in our case, the feasibility correspondence is nonempty and constant. Hence, it is closed and continuous, and thus also upper hemicontinuous. It is compact-valued by Lemma 15. □

To apply Proposition 4, we show the value function of the voter learning problem,

$$\mathbb{E}_\nu\Big[\max\Big\{\mathbb{E}_\pi[u(x_a,\theta)],\mathbb{E}_\pi[u(x_b,\theta)]+\nu\Big\}\Big]-\kappa D(\pi||\mu) \tag{1.33}$$

is indeed upper semicontinuous and bounded from above. The only property of the information cost that this result uses is lower semicontinuity of the Kullback-Leibler divergence.

Lemma 16. *The value function, (1.33), is upper semicontinuous in π and bounded from above.*

Proof. We separately show upper semicontinuity and boundedness from above for the information cost $-\kappa D(\pi||\mu)$ and for the instrumental value, $\mathbb{E}_\nu[\max\{\mathbb{E}_\pi[u(x_a,\theta)],\mathbb{E}_\pi[u(x_b,\theta)]+\nu\}]$, of information. Then, the their sum is upper semicontinuous and bounded from above.

The divergence $D_{\mathrm{KL}}(\cdot||\mu)$ is bounded from below as it is non-negative and lower semicontinuous by E. Posner (1975), Theorem 1. Thus, $-D_{\mathrm{KL}}(\pi||\mu)$ is upper semicontinuous and bounded from above.

Define the instrumental value of a posterior $V(\pi)$ as

$$V(\pi)=\mathbb{E}_\nu\Big[\max\Big\{\mathbb{E}_\pi[u(x_a,\theta)],\mathbb{E}_\pi[u(x_b,\theta)]+\nu\Big\}\Big].$$

The utility $u(x,\theta)=-(x-\theta)^\top A(x-\theta)$ is continuous (and hence upper semicontinuous) in θ and bounded from above by zero. By Villani (2009), Lemma 4.3, the function $\pi\mapsto\mathbb{E}_\pi[u(x,\theta)]=\int u(x,\theta)d\pi(\theta)$ is upper semicontinuous in π, and it is bounded from above by zero. Further,

$$\mathbb{E}_\nu[\max\{l,r+\nu\}]=lF_\nu(l-r)+r(1-F_\nu(l-r))+\int_{l-r}^{\infty}sdF_\nu(s).$$

Because valance ν has a continuous distribution, F_ν is differentiable. Thus, $\mathbb{E}_\nu[\max\{l,r+\nu\}]$ is continuous in l and r. Thus, $V(\pi)$ is upper semicontinuous.

Further,

$$lF_\nu(l-r) + r(1-F_\nu(l-r)) + \int_{l-r}^{\infty} sdF_\nu(s) \leq \max\{l,r\} + \frac{1}{2}\int |s|dF_\nu(s),$$

so by boundedness of $\mathbb{E}_\pi[u(x,\theta)]$ from above and ν having a finite first absolute moment, the instrumental value $V(\pi)$ is bounded from above. □

Corollary 5. *A solution to the voter learning problem (P) exists.*

Proof. Follows immediately from Proposition 4 and Lemma 16. □

Finally, for our proof of Theorem 2, we show the following result.

Lemma 17. *The value function, (1.33), converges uniformly as the valence shock converges in mean to zero.*

Proof. Let $\pi \in \Delta(\mathbb{R}^n)$. The information cost $-\kappa D(\pi||\mu)$ does not depend on the valence shock and can thus be ignored. The other component of the value function is the instrumental value, which we write as $V(\pi, F_\nu)$ as a function of the posterior π and CDF of ν,

$$V(\pi, F_\nu) = \int_{-\infty}^{\infty} \max\{\mathbb{E}_\pi[u(x_a,\theta)], \mathbb{E}_\pi[u(x_b,\theta)] + s\}dF_\nu(s).$$

Write $l := \mathbb{E}_\pi[u(x_a,\theta)]$ and $r := \mathbb{E}_\pi[u(x_b,\theta)]$, so $V(\pi, F_\nu) = \int_{-\infty}^{\infty} \max\{l, r+s\}dF_\nu(s)$. For ν degenerate at 0, that is $F_\nu(\nu) = \mathbb{1}_{\{\nu \geq 0\}}$, we have $V(\pi, \mathbb{1}_{\{\nu\geq 0\}}) = \max\{l,r\}$. We show $V(\pi, F_\nu)$ converges to $V(\pi, \mathbb{1}_{\{\nu\geq 0\}})$ as ν converges to zero in mean.

First, consider the case that $l \geq r$, so $V(\pi, \mathbb{1}_{\{\nu \geq 0\}}) = l$. Then,

$$\begin{aligned} l &= \int_{-\infty}^{\infty} l dF_\nu(s) \leq \int_{-\infty}^{\infty} \max\{l, r+s\} dF_\nu(s) \\ &= V(\pi, F_\nu) = l + \int_{-\infty}^{\infty} \max\{0, r-l+s\} dF_\nu(s) \\ &\leq l + \int_{0}^{\infty} \max\{0, s\} dF_\nu(s) = l + \frac{1}{2}\int_{-\infty}^{\infty} |s| dF_\nu(s) \\ \Rightarrow l &\leq V(\pi, F_\nu) \leq l + \frac{1}{2}\int_{-\infty}^{\infty} |s| dF_\nu(s), \end{aligned}$$

where we have used the symmetry of ν in the third line.

Second, consider the case that $r > l$, so $V(\pi, \mathbb{1}_{\{\nu \geq 0\}}) = r$. By symmetry of the density of ν,

$$\begin{aligned} r &= \int_{-\infty}^{\infty} r dF_\nu(s) = \int_{-\infty}^{\infty} (r+s) dF_\nu(s) \leq \int_{-\infty}^{\infty} \max\{l, r+s\} dF_\nu(s) \\ &= V(\pi, F_\nu) = r + \int_{-\infty}^{\infty} \max\{l-r, s\} dF_\nu(s) \\ &\leq r + \int_{-\infty}^{\infty} \max\{0, s\} dF_\nu(s) = r + \frac{1}{2}\int_{-\infty}^{\infty} |s| dF_\nu(s) \\ \Rightarrow r &\leq V(\pi, F_\nu) \leq r + \frac{1}{2}\int_{-\infty}^{\infty} |s| dF_\nu(s). \end{aligned}$$

Together, we have

$$|V(\pi, F_\nu) - V(\pi, \mathbb{1}_{\{\nu \geq 0\}})| = |V(\pi, F_\nu) - \max\{l, r\}| \leq \frac{1}{2}\int_{-\infty}^{\infty} |s| dF_\nu(s).$$

Thus, $V(\pi, F_\nu)$ converges uniformly to $V(\pi, \mathbb{1}_{\{\nu \geq 0\}})$ as ν converges in mean to zero. □

Distance-Based Information Costs The proof of Theorem 1 uses only Blackwell monotonicity (step 3), posterior separability (step 2), and a notion of reflection invariance (step 1) of the information cost. Thus, the result holds for all information costs that satisfy these conditions. More precisely, step 3 of the proof uses that the Kullback-Leibler divergence is strictly convex in its first argument for posteriors with finite divergence, to show a strict

mean-preserving contraction in posterior space strictly lowers the information cost. Step 2 of the proof uses linearity of the information cost under mixing between distributions over posteriors, which follows from posterior separability of the information cost. Step 1 uses that the Kullback-Leibler divergence is invariant under the constructed reflection Ref, $D_{\mathrm{KL}}(\mathrm{Ref}(\pi)||\,\mathrm{Ref}(\mu)) = D_{\mathrm{KL}}(\pi||\mu)$. This holds for all *invariant* divergences (Amari, 2016), which remain unchanged under any transformation of the state space. In fact, for invariant divergences, we sketch a somewhat shorter proof of Theorem 1 in the proof of Lemma 18. Other notable examples of invariant divergences, besides the Kullback-Leibler divergence, are the Rényi divergences, which have a foundation based on Blackwell dominance under repeated observation (Mu et al., 2021).

Furthermore, the proof of Theorem 1 generalizes beyond invariant divergences to certain distance-based divergences. Information costs based on invariant divergences have been criticized because they imply that any two states are equally costly to distinguish. The literature has, inspired by evidence from perceptual experiments, proposed distance-based information costs that make it more costly to distinguish between closer states (Hébert and Woodford, 2021; Pomatto et al., 2023). Recall that in step 1 of our proof we use that the divergence D satisfies

$$D(\mathrm{Ref}(\pi)||\,\mathrm{Ref}(\mu)) = D(\pi||\mu)$$

where Ref *is a reflection across a line that preserves the prior* μ. Because the prior μ is elliptical with covariance matrix Σ, the prior is preserved by members of the orthogonal group of inner product Σ^{-1}, which is $\{Q \in \mathbb{R}^{n\times n} | Q^\top \Sigma^{-1} Q = \Sigma^{-1}\}$. Reflections across a line are those members of the orthogonal group that deliver the identity function when applied twice and that have a line as the subset of the space that is invariant under the mapping. If the divergence D is invariant under all such reflections, then step 1 of our proof works for it. We proceed to show this condition is satisfied for certain distance-based divergences. The upshot will be that the information cost needs to be based on a distance that is compatible with inner product Σ^{-1}.

While there is no generally agreed upon definition of distance-based costs, we assume that if a posterior-separable information cost is based on distance $d\colon \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$, then it should be invariant under *isometries* of d. Recall that a bijection $g\colon \mathbb{R}^n \to \mathbb{R}^n$ is an isometry of $d\colon \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ if $\forall v, w \in \mathbb{R}^n\colon d(v,w) = d(g(v), g(w))$. Intuitively, if we relabel the states such that the distance d is preserved, the cost of information should not change. This should be seen as a minimal implication for an information cost to be based on distance d, which makes our results stronger than had we imposed a stronger requirement. Below we show that versions of recent proposals for distance-based information costs satisfy this condition.

Definition 3 (Distance-Compatible Information Cost)**.** *A posterior-separable information cost $c(\tau) = \mathbb{E}_\tau[D(\pi||\mu)]$ on state space $\mathbb{R}^n$ is compatible with distance $d\colon \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ if for all isometries $g\colon \mathbb{R}^n \to \mathbb{R}^n$ of d*

$$D(\pi||\mu) = D(g(\pi)||g(\mu)),$$

where $g\colon \Delta(\mathbb{R}^n) \to \Delta(\mathbb{R}^n)$ is the pushforward induced by g.

Particularly important in our context are the standard Euclidean distance on $\mathbb{R}^n$, $d(v,w) = \big((v-w)^\top (v-w)\big)^{1/2}$, and the (non-standard) Euclidean distance induced by a symmetric, positive definite matrix C, $d(v,w) = (v-w)^\top C(v-w)$. That is because we need the information cost to be compatible with a metric d of the state space $\mathbb{R}^n$ such that the reflection Ref in our proof is an isometry of d. Using the fact that Ref is an isometry of the distance induced by Σ^{-1}, we can easily prove the following corollary.

Corollary 6. *Let the information cost be posterior separable, $c(\tau) = \mathbb{E}_\tau[D(\pi||\mu)]$, consistent with the standard Euclidean distance, the divergence D be strictly convex in its first argument, and the prior be spherical, $\Sigma = I_n$. Then, the induced posterior means are on the line through the prior mean and $A(x_b - x_a)$ under any optimal information τ.*

More generally, maintaining posterior separability, and convexity of the divergence D, let the information cost be consistent with the Euclidean distance induced by symmetric, positive definite matrix C and the prior be spherical in an orthonormal basis of C, that is Σ^{-1} is a multiple of C. Then, the posterior means are on a line through the prior mean and $\Sigma A(x_b - x_a)$ under any optimal information τ.

Proof. Ref is a reflection with respect to Σ^{-1} and thus an isometry of the distance induced by Σ^{-1}. By $\Sigma^{-1} = kC$ with $k \in \mathbb{R}$, Ref is also an isometry of the distance induced by C. Thus, $D(\pi||\mu) = D(\mathrm{Ref}(\pi)||\,\mathrm{Ref}(\mu))$ and our proof of Theorem 1 applies. □

Intuitively, our proof of Theorem 1 requires that a reflection that preserves the prior (so the reflection defines a Bayes-consistent distribution over posteriors) to preserve the information cost. The elliptical prior with covariance matrix Σ is preserved by reflections with respect to the inner product induced by Σ^{-1}. Thus, we need the information-cost distance to induce the same geometry as Σ^{-1}. This is the case if the information-cost distance is based on an inner product induced by matrix C where C is a multiple of Σ^{-1}.

Finally, we present a few examples of information costs that are compatible with the Euclidean distance induced by a symmetric, positive definite matrix C. Strict convexity of these divergences, for posteriors with finite divergence, is either known or can be shown. (The first example is strictly convex only for posteriors that do not share the same mean, which suffices for the proof of Theorem 1.)

Example 1: Posterior-Variance Cost A multidimensional version of the posterior-variance cost can be defined as the divergence D_{Var} being a second central moment (a generalization of variance to arbitrary metric spaces),

$$D_{\mathrm{Var}}(\pi||\mu) = \mathbb{E}_\pi\Big[- d\big(\theta, \mathbb{E}_\pi[\theta]\big)^2\Big] = \int -d\big(\theta, \mathbb{E}_\pi[\theta]\big)^2 d\pi(\theta),$$

where $d(v,w) = \sqrt{v^\top C w}$ and C is a symmetric, positive definite $n \times n$-matrix. Under any bijection g on $\mathbb{R}^n$ that preserves the inner product $\langle \cdot, \cdot \rangle_C$, the information cost is preserved, that is $D(g(\pi)||g(\mu)) = D(\pi||\mu)$, which can be seen by

$$\begin{aligned} D_{\mathrm{Var}}(g(\pi)||g(\mu)) &= \int -d\big(\theta, \mathbb{E}_\pi[\theta]\big)^2 dg(\pi)(\theta) \\ &= \int -d\big(g(\theta), \mathbb{E}_\pi[g(\theta)]\big)^2 d\pi(\theta) \\ &= \int -d\big(g(\theta), g(\mathbb{E}_\pi[\theta])\big)^2 dg(\pi)(\theta) \\ &= \int -d\big(\theta), \mathbb{E}_\pi[\theta]\big)^2 d\pi(\theta) = D_{\mathrm{Var}}(\pi||\mu). \end{aligned}$$

Here, we have used that a bijection that preserves an inner product is linear and hence commutes with the expectation operator. That Examples 2 and 3 below are compatible with the Euclidean distance can shown in a similar fashion.

Example 2: Neighborhood-based Cost For state space $\mathbb{R}^n$, Hébert and Woodford (2021) propose the Fisher information cost, based on divergence

$$D(\pi||\mu) = \int_{\mathrm{supp}(\pi)} c(\theta) \frac{|\nabla f_\pi(\theta)|^2}{f_\pi(\theta)} d\theta$$

for posteriors π with density f_π (and infinite divergence if the posterior is not absolutely continuous with respect to the Lebesgue measure) and where $c(\theta)$ captures how costly it is locally to differentiate between states. If $c(\theta) = c$ is constant, one can show similar to above that the cost is based the standard Euclidean distance.

Example 3: Log-Likelihood Ratio Cost The log-likelihood ratio (LLR) cost, introduced and axiomatized by Pomatto et al. (2023) is defined for finite state spaces Θ. Pomatto et al. (2023) show, given a full-support prior, the LLR cost is posterior-separable

with divergence

$$D_{LLR}(\pi||\mu) = \sum_{\theta,\theta' \in \Theta} \beta(\theta,\theta') \frac{\pi(\theta)}{\mu(\theta)} \log\left(\frac{\pi(\theta)}{\pi(\theta')}\right),$$

if $\pi(\theta) > 0$ for all $\theta \in \Theta$ and infinite otherwise. The coefficients $\beta(\theta,\theta')$ capture how hard distinguishing between states θ and θ' is. If $\Theta \subset \mathbb{R}^n$ and $\beta(\theta,\theta') = f\left((\theta-\theta')^\top C(\theta-\theta')\right)$ for some function f, then it can be easily shown that the cost function is based on the Euclidean distance induced by C. While finite state space is not appropriate for our analysis since we assume the the prior μ is elliptical, we conjecture that under an appropriate generalization to infinite state spaces, the resulting LLR cost remains distance-based in our sense.

Corollary 1

Proof. We denote the k party platforms by x_j, $j \in \{1,...,k\}$ and assume the utility of voting for candidate j under ideal point θ is $u(x_j,\theta)+\nu_j$. Define the valence vector $\nu := (\nu_j)_{j\in\{1,...,k\}}$, which can have an arbitrary distribution.

Analogously to the proof of Theorem 1, we first show the instrumental value of information τ depends only on the projection of the induced distribution over posterior means on the $(k-1)$-dimensional subspace spanned by $\{x_j - x_1\}_{j=2,...,k}$.

$$\begin{aligned}
&\mathbb{E}_\tau\left[\mathbb{E}_\nu\left[\max_j\left\{\mathbb{E}_\pi\left[-\langle\theta-x_j,\theta-x_j\rangle_A\right]+\nu_j\right\}\right]\right]\\
=&\mathbb{E}_\tau\left[\mathbb{E}_\nu\left[\max_j\left\{-\langle x_j,x_j\rangle_A+2\langle x_j,\mathbb{E}_\pi[\theta]\rangle_A+\nu_j\right\}\right]\right]-\mathbb{E}_\mu\left[\langle\theta,\theta\rangle_A\right]\\
=&\mathbb{E}_\tau\left[\mathbb{E}_\nu\left[\max_j\left\{-\langle x_j,x_j\rangle_A+2\langle x_j-x_1,\mathbb{E}_\pi[\theta]\rangle_A+\nu_j\right\}\right]+2\langle x_1,\mathbb{E}_\pi[\theta]\rangle_A\right]-C_1\\
=&\mathbb{E}_\tau\left[\mathbb{E}_\nu\left[\max_j\left\{-\langle x_j,x_j\rangle_A+2\langle x_j-x_1,\mathbb{E}_\pi[\theta]\rangle_A+\nu_j\right\}\right]\right]-C_1+C_2
\end{aligned}$$

where $C_1 = \mathbb{E}_\mu\left[\langle\theta,\theta\rangle_A\right]$ and $C_2 = 2\langle x_1,\mathbb{E}_\mu[\theta]\rangle_A$ are constants and in fact zero under our prior. In the second line, we used the law of iterated expectations. In the last line, we used that the inner product is linear, to apply another law of iterated expectations. Thus, only the projection of $\mathbb{E}_\pi[\theta]$ on $\{x_j - x_1\}_{j=2,...,k}$ is payoff-relevant.

To replicate the second part of the proof of Theorem 1, we need to again define an appropriate reflection that preserves the instrumental value of information as well as the prior. Analogously to above, we let $\Delta\hat{x}_j := \Sigma A(x_j - x_1)$ and define the reflection as

$$\mathrm{Ref}_k(\theta) = 2\sum_{j=1}^{k} \frac{\langle \Delta\hat{x}_j, \theta\rangle_{\Sigma^{-1}}}{\langle \Delta\hat{x}_j, \Delta\hat{x}_j\rangle_{\Sigma^{-1}}}\Delta\hat{x}_j - \theta.$$

If $\Sigma = A = I_n$, this is just the standard reflection across the space spanned by $\{x_j - x_1\}_{j=2,\dots,k}$. In general, it is the suitable reflection that preserves the A-projection on this subspace as well as the prior (since it is a Σ^{-1}-reflection). The second part of the proof of Theorem 1 applies using the reflection Ref_k instead of Ref. □

Proof of Proposition 1

Proof. First, we show the conclusion of Theorem 1 still holds under the restriction to normal distributions. This implies voters' candidates for optimal signal structures are one-dimensional and normal. Hence, the candidate signal structures are completely Blackwell-ordered and thereby ordered by information cost.

Lemma 18. *Restrict the prior μ and feasible signal structures to be normal. The conclusion of Theorem 1 still holds; that is, revealed voter ideal points are on the line through the prior mean with direction $\Sigma A(x_b - x_a)$.*

Proof. Under the restriction to normal signal structures, the reflection argument underlying our proof of Theorem 1 does not hold anymore because the better signal structure, constructed by that proof, need not be normal. Instead, we apply an argument based on a so-called pre-garbling, which shows for all invariant information costs such as mutual information that agents learn only about the partition of payoff-equivalent states (Amari, 2016; Caplin et al., 2022). If the acquired signal structure was not measurable with respect to the partition of payoff-equivalent states, one can construct a better signal structure based on a pre-garbling, that is by, for each state, obtaining the average distribution over signals

conditional on the partition element of the state (Caplin et al., 2022).[38] The resulting signal structure does not distinguish between payoff-equivalent states and it is better because it is cheaper and equally instrumentally valuable. It is not hard to see that such a pre-garbling maintains normality of the signal structure. Thus, also under the restriction to normal signal structures, the optimal signal structure is measurable with respect to the partition of payoff-equivalent states, which in our case are those states θ that have the same A-projection $\langle x_b - x_a, \theta \rangle_A$ on the platform difference $x_b - x_a$. To obtain the result of Theorem 1, two additional steps are necessary. First, suppose the voter learns the A-projection $S = \langle x_b - x_a, \theta \rangle_A$ of the state θ on $x_b - x_a$ perfectly. Upon learning $S = s$, by joint normality, the posterior mean would be

$$\mathbb{E}[\theta|S = s] = \mathbb{E}[\theta] + (s - \mathbb{E}[S]) \cdot \frac{\mathrm{Cov}(\theta, S)}{\mathrm{Var}(S)} = c(s)\,\mathrm{Cov}(\theta, S) = c(s)\Sigma A(x_b - x_a),$$

where $c(s)$ is a scalar and we have used the normalization $\mathbb{E}[\theta] = 0$. Thus, the posterior means induced by S are on the line characterized by Theorem 1. Second, given that the voter actually acquires some garbling of S, the induced distribution over posterior means is a mean-preserving contraction of the one induced by S. Hence, the resulting distribution over posteriors means is also supported on the line characterized by Theorem 1. □

Second, we show the comparative statics regarding the cost parameter κ. The voter's objective is supermodular in κ and in the cost of information $c(\tau)$. Thus, a smaller κ implies a greater cost of information $c(\tau)$ in the strong order. The one-dimensional normal

[38] To be more specific, let S be a normal signal (modeled as a random vector), that is (S, θ) are jointly normal. Let $P(\theta) := \langle x_b - x_a, \theta \rangle_A = (A(x_b - x_a))^\top \theta$. The distribution of the pre-garbling $\tilde{S}$ conditional on some state θ should be identical to the distribution of S conditioned on the partition of payoff-equivalent states of θ, $\{\theta' \in \mathbb{R}^n | P(\theta') = P(\theta)\}$. Formally, we have

$$S|P(\theta) \sim \mathcal{N}\left(\mathbb{E}[S] + \Sigma_{SP}\Sigma_P^{-1}(P(\theta) - \mathbb{E}[P(\theta)]), \Sigma_S - \Sigma_{SP}\Sigma_P^{-1}\Sigma_{PS}\right).$$

where Σ_{SP}, Σ_P, and Σ_S are the relevant (cross-)covariance matrices of $P(\theta)$ and S. We can define the pre-garbling $\tilde{S}$ via

$$\tilde{S} = \mathbb{E}[S|P(\theta)] + \varepsilon$$

where $\varepsilon \sim \mathcal{N}(0, \Sigma_S - \Sigma_{SP}\Sigma_P^{-1}\Sigma_{PS})$ is independent of θ. Signal and state $(\tilde{S}, \theta)$ are jointly normal.

distributions over posterior means are completely ordered by the mean-preserving spread relation, which coincides with the ordering by variance. Thus, a greater cost of information implies a greater variance.

Third, we show the comparative statics regarding the degree of platform polarization α. By Lemma 18, the distribution ρ of revealed ideology can be written as $\rho = X\mathcal{N}(0, \sigma_\rho^2)$ with $X := \frac{\Sigma X(x_b - x_a)}{||\Sigma X(x_b - x_a)||} = \frac{\Sigma X(x_b^* - x_a^*)}{||\Sigma X(x_b^* - x_a^*)||}$. We show the value of information is supermodular in the standard deviation of revealed ideology σ_ρ and the degree of platform polarization α. This implies that the optimal variance is increasing in the strong set order in α. Because the cost of information does not depend on α, it is sufficient to show the instrumental value of information is supermodular in σ_ρ and α. By $x_b^\top A x_b = x_a^\top A x_a$, we have $\langle x_b - x_a, \frac{x_a + x_b}{2} \rangle_A = 0$. Using this and (1.11), the instrumental value of information as a function of σ_ρ and α, $V(\sigma_\rho, \alpha)$, can be written as an expectation of the one-dimensional variable $\hat{\theta}$,

$$\begin{aligned} V(\sigma_\rho, \alpha) &= \mathbb{E}_{\theta \sim X\mathcal{N}(0,\sigma_\rho^2)}\big[\mathbb{E}_\nu\big[\max\big\{-\langle \alpha(x_b^* - x_a^*), \theta\rangle_A, \langle \alpha(x_b^* - x_a^*), \theta\rangle_A + \nu\big\}\big]\big] \\ &= \mathbb{E}_{\hat{\theta} \sim \mathcal{N}(0,\sigma_\rho^2)}\Big[v(\hat{\theta}, \alpha)\Big] = \mathbb{E}_{Z \sim \mathcal{N}(0,1)}\Big[v(\sigma_\rho Z, \alpha)\Big], \end{aligned}$$

with

$$v(\hat{\theta}, \alpha) := \mathbb{E}_\nu\big[\max\big\{-\alpha \underbrace{\langle x_b^* - x_a^*, X\rangle_A}_{=:\Delta} \hat{\theta}, \alpha \underbrace{\langle x_b^* - x_a^*, X\rangle_A}_{=\Delta} \hat{\theta} + \nu\big\}\big].$$

For illustration, we graph the value function $v(\hat{\theta}, \alpha)$ as a function of $\hat{\theta}$ for different values of α in Figure 1.8. To show the instrumental value $V(\sigma_\rho, \alpha)$ is supermodular in σ_ρ and α, we show $\frac{d^2}{d\alpha d\sigma_\rho} V(\sigma_\rho, \alpha) > 0$. We have

$$\frac{d^2}{d\alpha d\sigma_\rho} V(\sigma_\rho, \alpha) = \mathbb{E}_{Z \sim \mathcal{N}(0,1)}\big[Z \frac{d^2}{d\alpha d\hat{\theta}} v(\sigma_\rho Z, \alpha)\big]. \tag{1.34}$$

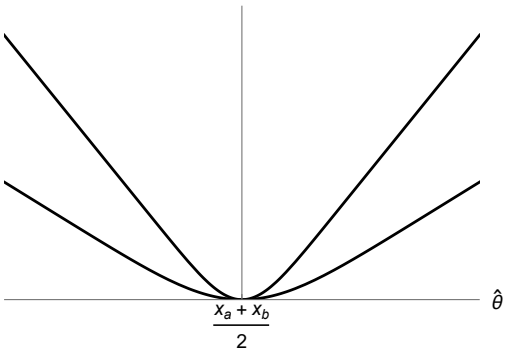


Figure 1.8: The value function $v(\hat{\theta}, \alpha)$ as a function of $\hat{\theta}$ for two values of α. The higher alpha corresponds to the steeper value function, which offers a higher return to information.

Assuming $\Delta > 0$ (the other case is analogous), the derivative of the value function $v(\hat{\theta}, \alpha)$ in $\hat{\theta}$ is, using the envelope theorem,

$$\frac{d}{d\hat{\theta}} v(\hat{\theta}, \alpha) = -\alpha\Delta + 2\alpha\Delta F_\nu(\alpha\Delta\hat{\theta}) = 2\alpha\Delta \left(F_\nu(\alpha\Delta\hat{\theta}) - \frac{1}{2} \right).$$

By symmetry of ν, the factor in brackets has the same sign as $\hat{\theta}$. Both factors are increasing in absolute value in α. Thus, this term is positive and increasing in α if $\hat{\theta} > 0$, zero if $\hat{\theta} = 0$, and negative and decreasing in α if $\hat{\theta} < 0$. Because (1.34) includes Z as an additional factor, the cross-derivative is positive. □

Finally, we argue that our result is, under some conditions, robust to a common component of ideal point. Suppose there is both a common component and an idiosyncratic component of the ideal point, $\theta_i = \omega + \delta_i$ where $\omega \sim \mathcal{N}(0, \Sigma_\omega)$, $\delta_i \sim \mathcal{N}(0, \Sigma_\delta)$, and ω and all δ_i are mutually independent. By the proof of Lemma 18, voters acquire noisy signals about the A-projection of their ideal point on the platform difference, $\hat{\theta} := (A(x_b - x_a))^\top \theta$. Defining $\hat{\omega} := (A(x_b - x_a))^\top \omega$ and $\hat{\delta} := (A(x_b - x_a))^\top \delta$ analogously, we can treat the voter's learning problem as a one-dimensional problem where they learn about $\hat{\theta} = \hat{\omega} + \hat{\delta}$ where $\hat{\omega} \sim \mathcal{N}(0, \sigma^2_{\hat{\omega}})$ and $\hat{\delta} \sim \mathcal{N}(0, \sigma^2_{\hat{\delta}})$, with $\sigma^2_{\hat{\omega}} = (A(x_b - x_a))^\top \Sigma_\omega A(x_b - x_a)$ and $\sigma^2_{\hat{\delta}} = (A(x_b - x_a))^\top \Sigma_\delta A(x_b - x_a)$. It remains to show the variance of the idiosyncratic uncertainty in (**??**) is increasing in the informativeness of the signal, or decreasing in the noise variance σ^2_ε of the normal signal

$S = \hat{\theta} + \varepsilon$. This is the case if and only if

$$\frac{d}{d\sigma_\varepsilon^2}\left(\frac{\sigma_{\hat{\omega}}^2 + \sigma_{\hat{\delta}}^2}{\sigma_{\hat{\omega}}^2 + \sigma_{\hat{\delta}}^2 + \sigma_\varepsilon^2}\right)^2 \left(\sigma_{\hat{\delta}}^2 + \sigma_\varepsilon^2\right) < 0 \Leftrightarrow \frac{d}{d\sigma_\varepsilon^2}\frac{\sigma_{\hat{\delta}}^2 + \sigma_\varepsilon^2}{\left(\sigma_{\hat{\omega}}^2 + \sigma_{\hat{\delta}}^2 + \sigma_\varepsilon^2\right)^2} < 0$$

$$\Leftrightarrow \left(\sigma_{\hat{\omega}}^2 + \sigma_{\hat{\delta}}^2 + \sigma_\varepsilon^2\right)^2 - 2\left(\sigma_{\hat{\delta}}^2 + \sigma_\varepsilon^2\right)\left(\sigma_{\hat{\omega}}^2 + \sigma_{\hat{\delta}}^2 + \sigma_\varepsilon^2\right) < 0 \Leftrightarrow \sigma_{\hat{\omega}}^2 < \sigma_{\hat{\delta}}^2 + \sigma_\varepsilon^2,$$

which is implied by $\sigma_{\hat{\delta}}^2 > \sigma_{\hat{\omega}}^2$. That is, the monotone comparative statics holds as long as the variance $\sigma_{\hat{\delta}}^2$ of the idiosyncratic component (when A-projected on the platform difference $x_b - x_a$) is greater than the variance $\sigma_{\hat{\omega}}^2$ of the common component (when A-projected on the platform difference $x_b - x_a$).

Proof of Proposition 2 For brevity, we refer by property L to the property that ideal points are on a line when projected onto the space spanned by the survey questions, as defined in the main text.

It is known (e.g. Ladha, 1991) that the one-dimensional spatial model with quadratic utility is equivalent to the one-dimensional item-response theory (IRT) model, which the empirical papers we referred to estimate (Jessee, 2009; Jessee, 2012; Tausanovitch and Warshaw, 2012; Shor and Rogowski, 2018; Fowler et al., 2022).

The IRT model is given as follows. Let y_{ij} denote the response of individual i to question j, which can be either 1 or 0. Under a one-dimensional IRT model, the likelihood is given by

$$\Pr(y_{ij} = 1) = \Phi(\alpha_j + \beta_j p_i),$$

where Φ is the logistic or the normal cumulative distribution function, $\alpha_j, \beta_j \in \mathbb{R}$ are question-specific parameters, and $p_i \in \mathbb{R}$ are individual-specific parameters.

In our proof, we show (Part I) if a multidimensional spatial model has property L, then there is a one-dimensional IRT model that is observationally equivalent, and (Part II) if there is a one-dimensional IRT model that is observationally equivalent to a given multidimensional spatial model, then the spatial model satisfies property L. Because of the equivalent to of

one-dimensional IRT models to one-dimensional spatial models with quadratic utility, this establishes our result.

Proof. First, in the spatial model, we have

$$\langle x_{j1} - \theta_i \rangle - \langle x_{j2} - \theta_i \rangle = \langle \Delta x_j, \theta_i - \overline{x}_j \rangle$$

where $\Delta x_j := x_{j1} - x_{j2}$ and $\overline{x}_j := \frac{x_{j1}+x_{j2}}{2}$. We suppose that the same response item is not part of two different questions, which seems to be satisfied in practice, so there are no restrictions that x_{j1} or x_{j2} for different j are the same.[39] Hence, there are no restrictions on Δx_j and $\overline{x}_j$, and we can reparametrize the model through $\{\Delta x_j, \overline{x}_j, \theta_i\}$ instead of $\{x_{j1}, x_{j2}, \theta_i\}$.

The multidimensional spatial model with parameters $\{\Delta x_j, \overline{x}_j, \theta_i\}$ is observationally equivalent to the one-dimensional IRT model with parameters $\{\alpha_j, \beta_j, p_i\}$ if and only if

$$\forall i, j \colon \langle \Delta x_j, \theta_i - \overline{x}_j \rangle = \alpha_j + \beta_j p_i \tag{1.35}$$

Part I Suppose the multidimensional spatial model satisfies property L. Then, define

$$\begin{aligned} p_i &:= \lambda_i \\ \beta_j &:= \langle \Delta x_j, \Delta\theta \rangle \\ \alpha_j &:= \langle \Delta x_j, \theta_1 - \overline{x}_j \rangle \end{aligned}$$

Then, (1.35) holds by

$$\langle \Delta x_j, \theta_j - \overline{x}_j \rangle = \langle \Delta x_j, \theta_1 - \overline{x}_j + \lambda_i \Delta\theta + \theta_i^{\perp} \rangle = \alpha_j + \beta_j \lambda_i$$

and the IRT model with parameters $\{\alpha_j, \beta_j, p_i\}$ is observationally equivalent.

[39] This condition is sufficient but not necessary for our proof.

Part II If there is a one-dimensional model that is observationally equivalent to the multidimensional spatial model, then (1.35) holds, which implies

$$\langle \Delta x_j, \theta_i - \theta_1 \rangle = \beta_j (p_i - p_1).$$

Take $i = 2$, then the projection of $\theta_2 - \theta_1$ on all Δx_j is given, so $\theta_2 - \theta_1$ is uniquely pinned down in the space spanned by $\{\Delta x_j\}$. For any other $i > 2$, $\theta_i - \theta_1$ is $\beta_j(p_i - p_1) = \frac{p_i - p_1}{p_2 - p_1} \langle \Delta x_j, \theta_i - \theta_1 \rangle$. Thus, the projection of $\theta_i - \theta_1$ on the space spanned by $\{\Delta x_j\}$ is a multiple of the one of $\theta_2 - \theta_1$. Thus, property L holds. □

Sufficiency of First-Order Conditions for Equilibrium Platforms Recall that in the context of Theorem 3, all our *equilibrium candidates*, that is pairs of platforms (x_a, x_b) that satisfy the necessary first-order conditions of optimality, are of the form

$$(x_a, x_b) = \alpha(x_a^*, x_b^*)$$

with $\alpha \in [0, 1]$. It remains to show these platforms are indeed best responses to each other, that is, the first-order conditions are sufficient for optimality in these cases.

First, we show in Lemma 19 that it is sufficient for equilibrium candidates to be equilibria that party objectives are quasi-concave on certain compact subsets of $\mathbb{R}^n$. Second, we show in Lemma 20 that when the weight on vote share m is small enough or the valence shock ν is large enough, then this condition for quasi-concavity is satisfied. By comparison to existing results (Lindbeck and Weibull, 1987; Enelow and Hinich, 1989), our proofs are complicated by the fact that voter ideal points are not bounded because we assume a normal distribution in Theorem 3.

For the first part, we make use of the fact that any platform choice x_j outside the ellipse defined by $u(x_j, x_j^*) \geq -m$ is suboptimal, as we have shown at the beginning of the proof of Lemma 1. To formulate the sufficient condition for equilibrium candidates to be equilibria,

define

$$\mathcal{E}_a := \{x \in \mathbb{R}^n | (x - x_a^*)^\top A(x - x_a^*) \leq m\}$$
$$\mathcal{E}_b := \{x \in \mathbb{R}^n | (x - x_b^*)^\top A(x - x_b^*) \leq m\}.$$

Further, recall that because of the normal prior $\mu = \mathcal{N}(0, \Sigma)$ and the restriction to normal signals, the distribution ρ of posterior means is necessarily normal. Also, it is supported on the line through the origin with direction $\Sigma A(x_b - x_a)$. By the law of total variance, the variance of the normal distribution is bounded by the variance of the prior in that direction. Let R the set of all distributions ρ satisfying these three requirements.

With these definitions, we can state the sufficient (but not necessary) condition for equilibrium candidates to be equilibria.

Condition 1. *The following holds.*

- $U_a(x_a, x_b, \rho)$ *is quasi-concave in* x_a *on* $\mathcal{E}_a$ *for all* $x_b = \alpha x_b^*$ *with* $\alpha \in [0, 1]$ *and* $\rho \in R$.
- $U_b(x_a, x_b, \rho)$ *is quasi-concave in* x_b *on* $\mathcal{E}_b$ *for all* $x_a = \alpha x_a^*$ *with* $\alpha \in [0, 1]$ *and* $\rho \in R$.

Lemma 19. *If Condition 1 holds, then all equilibrium candidates are equilibria.*

Proof. For an equilibrium candidate, $(x_a, x_b) = \alpha(x_a^*, x_b^*)$ with $\alpha \in [0, 1]$, party a's best response is in the ellipse $\mathcal{E}_a$, as argued above. Quasi-concave utility over $\mathcal{E}_a$ under the equilibrium x_b and any feasible ρ implies that the first-order condition is sufficient for optimality of x_a. The same holds for x_b. Thus, the equilibrium candidate (x_a, x_b, ρ) is an equilibrium as both platforms are best responses. □

Next, we give assumptions that ensure that Condition 1 holds. Lemma 20 uses the following assumption on the density of the valence shock, which is satisfied for example by the normal density and the Laplace or double exponential density. This assumption is far from necessary but it suffices to show certain terms vanish faster than a polynomial term diverges, which we use in our proof.

Assumption 1. *The density of the valence shock $f_\nu(x)$ is proportional to* $\exp\{-g(x)\}$ *where* $g(x)$ *is of the form*

$$g(x) = c_0 - c_1|x| - c_2x^2 - c_3|x|^3 - ... - c_m \cdot \begin{cases} |x|^m & \text{if } m \text{ odd} \\ x^m & \text{if } m \text{ even} \end{cases}$$

with $c_1, ..., c_m \geq 0$.

In this definition we take the absolute value of the odd polynomial terms, as present for example in the Laplace density, to ensure that the valence shock is symmetric. We assume that the constants c_1 through c_m are positive to ensure that the density is quasi-concave.

Lemma 20. *Under Assumption 1, if the weight on vote share m is small enough or if the valence shock ν is large enough, then Condition 1 holds. Formally, there exist $\underline{m} > 0$, such that for all m with $0 < m < \underline{m}$, Condition 1 holds. Given a valence shock ν that satisfies our assumptions, there exist $K > 0$, such that for all $k > K$, Condition 1 holds under valence shock $k\nu$.*

Proof. We show the Hessian of the party objective $U_a(x_a, x_b, \rho)$ in x_a is negative definite over $\mathcal{E}_a$ when $x_b = \alpha x_b^*$ with $\alpha \in [0, 1]$ and $\rho = \mathcal{N}(0, \Sigma_\rho)$ with $\Sigma_\rho \leq \Sigma$. This implies concavity and therefore quasi-concavity. The proof for concavity of $U_b(x_a, x_b, \rho)$ in x_b is analogous by symmetry.

Let ∇ denote the gradient with respect to x_a. The Hessian $H_a(x_a, x_b, \rho, m)$ of party a's objective with respect to x_a is

$$\begin{aligned} H_a(x_a, x_b, \rho, m) &:= \nabla^2 U_a(x_a, x_b, \rho) = m \int \nabla^2 F_\nu(\Delta u(\theta, x_a, x_b)) d\rho(\theta) + \nabla^2 u(x_a, x_a^*) \\ &= m \int \nabla\big(\nabla u(x_a, \theta) f_\nu(\Delta u(\theta, x_a, x_b))\big) d\rho(\theta) - I_n \\ &= m \int \left(-I_n f_\nu(\Delta u(\theta, x_a, x_b)) + \nabla u(x_a, \theta) \nabla u(x_a, \theta)^\top f'_\nu(\Delta u(\theta, x_a, x_b))\right) d\rho(\theta) - I_n \\ &= m \int \left(-I_n f_\nu(\Delta u(\theta, x_a, x_b)) + 4(\theta - x_a)(\theta - x_a)^\top f'_\nu(\Delta u(\theta, x_a, x_b))\right) d\rho(\theta) - I_n \end{aligned}$$

where

$$\Delta u(\theta, x_a, x_b) := u(x_a, \theta) - u(x_b, \theta).$$

We were able to exchange integration and differentiation because the derivative and second derivative are bounded (componentwise) by a constant, which is integrable under the probability measure ρ. We can bound both derivatives by constants because they consist of polynomial terms multiplied by an exponential function (with a decreasing polynomial exponent), so the integrands are eventually radially decreasing by Assumption 1. Thus, the supremum of the integrand is obtained on a compact sphere, on which the integrand obtains its finite maximum by continuity.

If $m = 0$, the Hessian $H_a(x_a, x_b, \rho, m)$ is $-A$, which is negative definite, so the objective is concave. To show for m small enough, the objective is concave, we first prove that the Hessian $H_a(x_a, x_b, \rho, m)$ is continuous in (x_a, x_b, ρ, m). For that we endow the domain of ρ, $\Delta(\mathbb{R}^n)$, with the weak topology on $\Delta(\mathbb{R}^n)$.

First, we show continuity of the Hessian in (x_a, x_b, ρ). Let the sequence $(x_a^n, x_b^n, \rho^n, m)_{n\in\mathbf{N}}$ converge to (x_a, x_b, ρ, m). We show that

$$\begin{aligned}
&\lim_{n\to\infty} \left(H_a(x_a, x_b, \rho, m) - H_a(x_a^n, x_b^n, \rho^n, m)\right) \\
&= \lim_{n\to\infty} \left(H_a(x_a, x_b, \rho, m) - H_a(x_a, x_b, \rho^n, m)\right) + \lim_{n\to\infty} \left(H_a(x_a, x_b, \rho^n, m) - H_a(x_a^n, x_b^n, \rho^n, m)\right) \qquad (1.36) \\
&= 0 + 0 = 0.
\end{aligned}$$

The first term of (1.36) is the difference between the integral of

$$m\left(-I_n f_\nu(\Delta u(\theta, x_a, x_b)) + 4(\theta - x_a)(\theta - x_a)^\top f'(\Delta u(\theta, x_a, x_b))\right)$$

with respect to ρ^n and with respect to ρ as $n \to \infty$. The limit is zero componentwise because the integrand is bounded (as we argued above), the integrand is continuous in θ, and ρ^n converges to ρ in the weak topology.

The limit of the second term of (1.36) is zero because the integrand is Lipschitz continuous in (x_a, x_b) with respect to the distance d, say induced by the L^1-norm. Lipschitz continuity with Lipschitz constant C implies that we can bound the term by

$$\lim_{n\to\infty} m \int Cd((x_a^n, x_b^n), (x_a, x_b))d\rho^n = \lim_{n\to\infty} mCd((x_a^n, x_b^n), (x_a, x_b)) = 0.$$

Lipschitz continuity follows from the gradient of the integrand in (x_a, x_b) being bounded componentwise, which follows analogously to how we showed above that the integrand is bounded.

If $H_a(x_a, x_b, \rho, m)$ is continuous in (x_a, x_b, ρ), then it is jointly continuous in (x_a, x_b, ρ, m) as m simply multiplies the integrand. Then, $v^\top H_a(x_a, x_b, \rho, m)v$ is jointly continuous in those variables and v.

Using continuity of the Hessian, we show for m small enough, the Hessian is negative definite for $x_a, x_b \in D$ and $\rho \in R$. Recall that a matrix $H \in \mathbb{R}^{n\times n}$ is negative definite if for all $v \in \mathbb{R}^n$, $v^\top Hv \leq 0$. Given m, choose the $x_a, x_b \in D$, σ_ρ^2 with $\sigma_\rho^2 \leq 1$, and $v \in \mathbb{R}^n$ with $\langle v, v\rangle = 1$ to maximize $v^\top H_a(x_a, x_b, \rho, m)v$. By the above, $v^\top H_a(x_a, x_b, \rho, m)v$ is continuous in $(x_a, x_b, \sigma_\rho^2, v, m)$ and the choice set is compact. At $m = 0$, we have $H_a(x_a, x_b, \rho, m) = -I_n$, so the value is -1 irrespective of the choice of (x_a, x_b, ρ, v), so the maximum is also -1. By Berge's maximum theorem, the value function is continuous, so for some $\underline{m}$ the value function crosses zero for the last time and is negative for $m < \underline{m}$. This implies that for $m < \underline{m}$, $v^\top H_a(x_a, x_b, \rho, m)v < 0$ for all $v \in \mathbb{R}^n$, so the Hessian is negative definite for all $x_a, x_b \in D$ and $\rho \in R$. This implies concavity of $U_a(x_a, x_b, \rho)$ in x_a for all $x_a, x_b \in D$ and all $\rho \in R$.

A similar argument shows that scaling up the valence shock ν by a large enough factor k makes the party objective concave. Valence shock $k\nu$ has the density $\frac{1}{k}f_\nu(\frac{x}{k})$. Reparametrizing by $c = 1/k$, we get the density $cf_\nu(cx)$ and the derivative of the density being $c^2 f'_\nu(cx)$. At $c = 0$, the Hessian is thus $-I_n$, which is negative definite, for all x_a, x_b, ρ. To show for c small enough (or, equivalently, k large enough), the Hessian is negative definite, again, let a fictitious adversarial agent choose $x_a, x_b \in D$, $\rho \in R$, and $v \in \mathbb{R}^n$ with $\langle v, v\rangle = 1$ to maximize $v^\top H_a(x_a, x_b, \rho)v$. Again, by continuity we can apply Berge's maximum theorem to obtain that the Hessian is negative definite for all $x_a, x_b \in D$ and $\rho \in R$ for c small enough.[40] □

[40]The arguments above use that for m small enough or ν large enough, the party objective becomes dominated by the ideological motive, which is concave, while the vote share motive becomes arbitrarily small. Additionally, one could show for large enough valence shock ν, the vote share alone becomes concave. This follows from the fact that in the preceding paragraph, the density scales by c while the derivative of the density scales by c^2. Thus, the negative definite density term in the integrand dominates the integrand for small enough c.

Chapter 2

Attention and Regret

> *[Regret] may also depend on the extent to which the individual blames himself for his original decision. [...] The neglect of this dimension of regret—although a useful simplifying assumption for many problems—is a serious obstacle to the development and generalisation of regret theory.* - Sugden, 1986

2.1 Introduction

Decisions under uncertainty can trigger regret and rejoicing, when it is learned that another action would have done better or worse. These emotions are often considered futile "crying over spilled milk" and their anticipation has been shown to distort choices, posing an puzzling question: Why do such emotions exist if only the actual outcomes should matter to an agent, particularly from an evolutionary perspective?[1]

Although the anticipation of regret and rejoicing can distort decision-making, we show that these emotions can nevertheless improve the quality of decisions if we take into account *attention*. When facing uncertainty, individuals can improve decision-making by paying costly attention, which we understand broadly as encompassing both external information

[1]Bell (1982) and Loomes and Sugden (1982) study the behavioral implications of regret and rejoicing theoretically. Evidence that their anticipation affects choices is reviewed in Zeelenberg (1999a).

acquisition and internal information processing such as deliberation, planning, or memory search. The anticipation of regret and rejoicing can motivate individuals to pay more attention to avoid regret and seek rejoicing. Such an additional motive to pay attention is beneficial, if the agent is biased toward insufficient attention due to the temptation to avoid attention effort or due to present bias, as benefits are only realized after the cost of attention is incurred.

Through a model, we show that, among all conceivable emotions as a function of the decision outcome, regret and rejoicing are indeed the optimal ones to counteract a bias for insufficient attention. We take a principal-agent approach (e.g. Robson, 2001; Samuelson and Swinkels, 2006; Rayo and Becker, 2007) to study the optimal emotions that evolution, whether natural or cultural, would have endowed us with for that purpose. We allow the principal to choose *any* emotional payoffs as a function of the chosen action and ex-post feedback about the realized state of the world. Under complete feedback, the model endogenizes regret theory, including disproportionate aversion to large regrets. Intuitively, regret can be seen as a form of *self-reproach*, punishing the agent for shirking on attention, and rejoicing as self-congratulation for the converse. Coherent with this interpretation, our model predicts that regret is stronger following a mistake in simple decision problems than in harder ones. The model also explains why the optimal mix of self-punishments and rewards is tilted towards the former, that is, why regret is stronger than rejoicing. Finally, we show that under incomplete feedback, disappointment and elation are additionally used as second-best motivators. Disappointment and elation also arise from counterfactual comparisons, but, in contrast to regret and rejoicing, they concern counterfactual outcomes had another *state* realized.

In our model, the principal, whether nature or a parent, delegates decision-making including attention (or information acquisition, depending on the interpretation) to an agent. Both parties care about the attention cost as well as the decision outcome but the agent overweights the attention cost, which we refer to as the agent's bias. The principal can incen-

tivize the agent by endowing them with emotions, which enter the agent's utility. Emotions can be chosen as a function of the chosen action and feedback on the realized state, which captures two constraints (discussed in more detail in section 2.2). First, emotions do not depend on the amount of attention paid, which we assume is difficult to verify, inducing a moral hazard concern. Second, we allow for an uncertain bias of the agent, on which emotions cannot condition and study the resulting effects of adverse selection. Such uncertainty may be due to the stochastic nature of temptation or due to heterogeneity in present bias. To investigate the incentive problem pertaining not just the amount but also the type of obtained information, we allow the agent to pay attention by acquiring *any* signal structure subject to an entropy-based cost as in the rational-inattention literature (Sims, 2003).

The principal-agent problem we study is complicated by the presence of both adverse selection and moral hazard, alongside the high dimensionality of the agent's strategy due to flexible information acquisition. To render the problem tractable, we impose two symmetry assumptions. First, we assume that the delegated decision problem is symmetric, that is, invariant under relabeling of actions. Second, we confine ourselves to mechanisms where the emotional payoffs preserve this symmetry of the decision problem.

We find that the optimal emotions are *regret* and *rejoicing*. Both arise from counterfactual thinking, as they result from comparisons to what would have happened, had one chosen differently. Formally, let $v(a, \omega)$ denote the payoff under action a and state ω. Regret or rejoicing experienced in state ω and under action a is a function of the vector of payoff differences $v(a, \omega) - v(a', \omega)$ to alternative actions a' under the realized state ω. The anticipation of regret and rejoicing increases the agent's incentive to pay attention, as the agent aims not only to achieve a better decision outcome but also to avoid regret and seek rejoicing. While various emotions can motivate attention by increasing the stakes of the decision problem, regret and rejoicing *emerge as optimal* in our model. This is because how much attention the agent pays to a certain state is determined by the utility differences between actions in that state. Under an unknown bias, it turns out how much the principal wishes

to amplify these utility differences, depends on the preexisting payoff differences, exclusive of emotions. Regret and rejoicing, as defined above, achieve precisely this amplification by being functions of the payoff differences to counterfactual actions.

Focusing on decision problems with two actions, which have received significant attention in the regret literature, we show our main results. The optimal mechanism *endogenizes* classical regret theory (Bell, 1982; Loomes and Sugden, 1982) including disproportionate aversion to large regrets. This aversion, crucial to regret theory, leads to common deviations from expected utility theory, such as the Allais paradox, the common ratio effect, and simultaneous gambling and insurance. Our model predicts a more general version of this property than assumed in aforementioned studies. However, we demonstrate that our more general version is equivalent to the conjunction of above-mentioned deviations from expected utility. Restricting ourselves to symmetric decision problems, we show that while these choice distortions do not come into play, attention is distorted. We show that disproportionate aversion to large regrets skews attention towards extreme events, or states where the payoff difference between the actions is greater.

Our model advances regret theories in several dimensions. First, it addresses the self-recrimination aspect of regret, which is inadequately captured in classical regret theory, as argued by one of its pioneers (Sugden, 1986). One implication is that a realistic theory of regret should take into account whether one could have known that another action would turn out better (ibid.). In our framework, regret and rejoicing can be seen as punishments and rewards for paying (insufficient) attention. Indeed, the model predicts that when it was less costly to find out the right action, mistakes are more severely punished by regret. Second, existing models typically focus solely on regret, overlooking rejoicing, or are agnostic about the relative intensity of both emotions (e.g. Bell, 1982; Loomes and Sugden, 1982; Quiggin, 1994; Sarver, 2008). In contrast, our model predicts both regret *and* rejoicing, while

confirming the intuition that regret is stronger than rejoicing.[2] Third, our model extends regret theory to more than two actions, which has proven difficult for classical regret theory (Loomes and Sugden, 1982; Loomes and Sugden, 1987). A frequently used model for general choice sets considers only regret to the ex-post best action (Quiggin, 1994). Our model makes the psychologically plausible prediction that regret and rejoicing are affected by all counterfactual actions. In particular, under a known bias, each counterfactual action affects regret and rejoicing in proportion to the probability with which the action should have been taken according to the first-best attention strategy (Theorem 10). This probability provides a notion of choiceworthiness and addresses the other main point in Sugden (1986): self-blame is affected by another action to the "extent to which an action is a serious candidate for choice". Finally, we extend regret theory to explicitly incorporate the role of feedback, to which we turn next.

In an extension of the model, emotions can be chosen as any function of the chosen action and *incomplete feedback* about the state. With complete feedback, the principal can identify not only the realized payoff but also counterfactual payoffs from alternative actions, which regret relies on. However, this is not always plausible. For instance, if the agent is selecting between different economic projects, only the return of chosen project and general market conditions may be observable ex post, but not the return of alternative projects. To model such situations, we assume the principal can condition emotions only on the payoff of the chosen action, as well as on a partitional signal of the realized state, such as market conditions.

Under incomplete feedback and an unknown bias, all emotional payoffs have to be determined jointly, making the solution analytically intractable. Therefore, we simplify by assuming a known bias, which can be perfectly offset by a combination of regret and rejoicing with another set of emotions, namely *disappointment* and *elation*. Our prediction aligns

[2]Although it turns out that the distinction between regret and rejoicing does not matter for choice behavior (e.g. Loomes and Sugden, 1982), it has implications for welfare comparisons and menu choice (Sarver, 2008).

with empirical evidence, that regret and rejoicing are predominant when there is complete feedback post-decision, while disappointment and elation gain prominence when only the realized payoff is known (Zeelenberg, 1999a; Camille et al., 2004). We are to our knowledge the first to incorporate both sets of emotions as a function of the feedback structure. Furthermore, under a known bias, the result is robust to discarding the above-mentioned symmetry assumptions and to generalizing the entropy-based attention cost to any posterior-separable cost function.

This paper also has implications for how to evaluate the emotions of regret and disappointment. Knowing their benefits and costs is useful, especially when we can influence these emotions through socialization or therapy. In this context, our work contributes to a broader debate on whether emotions are adaptive (Keltner and Gross, 1999). More specifically, Camille et al. (2004) find that patients who do not experience regret as the result of a damage to their orbitofrontal cortex make worse decisions under risk, in line with the hypothesis that regret is an adaptive emotion to improve decision-making (e.g. Janis and Mann, 1977; Zeelenberg, 1999b; O'Connor et al., 2012). The existing well-developed hypothesis as to *how* regret and disappointment improve decision-making is that they serve as heuristics for learning. When an agent faces a decision under uncertainty repeatedly, reinforcement learning with regret (Foster and Vohra, 1999) or disappointment and elation (Schultz et al., 1997) has desirable asymptotic properties. By contrast, we maintain the assumption of Bayesian learning and show that the *anticipation* of regret and disappointment can serve as optimal motivators for costly learning. While we believe that these two approaches can be complementary, our approach is able to explain more properties of regret, such as disproportionate aversion to large regrets and the relevance of the difficulty of the decision problem as well as of feedback.

Related Literature Regret and disappointment are the primary emotions studied in the context of decision-making under uncertainty (Zeelenberg et al., 2000; Loewenstein and Lerner, 2003). Regret theory, formalized by Bell (1982) and Loomes and Sugden (1982) for

decision problems with two actions, is endogenized in our paper. Bleichrodt and Wakker (2015) review evidence supporting regret theory and its applications to financial and health-related decisions. In contrast to the above-mentioned papers, we endogenize regret and disappointment from their hypothesized function as internal self-control mechanisms.

At a broader level, our work contributes to a large body of research, initiated by Demski and Sappington (1987), that explores how a principal should incentivize an expert to pay attention, or, in the terminology of that literature, to acquire information. Our model distinguishes itself by incorporating flexible information acquisition and both moral hazard, due to the acquired information not being contractible, and adverse selection, due to the agent's bias being private information of the agent.[3]

While Osband (1989) also addresses both moral hazard and adverse selection, his model is confined to a more specific setup. The principal employs an expert to predict a variable under a quadratic loss, with the expert drawing independent observations at a cost unknown to the principal. He concludes that the optimal contract is quadratic in prediction, linear in realization, and that self-screening offers little value. Our model imposes fewer constraints on the decision problem and allows the agent to acquire any signal structure at a posterior-separable cost, as seen in more recent contributions to this field (Rappoport and Somma, 2017; Carroll, 2019; Lindbeck and Weibull, 2020; Yoder, 2022; Whitmeyer and Zhang, 2022). In this regard, we leverage advances in rational inattention and information design, particularly those of the posterior approach (Caplin and Dean, 2013; Caplin et al., 2022), using concavification methods from Bayesian persuasion (Kamenica and Gentzkow, 2011), as well as the state-dependent choice probabilities approach (Matějka and McKay, 2015).

Yoder (2022) studies adverse selection and flexible information acquisition in a model with two states, but he assumes that the information-acquisition strategy, or at least the realized signal, is contractible. In contrast, our model only considers the chosen action and

[3]This type of adverse selection at the outset is to be distinguished from adverse selection in the interim, resulting from the realization of information acquisition being private information of the agent. The latter type of adverse selection, which is the focus of Thereze (2022a), is also present in our model.

realized state as contractible. We show that the principal cannot make use of self-screening and hence our solution methods and the resulting distortions differ: in his model, there is no distortion at the top, whereas in our case, information acquisition is skewed towards extreme states for all types.

Another difference from above-mentioned papers is that in our model the agent has a *direct interest* in the decision, albeit being biased towards insufficient information acquisition. This feature is shared with Szalay (2005), but in his paper the bias is known and the main result concerns the optimal mechanism when the principal has no access to transfers. He shows that, absent transfers, the principal can incentivize the agent by restricting the choice set and allowing only extreme actions.

This paper is part of a larger project that relates the optimal mechanism in the face of an uncertain bias to non-standard preferences. In Vaeth (2022), we investigate the optimal emotions when the agent's bias does not affect attention but the decision-making stage, for instance, because some actions demand more effort. We give conditions for when the induced utility function replicates loss-aversion around a reference point or goal and when mental accounting is used as an incentive instrument. Our work has relevance to the literature on contract theory with behavioral agents (Kőszegi, 2014), specifically O'Donoghue and Rabin (2006), Lockwood (2020) and Beshears et al. (2020), who explore optimal mechanisms to disincentivize overconsumption and incentivize effort and savings, respectively, for agents with heterogeneous present bias.

The paper is organized as follows. Section 2.2 introduces and discusses the model. Section 2.3 solves the model when the principal has complete feedback about the realized state and the bias is unknown. Section 2.7 extends the model to incomplete feedback about the state, when the bias is known. Section 2.5 concludes.

2.2 Model

Decision Problem The principal (evolution, she) faces a *decision problem*, which is a tuple (Ω, μ, A, v). It comprises a finite state space Ω, a full-support prior $\mu \in \Delta(\Omega)$ on the state space, a finite set of actions A, and a utility function $v\colon A \times \Omega \to \mathbb{R}$, which we refer to as the payoff.[4]

Attention The principal delegates a two-stage problem to the agent (he). First, he acquires a signal structure about the state (the attention stage) and, second, he chooses an action from A conditional on the realized signal (the decision-making stage). A signal structure is a pair (S, σ) where S is a finite signal space and $\sigma\colon \Omega \to \Delta(S)$ is the state-dependent distribution over signals. Following the rational-inattention approach (Sims, 2003), we assume that the agent can select *any* signal structure subject to the *mutual information* cost, defined below. It is well known (e.g. Matějka and McKay, 2015) that the agent's signal can be identified with an action recommendation, reducing the agent's two-stage problem to the selection of state-dependent choice probabilities $p\colon \Omega \to \Delta(A)$, written $p(a|\omega)$. The state-dependent choice probabilities p, which we also refer to as an *attention strategy*, entail an attention cost $c(p)$ that corresponds to the mutual information of the action and state,

$$c(p) := \sum_{a \in A} -p(a) \log p(a) - \sum_{\omega \in \Omega} \sum_{a \in A} -\mu(\omega) p(a|\omega) \log p(a|\omega).$$

Here, $p(a) := \sum_{\omega \in \Omega} \mu(\omega) p(a|\omega)$ is the unconditional probability of a and log refers to the natural logarithm. The first sum is the entropy of the unconditional choice probabilities, while the second is the expected entropy of the conditional choice probabilities. Intuitively, entropy quantifies the uncertainty of a distribution and mutual information measures how much knowing the state reduces the uncertainty about the action. Hence, the attention cost reflects the extent to which the action is tailored to the state.

[4] We assume a finite state space to circumvent measure theoretical issues. We conjecture that results carry over to infinite state spaces under the appropriate technical assumptions.

Preferences The principal's utility V derived from an attention strategy p is

$$V(p) = -\lambda c(p) + \sum_{\omega\in\Omega}\sum_{a\in A}\mu(\omega)p(a|\omega)v(a,\omega).$$

That is, she internalizes the attention cost, weighted by λ, and derives utility $v(a,\omega)$ from action a under state ω. We interpret λ as the difficulty of the decision problem, as it scales the cost of learning about the state. This parameter will be used for performing comparative statics.

As in the literature on delegation, the agent has a bias, represented here as a discount factor β between 0 and 1 applied to the utility $v(a,\omega)$. Equivalently, this bias can be viewed as an overweighting of the attention cost by a factor of $1/\beta$. In our main interpretation, the bias captures a present bias, or the temptation to avoid the attention cost.

We depart from standard delegation models in two respects. First, the bias β is uncertain and private information of the agent, with cumulative distribution function F. This uncertainty can be due to incomplete knowledge of the agent's fixed bias as in models of self-signaling (Bénabou and Tirole, 2004) or due to the stochastic nature of the bias, as in random-indulgence models (Chatterjee and Krishna, 2009; Duflo et al., 2011; Dekel and Lipman, 2012).[5] Second, the principal can use action- and state-contingent payoff-transfers $m(a,\omega)$ to incentivize the agent. We interpret these transfers as *emotions*, which can depend on the chosen action and the realized state.[6] The agent with bias β derives utility U_β from the attention strategy p,

$$U_\beta(p) = -\lambda c(p) + \beta\sum_{\omega\in\Omega}\sum_{a\in A}\mu(\omega)p(a|\omega)\big(v(a,\omega)+m(a,\omega)\big).$$

[5] In the terminology Dekel and Lipman (2012), our agent is a random Strotz type (as opposed to a random Gul-Pesendorfer type).

[6] As common in many models, we assume that emotions are additively separable in utility (e.g. Bell, 1982, 1985; Loomes and Sugden, 1982, 1986). We also make this assumption because it allows the principal to implement any emotion-inclusive utility function $u(a,\omega) = v(a,\omega)+m(a,\omega)$. Up to determining a state-wise constant, this is assumption is immaterial for the model's prediction of u.

Equivalently, the agent maximizes

$$\frac{U_\beta(p)}{\beta} = -\frac{\lambda c(p)}{\beta} + \sum_{\omega\in\Omega}\sum_{a\in A}\mu(\omega)p(a|\omega)\big(v(a,\omega)+m(a,\omega)\big).$$

Principal's Problem and Tie-Breaker Thus far, the model can only identify optimal emotions up to a state-wise constant. Given that the agent is an expected-utility maximizer, shifting all emotional payoffs in a given state by a constant does not alter his behavior.[7] Moreover, emotions do not enter the principal's utility directly because, as usual, evolution just cares about the induced behavior. Thus, the principal is indifferent between any shifts of emotions by state-wise constants. Yet, while such constants are not relevant for the agent's behavior, they do matter for welfare comparisons.

To resolve such indifferences, we introduce a refinement based of an implementation cost $\psi\colon \mathbb{R} \mapsto \mathbb{R}$ for positive as well as negative emotions. From the possible solutions to the principal's problem, we refine the principal's choice to the mechanism that minimizes the expected cost of emotion $\Psi(m)$,

$$\Psi(m) := \int \left(\sum_{\omega\in\Omega}\sum_{a\in A}\mu(\omega)p_\beta(a|\omega)\psi(m(a,\omega))\right) dF(\beta)$$

where ψ is strictly convex, differentiable and symmetric around 0.[8,9] The cost of emotions $\psi(m)$ could be due to the energy used to induce emotions in the brain or due to the opportunity cost from diminished capacity for other cognitive processes. Under the cultural-evolution interpretation, the cost can be interpreted as the cost of inculcating emotions through parenting and socialization, similar to Kaplow and Shavell (2007).

[7]This is a subtle issue. As in regret theory, given emotions, the agent is an expected-utility maximizer. However, since emotions can be a function of the choice set, the agent's choice behavior can violate expected-utility axioms.

[8]Instead of a cost of emotions, one could also assume a budget constraint that emotions have to be 0 in expectation, perhaps because of hedonic adaptation. It turns out that this budget constraint is implied by assuming a quadratic cost of emotions $\psi(m) = m^2$.

[9]If one assumed instead that the principal internalizes the agent's emotions positively, equivalent to setting $\psi(m) = -m$, then there would be no solution because the principal would prefer infinitely positive emotions.

Assumption 2. *The cost of emotions $\psi\colon \mathbb{R} \to \mathbb{R}$ is strictly convex, differentiable, and symmetric around 0.*

Note that under strict convexity, symmetry around 0 implies that the cost ψ is minimal at 0. Except for Corollary 8, which shows that regret is stronger than rejoicing, all our results hold assuming the cost is minimal at 0 without assuming symmetry.

To ensure that the optimal incentives for attention shape the payoff structure rather than the implementation cost, we treat the minimization of the cost of emotions as a tie-breaker, rather than directly incorporating it into the principal's utility function.

Interpretation and Discussion Our model aligns with the principal-agent approach to the evolution of preferences (e.g. Robson, 2001; Samuelson and Swinkels, 2006; Rayo and Becker, 2007). This approach likens evolution to a principal that chooses preferences to incentivize fitness-maximizing behavior while operating under certain contracting constraints. In our model, these constraints include: (1) the agent is inherently biased, and (2) the degree of this bias and attention are private information of the agent.

Regarding the first constraint, we assume that evolution cannot directly mitigate the bias, such as temptation. Instead, emotions like regret and disappointment function as kludges to improve decision-making. This assumption aligns with evidence for an incremental evolution of the mind, which adds new structures while leaving old structures intact (Linden, 2007; Marcus, 2008).[10] Furthermore, this view is consistent with standard views of self-control issues as manifestations of a divided self. Dual-self models, also known as planner-doer models or dual-system models of the brain, conceptualize the organism as a far-sighted self interacting with a short-sighted self (Thaler and Shefrin, 1981; Bénabou and Pycia, 2002; Bernheim and Rangel, 2004; Fudenberg and Levine, 2006, 2012; Brocas and Carrillo, 2008). These two selfs can be identified with evolutionary newer and older brain structures, such as the prefrontal cortex and the amygdala, which aligns with neuroscientific evidence of the prefrontal cortex exerting control over other brain structures (Miller and Cohen, 2001;

[10] Ely (2011) presents a theoretical foundation for why an organism might be kludged.

see also Brocas and Carrillo, 2008).[11] Compatible with this, neuroscientific research points toward a pivotal role of the prefrontal cortex in experiencing regret (Camille et al., 2004; Coricelli et al., 2005).

Regarding the second constraint, we assume that evolution cannot condition emotions on the acquired signal structure or on the stochastic bias β. This might stem from the fact that once the outcome materializes, it is easier to recall and verify actions taken than the intensity of temptation or the amount of attention paid before decision-making, as is assumed in models of self-signalling (Bodner and Prelec, 2003; Bénabou and Tirole, 2004).

Another, possibly complementary, interpretation of the principal is as cultural evolution or parents. Emotions may be culturally transmitted, for instance, through parental praise or reprimands for children's decisions, leading to the internalization of such feedback as self-congratulatory (rejoicing) or self-blaming (regret) emotions. Under this interpretation, the first constraint above amounts to assuming that culture or parents cannot affect an underlying present bias or temptation directly. The second constraint means that the degree of the bias and attention are private information of the child and parental feedback can condition only on the feedback about the decision outcome.

Finally, our model can be applied to contractual relationships, where information acquisition is delegated to an expert and payoff-transfers are monetary instead of emotions. For example, in corporate settings, shareholders entrust CEOs with the responsibility of selecting viable projects, and in financial markets, investors delegate the choice of investment portfolios to fund managers. For risk-neutral parties, we can interpret our transfers as monetary, subject to an ex-ante budget constraint. This constraint allows the principal to deviate from the budget as long as the the *expected* transfers do not exceed it. Among the contracts that optimize the agent's behavior and comply with the ex-ante budget constraint,

[11] While our main interpretation of the principal is evolution creating brain structures that induce regret, a dual-self interpretation is also possible. Novel compared to existing dual-self models is that our principal can use emotions as rewards or penalties to incentivize the agent. An exception is papers on goal-setting (Koch and Nafziger, 2011; Hsiaw, 2013), however, they constrain the choice of a utility function to specific functional forms.

we assume that the principal chooses the one that minimizes the variance of the ex-post budget imbalance. Under this interpretation, the agent's bias need not denote a behavioral bias but can also capture a divergence of interests between the principal and the agent. We discuss these two interpretations further in Appendix 2.4.

The reader who prefers a more standard economic interpretation of the model, is invited to reinterpret emotions in the following as monetary transfers.

2.3 Results

In this section, we maintain the assumption that the principal can condition emotions on the chosen action a and the realized state ω. The next section analyzes the contracting problem when only a partition of the state is contractible.

We assume that the distribution of the agent's type β is non-degenerate, so there is some uncertainty about the bias of the agent. Let $\text{supp}(\beta)$ denote the support of the distribution of the agent's type. We make the assumption that $\inf(\text{supp}(\beta)) > 0$, so we can show later that finite emotions are optimal. Further, we assume for convenience that the distribution of β is discrete, which simplifies some parts of the proof of Lemma 21 and Theorem 9.

Assumption 3. *The bias β non-degenerate, discrete distribution with* $\text{supp}(\beta) \subseteq [\underline{\beta}, 1]$ *where* $\underline{\beta} > 0$.

Due to the presence of both moral hazard and adverse selection in the principal's problem, the problem is very complex in general. In particular, the agent's choice of an attention strategy is an high-dimensional problem and generally does not have analytic solutions. To render the problem tractable, we make two symmetry assumptions. First, we impose that relative to the choice set A and prior μ, the utility v is *exchangeable* (Matějka and McKay, 2015). Second, we restrict ourselves to mechanisms where emotional payoffs maintain exchangeability of the utility. Exchangeable utility captures that the decision problem invariant under relabeling the actions. It helps to characterize the agent's attention strategy, which is

given by a multinomial logit formula (Matějka and McKay, 2015). This allows us to reduce the principal's problem to a state-wise optimization problem.

Definition 4 (Exchangeability)**.** *A utility function $u\colon A \times \Omega \to \mathbb{R}$ is exchangeable if the distribution over payoff vectors $(u(a,\omega))_{a\in A} \in \mathbb{R}^A$ induced by prior μ is invariant under permutation of the coordinates.*

As an example, the utility function according to the payoff-matrix in Table 2.1 is exchangeable if $\mu(\omega_1) = \mu(\omega_2)$. That is because the payoff vector $(0,1)$ is as likely as the payoff vector $(1,0)$. If either $\mu(\omega_1) \neq \mu(\omega_2)$ or if we changed one of the payoffs, then the utility would not be exchangeable anymore.

	ω_1	ω_2
a_1	1	0
a_2	0	1

Table 2.1: Exemplary payoff-matrix

Assumption 4. *The utility v is exchangeable.*

This symmetry assumption is natural, for example, when before any attention the actions are indistinguishable and arguably should hold at a sufficiently ex-ante stage. Under our main interpretation, this assumption is plausible if emotions are adapted to the general attention problem that the agent faces and not to specific contexts.

For tractability, we focus on emotions m that maintain exchangeability of the decision problem. Thus, the emotions do not depend on the labeling of actions. This symmetry assumption is natural because the underlying decision problem is symmetric in the actions, so the first-best attention strategy is symmetric in the actions.[12]

[12]Simulations under two states suggest that inducing asymmetries through emotions cannot benefit the principal. For a quadratic posterior-separable cost under two states, we can show that the asymmetric emotions are suboptimal (proof upon request).

We study the following principal's problem (P).

$$\max_{m,(p_\beta)_\beta} \int \left(-c(p_\beta) + \sum_{\omega\in\Omega} \sum_{a\in A} \mu(\omega) p_\beta(a|\omega) v(a,\omega) \right) dF(\beta) \tag{P}$$

s.t.

$$\forall \beta \colon p_\beta \in \arg\max_p -c(p) + \beta \sum_{\omega\in\Omega} \sum_{a\in A} \mu(\omega) p(a|\omega) \big(v(a,\omega) + m(a,\omega) \big) \tag{IC}$$

$$v + m \text{ is exchangeable.} \tag{EX}$$

Out of the solutions to (P), the principal chooses the one that minimizes the cost of emotions $\Psi(m)$.

If the agent's bias β was known, there would be a simple way to offset it through emotions. Any emotions that raise the stakes of the decision problem by $1/\beta$, for example $m(a,\omega) = \left(\frac{1}{\beta} - 1\right) v(a,\omega)$, perfectly align the interest of the agent and principal, since

$$\beta \big(v(a,\omega) + \left(\frac{1}{\beta} - 1\right) v(a,\omega) \big) = v(a,\omega).$$

Thus, different types require differently strong incentives to offset their bias.

Therefore, one might wonder whether by offering a menu of emotion schemes, the principal can tailor incentives to the agent's type.[13] Even more generally, the principal could elicit a message from the agent before and after attention prior to assigning an emotion scheme, to screen the agent. Assuming that every emotion scheme has to induce an exchangeable utility, we show that under two states, the principal cannot do better than offering a single emotion scheme, so the formulation of (P) above is without loss.[14]

Lemma 21. *Let $|\Omega| = 2$. Under Assumptions 1 to 3, any mechanism is weakly dominated by a single emotion scheme and the formulation* (P) *is without loss.*

[13] In analogy to the term "transfer schemes", we also use "emotion scheme" to refer to the function m.

[14] We conjecture that this Proposition generalizes to any (finite) state space.

All proofs are relegated to the Appendix. The proof shows that offering a menu makes types sort into emotion schemes to the principal's disadvantage. Ideally, the principal would want to offer stronger incentives to types with lower β than to types with higher β because the former group is more biased. However, higher types have a higher return to attention and pick stronger incentives than lower types, contrary to the principal's aim. This leads to higher types paying more attention and lower types less than under a single emotion scheme. Hence, offering a menu increases the variance of attention, decreasing the principal's utility.

First, we characterize the solution of (P) for decision problems containing two actions and then turn to the case of $n > 2$ actions.

2.3.1 Two Actions

Let $A = \{a_1, a_2\}$. To see whether the optimal emotions can be interpreted as regret and rejoicing, we define a regret solution of the principal's problem, which is based on classical regret theory in Bell (1982) and Loomes and Sugden (1982). For a_i, $i = 1, 2$, denote by a_{-i} the other action in A, so $a_{-1} = a_2$ and $a_{-2} = a_1$.

Definition 5 (Regret Solution)**.** *Under two actions, $A = \{a_1, a_2\}$, a solution $(m, (p_\beta)_\beta)$ to the principal's problem* (P) *is a regret solution if there exists a function $R\colon \mathbb{R} \to \mathbb{R}$ such that for all $(a_i, \omega) \in A \times \Omega$:*

$$m(a_i, \omega) = R\big(v(a_i, \omega) - v(a_{-i}, \omega)\big)$$

and

(1) $R(x) \gtreqless 0 \Leftrightarrow x \gtreqless 0$,

(2) $\frac{R(x)-R(-x)}{x}$ *is strictly increasing on* $\mathbb{R}_{>0}$.

Note that a regret solution does not only require that the agent's preferences can be represented by a utility function as in regret theory, but further that the emotional utility m itself can be written as the regret utility term in Loomes and Sugden (1982) and Bell

(1982). That is because, like the emotional utility term in regret theory, m should be interpretable as experienced utility. The function R is called a regret-rejoicing function in regret theory and captures the emotions of regret (if $v(a_i, \omega) - v(a_{-i}, \omega) < 0$) and rejoicing (if $v(a_i, \omega) - v(a_{-i}, \omega) > 0$). Thus, in regret theory the agent cares about the outcomes of actions but also anticipates and cares about the emotional consequences, which enter the utility function as an additive component.

Property 1 states that if the actual and the counterfactual payoff are the same ($x = 0$), the agent experiences neither regret nor rejoicing ($R(x) = 0$). If the actual payoff is better than the counterfactual payoff ($x > 0$), then the emotion is positive ($R(x) > 0$) and if the counterfactual payoff is better ($x < 0$), then the emotion is negative ($R(x) < 0$).[15]

Property 2 requires that the sum of regret and rejoicing for some payoff difference is increasing more than proportionally. It is often called disproportionate aversion to large regrets and captures that large payoff differences get overweighted by regret, which leads to violations of expected-utility axioms. In fact, Property 2 *captures all the behavioral consequences of regret* for attention and decision-making. To see that, first, note that Property 2 is equivalent to requiring that

$$Q(x) := x + R(x) - R(-x)$$

satisfies $\frac{Q(x)}{x}$ is strictly increasing on $\mathbb{R}_{\geq 0}$, by virtue of $\frac{Q(x)}{x} = 1 + \frac{R(x)-R(-x)}{x}$. Further, as the agent is an expected-utility maximizer, adding a state-dependent constant to utility does not change incentives for attention or decision-making. Hence, behavior is determined by the *stakes* in each state, that is the state-wise utility difference between the two actions.

[15] Loomes and Sugden (1982) make the stronger assumption that $R(x)$ is strictly increasing in x, that is they assume the better the actual compared to the counterfactual outcome, the better the emotion. The model of this paper does predict this property for negative x, but for large enough $x > 0$, $R(x)$ can be decreasing. In the context of the model, this can be interpreted as saying that for large utility differences, there is nothing to congratulate oneself for because any type would have paid enough attention to choose the right action with stakes that large.

Denoting the *emotion-inclusive utility*

$$u(a,\omega) := v(a,\omega) + m(a,\omega),$$

Q is exactly the function that maps the stakes without emotions,

$$\Delta v(a_i,\omega) := v(a_i,\omega) - v(a_{-i},\omega),$$

into the stakes including emotions,

$$\Delta u(a_i,\omega) := u(a_i,\omega) - u(a_{-i},\omega)$$

by

$$\begin{aligned}\Delta u(a_i,\omega) &= \Big(v(a_i,\omega) + R\big(v(a_i,\omega) - v(a_{-i},\omega)\big)\Big) - \Big(v(a_{-i},\omega) + R\big(v(a_{-i},\omega) - v(a_i,\omega)\big)\Big)\\ &= \Delta v(a_i,\omega) + R\big(\Delta v(a_i,\omega)\big) - R\big(-\Delta v(a_i,\omega)\big)\\ &= Q(\Delta v(a_i,\omega)).\end{aligned}$$

Property 2 is more general than the assumption in Bell (1982) and Loomes and Sugden (1982) that Q is strictly convex, which implies that $\frac{Q(x)}{x}$ is strictly increasing on $\mathbb{R}_{>0}$.

Despite Property 1 and 2 being less restrictive than the assumptions in Bell (1982) and Loomes and Sugden (1982), we show in the Appendix 2.6, Proposition 5, that our weaker assumptions are still sufficient for many predictions of regret theory, notably the Allais paradox, the common-ratio effect, the isolation effect, and simultaneous gambling and insurance. As we show, Property 2 is in fact equivalent to the conjunction of these effects in the context of regret theory.

The following theorem shows that principal's problem indeed has a regret solution. This is remarkable, since one could expect that the optimal emotions under some outcome depend on the whole decision problem. That is they could depend on the the payoff v of every

outcome and on the prior. However, they follow the simple regret formula, which means that emotions depend only on the difference of the realized payoff to the payoff had one taken the other action, and they do so following a regret-rejoicing function, which does not depend on the state nor on the prior. However, the regret-rejoicing function depends on λ and on the distribution of β, which we use for comparative statics below.

Theorem 8. *Under a binary decision problem and Assumptions 1 to 3, the principal's problem* (P) *has a regret solution. The regret-rejoicing function R of such a solution is unique almost everywhere on $\mathbb{R}$.*

The intuition for why regret and rejoicing are optimal is as follows. One can show that under exchangeability, the principal's problem separates by state, so the optimal emotions depend only on the payoffs in the realized state. In any state, the agent's stochastic choice probabilities are given by a multinomial logit function (Matějka and McKay, 2015):

$$p_\beta(a_i|\omega) = \frac{e^{\frac{\beta}{\lambda}\Delta u(a_i,\omega)}}{1 + e^{\frac{\beta}{\lambda}\Delta u(a_i,\omega)}} \tag{2.1}$$

This illustrates that the optimal attention strategy is such that the probability with which the agent picks the right action in some state depends on the (emotion-inclusive) stakes in that state. The higher the stakes, the more attention the agent pays to make sure that he takes the right action under this state. What regret and rejoicing do is exactly to increase the stakes: regret penalizes the wrong action and rejoicing rewards the right action, state by state. This explains Property 1.

The solution endogenously satisfies Property 2, or disproportionate aversion to large regrets, by the following. As shown above, Property 2 is equivalent to $Q(x)/x$ being strictly increasing on $\mathbb{R}_{>0}$, where Q maps, in each state, the stakes without emotions Δv into the emotion-inclusive stakes $\Delta u = Q(\Delta v)$. If β was known, the optimal solution would be

$$Q(\Delta v) = \frac{\Delta v}{\beta},$$

since it perfectly offsets the agent's discounting of Δv. Under unknown β, the optimal Q increases the stakes in some state to offset a *weighted average* of the bias. The average is taken with respect to the distribution of types weighted by how much these types' choice probabilities are *sensitive* to the stakes. Formally, the optimal $Q(\Delta v)$ is implicitly characterized by the ratio of v and an endogenous weighted average of the bias,

$$Q(\Delta v) = \frac{\Delta v}{\frac{\int p'_\beta(Q(\Delta v))\beta dF(\beta)}{\int p'_\beta(Q(\Delta v))dF(\beta)}}, \tag{2.2}$$

where

$$p'_\beta(Q(\Delta v)) = \frac{d}{dQ(\Delta v)} \frac{e^{\frac{\beta}{\lambda}Q(\Delta v)}}{1 + e^{\frac{\beta}{\lambda}Q(\Delta v)}},$$

For higher stakes, higher types are comparatively less sensitive to the stakes because they already take the right action with probability close to 1. Thus, the emotions are more targeted to the low types, which require proportionally higher emotions to offset their bias. This leads to disproportionate aversion to large regrets.

In general, equation (2.2) can have multiple solutions. We show that the solutions are increasing in the strong set order and thus wherever Q is non-unique it has an upward jump. The non-decreasing function Q can have only countably many jump discontinuities, from which we show that R is unique up to at most countably many points.

Finally, the principal's problem can have solutions that are not regret solutions. When (2.2) has multiple solutions and there are duplicate states, she could pick from the solution set as a function of the state. Then, emotions are not a state-independent function of the set of payoff-differences anymore violating our regret solution definition. However, the regret solution is arguably the simplest type of solution because the principal can choose emotions only as a function of the vector of state-wise payoffs. In fact, she can employ the same function R *independent of the decision problem.*

Turning to the behavior induced by the optimal contract, it is useful to recall the distinction between the attention stage (choice of a signal structure) and the decision-making

stage (choice of an action conditional on the signal). Corollary 7 shows that the principal only distorts behavior in the former stage.

Corollary 7. *The agent's behavior at the decision-making stage is not distorted. The agent's attention is skewed towards more extreme events.*

A common theme in the principal-expert literature is that the principal might want to distort the expert's incentives at the decision-making stage to incentivize more information acquisition, as already pointed out by Demski and Sappington (1987). In the current model, that is not necessary. Because of the symmetry implied by exchangeability, under any posterior that any type of the agent will optimally acquire, the action he takes is the same as the principal's preferred action, as we show in the Appendix. However, under other posteriors, regret and rejoicing can distort behavior, as is known from regret theory (Bell, 1982; Loomes and Sugden, 1982).

Regarding attention, Property 2 implies that the agent pays relatively more attention to more extreme states compared to the first-best. We measure extremeness of a state ω by the stakes $|v(a_1,\omega) - v(a_2,\omega)| = |\Delta v(a_1,\omega)|$. How much attention the type β agent pays to some state ω can be measured by the log-likelihood ratio of the signal under ω. Because the signal is identical to the action, equation (2.1) implies that the log-likelihood ratio equals

$$\log\left(\frac{p_\beta(a_1|\omega)}{p_\beta(a_2|\omega)}\right) = \log\left(e^{\frac{\beta}{\lambda}Q(\Delta v(a_1,\omega))}\right) = \frac{\beta}{\lambda}Q\big(\Delta v(a_1,\omega)\big).$$

By contrast, the log-likelihood ratio with respect to the first-best attention strategy p^* is

$$\log\left(\frac{p^*(a_1|\omega)}{p^*(a_2|\omega)}\right) = \log\left(e^{\frac{1}{\lambda}\Delta v(a_1,\omega)}\right) = \frac{1}{\lambda}\Delta v(a_1,\omega).$$

The ratio of those two log-likelihood ratios, whether for a fixed type β or in expectation over β, grows with the extremeness of the state, since $Q\big(\Delta v(a_1,\omega)\big)/\Delta v(a_1,\omega)$ is increasing as

$|\Delta v(a_1, \omega)|$ increases by Property 2. Thus, the model predicts that the agent pays relatively more attention to more extreme events.

The prediction of disproportionate attention to states with a high payoff difference is similar in spirit to models of salience or focusing based on contrast (Bordalo et al., 2012; Kőszegi and Szeidl, 2013). Salience effects or focusing are often understood as a feature of automatic bottom-up attention instead of deliberate top-down attention, such as rational inattention (Bordalo et al., 2022). In contrast, we show that such focusing effects are not only consistent with a model of rational (in)attention, but that they are in fact optimal distortions to motivate rational inattention.[16]

The model makes two additional predictions regarding emotions that we relate to regret theory. First, while not imposed in standard regret theory, it is widely assumed that regret is stronger than rejoicing (e.g. Larrick and Boles, 1995; Humphrey, 2004), for which there is evidence based on self-reported emotions (Beattie et al., 1994; Mellers et al., 1999). Moreover, models of regret often do not include rejoicing at all (Quiggin, 1994; Sarver, 2008), suggesting that regret is considered more prominent than rejoicing.

Corollary 8. *Regret is stronger than rejoicing, that is,* $\forall x > 0 \colon |R(-x)| > |R(x)|$.

This result stems from the principal's motive to minimize the cost of inducing emotions Ψ. Rejoicing is costly for the principal when the agent takes the right action, and regret is costly for the principal when the agent takes the wrong action. By the formula for conditional choice probabilities (2.1), the agent takes, in any state, the right action with higher probability than the wrong action. Thus, the principal finds it cheaper to use regret rather than rejoicing. This motive is present in the analysis of optimal moral systems by Kaplow and Shavell (2007), who assume an inculcation cost for guilt and virtue. Further, it

[16] Lieder et al. (2018) offer a different explanation for overrepresentation of extreme events based on the resource-rationality approach.

echoes similar motives elsewhere in game theory as neatly expressed by Schelling (1980), "a promise is costly when it succeeds, and a threat is costly when it fails."[17]

Second, in contrast to existing models of regret, our model endogenizes the optimal mechanism as a function of the parameter λ, which can be interpreted as the difficulty of the decision problem. The higher λ the more costly it is to learn about the state and find the right action. Sugden (1986) argues that a short-coming of standard regret theory is that it does not incorporate that regret depends on how easy it was that one could have known the right action. He relates this to what he calls the self-recrimination aspect of regret. The following corollary shows that our model predicts stronger regret if it was cheaper to find out the optimal choice.

Corollary 9. *A smaller λ leads to larger regret from taking the wrong action, that is higher $|R(-x)|$ for all $x > 0$. Further, it leads to higher emotion-inclusive stakes, that is higher $Q(x)$ for all $x > 0$.*

Two mechanisms contribute to this effect. First, a smaller λ makes the types with lower β comparatively more responsive to changes in stakes, which increases $Q(x)$ and second, a smaller λ makes it more likely that the agent takes the right action, which makes regret comparatively cheaper than rejoicing for the principal.

Corollary 9 can also be related to inter-personal differences. The parameter λ not only captures properties of the decision problem but also of the agent, namely how costly it is for them to process information. A small λ corresponds to decision-makers for whom it is relatively cheap to process information or, equivalently, who have a greater information processing capacity, for example due to a smaller cognitive load. The model predicts that they will experience stronger regret when taking the wrong action. This can help explain why Schwartz et al. (2002) and Iyengar et al. (2006) find that "maximizers" are less satisfied with their choices than "satisficers" despite the first group making better choices on average.

[17]Note that Corollary 8 does not necessarily imply that emotions are negative in expectation. In fact, for a quadratic cost of emotions $\psi(x) = x^2$, one can show that emotions are 0 in expectation.

Relation to Classic Regret Theory

Our Property 2 is slightly weaker than the standard convexity assumption of Q, or equivalently the assumption that R is decreasingly concave (Bell, 1982; Loomes and Sugden, 1982). One can see how they relate using the terminology of production functions. Viewing Q as a production function for the sake of the analogy, convexity corresponds to increasing *marginal* product, while Property 2 corresponds to increasing *average* product. Under differentiability, increasing *average* product is equivalent to the marginal product being greater than the average product, in contrast to convexity which requires that the marginal product at x is greater than *any* marginal product up to x. Loomes and Sugden (1982) motivate the *convex* regret-rejoicing function Q through its implication of the Allais paradox, the common-ratio effect, the isolation effect, and under linear utility over money, the reflection effect, and simultaneous gambling and insurance. They show that the reflection effect holds under regret theory independent of assumptions on Q and that the isolation effect is equivalent to the common-ratio effect. Thus, we focus on the (a) Allais paradox, (b) the common-ratio effect, and (c) simultaneous gambling and insurance below.[18] Proposition 5 shows that Property 2 is in fact equivalent to the conjunction of (a) to (c) in the context of regret theory.

Note that regret theory was developed for decisions without taking into account attention. In particular, it has been applied to decisions between lotteries, where the nature of uncertainty is such that the state cannot be learned by paying attention. One way to reconcile this with our model is to assume that the agent's internal self-control mechanism treats those decisions *as if* the state could be learned and triggers regret and rejoicing regardless.[19]

The above-mentioned behavioral effects are formulated in terms of lotteries. To identify lotteries with actions in our present framework, let $x(a, \omega)$ denote the monetary outcome

[18]Bell (1982) observes that Property 2 is sufficient for simultaneous gambling and insurance and for rejection of probabilistic insurance. Loomes and Sugden (1987) demonstrate that superadditivity of regret, which we show below follows from Property 2, predicts preferences reversals.

[19]Regret and rejoicing can be understood as rational *rules*, that perform well in most environments, in which arguably some learning about the state is possible. See for example Heiner (1983), Aumann (1997), and Halevy and Feltkamp (2005) for related discussions on rule rationality.

of action a under state ω. Given the prior μ over states, an action induces a lottery over monetary amounts $x(a, \omega)$. Following Loomes and Sugden (1982), lotteries can be identified with any action that induces such lottery given the belief μ. Two independent lotteries can be identified with two actions such that the two distributions over monetary amounts induced by the prior are independent. Let $(x_1, p; x_2, 1-p)$ denote a lottery that gives monetary outcome x_1 with probability p and monetary outcome x_2 with probability $1-p$. Further, let (x, p) be shorthand for $(x, p; 0, 1-p)$.

Like Loomes and Sugden (1982), we consider choice sets of two actions, $A = \{a_1, a_2\}$. Let $v\colon \mathbb{R} \to \mathbb{R}$ be the decision-maker's increasing utility function over money. According to regret theory, the decision-maker's utility of action $a_i \in A$ is given by

$$\sum_{\omega \in \Omega} v(x(a_i, \omega)) + R\big(v(x(a_i, \omega)) - v(x(a_{-i}, \omega))\big)$$

for some regret-rejoicing function $R\colon \mathbb{R} \to \mathbb{R}$. Equivalently, let $Q(x) = x + R(x) - R(-x)$ and preferences between actions be given by

$$a_1 \succsim a_2 \Leftrightarrow \sum_{\omega \in \Omega} \mu(\omega) Q\big(v(x(a_1, \omega)) - v(x(a_2, \omega))\big) \geq 0 \tag{2.3}$$

where $Q(x) = x + R(x) - R(-x)$. At this point, we impose no assumptions on R or Q.

The following definitions are taken from Loomes and Sugden (1982).

Definition 6 (Allais Paradox)**.** *Let $X_1 = (x_1, p_1; x_2, \alpha)$ and $X_2 = (x_2, p_2 + \alpha)$ be independent lotteries where $1 \geq p_2 > p_1 > 0$ and $(1 - p_2) \geq \alpha \geq 0$. If there exists some probability $\overline{\alpha}$ such that $X_1 \sim X_2$ when $\alpha = \overline{\alpha}$, then (i) if $x_1 > x_2 > 0$, then $\alpha < \overline{\alpha} \Rightarrow X_1 \succ X_2$ and $\alpha > \overline{\alpha} \Rightarrow X_1 \prec X_2$ and (ii) if $0 > x_2 > x_1$, then $\alpha < \overline{\alpha} \Rightarrow X_1 \prec X_2$ and $\alpha > \overline{\alpha} \Rightarrow X_1 \succ X_2$.*

Definition 7 (Common Ratio Effect)**.** *Let $X_1 = (x_1, \lambda p)$ and $X_2 = (x_2, p)$ be independent lotteries, where $1 \geq p > 0$ and $1 > \lambda > 0$. If there exists some probability $\overline{p}$ such that*

$X_1 \sim X_2$ *when* $p = \overline{p}$*, then (i) if* $x_1 > x_2 > 0$*, then* $p < \overline{p} \Rightarrow X_1 \succ X_2$ *and* $p > \overline{p} \Rightarrow X_1 \prec X_2$ *and (ii) if* $0 > x_2 > x_1$*, then* $p < \overline{p} \Rightarrow X_1 \prec X_2$ *and* $p > \overline{p} \Rightarrow X_1 \succ X_2$*.*

Definition 8 (Simultaneous Gambling and Insurance)**.** *Let* v *be linear in money. Let* $X_1 = (0, 1)$ *and* $X_2(x, p; -px/(1-p), 1-p)$ *be independent lotteries where* $0 < p < 1$ *and* $x > 0$*. (i) For* $p > 0.5$*,* $X_1 \succ X_2$*, (ii) for* $p < 0.5$*,* $X_1 \prec X_2$*, and for (iii) for* $p = 0.5$*,* $X_1 \sim X_2$*.*

Proposition 5. *Suppose the decision-maker has preferences between lotteries given by (2.3). The decision-maker's preferences exhibit the Allais paradox, common-ratio effect, and simultaneous gambling and insurance if and only if* $Q(x)/x$ *is strictly increasing on* $\mathbb{R}_{>0}$*.*

2.3.2 More than Two Actions

As stressed in the literature, it is not trivial how to generalize regret theory to more than two actions (Loomes and Sugden, 1982, Loomes and Sugden, 1987, Quiggin, 1994). We define a *general regret solution* based on the utility representation in Loomes and Sugden (1987) and Sugden (1993), which to our knowledge accommodates all proposals. They define a representation of preferences where the utility depends on the realized outcome and on the set of counterfactual outcomes under the same state. They do not impose additional assumptions that would correspond to generalizations of Properties 1 and 2. In fact, it is not trivial how these would generalize to more than two actions. This representation is more general than another widely-used representation due to Quiggin (1994), according to which regret depends on the difference of the actual to the highest attainable payoff under the realized state. As we will show, that form of regret is in general not optimal in our model.

Definition 9 (General Regret Solution)**.** *Under a finite choice set* A*, a solution* $(m, (p_\beta)_\beta)$ *to the principal's problem* (P) *is a general regret solution if there exists a function* R *such that for all* $(a, \omega) \in A \times \Omega$*:*

$$m(a, \omega) = R\big(\{v(a, \omega) - v(a', \omega)\}_{a' \in A \setminus \{a\}}\big),$$

and for all $a, a' \in A, \omega \in \Omega$:

$$v(a,\omega) > v(a',\omega) \Leftrightarrow v(a,\omega) + m(a,\omega) > v(a',\omega) + m(a',\omega).$$

The definition states that emotions are a state-independent function of the *set* of payoff differences between the actual and the counterfactual payoffs under the realized state. In particular, a permutation of the counterfactual payoffs does not change the emotion as well as shifting all payoffs under some state by a constant. Moreover, adding emotions preserves the strict ordering of payoffs, state-wise. In the Appendix 2.6, we comment on how our representation relates exactly to the one in Loomes and Sugden (1987).

Finally, assume that the principal has strict preferences over actions, given the state. This can be interpreted as a genericity assumption or as the state being defined in a sufficiently fine-grained way. We use this assumption to prove that state-wise permutations of payoffs translate to corresponding permutations of the optimal emotions.

Assumption 5. *For all* $\omega \in \Omega$*, for all* $a, a' \in A$*:* $(a \neq a' \rightarrow v(a,\omega) \neq v(a',\omega))$.

Under these assumptions, we prove the following theorem.

Theorem 9. *Under finite A and Assumptions 1 to 3, the principal's problem* (P) *has a general regret solution.*

The main part of the proof is to show that the principal's problem separates by state and that state-wise optimal emotions exist, a result which we also use for proving Theorem 8. Moreover, we show that state-wise optimal emotions are invariant under state-wise permutations or shifts of payoffs, and preserve strict orders of payoffs.[20]

Our model does not only predict that the optimal emotions are a general regret solution. Given λ and the distribution of β, the model can be used to compute the precise optimal

[20]Solutions that are not general regret solutions exist for the same reason as in the case with two actions. Further, one can show analogously to the binary action case that the principal distorts behavior only at the attention stage and not at the decision-making stage.

emotional payoffs. While optimal emotions could, again, depend on the whole decision problem, they boil down to a state-wise $|A|$-variable optimization problem. That is, the optimal emotions are a function of the vector of payoffs in the realized state. They are given by the variance-of-emotions-minimizing solution to

$$\max_{u(\cdot,\omega)\colon A\to\mathbb{R}} \int \left(\lambda H\big(p_\beta(u(\omega))\big) + \sum_{a\in A} p_\beta(a|u(\omega))v(a,\omega) \right) dF(\beta),$$

where $p_\beta(a|u(\omega))$ is given by the multinomial logit formula

$$p_\beta(a|u(\omega)) = \frac{e^{\frac{\beta}{\lambda}u(a,\omega)}}{\sum_{a'\in A} e^{\frac{\beta}{\lambda}u(a',\omega)}},$$

and $H(p_\beta(u(\omega)))$ is the entropy of the probability vector $p_\beta(u(\omega)) = (p_\beta(a|u(\omega)))_{a\in A}$. Unfortunately, the problem does not allow for a simple analytical solution or for comparative statics as in binary decision problems. However, once the the problem is solved for a vector of payoffs, it applies to payoff-equivalent states in *any decision problem*. Further, when the bias is known, the next section shows that there is a simple analytic and intuitive solution (Theorem 10).

2.4 Interpretation as Contracting with an Expert

As mentioned in the main text, the model can be interpreted in a more conventional way as a principal hiring an agent (expert) to acquire information. Of particular relevance are CEO and fund manager compensation.

The principal reaps utility $v(a,\omega)$ from action a under state ω and pays monetary transfers (wages) $m(a,\omega)$ to the agent, which are valued quasi-linearly by the agent. The bias of the expert can capture both a conflict of interest between principal and agent, as well as a behavioral bias of the expert. We discuss both cases in turn.

2.4.1 Hiring an Expert with Conflicting Interests

To some degree our model can be interpreted as the optimal monetary contract between a principal and an expert with a conflict of interest of how much information to acquire. This is the case if both principal and agent care about the information cost as well as the decision outcome but differ in how they weigh these considerations. For instance, investors or shareholders may internalize the time cost of information gathering to some degree due to the resulting decision delay or because it diverts the agent from other valuable activities. CEOs or fund managers may care about the decision outcome out of intrinsic motivation or reputational concerns. However, from the principal's perspective, the agent might overemphasize the information cost in the absence of financial incentives, potentially leading to suboptimal information acquisition. In these contexts, it is realistic to assume that information acquisition is unverifiable and only the outcome from decision-making is contractible. Information acquisition often does not produce tangible evidence and even when it does, the evidence might be hard to communicate or verify without the necessary expertise on the principal's side (Lindbeck and Weibull, 2020). Finally, we take serious the possibility that the principal is uncertain about how much the agent overweights the information cost, by incorporating adverse selection.[21]

Ex-Ante Budget Constraint Under a quadratic implementation cost $\psi(x) = x^2$, our cost-of-emotions refinement is equivalent to an ex-ante budget constraint plus the refinement of reducing the variance of emotions. Note that the mean squared emotions equal the squared mean emotions plus the variance of emotions, as for any random variable X, $\mathbb{E}[X^2] = \mathbb{E}[X]^2 + \text{Var}[X]$. Thus, minimizing the quadratic implementation cost is equivalent to minimizing the sum of the squared expected emotions and the variance of emotions. The principal can add a constant to emotions uniformly to reduce the expected emotions to 0, without affecting the variance of emotions. Further, this does not affect the agent's behavior. Thus, minimizing

[21] A shortcoming of this application is that our model does not feature a limited liability constraint, which binds if the budget b is low enough.

the quadratic implementation cost is equivalent to an ex-ante budget constraint

$$\int \left(\sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) m(a,\omega) \right) dF(\beta) = 0 \tag{BC}$$

with $b = 0$ and the refinement of minimizing the variance of emotions,

$$\int \sum_{\omega \in \Omega} \sum_{a \in A} \mu(\omega) p_\beta(a|\omega) m_\beta(a,\omega)^2 dF(\beta).$$

Under an interpersonal interpretation with m as monetary transfers, Ψ can be thought of as a implementation cost or as a cost of deviating from an ex-ante budget. Further, when interpreting the principal's problem in organizational contexts where monetary transfers are not used, we can interpret transfers as the reduced-form utility of reward- and punishment-actions taken by the principal. Rewards, such as praise and gifts, as well as punishments, such as making the agent perform unfavorable tasks (which require monitoring), might both be costly for the principal, as captured by ψ.[22]

Hiring a Dynamically-Inconsistent Expert If the bias does, however, denote a behavioral bias of the agent, then our optimal contract can be readily interpreted as the *welfare-maximizing* commitment contract for the agent, subject to an ex-ante *budget balance constraint*.[23]

Furthermore, as demonstrated below, our model also applies to a *profit-maximizing* principal who hires a dynamically-inconsistent expert, subject to a *participation constraint*. Here, the principal values the decision outcome net of transfers paid to the agent, but does not care about the agent's attention cost directly. The agent, on the other hand, cares only about the transfers received net of the attention cost. Instead of a budget constraint, the principal faces an ex-ante participation constraint that guarantees the agent a certain expected

[22] See for example Andreoni et al. (2003) for more examples of rewards and punishments in different social and institutional contexts.

[23] At first glance, a welfare-maximizing principal would have a different objective than our principal because she would care about the agent's transfers. However, this is immaterial as transfers cancel out in expectation due to the ex-ante budget constraint.

utility before his bias realizes. The participation constraint effectively makes the principal internalize the agent's attention cost, rendering the problem equivalent to ours.

In these contexts, our work contribute to the understanding of compensation schemes, specifically those that assess the agent's performance relative to a peer benchmark. Relative performance evaluation is common in compensation schemes for fund managers (Ma et al., 2019; Evans et al., 2023) and, to some degree, for CEOs (Edmans et al., 2017).[24] Such benchmarks can be justified by by Holmström's (1979) informativeness criterion in standard moral-hazard settings, because subtracting exogenous common shocks to the return of the industry increases the informativeness of the performance evaluation. Our model offers a different explanation based on incentives for information acquisition. Peer comparisons can be interpreted as counterfactual comparisons, since payoffs of relevant peers inform the principal about the payoff that alternative actions would have lead to. Indeed, Theorem 10 shows that under a known bias, transfers proportional to the difference between the realized return and the expected return under the realized market conditions are optimal.

Consider a principal who hires a dynamically-inconsistent agent (expert) to acquire information.[25] The principal reaps utility $v(a, \omega)$ from action a under state ω and pays monetary transfers (wages) $m(a, \omega)$ to the agent. The agent maximizes transfers paid to them. Both have quasi-linear utility. After signing the contract, the agent faces a stochastic temptation shock to his preferences, as in random-indulgence models. Accordingly, the individual-

[24] Evans et al. (2023) find that 71% of mutual funds in the US use peer-benchmarked compensation.
[25] See Kaur et al. (2015) for evidence on self-control issues at work.

rationality constraint needs to hold in expectation. The principal's problem is

$$\max_{m,(p_\beta)_\beta} \int \left(-c(p_\beta) + \sum_{\omega\in\Omega} \sum_{a\in A} \mu(\omega) p_\beta(a|\omega) \big(v(a,\omega) - m(a,\omega) \big) \right) dF(\beta)$$

s.t.

$$\forall \beta\colon \ p_\beta \in \arg\max_p -c(p) + \beta \sum_{\omega\in\Omega} \sum_{a\in A} \mu(\omega) p(a|\omega) m(a,\omega) \tag{IC}$$

$$\int \left(-c(p_\beta) + \beta \sum_{\omega\in\Omega} \sum_{a\in A} \mu(\omega) p_\beta(a|\omega) m(a,\omega) \right) dF(\beta) \geq U_0 \tag{IR}$$

If the agent was dynamically consistent ($\beta \equiv 1$), the principal would achieve the first-best by selling the firm to the agent. In the presence of a bias, the principal can be seen as maintaining control of the firm in order to provide a self-control mechanism to the agent for their mutual benefit.

By standard arguments, the (IR) constraint is binding. Inserting the (IR) constraint into the objective, we obtain the following equivalent formulation.

$$\max_{m,(p_\beta)_\beta} \int \left(-c(p_\beta) + \sum_{\omega\in\Omega} \sum_{a\in A} \mu(\omega) p_\beta(a|\omega) v(a,\omega) \right) dF(\beta) - U_0$$

s.t.

$$\forall \beta\colon \ p_\beta \in \arg\max_p -c(p) + \beta \sum_{\omega\in\Omega} \sum_{a\in A} \mu(\omega) p(a|\omega) m(a,\omega) \tag{IC}$$

$$\int \left(-c(p_\beta) + \beta \sum_{\omega\in\Omega} \sum_{a\in A} \mu(\omega) p_\beta(a|\omega) m(a,\omega) \right) dF(\beta) = U_0 \tag{IR'}$$

The constant $-U_0$ in the principal's objective can be ignored. Identifying m with $v + m$ in the main text, the objective and (IC) constraint are identical. The (IR') constraint can be satisfied by adding a uniform constant to emotions, so (IR') acts in effect as an ex-ante budget constraint. Adding the refinement that the principal chooses the mechanism with lowest variance of transfers, renders the problem essentially equivalent to the principal's problem

in the main text.[26] The only difference is that the optimal transfers emulate the payoff v in addition to the emotions in the main formulation. That is because in this formulation the agent has no direct interest in the decision. Additionally, the optimal transfers are uniformly shifted by a constant, which depends on the outside option U_0.

2.5 Conclusion

While counterfactual thinking is often dismissed as irrational (e.g. Howard, 1992)[27] and seen as distortions to rational decision-making, the present model shows that emotions based on counterfactual reasoning are actually optimal to motivate better decision-making. This paper uses a rational-inattention model, but any model that trades off decision-making costs with better decision outcomes would arguably predict a useful role for regret and disappointment to amplify the stakes of the decision problem. Although the model allows for any emotions as a function of ex-post feedback, simple and well-known emotions arise endogenously. Under complete feedback about the state of the world, regret and rejoicing are optimal. In a binary decision problem, these emotions feature disproportionate aversion to large regrets and imply that attention is skewed towards extreme states. When feedback about the state is incomplete, disappointment and elation supplement regret and rejoicing as a form of second-best emotions. Extending standard regret theory, the model predicts stronger regret the easier one could have known the right action, consistent with the psychology of a self-blaming aspect of regret.

Many open questions remain. Our analysis is for tractability and conciseness restricted by symmetry assumptions. Can the principal profit from asymmetric incentive schemes in symmetric decision problems? What are the optimal incentives when the underlying decision problem is asymmetric? When delegating asymmetric decision problems, motivating attention may come only at the cost of distorting decision-making. Analyzing the optimal

[26] The equivalence also holds under an appropriate exchangeability constraint or when the principal chooses a menu of transfer schemes.

[27] However, see Bradley and Stefánsson (2017) for a defense of counterfactual comparisons as rational.

incentives could help shed further light on the resulting trade-off and thereby unite themes from the current paper, which shows that regret is optimal to motivate attention absent distortions of choice, and studies in regret theory, which focus on how regret distorts decision-making.

On a broader scale, our research underscores parallels between internal psychological mechanisms, such as counterfactual comparisons, and more traditional principal-agent problems. We believe that such a perspective has the potential to profit both behavioral economics and mechanism design. Behavioral economics might glean insights from mechanism design to better understand the function and benefits of non-standard preferences. Similarly, mechanism design can turn to human psychology as a source of innovative solutions for agency problems, which evolution has already solved.

2.6 Appendix

Proof of Lemma 21 A general mechanism elicits a message from the agent *before* as well as *after* attention. Conditional on the messages, an action recommendation is sent and emotion-transfers are assigned conditional on the messages and the outcome. By the revelation principle it is sufficient to restrict attention to mechanisms, where the first message corresponds to the agent's type. We maintain the assumption that any emotion scheme induces an exchangeable decision problem for the agent. First, we analyze the case when the mechanism elicits a message only before attention. At the end of the proof, we argue that eliciting a message after attention can be reduced to a mechanism of the former case.

The principal chooses a vector of emotion schemes $m_\beta \colon A \times \Omega \to \mathbb{R}$ and attention strategies $p_\beta \colon \Omega \to \Delta(A)$, indexed by the agent's type, each, subject to incentive compatibility, budget balance, and exchangeability. The incentive-compatibility constraint below captures both the truthful reporting of one's type as well as the subsequent attention and decision-

making, accounting for the possibility of double deviations. The principal's problem is

$$\max_{(m_\beta,p_\beta)_\beta} \int \left(-c(p_\beta) + \sum_{\omega\in\Omega}\sum_{a\in A} \mu(\omega)p_\beta(a|\omega)v(a,\omega) \right) dF(\beta)$$

s.t.

$$\forall\beta\colon\ (\beta,p_\beta) \in \arg\max_{\hat\beta,\hat p} -c(\hat p) + \beta \sum_{\omega\in\Omega}\sum_{a\in A} \mu(\omega)\hat p(a|\omega)\big(v(a,\omega) + m_{\hat\beta}(a,\omega)\big) \tag{IC}$$

$$\int \left(\sum_{\omega\in\Omega}\sum_{a\in A} \mu(\omega)p_\beta(a|\omega)m_\beta(a,\omega) \right) dF(\beta) = 0 \tag{BC}$$

$$\forall\beta\colon\ v + m_\beta \text{ is exchangeable.} \tag{EX}$$

Out of the solutions $(m_\beta, p_\beta)_\beta$, the principal chooses the one that minimizes the variance of emotions,

$$\int \sum_{\omega\in\Omega}\sum_{a\in A} \mu(\omega)p_\beta(a|\omega)m_\beta(a,\omega)^2 dF(\beta).$$

Osband (1989) shows that in their model offering a menu of emotion schemes can in principle benefit the principal, but using numerical simulations he shows that these gains are very small. In contrast to Osband, in our setup, heterogeneity concerns the agent's bias instead of his cost parameter. Thus, the first-best attention strategy is the same for every agent, which the proof relies on. In Yoder (2022), the principal can effectively screen agents when the acquired posterior is contractible. We show that the moral hazard constraint implies that the principal can induce only gross value functions that are convex in the posterior. Offering a menu makes the gross value function more convex, leading to further distortions at the top and bottom. Our proof does not rely on the specific shape of mutual information and we show the proposition for any posterior-separable attention cost (see section 2.7 for a definition) with strictly concave and differentiable H that is *symmetric in the states*.

Proof. Under exchangeability of v, if there are two states, $\Omega = \{\omega_1, \omega_2\}$, then one of the following needs to be the case. The payoff v is state-wise constant, or $\mu(\omega_1) = \mu(\omega_2) = \frac{1}{2}$

and there are two actions, $A = \{a_1, a_2\}$, with exchangeable payoffs, $v(a_1, \omega_1) = v(a_2, \omega_2)$ and $v(a_1, \omega_2) = v(a_2, \omega_1)$. If there were more than two actions and v is not state-wise constant, then more than two states are necessary to have all permutations of $(v(a, \omega))_{a \in A}$ be equiprobable.

If the payoffs are state-wise constant, then the optimal attention strategy is to pay zero attention, which is achieved at the smallest implementation cost by setting the emotions to 0 for all types. Thus, all u_β are the same.

Consider, then, the other case that or $\mu(\omega_1) = \mu(\omega_2) = \frac{1}{2}$ and there are two actions, $A = \{a_1, a_2\}$, with exchangeable payoffs, $v(a_1, \omega_1) = v(a_2, \omega_2)$ and $v(a_1, \omega_2) = v(a_2, \omega_1)$. Without loss, the payoff matrix for the principal can be written as

	ω_1	ω_2
a_1	v	0
a_2	0	v

where $v > 0$.

Agent's Behavior Turning to the agent, we first analyze the agent's behavior given some $u \colon A \times \Omega \to \mathbb{R}$ in the menu of utility functions $(u_\beta)_\beta = (v + m_\beta)_\beta$ that the principal offers.[28] By the exchangeability and an argument analogous to the above, the payoff matrix of any such u,

	ω_1	ω_2
a_1	$u(a_1, \omega_1)$	$u(a_1, \omega_2)$
a_2	$u(a_2, \omega_1)$	$u(a_2, \omega_2)$

must satisfy one of the following.

Either u is state-wise constant, that is $u(a_1, \omega_1) = u(a_2, \omega_1)$ and $u(a_2, \omega_1) = u(a_1, \omega_1)$, or $u(a_1, \omega_1) = u(a_2, \omega_2)$ and $u(a_1, \omega_2) = u(a_2, \omega_1)$. In the first case, we can make u constant over all outcomes without changing the agent's behavior (he pays no attention in either case), so

[28] By a change of variables, the principal chooses u directly instead of m.

it is captured by the second case, which we analyze below. We can write the agent's payoff matrix as

	ω_1	ω_2
a_1	$\underline{u} + \Delta u$	$\underline{u}$
a_2	$\underline{u}$	$\underline{u} + \Delta u$

where $\underline{u} = u(a_1, \omega_2) = u(a_2, \omega_1)$ and $\Delta u = u(a_1, \omega_1) - u(a_2, \omega_1) = u(a_2, \omega_2) - u(a_1, \omega_2)$.

We characterize the optimal attention strategy given some chosen utility function u. Following Caplin and Dean (2013), we solve for the optimal distribution over posteriors, induced by the attention strategy, analogous to Bayesian persuasion models (Kamenica and Gentzkow, 2011). Under two states, we can identify posteriors $\gamma \in \Delta(\Omega)$ with the probability in $[0, 1]$ they assign to state ω_1, which makes this posterior approach especially attractive. The problem of agent β given u is to pick a Bayes-consistent distribution over posteriors to maximize the expectation of the *net value* function N_β (Caplin and Dean, 2013):

$$\max_{\tau \in \Delta[0,1]} \mathbb{E}_\tau[N_\beta(\gamma|u)]$$

$$\text{s.t.}$$

$$\mathbb{E}_\tau[\gamma] = \frac{1}{2}$$

Here, $N_\beta(\gamma|u)$ is the net value of posterior γ for the agent of type β given u, which is

$$N_\beta(\gamma|u) = \beta(\underline{u} + \max\{\gamma\Delta u, (1-\gamma)\Delta u\}) + \lambda H(\gamma).$$

The first summand of the net value $N_\beta(\cdot|u)$ captures the *gross value*, or instrumental value, of attention and the second the attention cost. The gross value is continuous, convex, symmetric around $1/2$, and piecewise linear as the maximum of linear functions. The attention cost is mutual information, $H(\gamma) = \gamma \log(\gamma) + (1-\gamma)\log(1-\gamma)$, which is strictly concave,

differentiable and symmetric around 1/2. These are the only properties of H that we make use of for the proof.

As known from the literature on Bayesian persuasion, the solution can be obtained by concavifying the net value function $N_\beta(\cdot|u)$. As we show next, this Bayesian persuasion problem takes a particularly simple form due to symmetry.

The net value function has one global maximum if $\Delta u = 0$ and two global maxima otherwise. Restricted to $[1/2, 1]$, the net value function attains its maximum by continuity. On $[1/2, 1]$, the net value function is the sum of the linear gross value and the strictly concave attention cost, so the maximum must be unique. By symmetry and differentiability of H, we have $H'(1/2) = 0$, so the maximum is at 1/2 only if $\Delta u = 0$. In fact, by differentiability the maximum is characterized by $-\lambda H'(\gamma) = \beta \Delta u$ if $-\lambda H'(1) > \beta \Delta u$ and $\gamma = 1$ otherwise. By symmetry of H around 1/2, the net value function $N_\beta(\cdot|u)$ is symmetric around 1/2, so the net value function attains another maximum at the reflection point around 1/2.

Because the maxima are symmetric around the prior probability 1/2, there is a unique Bayes-consistent τ whose support is the set of global maxima. This is the unique τ that attains the highest expected posterior net value because each posterior in its support maximizes the net value function. This τ is either degenerate at 1/2 or binary and symmetric around 1/2. The agent selects action a_1 when the realized posterior is greater than 1/2 and a_2 when the realized posterior is less than 1/2, if $\Delta u > 0$ and vice versa if $\Delta u < 0$. Thus, given $v > 0$, the principal never offers an emotion-inclusive utility with $\Delta u < 0$, which would induce the agent to pay costly attention only to select the wrong action.

Thus, while the agent chooses a high-dimensional distribution over posteriors, we can characterize his choice by the one-dimensional maximizer of the net value function on $[1/2, 1]$, which corresponds to the higher posterior the agent acquires. As written above, the maximizer is characterized by $\lambda - H'(\gamma) = \beta \Delta u$ if $-\lambda H'(1) > \beta \Delta u$ and $\gamma = 1$ otherwise. Because H is strictly concave and symmetric, $-H'$ is strictly increasing on $[1/2, 1]$ with $-H'(1/2) = 0$.

Thus, the maximum is strictly increasing in Δu as long as the optimal γ is less than 1 and constant afterwards.

The complete strategy of the agent consists of choosing an emotion scheme and subsequently an attention strategy. By the above, the agent's behavior can be described as choosing a net value function and subsequently the maximum of the net value function on $[1/2, 1]$. This is equivalent to choosing the maximum on $[1/2, 1]$ of the pointwise maximum $N_\beta(\cdot|(u_\beta)_\beta)$ of the collection of net value functions,

$$\begin{aligned} N_\beta(\gamma|(u_\beta)_\beta) =& \max_{\hat{\beta}}\{N_\beta(\gamma, u_{\hat{\beta}})\} \\ =& \beta \max_{\hat{\beta}}\{\underline{u}_{\hat{\beta}} + \gamma\Delta u_{\hat{\beta}}\} + \lambda H(\gamma), \text{ for } \gamma \geq 1/2. \end{aligned}$$

We call $N_\beta(\gamma|(u_\beta)_\beta)$ the *effective* value function. By $\Delta u_\beta > 0$ for all β, the effective value function is supermodular on $[1/2, 1]$ in β and γ and thus the optimal γ is non-decreasing in β (Topkis, 1978). The effective value function on $[1/2, 1]$ is the sum of the maximum of increasing, linear functions and a concave attention cost. The maximum of increasing, linear functions is increasing and convex. Thus, non-decreasing γ implies a non-decreasing choice of the slope $\Delta u_{\hat{\beta}}$.

To summarize, for any type of the agent, the menu of emotion-inclusive utility functions induces an effective value function and the type's acquired signal structure can be described as picking the posterior in $[1/2, 1]$ that maximizes this net value function. We have shown that (1) a higher β implies a weakly more extreme posterior, and (2) a weakly higher slope Δu_β of the gross value function. Given the choice of an emotion-inclusive utility, (3) the agent's posterior is a strictly increasing function of the slope of the gross value function, if the posterior is below 1.

Principal's Problem Turning back to the principal, her net value function is

$$N_P(\gamma) = \max\{\gamma v, (1-\gamma)v\} + \lambda H(\gamma).$$

Her value from the symmetric distribution with support $\{\gamma, 1-\gamma\}$ is just $N_P(\gamma)$ by symmetry of N_P. The principal's value function N_P is strictly concave on $[1/2, 1]$ because it is the sum of the linear gross value and the strictly concave attention cost. Strict concavity implies strict quasi-concavity, which we are using below.

Suppose for the sake of contradiction that there are types β, β' such that $u_\beta \neq u_{\beta'}$. We show that the principal strictly benefits from offering only the emotion-inclusive utility u_{β^*} of a type β^* that selects the most favorable posterior to the principal,

$$\beta^* \in \arg\max_{\beta \in \text{supp}(\beta)} N_P(\gamma_\beta).$$

Such a type β^* exists by $\text{supp}(\beta)$ being finite.

By (2), greater types than β^* select a weakly higher slope and smaller types a weakly smaller slope of the gross value function. Offering only the emotion-inclusive utility u_{β^*}, by (3) weakly increases the posterior for type less than β^*, weakly decreases the posterior for types greater than β^*, and does not change the posterior of β^*. Because the ordering of posteriors, (1), still holds due to supermodularity of the new value function, each posterior moves weakly towards γ_{β^*}. By strict quasi-concavity of the principal's value function on $[1/2, 1]$ and γ_{β^*} yielding a weakly higher value to the principal, the principal benefits weakly from the induced changes and benefits strictly from strict changes. Thus, the principal is weakly better off offering only one emotion scheme.

It remains to show that in case of indifferences, the principal cannot decrease the implementation cost by offering different emotion schemes. By the above, the principal is strictly better off if some posterior moved strictly. The only case that no posterior changed, is when the only types whose emotion-inclusive utility changed selects $\gamma = 1$ at their original and at the new emotion-inclusive utility. By construction, the new value function is below the original induced value function at $\gamma = 1$, so it entails a smaller emotional payoff after selecting the better action under the realized state. Because $\beta \leq 1$, optimality entails that the slope

of the constructed gross value function is at least v. Then, the implementation cost being minimized entails positive emotions at selecting the better action under the realized state. Thus, the new emotion-inclusive utility assigning smaller emotions implies a reduction of the implementation cost.

General Mechanism As mentioned above, the principal could use an even more general mechanism. That is, she could elicit an additional message from the agent after attention and condition the emotion scheme on both the message before and after attention. We maintain the restriction that any assigned emotion scheme m is such that $v+m$ is exchangeable. Here, we sketch how the proof can be extended to account for such mechanisms. We show that by starting with a general mechanism, we can construct a mechanism *without* a message after attention that is equally good for the principal. Then, the above proof can be applied. The argument below uses that if the agent acquires symmetric posteriors, there is no difference between eliciting a message before and after attention, as we have shown above for symmetric value functions. Using this, we show that the principal might as well offer only one emotion scheme given the first message from the agent (thus making the second message meaningless). The argument proceeds in three steps.

First, given an incentive-compatible general mechanism, fix an arbitrary type β of the agent and consider the agent's attention problem following the first message. We can, again, by direct revelation identify the first message with the agent's type, which the agent truthfully reveals. The agent's attention problem still consists of concavifying a symmetric value function. That is because analogous to the argument above, the agent faces an effective value function, denoted by $\hat{N}_\beta$, given by the maximum of the value functions induced by the emotion schemes, which the agent can choose from after attention. Because each emotion scheme induces a symmetric value function, the maximum of the value functions is symmetric, too. This implies that the set of maximizers M_β of the value function $\hat{N}_\beta$ is symmetric around $1/2$. Any Bayes-consistent distribution with support at only the maximizers of the $\hat{N}_\beta$ is a solution to the attention problem of type β.

Second, there is a symmetric, principal's preferred solution to the attention problem of type β. The principal's value function is symmetric around 1/2, too. Hence, the principal prefers, out of the solutions to the agent's problem, those whose support includes only posteriors that give the highest value to the principal. Formally, let $M_{P,\beta} := \arg\max_{\gamma \in M_\beta} N_P(\gamma)$, which is symmetric around 1/2. The principal's preferred solutions to the attention problem of type β are the Bayes-consistent distributions with support in $M_{P,\beta}$. Because $M_{P,\beta}$ is symmetric and the prior is at 1/2, there is a Bayes-consistent, binary, symmetric distribution with support in $M_{P,\beta}$.

Third, starting with a mechanism that elicits a message before and after attention, we can construct an equally good mechanism that elicits a message only before attention. As we are allowing the principal to pick her preferred solution when the agent is indifferent, for any type we can select a principal's preferred solution to his attention problem that is moreover symmetric and binary by the above. By symmetry, this solution makes any agent pick the same emotion scheme (or, send the same second message) after each acquired posterior. This solution can also be implemented by offering to every agent *only this one emotion scheme*, after the first message. Offering less choice to the agent does not not violate any other incentive-compatibility constraints. Finally, we can apply the proof above to the resulting mechanism. □

Proof of Theorem 9 Before proving Theorem 9, which we use below to prove Theorem 8, we comment on how our general regret representation relates to the one in Loomes and Sugden (1987).

To introduce the general regret representation in Loomes and Sugden (1987), let us define $o(a, \omega)$ as the outcome from action a under state ω. Their representation requires that the agent maximizes a utility function

$$u(a, \omega) = \Phi(o(a, \omega), \{o(a', \omega) | a \in A \setminus \{a\}\}),$$

that is a function of the actual outcome and the set of counterfactual outcomes.

Our general regret solution

$$u(a,\omega) = v(a,\omega) + R\big(\{V(a,\omega) - v(a',\omega) | a' \in A \setminus \{a\}\}\big)$$

implies their representation but makes several additional assumptions. First, the emotion-inclusive utility $u(a,\omega)$ depends on the set of state-wise *payoffs* instead of outcomes. Thus, payoff-equivalent states are treated the same way. Second, the additive separability together with the dependence of emotions on the set of payoff *differences* additionally imposes that shifting all payoffs in some state by a constant does not change the emotions. Third, we assume that emotions do not change, in each state, the strict ordering of utilities, as in Sarver (2008). These three additional assumptions seem plausible in the context of regret theory. We now turn to the proof of Theorem 9, which we show for a general implementation cost. By Appendix 2.4, this captures as a special case the budget constraint with the refinement of minimal variance of emotions.

Proof: Let $A := \{a_1, ..., a_n\}$. Because the principal is constrained to choose emotions that induce an exchangeable emotion-inclusive utility u, the state-dependent choice probabilities of any type of the agent can be written as (Matějka and McKay, 2015)

$$p_\beta(a|\omega) = \frac{e^{\frac{\beta}{\lambda}u(a,\omega)}}{\sum_{a' \in A} e^{\frac{\beta}{\lambda}u(a',\omega)}} = \frac{e^{\frac{\beta}{\lambda}\Delta u(a,\omega)}}{1 + \sum_{a' \neq a_1} e^{\frac{\beta}{\lambda}\Delta u(a',\omega)}}. \tag{2.4}$$

where $\Delta u(a,\omega) := u(a,\omega) - u(a_1,\omega)$. Further, we define $u(\omega) := (u(a_1,\omega), ..., u(a_n,\omega))$ and $p_\beta(\omega) := (p_\beta(a_1|\omega), ..., p_\beta(a_n|\omega))$.

By (2.4), the choice probabilities only depend on the state-wise utilities and even further only on state-wise utility differences with respect to some reference action. Using this, we show that the principal's problem is separable by state.

By a change of variables, the principal chooses directly the emotion-inclusive utility u instead of emotions m. The principal's maximization problem consists of an outer minimiza-

tion problem regarding the implementation cost $\psi(u(a,\omega)-v(a,\omega))$ and an inner problem to maximize the payoff net of attention cost, which is subject to the exchangeability constraint:

$$\min_{u\colon A\times\Omega\to\mathbb{R}} \int \left(\sum_{\omega\in\Omega}\sum_{a\in A} \mu(\omega) p_\beta(a|\omega) \psi\big(u(a,\omega) - v(a,\omega)\big) \right) dF(\beta)$$

s.t.

$$u \in \underset{u\colon A\times\Omega\to\mathbb{R}}{\arg\max} \int \sum_{\omega\in\Omega} \mu(\omega) \left(\lambda \log(1/n) + \lambda H(p_\beta(\omega)) + \sum_{a\in A} p_\beta(a|\omega) v(a,\omega) \right) dF(\beta) \qquad (2.5)$$

s.t. $\{a_1, ..., a_n\}$ is exchangeable with respect to u.

where $p_\beta(a|\omega)$ is given by (2.4) and $H(p_\beta(\omega)) = \sum_{a\in A} -p_\beta(a|\omega)\log(p_\beta(a|\omega))$ is the entropy of the conditional choice probabilities. We have used that under exchangeable payoffs, the unconditional choice probabilities are all the same, and thus equal $1/n$. The constant entropy of the unconditional choice probabilities, $\lambda\log(1/n)$, can be dropped.

Inner Maximization Problem Turning to the inner maximization problem first, we can exchange the order of integration and summation in (2.5). We first relax the maximization problem by dropping the exchangeability constraint. Then, the inner objective is separable by state as $p_\beta(\omega)$ depends only on the vector of state-wise utilities, $u(\omega)$. Below, we show that any solution to the state-wise problem

$$\max_{u(\cdot,\omega)\colon A\to\mathbb{R}} \int \left(\lambda H(p_\beta(\omega)) + \sum_{a\in A} p_\beta(a|\omega) v(a,\omega) \right) dF(\beta) \qquad (2.6)$$

preserves the strict order of utilities and that a solution exists. Then, we go on to show that there is an exchangeable utility function that maximizes (2.6) state-by-state. By every state having positive probability, this implies that any solution to the constrained inner-maximization problem (including the exchangeability constraint) must maximize the state-wise problem in every state.

Lemma 22. *Any solution to the state-wise maximization problem (2.6) satisfies*

$$\forall a, a' \in A\colon v(a,\omega) > v(a',\omega) \Rightarrow u(a,\omega) > u(a',\omega).$$

Proof. If $v(a,\omega) > v(a',\omega)$ but $u(a,\omega) < u(a',\omega)$, then we can improve the principal's utility by exchanging $u(a,\omega)$ and $u(a',\omega)$, that is setting $u(a,\omega)$ to the original value of $u(a',\omega)$ and vice versa. For any type β, this would not change the entropy term while improving the expected payoff v. That is because for any type β, exchanging these two u's only exchanges the probabilities $p_\beta(a|\omega)$ and $p_\beta(a'|\omega)$ due to (2.4). The entropy $H(p_\beta(u(\omega)))$ remains unchanged for any type β because it is invariant under permutations of the probabilities, which is immediate from the definition of entropy. The expected payoff $\sum_{\hat{a}\in A} p_\beta(\hat{a}|\omega)v(\hat{a},\omega)$ would be improved because the conditional choice probabilities of $v(a,\omega)$ and $v(a',\omega)$ are exchanged, giving higher probability to the higher payoff. Hence, at a maximum, $v(a,\omega) > v(a',\omega)$ implies $u(a,\omega) \geq u(a',\omega)$.

Suppose for the sake of contraction that $v(a,\omega) > v(a',\omega)$ and $u(a,\omega) = u(a',\omega)$. Set $\tilde{u}(a,\omega) = u(a,\omega) + \varepsilon$ and $\tilde{u}(a',\omega) = u(a',\omega) - \varepsilon$ and $\tilde{u}(\hat{a},\omega) = u(\hat{a},\omega)$ for all $\hat{a} \neq a, a'$. Fix any type $\beta \in (0,1]$. At $\varepsilon = 0$, $p_\beta(a|u(\omega)) = p_\beta(a'|u(\omega))$ by (2.4). By symmetry of (2.4) in the u's, at $\varepsilon = 0$, for all $\hat{a} \neq a, a'$: $\frac{dp_\beta(\hat{a}|\omega)}{du(a,\omega)} = \frac{dp_\beta(\hat{a}|\omega)}{du(a',\omega)}$ and hence $\frac{dp_\beta(\hat{a}|\omega)}{d\varepsilon} = \frac{dp_\beta(\hat{a}|\omega)}{du(a,\omega)} - \frac{dp_\beta(\hat{a}|\omega)}{du(a',\omega)} = 0$ and $\frac{dp_\beta(a|\omega)}{d\varepsilon} = -\frac{dp_\beta(a'|\omega)}{d\varepsilon}$. By symmetry of entropy in the conditional choice probabilities, at $\varepsilon = 0$, $\frac{dH}{dp_\beta(a|\omega)} = \frac{dH}{dp_\beta(a'|\omega)}$ and hence $\frac{dH}{d\varepsilon} = \frac{dp_\beta(a'|\omega)}{d\varepsilon}\frac{dH}{dp_\beta(a'|\omega)} + \frac{dp_\beta(a|\omega)}{d\varepsilon}\frac{dH}{dp_\beta(a|\omega)} = 0$ by the equalities above. However, $v(a,\omega) - v(a',\omega) > 0$, so at $\varepsilon = 0$, the derivative of the expected payoff is positive by $\forall \hat{a} \neq a, a'$: $\frac{dp_\beta(\hat{a}|\omega)}{d\varepsilon} = 0$, $\frac{dp_\beta(a|\omega)}{d\varepsilon} > 0$ and $\frac{dp_\beta(a'|\omega)}{d\varepsilon} < 0$. Hence, increasing ε by some positive amount improves the principal's utility for all types β, which is why we must have $u(a,\omega) > u(a',\omega)$. □

Because of the genericity Assumption 5, the Lemma implies that

$$v(a,\omega) > v(a',\omega) \Leftrightarrow u(a,\omega) > u(a',\omega).$$

Before we show existence, we insert the agent's state-dependent choice probabilities (2.4) to simplify the integrand of the state-wise objective:

$$\begin{aligned}
&\sum_{a\in A}\Big(p_\beta(a|\omega)v(a,\omega)\Big) + \lambda H(p_\beta(u(\omega)))\\
=&\sum_{a\in A}\Big(p_\beta(a|\omega)v(a,\omega) - \lambda p_\beta(a|\omega)\log(p_\beta(a|\omega))\Big)\\
=&\sum_{a\in A}\left(p_\beta(a|\omega)v(a,\omega) - \lambda p_\beta(a|\omega)\log\left(\frac{e^{\frac{\beta}{\lambda}u(a,\omega)}}{\sum_{a'\in A}e^{\frac{\beta}{\lambda}u(a',\omega)}}\right)\right)\\
=&\sum_{a\in A}\left(p_\beta(a|\omega)\big(v(a,\omega)-\beta u(a,\omega)\big) + \lambda\log\left(\sum_{a'\in A}e^{\frac{\beta}{\lambda}u(a',\omega)}\right)\right)
\end{aligned}$$

Thus, the state-wise objective, which we denote by $G(u(\omega))$, is

$$G(u(\omega)) = \int\left(\sum_{a\in A}p_\beta(a|u(\omega))\big(v(a,\omega)-\beta u(a,\omega)\big) + \lambda\log\left(\sum_{a\in A}e^{\frac{\beta}{\lambda}u(a,\omega)}\right)\right)dF(\beta). \quad (2.7)$$

It will be useful to know the derivatives of G and $p_\beta(a|\omega)$ with respect to the components of u, which have simple expressions. Using

$$p_\beta(a|\omega) = \frac{d}{du(a,\omega)}\log\left(\sum_{a'\in A}e^{\frac{\beta}{\lambda}u(a',\omega)}\right),$$

which follows from (2.4), one can show by straightforward differentiation

$$\frac{dG(u(\omega))}{du(a,\omega)} = \int\left(\sum_{a'\in A}\frac{dp_\beta(a'|\omega)}{du(a,\omega)}\big(v(a',\omega)-\beta u(a',\omega)\big)\right)dF(\beta) \quad (2.8)$$

$$\frac{dp_\beta(a|\omega)}{du(a,\omega)} = \frac{\beta}{\lambda}p_\beta(a|\omega)(1-p_\beta(a|\omega)) = \frac{\beta}{\lambda}p_\beta(a|\omega)\sum_{a'\neq a}p_\beta(a'|\omega) \quad (2.9)$$

$$\frac{dp_\beta(a'|\omega)}{du(a,\omega)} = -\frac{\beta}{\lambda}p_\beta(a'|\omega)p_\beta(a|\omega),\ \text{for } a\neq a'. \quad (2.10)$$

Lemma 23. *The state-wise maximization problem (2.6) has a solution.*

Proof. We prove the statement by showing that any $u(\cdot,\omega) \in \mathbb{R}^A$ is weakly dominated (in the sense of achieving a lesser or equal objective) by some $\tilde{u}(\cdot,\omega) \in \mathbb{R}^A$ in a fixed compact subset of $\mathbb{R}^A$. Thus, we can restrict attention to this compact set, on which a maximum (which must be a global maximum) is obtained by continuity of the objective.

Without loss, assume $v(a_1,\omega) \leq ... \leq v(a_n,\omega)$. We can also assume without loss that $v(a_1,\omega) = u(a_1,\omega) = 0$. That is because from the state-wise objective (2.6) it is clear that shifting all $u(a,\omega)$ by the same constant does not affect choice probabilities and hence also not the objective. Similarly, shifting all $v(a,\omega)$ by some constant c changes the objective exactly by that constant since $\sum_{a\in A} p_\beta(a|\omega)(v(a,\omega)+c) = c + \sum_{a\in A} p_\beta(a|\omega)v(a,\omega)$. Hence, the constant does not affect the set of maximizers.

By Lemma 23, if for some a, a', $v(a,\omega) > v(a',\omega)$ and $u(a',\omega) > u(a,\omega)$, then we can exchange $u(a,\omega)$ and $u(a',\omega)$ and improve the principal's utility. If $v(a,\omega) = v(a',\omega)$ and $u(a',\omega) > u(a,\omega)$, then exchanging $u(a,\omega)$ and $u(a',\omega)$ does not change the principal's utility. Thus, every u that does not satisfy $u(a_1,\omega) \leq ... \leq u(a_n,\omega)$ is weakly dominated by a $\tilde{u}$ that does satisfy it. Thus, we can restrict attention to the set of utility-vectors u with $0 = u(a_1,\omega) \leq ... \leq u(a_n,\omega)$.

Finally, we bound $u(a_i,\omega)$ for $i = 2,...,n$ from above by $u(a_{i-1},\omega) + \frac{v(a_n,\omega)}{\inf\beta}$ ($\frac{1}{\inf\beta}$ exists by Assumption 3). Suppose for some $m \in \{2,...,m\}$, $u(a_m,\omega)$ violated the bound. Then, for all $i \geq m, k \leq m-1$:

$$u(a_i,\omega) \geq u(a_m,\omega) > u(a_{m-1},\omega) + \frac{v(a_n,\omega)}{\inf\beta} \geq u(a_k,\omega) + \frac{v(a_n,\omega)}{\inf\beta}. \tag{2.11}$$

Using (2.8) through (2.10), the derivative of decreasing $u(a_m, \omega)$ through $u(a_n, \omega)$ by the same amount is

$$\begin{aligned}
&\left.\frac{dG(u(a_1,\omega),...,u(a_{m-1},\omega),u(a_m,\omega)-x,...,u(a_n,\omega)-x)}{dx}\right|_{x=0} = -\sum_{i=m}^{n}\frac{dG(u)}{du(a_i,\omega)} \\
&= \sum_{i=m}^{n}\int \frac{\beta}{\lambda}p_\beta(a_i|\omega)\left(\sum_{\substack{k=1,\\k\neq i}}^{n} p_\beta(a_k|\omega)\Big[\big(v(a_k,\omega)-\beta u(a_k,\omega)\big)-\big(v(a_i,\omega)-\beta u(a_i,\omega)\big)\Big]\right)dF(\beta) \\
&= \sum_{i=m}^{n}\int \frac{\beta}{\lambda}p_\beta(a_i|\omega)\left(\sum_{k=1}^{m-1} p_\beta(a_k|\omega)\Big[\big(v(a_k,\omega)-\beta u(a_k,\omega)\big)-\big(v(a_i,\omega)-\beta u(a_i,\omega)\big)\Big]\right)dF(\beta) \\
&\geq \sum_{i=m}^{n}\int \frac{\beta}{\lambda}p_\beta(a_i|\omega)\left(\sum_{k=1}^{m-1} p_\beta(a_k|\omega)\Big[\beta\big(u(a_i,\omega)-u(a_k,\omega)\big)-v(a_n,\omega)\Big]\right)dF(\beta) \\
&= \sum_{i=m}^{n}\int \frac{\beta}{\lambda}p_\beta(a_i|\omega)\beta\left(\sum_{k=1}^{m-1} p_\beta(a_k|\omega)\Big[u(a_i,\omega)-u(a_k,\omega)-\frac{v(a_n,\omega)}{\beta}\Big]\right)dF(\beta),
\end{aligned}$$

which is strictly positive since the term in the brackets is strictly positive by (2.11). Thus, we can decrease $u(a_m, \omega)$ through $u(a_n, \omega)$ until $u(a_i, \omega) \leq u(a_{i-1}, \omega) + \frac{v(a_n,\omega)}{\inf\beta}$, while improving the objective. Thus, each $u(a_i, \omega)$ can be bounded by $(i-1)\frac{v(a_n,\omega)}{\inf\beta}$, which defines a compact set. □

Lemma 24. *There exists an exchangeable utility function $u\colon A \times \Omega \to \mathbb{R}$ that maximizes (2.6) state-by-state.*

Proof. We construct a utility function that maximizes (2.6) state-by-state and which inherits exchangeability from v. The state-wise objective (2.6) is symmetric in the components of $v(\cdot, \omega)$. Thus, permuting the components of $v(\cdot, \omega)$ leads to a permutation of the components of the solution candidates $u(\cdot, \omega)$. Define an equivalence relation on Ω: two states are equivalent if their payoff-vectors $v(\cdot, \omega)$ are a permutation of each other. For each equivalence class, pick a representative state and a solution $u(\cdot, \omega)$ of the state-wise objective given that state's payoff-vector. Then, for other states ω' in the equivalence class, assign the analogous permutation of $u(\cdot, \omega)$, which is a solution given $v(\cdot, \omega')$. This procedure is well-defined

because by Assumption 5, for any state ω, all components of $v(\cdot, \omega)$ are pairwise different. Thus, any permutation of $v(\cdot, \omega)$ can be obtained in a *unique* way by a sequence of pairwise permutations of components. The same holds for all components of any state-wise solution $u(\cdot, \omega)$, as by by Lemma 22 all components are pairwise different as well. On the equivalence class, any permutation of the state-wise solution of the representative state is thus as likely as any permutation by exchangeability of v. Doing the above procedure for every equivalence class, the resulting u is exchangeable. □

We have shown thus shown that the inner maximization problem separates into state-wise maximization of (2.6). Next, we show that the overall principal's problem separates by state.

State-Wise Principal's Problem Because the integral over states is a finite sum, we can exchange the order of integration in the implementation cost, which becomes

$$\sum_{\omega \in \Omega} \mu(\omega) \int \left(\sum_{a \in A} p_\beta(a|u(\omega)) \psi\big(u(a_i, \omega) - v(a_i, \omega)\big) \right) dF(\beta).$$

Because by constraint u is exchangeable, $p_\beta(a|u(\omega))$ depends only on the state-wise utility vector $u(\omega)$ by (2.4), so the outer objective separates by state. By the exchangeability assumption on v, one can make a symmetry argument that there is a state-wise minimizer u subject to it being among the state-wise maximizers of (2.6), such that the actions are exchangeable with respect to u. It follows that if a solution to the principal's problem exists, it must solve for every state $\omega \in \Omega$ the **state-wise principal's problem**:

$$\min_{u(\cdot,\omega)\colon A \to \mathbb{R}} \int \left(\sum_{a \in A} p_\beta(a|u(\omega)) \psi\big(u(a, \omega) - v(a, \omega)\big) \right) dF(\beta)$$

s.t.

$$u(\cdot, \omega) \in \operatorname*{arg\,max}_{u(\cdot,\omega)\colon A \to \mathbb{R}} \int \left(-\lambda H(p_\beta(u(\omega))) + \sum_{a \in A} p_\beta(a|u(\omega)) v(a, \omega) \right) dF(\beta)$$

Lemma 25. *A solution to the state-wise principal's problem exists. The solution is unique if the solution to the inner maximization problem is unique up to a common additive constant.*

Proof. We show again that we can restrict attention to a compact subset of $\mathbb{R}^A$, on which the minimum is obtained by continuity of the objective (ψ is convex implies that ψ is continuous).

First, we reparametrize $(u(a_1,\omega),...,u(a_n,\omega))$ by

$$(u_1, \Delta u_2, ..., \Delta u_n) := (u(a_1,\omega), u(a_2,\omega) - u(a_1,\omega), ..., u(a_n,\omega) - u(a_1,\omega)).$$

Using $u(a_i,\omega) = u_1 + \Delta u_i$, the outer objective can be rewritten as

$$\sum_{i=1}^{n} \psi\big(u_1 + \Delta u_i - v(a_i,\omega)\big) \int p_\beta(a_i|u(\omega))dF(\beta) \tag{2.12}$$

where $p_\beta(a_i|u(\omega))$ depends only on $(\Delta u_2, ..., \Delta u_n)$. By assumption, ψ is convex and minimized at 0. Thus, keeping $(\Delta u_2, ..., \Delta u_n)$ fixed, the objective is radially unbounded in u_1 as it is increasing for u_1 to the right of the interval

$$\big[- \max_{i=2,...,n} \{\Delta u_i - v(a_i,\omega)\}, - \min_{i=2,...,n} \{\Delta u_i - v(a_i,\omega)\}\big]$$

and decreasing when u_1 is to the left of the interval. Hence, because Δu_i is bounded by the proof in Lemma 23, this implies a bound on u_1. We know that u_1 does not affect the value of the inner objective because it does not affect the choice probabilities, so from the inner maximization problem, there is no other constraint on u_1. Further, by the proof of Lemma 23, $(\Delta u_2, ..., \Delta u_n)$ is in a compact set. By Berge's maximum theorem, the set of maximizers form a compact set themselves. Hence, we can restrict attention to a product of two compact sets, which is compact, on which the minimum is obtained.

For the second statement of the Lemma, if the solution to $(\Delta u_2, ..., \Delta u_n)$ from the inner maximization problem is unique, then the minimization problem above gives a unique solu-

tion for u_1 because equation (2.12) is a strictly convex function in u_1 keeping $(\Delta u_2, ..., \Delta u_n)$ fixed. □

For $c \in \mathbb{R}$, let $\mathbf{c} := (c, ..., c) \in \mathbb{R}^A$ (with slight abuse of notation).

Lemma 26. *$u(\cdot, \omega) \in \mathbb{R}^A$ solves the state-wise principal's problem given the payoff vector $v(\cdot, \omega) \in \mathbb{R}^A$ if and only if $u(\cdot, \omega) + \mathbf{c}$ does so given payoff vector $v(\cdot, \omega) + \mathbf{c}$.*

Proof. As noted in Lemma 23, shifting all components of $v(\cdot, \omega)$ or $u(\cdot, \omega)$ by constants does not affect the inner maximization problem. The solutions to the inner maximizaton problem given $v(\cdot, \omega)$ and $v(\cdot, \omega) + \mathbf{c}$ are the same as this change only moves the principal's utility by c independent of the agent's choice. And, if $u(\cdot, \omega)$ is a solution to the inner maximization problem, then $u(\cdot, \omega) + \mathbf{c}$ is, too, since shifting all components of $u(\cdot, \omega)$ by a constant does not affect the principal's utility as it does not affect the agent's choice probabilities. Thus, if $u(\cdot, \omega)$ is in the constraint set defined by the inner maximization problem given $v(\cdot, \omega)$, then $u(\cdot, \omega) + \mathbf{c}$ is in the said set given $v(\cdot, \omega) + \mathbf{c}$.

Second, the outer objective is unchanged if $\mathbf{c}$ is added to both $u(\cdot, \omega)$ and $v(\cdot, \omega)$ because the agent's choice probabilities are unaffected and $\psi(u(a_i, \omega) - v(a_i, \omega))$ is unaffected. If $u(\cdot, \omega)$ minimizes the implementation cost given $v(\cdot, \omega)$, then $u(\cdot, \omega) + \mathbf{c}$ minimizes it given $v(\cdot, \omega) + \mathbf{c}$ over the original constraint set plus $\mathbf{c}$. However, this is the actual constraint set, by the above. The other direction follows from shifting by $-\mathbf{c}$. □

Solution to Principal's Problem

Lemma 27. *There exists a solution to the principal's problem such that there is a function R, such that for all $(a, \omega) \in A \times \Omega$,*

$$u(a, \omega) = v(a, \omega) + R(\{v(a, \omega) - v(a', \omega)\}_{a' \in A}).$$

Proof. Define a partition $\mathbb{P}$ on the states based on the equivalence relation that the payoff vector $v(\cdot, \omega) \in \mathbb{R}^A$ of two states are the same up to permutations of its components, shifting

all components by a constant, and combinations of both. Formally, $\omega \sim \omega'$ if there exists a permutation $\tilde{v}(\cdot, \omega)$ of $v(\cdot, \omega)$ and a $c \in \mathbb{R}$ such that $v(\cdot, \omega) = \tilde{v}(\cdot, \omega) + \mathbf{c}$. We proceed by constructing $u(\cdot, \omega)$ on each partition element.

Let $P \in \mathbb{P}$ be some partition element and ω be some state in P. Define $u(\cdot, \omega)$ to be a solution to the state-wise principal's problem, which exists by Lemma 25. So, for each other state ω' in P, there exist a $c \in \mathbb{R}$ and a permutation $\tilde{v}(\cdot, \omega)$ of the payoff vector $v(\cdot, \omega)$, such that $v(\cdot, \omega') = \mathbf{c} + \tilde{v}$. Moreover, the constant c and the way $v(\cdot, \omega)$ is permuted are unique, the latter because by Assumption 5. Define $u(\cdot, \omega')$ to be the analogous permutation of $u(\cdot, \omega)$ plus $\mathbf{c}$. This gives a solution to the state-wise principal's problem under ω' because the state-wise principal's problem is the same up to this permutation and by Lemma 26. Because the actions are exchangeable with respect to $v(\cdot, \omega)$, each state in partition element in P that has $c = 0$ relative to ω (that is a "pure permutation") has the same probability. Hence, by construction $u(\cdot, \omega)$ is a pure permutation on these elements, so the actions are exchangeable with respect to $u(\cdot, \omega)$. Thus, the constructed $u(\cdot, \omega)$ solves the principal's problem.

Further, $u(\cdot, \omega)$ is a function only of the payoff vector $v(\cdot, \omega)$. Denote this function by $M \colon \mathbb{R}^A \to \mathbb{R}^n$ and by M_i the $i - th$ component of M. Define for all $(v_1, ..., v_{n-1}) \in \mathbb{R}^A$

$$\tilde{R}(v_1, ..., v_{n-1}) := M_n(v_1, ..., v_{n-1}, 0).$$

By construction, a permutation of the arguments of M permutes the image in an analogous way. Formally, for all pairwise different $i, j, k \in \{1, ..., n\}$

$$M_i(v_1, ..., v_i, ..., v_j, ..., v_n) = M_j(v_1, ..., v_j, ..., v_i, ..., v_n)$$
$$M_k(v_1, ..., v_i, ..., v_j, ..., v_n) = M_k(v_1, ..., v_j, ..., v_i, ..., v_n)$$

This implies that $\tilde{R}$ is invariant under relabeling of its components, so we can define

$$R(\{v_1, ..., v_n\}) := \tilde{R}(-v_1, ..., -v_n).$$

We have by Lemma 26,

$$\begin{aligned}
&u(a_i, \omega) = M_i(v(a_1, \omega), ..., v(a_2, \omega)) \\
&= v(a_i, \omega) + M_i(v(a_1, \omega) - v(a_i, \omega), ..., v(a_i, \omega) - v(a_i, \omega), ..., v(a_2, \omega) - v(a_i, \omega)) \\
&= v(a_i, \omega) + M_n(v(a_1, \omega) - v(a_i, \omega), ..., v(a_2, \omega) - v(a_i, \omega), 0) \\
&= v(a_i, \omega) + \tilde{R}(v(a_1, \omega) - v(a_i, \omega), ..., v(a_2, \omega) - v(a_i, \omega)) \\
&= v(a_i, \omega) + R(\{v(a_i, \omega) - v(a', \omega)\}_{a' \in A \setminus \{a_i\}})
\end{aligned}$$

□

Note that the constructed solution to the state-wise principal's problem depends only on the state-wise actual and counterfactual payoffs as well as the distribution over β and λ. Thus, the solution does not depend on the whole on the decision-problem that the principal delegates.

Proof of Theorem 8

Proof. By Theorem 9 there is a solution with

$$u(a, \omega) = v(a, \omega) + R(v(a, \omega) - v(a_-, \omega)), \tag{2.13}$$

if $v(a, \omega) - v(a_-, \omega) \neq 0$ for all states (because we have assumed strict preferences of the principal for that theorem). First, we show that this statement extends to cases where $v(a, \omega) - v(a_-, \omega) = 0$ for some states. Next, we show that any two such solutions have R and Q that are unique up to a measure zero set in $\mathbb{R}$ (with respect to the Lebesgue measure). Then, we show that any such solution satisfies Properties 1 and 2.

Extending Theorem 1 By Theorem 9, a solution to the principal's problem exists and every solution must, in every state, solve the state-wise principal's problem, which is

$$\min_{u(\cdot,\omega)\colon A\to\mathbb{R}} \int \left(\sum_{i=1}^{2} p_\beta(a_i|u(\omega))\psi\big(u(a_i,\omega) - v(a_i,\omega)\big) \right) dF(\beta)$$

s.t.

$$u(\cdot,\omega) \in \operatorname*{arg\,max}_{u(\cdot,\omega)\colon A\to\mathbb{R}} \int \left(\lambda H(p_\beta(u(\omega))) + \sum_{i=1}^{2} p_\beta(a_i|u(\omega))v(a_i,\omega) \right) dF(\beta)$$

If $v(a_1,\omega) = v(a_2,\omega)$, the payoff $\sum_{i=1}^{2} p_\beta(a_i|u(\omega))v(a_i,\omega)$ is unaffected by the conditional choice probabilities $p_\beta(a_i|u(\omega))$. The entropy term is maximized if

$$p_\beta(a_1|u(\omega)) = p_\beta(a_2|u(\omega)) = 1/2,$$

which is achieved for all types if and only if $u(a_1,\omega) = u(a_2,\omega)$. In that case, because the implementation cost is minimized at 0, the unique solution is

$$u(a_1,\omega) = u(a_2,\omega) = v(a_1,\omega) = v(a_2,\omega).$$

Take our solution from Theorem 9 for states where $v(a_1,\omega) \neq v(a_2,\omega)$, we can extend it for states where $v(a_1,\omega) = v(a_2,\omega)$ and maintain exchangeability, so it must be a solution.

Uniqueness By (2.7), the inner objective can be rewritten as

$$\int \left(\sum_{i=1}^{2} p_\beta(a_i|u(\omega))\big(v(a_i,\omega) - \beta u(a_i,\omega)\big) + \lambda \log\left(\sum_{i=1}^{2} e^{\frac{\beta}{\lambda}u(a_i,\omega)} \right) \right) dF(\beta).$$

Using $p_\beta(a_1|\omega) = 1 - p_\beta(a_2|\omega)$, we can rewrite the objective as

$$\int p_\beta(a_2|u(\omega))\Big[v(a_2,\omega) - v(a_1,\omega) - \beta\big(u(a_2,\omega) - u(a_1,\omega)\big)\Big] + v(a_1,\omega) - \beta u(a_1,\omega) + \lambda\left(\frac{\beta}{\lambda}u(a_1,\omega) + \log\left(1 + e^{\frac{\beta}{\lambda}(u(a_2,\omega)-u(a_1,\omega))}\right)\right) dF(\beta)$$

Defining $\Delta u(\omega) = u(a_2, \omega) - u(a_1, \omega)$ and $\Delta v(\omega) = v(a_2, \omega) - v(a_1, \omega)$, and subtracting the constant $v(a_1, \omega)$, we obtain that the objective is

$$\int \Big(p_\beta(a_2|u(\omega))\big(\Delta v(\omega) - \beta \Delta u(\omega)\big) + \lambda \log\Big(1 + e^{\frac{\beta}{\lambda}\Delta u(\omega)}\Big)\Big)\, dF(\beta) \tag{2.14}$$

Note that $p_\beta(a_2|u(\omega))$ depends only on $\Delta u(\omega)$ by (2.4), so with slight abuse of notation, we write $p_\beta(a_2|\Delta u(\omega))$. The first-order condition is

$$\int \left(\frac{dp_\beta(a_2|\Delta u(\omega))}{d\Delta u(\omega)} \big(\Delta v(\omega) - \beta \Delta u(\omega)\big) \right) dF(\beta) = 0, \tag{2.15}$$

which can be rewritten as

$$\Delta u(\omega) = \Delta v(\omega) \left(\frac{\int \frac{dp_\beta(a_2|\Delta u(\omega))}{d\Delta u(\omega)} \beta dF(\beta)}{\int \frac{dp_\beta(a_2|\Delta u(\omega))}{d\Delta u(\omega)} dF(\beta)} \right)^{-1}. \tag{2.16}$$

as used in the main text.

The solution $\Delta u(\omega)$ need not be unique. However, using the supermodularity of the objective in $\Delta u(\omega)$ and $\Delta v(\omega)$, we show that there is a unique maximizer $\Delta u(\omega)$ for almost all $\Delta v(\omega)$.

Lemma 28. *The set of solutions $\Delta u(\omega)$ as a function of $\Delta v(\omega)$ are increasing in the strong set order. For almost all $\Delta v(\omega)$, the solution $\Delta u(\omega)$ is unique.*

Proof. The objective (2.14) has a positive cross-derivative in $\Delta v(\omega)$ and $\Delta u(\omega)$, which, using (2.15), is

$$\int \frac{dp_\beta(a_2|\Delta u(\omega))}{d\Delta u(\omega)} dF(\beta) > 0.$$

It is positive because $p_\beta(a_2|\Delta u(\omega))$ is strictly increasing in $\Delta u(\omega)$ for all $\beta > 0$. Thus, by Topkis characterization theorem (Topkis, 1978), the objective is supermodular in $\Delta u(\omega)$ and $\Delta v(\omega)$. By Topkis theorem (ibd.), the set of $\Delta u(\omega)$ that maximize the objective as a function of $\Delta v(\omega)$ is increasing in the strong set order. Since the cross-partial is strictly

positive, an even stronger statement holds: if $\Delta u'' > \Delta u'$ are both maximizers at $\Delta v(\omega)$, then at any $\Delta v(\omega)' > \Delta v(\omega)$, $\Delta u''$ achieves a strictly higher payoff than $\Delta u'$. That is because the derivative of the objective between $\Delta u'$ and $\Delta u''$ is strictly higher at any point, so the integral over the derivative is now positive, where it was 0 before. Thus, any selection function of maximizers as a function of $\Delta v(\omega)$ must have a jump discontinuity wherever the set of maximizers is not a singleton. Since monotonic functions can have at most countably many jump discontinuities, there are at most countably many $\Delta v(\omega)$ for which there is not a unique maximizer. Thus, for almost all $\Delta v(\omega)$, the solution is unique. □

By Lemma 25 and 28, the solution to the outer minimization problem subject to the inner maximization problem is unique for almost all $\Delta v(\omega)$. Thus, $R(x)$ is unique for almost all $x \in \mathbb{R}$ and thereby also $Q(x) := x + R(x) - R(-x)$, which equals $\Delta u(\omega)$ given $\Delta v(\omega) = x$.

Property 1 Above, we have already argued that $R(0) = 0$. First, we show that $R(x) \geq 0$ for $x > 0$. Note that equation (2.16) show that $Q(x) - x > 0$ when $x > 0$. Thus, if $R(x)$ was negative then $R(x)$ and $R(-x)$ were negative by $R(x) - R(-x) = Q(x) - x$. Then, we could increase $R(x)$ and $R(-x)$ by the same small amount (thus maintaining $Q(x)$ being a solution to the state-wise maximization problem) and decrease the implementation cost by ψ being strictly convex and minimal at 0. Analogously, one can show that $R(x) \leq 0$ if $x < 0$.

If $R(x)$ was 0, then and $R(-x)$ would be negative by $R(x) - R(-x) = Q(x) - x$. Because ψ is differentiable and has a global minimum at 0, it must have derivative 0 at 0. Then, we could increase $R(x)$ and $R(-x)$ by the some small enough amount (thus maintaining $Q(x)$ being a solution to the state-wise maximization problem) and decrease the implementation cost, by ψ being strictly convex and having derivative 0 at 0. Analogously, one can show that $R(x) < 0$ if $x < 0$.

Property 2 As argued in the main text, Property 2 is equivalent to $Q(x)/x$ being strictly increasing on $\mathbb{R}_{>0}$. We show that any $Q(x) := x + R(x) - R(-x)$ satisfies this property.

Rearranging (2.16), we get

$$\frac{Q(x)}{x} = \left(\frac{\int \frac{dp_\beta(a_2|Q(x))}{dQ(x)} \beta dF(\beta)}{\int \frac{dp_\beta(a_2|Q(x))}{dQ(x)} dF(\beta)} \right)^{-1}. \tag{2.17}$$

Because β has a discrete distribution by Assumption 3, the integral $\int g(\beta)dF(\beta)$ over some measurable function $g(\beta)$ equals $\sum_{\beta \in \text{supp}(\beta)} g(\beta)f(\beta)$ with probability mass function f. The term in parentheses in (2.17) is the expectation of β with respect to the probability mass function

$$\frac{\frac{dp_\beta(a_2|Q(x))}{dQ(x)} f(\beta)}{\int \frac{dp_b(a_2|Q(x))}{dQ(x)} dF(b)}.$$

Since we are considering $x > 0$, we have $Q(x) > 0$. Lemma 29 below shows that this probability mass function undergoes a strict likelihood-ratio shift downwards as $Q(x)$ increases for $Q(x) > 0$. Because $Q(x)$ is strictly increasing by Lemma 28 and because strict likelihood-ratio dominance implies strict first-order stochastic dominance, this implies that the term in parentheses in (2.17) is strictly decreasing in x. Thus, $Q(x)/x$ is strictly increasing in x.

Lemma 29. *The probability mass function*

$$\frac{\frac{dp_\beta(a_2|Q(x))}{dQ(x)} f(\beta)}{\int \frac{dp_b(a_2|Q(x))}{dQ(x)} dF(b)}$$

undergoes a strict likelihood-ratio shift downwards as $Q(x)$ increases for $Q(x) > 0$.

Proof. We need to show that for all $1 \geq \alpha > \beta > 0$,

$$\frac{\left(\frac{\frac{dp_\alpha(a_2|Q(x))}{dQ(x)} f(\alpha)}{\int \frac{dp_b(a_2|Q(x))}{dQ(x)} dF(b)} \right)}{\left(\frac{\frac{dp_\beta(a_2|Q(x))}{dQ(x)} f(\beta)}{\int \frac{dp_b(a_2|Q(x))}{dQ(x)} dF(b)} \right)}$$

is decreasing in $Q(x)$. Canceling out the integration constants in the denominators, this term is equal to

$$\frac{\frac{dp_\alpha(a_2|Q(x))}{dQ(x)}f(\alpha)}{\frac{dp_\beta(a_2|Q(x))}{dQ(x)}f(\beta)} = \frac{\left(\frac{\frac{\alpha}{\lambda}e^{\frac{\alpha}{\lambda}Q(x)}f(\alpha)}{\left(1+e^{\frac{\alpha}{\lambda}Q(x)}\right)^2}\right)}{\left(\frac{\frac{\beta}{\lambda}e^{\frac{\beta}{\lambda}Q(x)}f(\beta)}{\left(1+e^{\frac{\beta}{\lambda}Q(x)}\right)^2}\right)} = \frac{\alpha f(\alpha)}{\beta f(\beta)}e^{\frac{\alpha-\beta}{\lambda}Q(x)}\left(\frac{1+e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\alpha}{\lambda}Q(x)}}\right)^2$$

Ignoring the constant $\frac{\beta f(\beta)}{\alpha f(\alpha)}$, the term is decreasing in $Q(x)$ if the logarithm of the term is decreasing in $Q(x)$.

$$\log\left(e^{\frac{\alpha-\beta}{\lambda}Q(x)}\left(\frac{1+e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\alpha}{\lambda}Q(x)}}\right)^2\right) = (\alpha-\beta)\frac{Q(x)}{\lambda} + 2\log\frac{1+e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\alpha}{\lambda}Q(x)}}$$

Differentiating with respect to $Q(x)$:

$$\begin{aligned}
&\frac{d}{dQ(x)}\left((\alpha-\beta)\frac{Q(x)}{\lambda} + 2\log\frac{1+e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\alpha Q(x)/\lambda}}\right) = \frac{\alpha-\beta}{\lambda} + 2\frac{\frac{d}{dQ(x)}\frac{1+e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\alpha}{\lambda}Q(x)}}}{\frac{1+e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\alpha}{\lambda}Q(x)}}} \\
&= \frac{\alpha-\beta}{\lambda} + 2\frac{\frac{\beta}{\lambda}e^{\frac{\beta}{\lambda}Q(x)}(1+e^{\frac{\alpha}{\lambda}Q(x)}) - \frac{\alpha}{\lambda}e^{\frac{\alpha}{\lambda}Q(x)}(1+e^{\frac{\beta}{\lambda}Q(x)})}{\left(1+e^{\frac{\alpha}{\lambda}Q(x)}\right)^2}\frac{1+e^{\frac{\alpha}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}} \\
&= \frac{\alpha-\beta}{\lambda} + 2\frac{\beta-\alpha}{\lambda}\frac{e^{\frac{\alpha+\beta}{\lambda}Q(x)}}{(1+e^{\frac{\beta}{\lambda}Q(x)})(1+e^{\frac{\alpha}{\lambda}Q(x)})} + 2\frac{\frac{\beta}{\lambda}e^{\frac{\beta}{\lambda}Q(x)} - \frac{\alpha}{\lambda}e^{\frac{\alpha}{\lambda}Q(x)}}{(1+e^{\frac{\alpha}{\lambda}Q(x)})(1+e^{\frac{\beta}{\lambda}Q(x)})}
\end{aligned}$$

Multiply by $\lambda(1+e^{\frac{\beta}{\lambda}Q(x)})(1+e^{\frac{\alpha}{\lambda}Q(x)}) > 0$ gives

$$\begin{aligned}
&(\alpha-\beta)(1+e^{\frac{\beta}{\lambda}Q(x)})(1+e^{\frac{\alpha}{\lambda}Q(x)}) + 2(\beta-\alpha)e^{\frac{\alpha+\beta}{\lambda}Q(x)} + 2(\beta e^{\frac{\beta}{\lambda}Q(x)} - \alpha e^{\frac{\alpha}{\lambda}Q(x)}) \\
&= (\alpha-\beta) + (\alpha+\beta)e^{\frac{\beta}{\lambda}Q(x)} - (\alpha+\beta)e^{\frac{\alpha}{\lambda}Q(x)} + (\beta-\alpha)e^{\frac{\alpha+\beta}{\lambda}Q(x)} \\
&= (\alpha+\beta)(e^{\frac{\beta}{\lambda}Q(x)} - e^{\frac{\alpha}{\lambda}Q(x)}) + (\alpha-\beta)(1-e^{\frac{\beta}{\lambda}Q(x)}e^{\frac{\alpha}{\lambda}Q(x)}).
\end{aligned}$$

By $\alpha, \beta > 0$ and $Q(x) > 0$, we have $e^{\frac{\alpha}{\lambda}Q(x)} > e^{\frac{\beta}{\lambda}Q(x)}$ and hence the first term is negative. By $\alpha > \beta$ and $Q(x) > 0$, we have $e^{\frac{\beta}{\lambda}Q(x)} > 1$, $e^{\frac{\alpha}{\lambda}Q(x)} > 1$, and $e^{\frac{\beta}{\lambda}Q(x)}e^{\frac{\alpha}{\lambda}Q(x)} > 1$. Hence, the second term is negative. This establishes the negative derivative in $Q(x)$. □

Thus, Q and R of the regret solution satisfy

$$Q(x) \in \arg\max_{y \in \mathbb{R}} \int \frac{e^{\beta y/\lambda}}{1 + e^{\beta y/\lambda}}(x - \beta y) + \lambda \log(1 + e^{\beta y/\lambda}) dF(\beta)$$
$$R(x) = \arg\min_{R \in \mathbb{R}} \left(\int p_\beta(Q(x)) dF(\beta) \right) \psi(R) + \left(1 - \int p_\beta(Q(x)) dF(\beta) \right) \psi(x + R - Q(x))$$

where $p_\beta(Q(x)) := \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}$. □

Proof of Proposition 5

Proof. Loomes and Sugden (1982) show that the Allais paradox and the common-ratio effect are equivalent to strict superadditivity,

$$\forall x_1, x_2 \colon 0 < x_2 < x_1 \to Q(x_1) - Q(x_1 - x_2) - Q(x_2) > 0, \tag{2.18}$$

and that simultaneous gambling and insurance is equivalent to

$$\forall x > 0 \colon Q(\frac{p}{1-p}x) \gtreqless \frac{p}{1-p}Q(x) \quad \text{if } p \gtreqless 0.5. \tag{2.19}$$

We show that (2.18) and (2.19) are implied by $Q(x)/x$ being strictly increasing on $\mathbb{R}_{>0}$ and that (2.19) implies that $Q(x)/x$ is strictly on $\mathbb{R}_{>0}$.

To show that (2.18) is implied by Property 2, let $x_1 > x_2 > 0$ and define $Q_1 := Q(x_1)$, $Q_2 := Q(x_2)$, and $Q_{12} := Q(x_1 - x_2)$. We can assume without loss $0 < x_1 - x_2 < x_2 < x_1$. That is because if $0 < x_2 < x_1 - x_2 < x_1$, we can simply exchange x_2 and $x_1 - x_2$ in (2.18).

By $Q(x)/x$ being strictly increasing on $\mathbb{R}_{>0}$:

$$\begin{aligned}
&\frac{Q_1}{x_1} > \frac{Q_2}{x_2} \\
\Rightarrow\ & Q_1 x_2 > Q_2 x_1 \\
\Rightarrow\ & Q_1 x_2 - Q_2 x_2 > Q_2 x_1 - Q_2 x_2 \\
\Rightarrow\ & (Q_1 - Q_2) x_2 > Q_2 (x_1 - x_2) \\
\Rightarrow\ & \frac{Q_1 - Q_2}{x_1 - x_2} > \frac{Q_2}{x_2}
\end{aligned}$$

Using again that $Q(x)/x$ is strictly increasing on $\mathbb{R}_{>0}$,

$$\frac{Q_2}{x_2} > \frac{Q_{12}}{x_1 - x_2}.$$

Together,

$$\frac{Q_1 - Q_2}{x_1 - x_2} > \frac{Q_2}{x_2} > \frac{Q_{12}}{x_1 - x_2} \Rightarrow Q_1 - Q_{12} - Q_2 > 0.$$

Equation (2.19) follows from $Q(x)/x$ strictly increasing on $\mathbb{R}_{>0}$ by

$$p > 0.5 \Rightarrow \frac{p}{1-p} > 1 \Rightarrow \frac{Q(\frac{p}{1-p}x)}{\frac{p}{1-p}x} > \frac{Q(x)}{x} \Rightarrow Q(\frac{p}{1-p}x) > \frac{p}{1-p}Q(x),$$

and analogously for $p < 0.5$. When $p = 0.5$, $\frac{p}{1-p} = 1$ and the statement trivially holds.

For the other direction, suppose $Q(x)/x$ was not strictly increasing on $\mathbb{R}_{>0}$. Then, there exists $0 < x < y$ such that $\frac{Q(y)}{y} \leq \frac{Q(x)}{x}$. Let $p = \frac{y/x}{1+y/x} \in (1/2, 1)$. Then, $\frac{p}{1-p} = \frac{y}{x}$ and thus

$$\frac{Q(\frac{p}{1-p}x)}{\frac{p}{1-p}x} \leq \frac{Q(x)}{x} \Rightarrow Q(\frac{p}{1-p}x) \leq \frac{p}{1-p}Q(x)$$

despite $p > 1/2$ violating (2.19). □

Proof of Corollary 7

Proof. Regarding the action-implementation stage, we need to establish that given the information the agent of type β acquires, he takes the principal's preferred action. That is, we show that the posterior of type β under some action $a_i \in A$ makes a_i the optimal choice for the principal. This posterior is proportional to the product of the prior and the likelihood, the latter of which is given by the state-dependent choice probabilities. Consider the posterior probability of some emotion-*exclusive* payoff vector $(v(a_i, \omega), v(a{-}i, \omega)) = (v_1, v_2)$ and its permutation (v_2, v_1) with $v_1 > v_2$. Both payoff vectors have the same prior probability by exchangeability. Because emotions preserve the state-wise ordering of payoffs, the conditional choice probability of any agent β of action a_i is higher than the one of a_{-i} by (2.4). Thus, the posterior probability of (v_1, v_2) conditional on a_i is higher than conditional on a_{-i}, contributing to a higher expected utility for the principal of a_i than a_{-i}. Integrating over all pairs of payoff vectors and their permutation, the expected utility of action a_i with respect to v is higher than the one of action a_{-i}. □

Proof of Corollary 8

Proof. Let $x > 0$ and denote $p_\beta(Q(x)) := \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}$. From the proof of Theorem 8, regret and rejoicing, $R(-x)$ and $R(x)$, minimize the implementation cost subject to being consistent with $Q(x)$:

$$\min_{R(x),R(-x)\in\mathbb{R}} \left(\int p_\beta(Q(x))dF(\beta)\right)\psi(R(x)) + \left(1 - \int p_\beta(Q(x))dF(\beta)\right)\psi(R(-x))$$

$$\text{s.t.}$$

$$x + R(x) - R(-x) = Q(x)$$

That x is positive implies $Q(x) > 0$, which in turn implies that for all $\beta > 0$:

$$p_\beta(Q(x)) = \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1 + e^{\frac{\beta}{\lambda}Q(x)}} > \frac{1}{2} > 1 - p_\beta(Q(x))$$
$$\Rightarrow \int p_\beta(Q(x))dF(\beta) > 1 - \int p_\beta(Q(x))dF(\beta)$$

Thus, in the implementation-cost minimization problem, the weight on $\psi(R(x))$ is greater than the weight on $\psi(R(-x))$. Suppose for the sake of contradiction that we had $R(x) > |R(-x)|$. By convexity of ψ and symmetry around 0, there is an $\varepsilon > 0$ such that the implementation cost decreases if $R(x)$ and $R(-x)$ are both lowered by ε, which maintains the constraint $x + R(x) - R(-x) = Q(x)$. □

Proof of Corollary 9

Proof. We show, first, that decreasing λ increases $Q(x)$ for all $x > 0$ and, second, that smaller λ and greater $Q(x)$ lead to a greater $R(-x)$.

First, decreasing λ has the same effect as increasing $Q(x)$ in Lemma 29 because both enter the density in (2.16) only through $\frac{Q(x)}{\lambda}$. Thus, decreasing λ leads to a likelihood-ratio shift downwards and by (2.16) to a greater $Q(x)$ for any $x > 0$.

Second, given $Q(x)$, $R(-x)$ is determined by the implementation-cost minimization problem. Inserting the constraint $x + R(x) + R(-x) = Q(x)$, we can rewrite the implementation cost Ψ with slight abuse of notation as a function of $R(-x), Q(x)$, and λ:

$$\Psi(R(-x), Q(x), \lambda) := \left(\int \frac{e^{\frac{\beta}{\lambda} Q(x)}}{1 + e^{\frac{\beta}{\lambda} Q(x)}} dF(\beta) \right) \psi(Q(x) - x + R(-x)) + \left(1 - \int \frac{e^{\frac{\beta}{\lambda} Q(x)}}{1 + e^{\frac{\beta}{\lambda} Q(x)}} dF(\beta) \right) \psi(R(-x))$$

We show that the Ψ satisfies strictly increasing differences in $(R(-x), (Q(x), -\lambda))$, which implies that the solution $R(-x)$ to the minimization problem is strictly smaller, when $Q(x)$ is greater and λ is smaller by Topkis Theorem. By Topkis Characterization Theorem, it is sufficient to show that $\frac{d^2\Psi}{dQ(x)dR(-x)} > 0$ and $-\frac{d^2C}{d\lambda dR(-x)} = \frac{d^2C}{d(-\lambda)dR(-x)} > 0$. We can restrict attention to $Q(x) > 0$, $R(-x) < 0$ and $R(x) = Q(x) - x + R(-x) > 0$, which are satisfied at

the optimum by Theorem 8 anyways.

$$\frac{d\Psi}{dR(-x)} = \left(\int \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}dF(\beta)\right)\psi'(Q(x)-x+R(-x)) + \left(1-\int \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}dF(\beta)\right)\psi'(R(-x))$$

$$\Rightarrow \frac{d^2\Psi}{dQ(x)dR(-x)} = \left(\int \frac{d}{dQ(x)}\frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}dF(\beta)\right)\psi'(Q(x)-x+R(-x)) + \left(\int \frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}dF(\beta)\right)\psi''(Q(x)-x+R(-x)) - \left(\int \frac{d}{dQ(x)}\frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}dF(\beta)\right)\psi'(R(-x))$$

By differentiability of the integrand in Q and boundedness of this derivative, we exchanged integration over β and differentiation by $Q(x)$ according to Leibniz rule. Since ψ is strictly convex and minimal at 0, we have $\psi'(x) \gtrless 0 \Leftrightarrow x \gtrless 0$ and $\psi''(x) > 0$. By the constraints on $Q(x)$ and $R(-x)$, we have $R(-x) < 0$ and $Q(x) - x + R(-x) > 0$, so $\psi'(R(-x)) < 0$ and $\psi'(Q(x) - x + R(-x)) > 0$. Finally, by

$$\frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}} > 0 \quad \text{and}$$
$$\frac{d}{dQ(x)}\frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}} > 0$$

for $Q(x) > 0$, we obtain that $\frac{d^2\Psi}{dQ(x)dR(-x)}$ is strictly positive.

Regarding λ,

$$\frac{d^2\Psi}{d\lambda dR(-x)} = \left(\int \frac{d}{d\lambda}\frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}dF(\beta)\right)\psi'(Q(x)-x+R(-x)) - \left(\int \frac{d}{d\lambda}\frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}}dF(\beta)\right)\psi'(R(-x))$$

By

$$\frac{d}{d\lambda}\frac{e^{\frac{\beta}{\lambda}Q(x)}}{1+e^{\frac{\beta}{\lambda}Q(x)}} < 0$$

and we obtain that $-\frac{d^2\Psi}{d\lambda dR(-x)} < 0$ and hence $\frac{d^2\Psi}{d(-\lambda)dR(-x)} > 0$. □

2.7 Online Appendix

Incomplete Feedback So far, we have assumed that the principal can choose emotions as an arbitrary function of the action and realized state. In many decision contexts, however, ex post only the payoff is revealed but not the full state. For example, if the agent pursues some economic project, the realized return might be observable but not the return had he pursued another project. In particular, when only the realized payoff $v(a, \omega)$ is learned but not the state ω, the principal cannot in general infer the payoffs of counterfactual actions. Thus, she cannot implement the general regret solution from the previous section. This section generalizes the principal's problem to account for incomplete feedback under the simplifying assumption of a known bias.

As noted earlier, under known bias it is possible to implement the first-best attention strategy by offering a contract that scales up the payoff-differences by $1/\beta$. However, to find the emotion scheme m with minimal variance, we need to determine *all* emotion schemes that implement the optimal attention strategy. Below, we argue how this exercise teaches us something about the nature of the optimal mechanism when the first-best attention strategy cannot be implemented due to an uncertain bias.

Ex-post feedback also plays an important role in the study of regret and disappointment. While classical regret theory (Bell, 1982; Loomes and Sugden, 1982) does not take into account feedback, several experiments have documented that regret affects decisions more if feedback is expected (Zeelenberg, 1999a). To accommodate for that evidence, Humphrey (2004) introduces a theory of feedback-conditional regret which employs two regret-rejoicing

functions, one for complete feedback and one for incomplete feedback. Gabillon (2020) proposes a regret theory for arbitrary feedback structures, where the regret-rejoicing function is applied to the difference between the actual payoff and the highest counterfactual payoff, given the posterior beliefs after feedback. Evidence from Camille et al. (2004) based on self-reported emotions, skin conductance responses, and choices suggests, however, that under partial feedback (only the payoff of the chosen action is learned) primarily disappointment and elation are experienced, while under complete feedback mainly regret and rejoicing are experienced. We are not aware of theories that combine regret and disappointment depending on the feedback structure.

To incorporate feedback into our model, we assume that the principal can choose emotions only as a function of the feedback that is received ex post. Formally, we assume the emotion scheme m is measurable with respect to a feedback structure $\mathcal{F}\colon A \times \Omega \to Z$ where Z is a message space, that is,

$$\forall a, a' \in A, \omega, \omega' \in \Omega\colon \mathcal{F}(a, \omega) = \mathcal{F}(a', \omega') \to m(a, \omega) = m(a', \omega').$$

We consider partitional feedback structures that reveal the payoff $v(a, \omega)$, the action a, and a partition $\mathbb{P}$ of the state space.[29] In one extreme case, the feedback structure reveals the state fully (the finest partition; complete feedback) and in the other extreme, only the payoff is revealed (the coarsest partition; partial feedback). In intermediate cases, some information about the state is revealed, such as general market conditions, while more specific information, such as how each particular economic project would have done, is not revealed.

In some decision-problems, however, the agent can back out more information about the state from the material payoff and the action. To rule out such cases, we make an additional assumption that the message space is a "connected" on each partition. To formulate this

[29]A partition of the state space is a set of mutually exclusive and exhaustive subsets $\mathbb{P} = \{P_1, ..., P_n\}$ of Ω. In the Appendix, we show that the resulting emotions are essentially of the same form, when we allow for other deterministic feedback structures that reveal at least the payoff. Further, none of the results below would change if the action was not revealed.

condition, we endow the message space, Z, with a symmetric relation $\sim$, which makes Z an undirected graph where $\sim$ stands for the edges of the graph. We define $z \sim z'$ if and only if there exists $\omega \in \Omega$ and $a, a' \in A$ such that $z = \mathcal{F}(a, \omega)$ and $z' = \mathcal{F}(a', \omega)$. Thus, two messages are directly connected by an edge if they are observed under the same state. Because the state-wise payoff-differences determine the incentives for attention, this is the relevant notion of connectedness that ensures that the emotions of the two messages have to be determined jointly. More specifically, the following condition ensure that the set of messages consistent with a partition cannot be divided into disjoint subsets such that emotions can be chosen independently on each subset.

Assumption 6 (Partitional Feedback Structure)**.** *There is a partition $\mathcal{P}$ of Ω such that*

$$\mathcal{F}(a, \omega) = (v(a, \omega), a, \mathcal{P}(\omega))$$

where $\mathcal{P}(\omega)$ is the partition cell that contains ω. We define Z as the image of $\mathcal{F}$ in $\mathbb{R} \times A \times \mathcal{P}$.

For each partition cell $P \in \mathcal{P}$, the subset of the message space that is consistent with this partition cell, $\{(v, a, P) \in Z | v \in \mathbb{R}, a \in A\}$, is a connected subgraph of Z.

For simplicity, we analyze the model under a known bias $\beta \in (0, 1)$. Thus, we consider a pure moral hazard problem without an adverse selection component.

Assumption 7. *The distribution of $\beta \in (0, 1)$ is degenerate.*

Under known β, we can generalize the model in two dimensions. First, we drop the exchangeability assumption on v and the restriction to mechanisms with exchangeable $v + m$. Second, we generalize the entropy-based attention cost to any posterior-separable cost function. To be coherent with the notation above, we continue to define the attention cost in terms of the state-dependent stochastic choice function.[30] For each state-dependent stochas-

[30]This is without loss if the function H below is strictly concave. As is well-known, if we defined the attention cost as a function of the acquired signal structure, we can without loss identify the signal with the action under strictly concave H.

tic choice function $p\colon \Omega \to \Delta(A)$, let $\tau(p) \in \Delta(\Delta(\Omega))$ denote the distribution over posteriors $\gamma \in \Delta(\Omega)$ induced by p via Bayesian updating.

Definition 10 (Posterior-Separable attention cost)**.** *The attention cost c is posterior-separable if*

$$c(p) = \mathbb{E}_{\tau(p)}[H(\mu) - H(\gamma)]$$

where $H\colon \Delta(\Omega) \to \mathbb{R}$ is strictly concave and differentiable on $\Delta(\Omega)$.

Posterior-separable attention costs arise naturally from several foundations, including information theory (Sims, 2003), sequential sampling (Morris and Strack, 2019; Hébert and Woodford, 2023), and constant marginal cost of experimentation (Pomatto et al., 2023). The function H can be understood as a measure of uncertainty of a distribution, with entropy as an example. Thus, under a posterior-separable cost, the agent incurs a cost proportional to the expected reduction of the uncertainty of their belief. That H is strictly concave ensures that more information (in the Blackwell sense) is more costly.

Assumption 8. *The attention cost $c(p)$ is posterior separable.*

Given incomplete feedback and the absence of adverse selection, the new principal's problem $\mathrm{P}^{\mathrm{inc}}$ is to pick the solution with minimal variance of emotions to:

$$\max_{m,p} -c(p) + \sum_{\omega\in\Omega}\sum_{a\in A} \mu(\omega)p(a|\omega)v(a,\omega) \tag{$\mathrm{P}^{\mathrm{inc}}$}$$

s.t.

$$p \in \arg\max_{\tilde{p}} -c(\tilde{p}) + \beta \sum_{\omega\in\Omega}\sum_{a\in A} \mu(\omega)\tilde{p}(a|\omega)\big(v(a,\omega) + m(a,\omega)\big) \tag{IC}$$

$$\sum_{\omega\in\Omega}\sum_{a\in A} \mu(\omega)p(a|\omega)m(a,\omega) = 0 \tag{BC}$$

$$\forall a, a' \in A, \omega, \omega' \in \Omega\colon \mathcal{F}(a,\omega) = \mathcal{F}(a',\omega') \to m(a,\omega) = m(a',\omega') \tag{F}$$

We introduce the following class of solutions, which we interpret below as a combination of regret and disappointment. To simplify the exposure, we assume that the first-best attention

strategy p, which maximizes the principal's objective, is unique. Further, we define the consideration set $S \subseteq A$ as the set of actions that have positive probability according to the first-best attention strategy p, formally $S = \{a \in A | \sum_{\omega \in \Omega} \mu(\omega) p(a|\omega) > 0\}$.

Definition 11 (Linear Regret-Disappointment Solution)**.** *A solution (m, p) to the principal's problem* $\mathrm{P}^{\mathrm{inc}}$ *is a linear regret-disappointment solution if for all $a \in S, \omega \in \Omega$:*

$$m(a, \omega) = \left(\frac{1}{\beta} - 1\right) \big(v(a, \omega) - v(\mathcal{P}(\omega))\big) \tag{2.20}$$

where $v(\mathcal{P}(\omega))$ is the expected payoff conditional on partition element $\mathcal{P}(\omega)$ with respect to the first-best attention strategy p,

$$v(\mathcal{P}(\omega)) = \frac{1}{\mu(\mathbb{P}(\omega))} \sum_{\omega' \in \mathbb{P}(\omega)} \sum_{a \in S} \mu(\omega') p(a|\omega') v(a, \omega').$$

Such a solution has the natural interpretation as setting for each partition cell a reference point or benchmark $v(\mathcal{P}(\omega))$, which corresponds to the expected payoff the agent should achieve – according to the first-best attention strategy – conditional on that partition cell.[31] Note that in contrast to other models of reference points (Kahneman and Tversky, 1979; Gul, 1991; Kőszegi and Rabin, 2006), the optimal gain-loss utility here is linear.

It is instructive to compare (2.20) to the regret theory by Gabillon (2020), where regret depends on the difference between the realized payoff and the highest payoff that could have been achieved through some action given the ex-post feedback about the state. As in their proposal, the reference point is updated on feedback about the state.[32] However, in our model, the reference point is not the highest attainable counterfactual payoff but instead the *expected* payoff conditional on ex-post feedback. Because this expected payoff is also affected by the payoffs of the chosen action *had different states realized*, one can interpret

[31] In the Appendix, we show that for non-quadratic costs of emotions, an analogous solution with a different reference point v is optimal.

[32] To be precise, in (2.20), the reference point is updated only on feedback about the state that is action-independent, namely what is learned through the partition of the state space. It is not updated on what is learned from the realized payoff $v(a, \omega)$.

the emotional payoff as a combination of regret and disappointment. Next, we make this intuition more clear.

Another, mathematically equivalent, formulation of (2.20) relates the emotional payoff to regret and disappointment. Let $\gamma\colon A \to \Delta(\Omega)$ map each action to the posterior associated with it via Bayesian updating, which we write $\gamma(\omega|a)$. We can define, *conditional on partition cell* $\mathbb{P}(\omega)$, the probability, $p(a|\mathbb{P}(\omega))$, and the expected payoff of, $v(a|\mathbb{P}(\omega))$, of a,

$$\begin{aligned} p(a|\mathbb{P}(\omega)) &:= \sum_{\omega' \in \mathcal{P}(\omega)} \frac{\mu(\omega')}{\mu(\mathbb{P}(\omega))} p(a|\omega') \\ v(a|\mathcal{P}(\omega)) &:= \sum_{\omega' \in \mathcal{P}(\omega)} \frac{\gamma(\omega|a)}{\gamma(\mathbb{P}(\omega)|a)} v(a,\omega). \end{aligned}$$

Given this notation, we can rearrange terms and rewrite equation (2.20) as

$$\begin{aligned} m(a,\omega) = \sum_{a \neq a'} p(a'|\mathcal{P}(\omega)) \left(\frac{1}{\beta} - 1\right) \big(v(a,\omega) - v(a'|\mathcal{P}(\omega))\big) \\ + p(a|\mathcal{P}(\omega)) \left(\frac{1}{\beta} - 1\right) \big(v(a,\omega) - v(a|\mathcal{P}(\omega))\big). \end{aligned} \tag{2.21}$$

The term in the first line can be interpreted as regret and rejoicing. The emotion depends on the weighted difference between the actual payoff and the expected payoffs of all other actions a', conditional on partition cell $\mathcal{P}(\omega)$. The expected payoff of each action is computed with respect to the first-best attention strategy. The regret-like comparisons to different actions are weighted by the probabilities of the respective actions, again conditional on the partition element $\mathcal{P}(\omega)$.

This is in contrast to a widely used theory of regret where the comparison is with respect to the ex-post optimal action (Quiggin, 1994; Sarver, 2008), which allows for regret only and no rejoicing. This proposal is arguably extreme when the agent has a large choice set and does not seem psychologically plausible when the ex-post optimal action was not even considered. Moreover, in some relevant applications as in finance, under no short-sale constraints the ex-post optimal return is infinite (Qin, 2020). In contrast, our prediction corresponds to the

proposal of Loomes and Sugden (1982) to generalize regret theory to multiple actions based on such a weighted sum of comparisons to *all* other actions, not just the ex-post optimal one. They highlight the importance of providing a theory of the appropriate action weights. Sugden (1986) writes that

> *"Any plausible theory of such weights must, I think, take some account of the extent to which an action is a serious candidate for choice or, to put it slightly differently, of the extent to which the individual could sensibly blame himself for not having chosen it."*

In our contexts, the action weights are naturally provided by the state-dependent choice probabilities under the first-best attention strategy. These choice probabilities capture with what probability the agent should have taken some action under the realized state. This implies that regret and rejoicing do not depend on actions outside of the consideration set, which are taken with probability 0. This captures that as Sugden (1986) writes, only actions that are "serious candidates" or "real options" should be sources of regret and rejoicing. While regret and rejoicing are linear here and thus predict no violations of expected utility axioms from choice, linear regret still has effects on menu choice (Sarver, 2008).

The term in the second line of (2.21) can be interpreted as disappointment and elation. The emotion depends on the different between the payoff $v(a, \omega)$ and the expected payoff of the chosen action. The expected payoff of the chosen action is computed with respect to the posterior that the agent holds when taking the action. However, in contrast to standard models of disappointment and elation (Bell, 1985; Loomes and Sugden, 1986; Gul, 1991), the expected payoff is conditioned on what is learned about the state. The weight on disappointment and elation is the probability of taking the actual action conditional on the partition cell.

Theorem 10. *Under Assumptions 6 to 6, any solution of* $\mathrm{P}^{\mathrm{inc}}$ *is a linear regret-disappointment solution and induces the first-best attention strategy.*

Because we assumed that the bias is known, the optimal mechanism can perfectly offset it by scaling up the utility differences by $1/\beta$. Using a result in Whitmeyer and Zhang (2022), which builds on the Lagrangian Lemma of Caplin et al. (2022), we show that such inflating of the utility differences is also *necessary*, in order to implement the first-best attention strategy. The remaining goal of the principal is to reduce the variance of emotions, while scaling up the utility differences. This goal is constrained by the feedback structure. Under complete feedback, the second line of (2.21) is 0, and only regret is used. However, for incomplete feedback structures, both lines of (2.21) are generally non-zero and accordingly the optimal emotions can be interpreted as a combination of regret and disappointment.

Disappointment and elation are used because, under incomplete feedback, the choice of all emotional utilities consistent with some partition element of states is interrelated (Appendix, Lemma 30). Thus, all payoffs under that partition element, including the payoffs of the chosen action *had another state realized*, affect the optimal emotions under some feedback signal.[33]

Thus, in our model, disappointment and elation can be thought of as *second-best* motivators for attention when complete feedback is not available. They are second-best because they lead to a higher variance of emotions than using regret and rejoicing.[34]

One might be surprised that disappointment and elation motivate more attention because they depend on the realization of the state, which the agent cannot affect. By contrast, regret and rejoicing seem to depend more directly on the agent's action, which perhaps more intuitively incentivizes attention. However, although the agent cannot choose the unconditional distribution of states, he can implicitly choose the distribution of states *conditional* on an action by paying attention. Indeed, this is the perspective of attention taken by the posterior approach, in which the agent acquires posteriors over states conditional on each taken action (Caplin and Dean, 2013). This goes to say that, like regret and rejoicing, the

[33]This statement holds also under unknown bias.

[34]Under an unknown bias, as we have shown under complete feedback, regret and rejoicing are not only optimal because they minimize the variance of emotions but also because they allow tailoring the incentives to the state.

emotions of disappointment and elation penalize bad outcomes and reward good outcomes and thus raise the return to attention. As argued above, regret and rejoicing raise the return to attention in the optimal way but under constraints through the feedback structure, disappointment and elation have a role to play.

Proof of Theorem 10 Because the feedback structure reveals the material payoff by Assumption 6, the principal can choose the utility $u(z)$ for each message $z = (v, a, P) \in Z$. For $z = (v, a, P)$, we write $v(z)$ and $a(z)$ for the first and third component, respectively. The following Lemma uses of the fact that the optimal attention strategy under posterior-separable attention costs depends on the state-wise utility differences as shown in Whitmeyer and Zhang (2022).

Lemma 30. *A first-best attention strategy is induced only if for all elements z_1, z_2 of Z that are directly connected by an edge and have $a(z_1), a(z_2) \in S$,*

$$u(z_1) - u(z_2) = \frac{1}{\beta}(v(z_1) - v(z_2)).$$

Proof. By Theorem 4.3 in Whitmeyer and Zhang (2022), to induce an attention strategy that is optimal for the principal, the relative incentives, such as $u(a, \omega) - u(a_1, \omega)$ for $a \in S, \omega \in \Omega$, are uniquely pinned down. Since the agent discounts the relative incentives by β, for the relative incentives to equal those of the principal, they need to be scaled up by $\frac{1}{\beta}$. This is achieved if and only if for all $\omega \in \Omega$, $u(a, \omega) - u(a_1, \omega) = \frac{1}{\beta}(v(a, \omega) - v(a_1, \omega))$. This in turn is equivalent to all elements m_1, m_2 of Z that are directly connected by an edge (that is, are achieved at the same state) to satisfy $u(m_1) - u(m_2) = \frac{1}{\beta}(v(m_1) - v(m_2))$. □

The above is an only-if statement. It ensures only that if we restricted the agent's choice to the consideration set S, a first-best attention strategy would be incentive-compatible (Whitmeyer and Zhang, 2022). To implement the first-best attention strategy when the whole decision problem A is feasible, it is sufficient to punish the actions in $A \setminus S$ with sufficiently negative emotions. As long as the principal deters the agent from choosing any

action in $A \setminus S$, she is indifferent with regards to how strongly she punishes these actions as it does not affect her utility nor the implementation cost of realized emotions.

The utility difference being scaled up by $\frac{1}{\beta}$ is a transitive relation that extends to all nodes on the largest connected subgraph. If z and $\tilde{z}$ are on the same connected subgraph of Z, then there is a finite chain of elements $(z_1, ..., z_k)$ of Z such that $z = z_1$, $\tilde{z} = z_k$, and each adjacent pair is directly connected. If each of those directly connected pairs has a utility difference scaled up by $\frac{1}{\beta}$, then

$$\begin{aligned} u(\tilde{z}) - u(z) &= (u(z_k) - u(z_{k-1})) + ... + (u(z_2) - u(z_1)) \\ &= \frac{1}{\beta}(v(z_k) - v(z_{k-1}) + ... + \frac{1}{\beta}(v(z_2) - v(z_1)) = \frac{1}{\beta}(v(\tilde{z}) - v(z)). \end{aligned}$$

Conversely, if the utility difference of all pairs of nodes on the same connected subgraph is scaled up by $1/\beta$, then trivially the utility difference of directly connected nodes is scaled up by $1/\beta$.

Thus, an optimal attention strategy of the principal is implemented only if the utility difference of all elements on the same connected subgraph are scaled up by $1/\beta$. By Assumption 6, the sets of all messages that are consistent with some partition element $P \in \mathcal{P}$ are connected. By the definition of edges, such sets for different partition elements are not directly connected. In other words, Z has one connected subgraph for each partition element $P \in \mathcal{P}$. This pins down the utility function u except for one constant per subgraph, which is determined by the implementation-cost minimization problem. More specifically, by the above, we can write

$$m(a, \omega) = \left(\frac{1}{\beta} - 1\right)\big(v(a, \omega) - v(\mathcal{P}(\omega))\big). \tag{2.22}$$

The remaining problem is to determine the "reference point" $v(\mathcal{P}(\omega))$ for each partition cell.

The optimal reference point is connected to the concept of central tendency from statistics. The central tendency $\mathcal{C}_f(X)$ of random variable X with respect to a strictly convex

function f is defined as

$$\mathcal{C}_f(X) := \arg\min_c \mathbb{E}[f(X-c)],$$

which is known to be unique under strict convexity of f. Here, we consider the central tendency $\mathcal{C}_\psi$ with respect to the strictly convex implementation cost ψ.

Lemma 31. *Under implementation cost $\psi(x) = |x|^k$ with $k > 1$, the implementation cost minimizing solution has reference point $v(P)$ equal to the central tendency of payoff v on the partition element P with respect to the optimal attention strategy.*

Proof. First, the centrality measure of emotions must be 0, otherwise we could decrease the implementation cost. Let $P \in \mathbb{P}$. The attention strategy that is implemented defines a distribution over emotions. The implementation cost on this partition is the expectation of $\psi(u(a,\omega) - v(a,\omega))$. Under the optimal emotions, $\psi(u(a,\omega) - v(a,\omega) - c)$ is minimized at $c = 0$, otherwise one could further reduce the implementation cost. This implies by definition that the central tendency of emotions is 0.

Next, we show two properties of the central tendency. First, $\mathcal{C}_\psi(X+b) = \mathcal{C}_\psi(X) + b$ due to

$$\arg\min_c \mathbb{E}[\psi(X+b-c)] = b + \arg\min_c \mathbb{E}[\psi(X-c)].$$

Second, under $\psi(x) = |x|^k$ with $k > 1$, $\mathcal{C}_\psi(aX) = a\mathcal{C}_\psi(X)$ for all $a \in \mathbb{R}$. First of all, note that $\psi(ax) = |ax|^k = (|a||x|)^k = |a|^k|x|^k = |a|^k\psi(x)$. Then,

$$\arg\min_c \mathbb{E}[\psi(aX-c)] = a\arg\min_c \mathbb{E}[\psi(a(X-c))] = a\arg\min_c \mathbb{E}[|a|^k\psi(X-c)] =$$
$$a\arg\min_c \left(|a|^k\mathbb{E}[\psi(X-c)]\right) = a\arg\min_c \mathbb{E}[\psi(X-c)].$$

Finally, using the two previous facts we show that the reference point of P is the central tendency of the payoff v, with respect to the distribution induced by the implemented attention strategy on P. The reference point $v(\mathbb{P}(\omega))$ is the (hypothetical) payoff $\tilde{v}$ that obtains

emotions 0 according to equation (2.22). By equation (2.22), emotions are a linear function of the payoff $v(a, \omega)$. Then, the realized payoff is also a linear function f of emotions. The payoff $\tilde{v}$ is thus the image of 0 under f, $\tilde{v} = f(0)$. By the above, $\mathcal{C}_\psi(aX + b) = a\mathcal{C}_\psi(X) + b$, so $\mathcal{C}_\psi$ commutes with linear functions. By 0 being the central tendency of emotions, $\tilde{v}$ is the central tendency of the payoff. □

As is known, the central tendency $C_\psi(X)$ of random variable X with respect to $\psi(x) = x^2$ equals the expectation of X. This concludes the proof.

Chapter 3

The Optimal Design of Public Recognition Schemes

3.1 Introduction

Economics typically focuses on *monetary* incentives. However, a growing body of evidence shows that *social image* is a powerful motivator beyond material rewards (Bursztyn and Jensen, 2017). Because individuals care about how they are perceived by others, public recognition schemes—which signal agents' characteristics through conveying information about their behavior—can be used to incentivize desired actions. For instance, organizations implement "employee of the month" programs that signal high performance and incentivize effort. Schools employ grading systems that signal ability to the job market, motivating human capital accumulation. Charities publicly acknowledge donors, signaling prosociality and incentivizing donations. This paper studies how public recognition schemes can be optimally designed to leverage these image concerns and promote desired behaviors.

While the optimal design of public recognition schemes has been studied across various contexts,[1] the existing literature leaves two fundamental questions open.

[1]Previous literature has studied the optimal design of: privacy under prosocial signaling (Bénabou and Tirole, 2006; Bénabou and Tirole, 2011; Ali and Bénabou, 2020), performance evaluation under worker career

First, what type of signal structures emerge as optimal when the principal is not constrained to specific forms? Previous work assumes particular forms of signal structures. These constraints may overlook other potentially optimal designs, so understanding the unconstrained optimal signal structures is crucial for both theory and practice.

Second, how do agents' preferences over their image affect the optimal design of recognition schemes? Specifically, prior studies assume that agents are risk neutral with respect to their image—they care linearly about the public's expectation about their type. However, if agents are risk averse regarding their image, then revealing information imposes a social welfare cost. Therefore, it is important to understand optimal public recognition schemes when agents exhibit risk aversion over their image.

In this paper, we make progress on these questions by leveraging majorization and information design tools. This approach allows us to study the optimal design problem without imposing any restrictions on the types of information structures available to the principal. By doing so, we provide a more general and robust framework that can accommodate a variety of settings and agent preferences.

To address these questions, we consider the following model. An agent has a private type θ—such as ability or degree of prosociality—drawn from a known distribution over the interval $\Theta = [\underline{\theta}, \overline{\theta}]$. Upon observing her realized type, the agent chooses an action $a \in \mathbb{R}_+$, incurring a cost that increases with higher actions and decreases with her type. The public, however, only receives a signal that is correlated with the agent's action. The agent values the public's belief about her type—her *image*. The principal designs the signal structure, referred to as a public recognition scheme, with the objective of incentivizing the agent to take higher actions. We investigate the optimal signal design when the signal is the sole instrument available to the principal for influencing the agent's behavior.

concerns (Holmström, 1999; Dewatripont et al., 1999), quality certification under endogenous quality provision (Albano and Lizzeri, 2001; Zapechelnyuk, 2020), grading schemes under job-market signaling (Dubey and Geanakoplos, 2014; Zubrickas, 2015; Onuchic and Ray, 2023), monopolistic pricing under conspicuous consumption (Rayo, 2013), taxation under welfare stigma (Blumkin et al., 2015), and status categories for contests (Moldovanu et al., 2007).

Our first result concerns the case when the agent is risk neutral with regard to image and cares about the public's expectation about her type. The resulting design problem is tractable despite its high dimensionality, and a simple class of signal structures—*monotone partitional signals*—emerges as optimal. Monotone partitional signals resemble grading or rating schemes and are often used in practice.

To understand why these signals emerge as optimal, we need to understand the set of feasible image payoffs within our model. The set of feasible image payoffs was characterized by Saeedi and Shourideh (2023) in the context of rating design under moral hazard. By employing the majorization tools from Kleiner et al. (2021), we provide a short alternative proof of this characterization, demonstrating that any second-order expectation can be generated by publicly randomizing over monotone partitional signals. Leveraging this connection to majorization, we generalize previous results from the literature on public recognition schemes. Specifically, we show that in quasi-linear settings, a monotone partitional signal structure is optimal—that is, the principal does not need to randomize—and we analyze when the optimal signal is fully revealing or involves pooling over certain intervals of types. Finally, our proof extends to higher-order expectations, showing they have the same characterization as second-order expectations.

In the second part of our paper, we relax the common assumption of risk neutrality, which dominates previous work. While other applications in behavioral economics have relaxed this assumption (e.g., Kőszegi, 2003), prior studies on public recognition schemes assume risk neutrality over image. They do so by presuming that agents value the public's expected belief about their type in a linear manner.

While risk neutrality is a natural starting point, there are two reasons to consider risk aversion over image. First, agents may have an intrinsic valuation of their image due to psychological factors, leading to risk-averse behavior. Butera et al. (2022) provide evidence for risk aversion over prosocial image, which supports this interpretation. Second, risk aversion over image can be induced from risk aversion over other goods when image has

instrumental value. For example, if an agent values the public's belief about her ability because her attainable wage equals her expected ability, then risk aversion over income translates to risk aversion over image.

Our second main result is that under risk aversion over image, full privacy may become optimal, leading the principal to forego the possibility of screening the agent. We illustrate this result in a two-type model and contrast it with monetary incentive schemes, where the principal always employs some degree of screening to provide incentives.

The intuition behind this result is as follows. First of all, risk aversion over image makes revealing information socially costly due to Jensen's inequality: disclosing information creates a mean-preserving spread in the distribution of beliefs, lowering aggregate welfare. The principal internalizes this cost either because they directly care about the agent's well-being or because of the agent's participation constraint. Consequently, risk aversion pushes towards greater privacy.

While this force resembles the familiar efficiency-equity trade-off in the optimal taxation literature following Mirrlees (1971), we demonstrate a qualitative difference between image and standard goods like money. In traditional screening problems involving monetary transfers, there is a trade-off between providing incentives (through unequal allocations) and mitigating welfare losses from risk aversion (favoring more equal allocations). Because, for small incentives, the loss from risk aversion is of second order, some incentive provision is always optimal; complete equality (egalitarianism) is suboptimal.

However, when it comes to image, the push towards equality can be so strong that full privacy—the only public recognition scheme that equalizes image across agents—becomes optimal. We show that the aggregate welfare loss and the gain from providing incentives are of the same order because to induce small incentives, the principal has to resort to random signal structures, and consequently random allocations of image. Consequently, it may be optimal not to reveal information at all because any level of information would impose a greater welfare loss than benefit through providing incentives.

This result has implications for the economics of privacy, reviewed in Acquisti et al. (2016). While some authors have argued for publicity as a means to incentivize socially desirable behavior and provide economically useful information (R. A. Posner, 1978), a lack of privacy can lead to data misuse by firms or governments (for the latter, see Tirole, 2021), wasteful signaling (Daughety and Reinganum, 2010), and suppression of learning about societal preferences (Ali and Bénabou, 2020). As pointed out by Butera et al. (2022), there is also a more immediate welfare cost of publicity stemming from risk aversion over social image. We show that this force is so strong that complete privacy may be optimal despite publicity incentivizing desirable behavior.

This paper is structured as follows. Section 3.2 introduces the model. Section 3.3 studies the model when agents are risk neutral over image. Section 3.4 studies the model when agents are risk averse over image.

3.2 Baseline Model

Let θ be an agent's private type drawn from $\Theta = [\underline{\theta}, \overline{\theta}]$ according to some CDF $F : \Theta \to \mathbb{R}$ with continuous and strictly positive derivative f. By θ we will denote a generic element of Θ and we use bold symbols, such as $\boldsymbol{\theta}$, for random variables. The agent incurs linear cost a/θ for action $a \in \mathbb{R}_{\geq 0}$, say effort; that is, higher types have smaller costs. The principal wants to maximize the aggregate action $\int a(\theta) dF(\theta)$. Without loss, we restrict attention to direct mechanisms.[2] The principal recommends an action schedule $a(\theta)$. To implement this action schedule, the principal does not use monetary transfers but instead makes use of image motives by committing to revealing a signal about the reported type.

As mentioned in the introduction, the model has a number of different interpretations, such as a firm designing a performance evaluation scheme to incentive effort when workers have career concerns, a school designing a grading schemes to incentive human capital ac-

[2]One can show that the usual taxation principle holds: we can let the agent choose the action and the principal chooses a signal structure about the chosen action. The principal can pool unchosen actions with the outside option, which gives the worst reputation, see Rayo (2013).

cumulation under job-market signaling, or, more abstractly, the optimal degree of privacy about agent's public good contributions when they engage in prosocial signaling.

Signal Structures and Image Payoff A version of this model was considered among others by Rayo (2013), who assumed that the principal is restricted to partitional signal structures. We do not impose any assumptions on the nature of signals and allow for arbitrary signal structures. A *signal structure* (S, σ) consists of the set of signals S, and a map $\sigma\colon \Theta \to \Delta(S)$ assigning a distribution of signals to each realized state.[3] Equivalently, a signal structure can be modeled as a Bayes-consistent distribution $\tau \in \Delta(\Delta(\Theta))$ over posteriors $\pi \in \Delta(\Theta)$ (Kamenica and Gentzkow, 2011). The distribution τ is Bayes consistent if $\mathbb{E}_\tau[\pi] = \mu$, where μ is the measure $\mu \in \Delta(\Theta)$ associated with CDF F. In general, agents can value the belief π held by the public as $v(\pi)$, $v\colon \Delta(\Theta) \to \mathbb{R}$. Because the agent knows their type θ, they need to condition on $\boldsymbol{\theta} = \theta$ to form their expected image payoff. We say the payoff $V : \Theta \to \mathbb{R}$ is *generated* by a signal structure $\tau \in \Delta(\Delta(\Theta))$ if $V(\theta)$ corresponds to the expected value $v(\pi)$ of the posterior π, conditioning on the true state being θ, that is,

$$V(\theta) = \mathbb{E}_\tau[v(\pi)|\boldsymbol{\theta} = \theta].$$

We say $V : \Theta \to \mathbb{R}$ is *feasible* if it is generated by some signal structure τ.

Agent's Utility The utility $\tilde{U}(\theta, \theta')$ of agent θ from reporting θ' is

$$\tilde{U}(\theta, \theta') = V(\theta') - \frac{1}{\theta} a(\theta').$$

Equivalently, we can define the agent θ as choosing θ' to maximize utility $U(\theta, \theta') = \theta\tilde{U}(\theta, \theta')$, that is,

$$U(\theta, \theta') = \theta V(\theta') - a(\theta').$$

[3]To ensure that conditional expectations are well-defined, we require that the set of signals S is a Polish space endowed with its Borel sigma-algebra and $\sigma\colon \Theta \to \Delta(S)$ is such that for all measurable $A \subseteq S$, $\sigma(A, \theta)$ is measurable in θ.

Instead of letting the principal choose a signal structure, we write their problem as choosing the second-order expectation $V(\theta)$ subject to feasibility, that is, $V(\theta)$ can be generated by some signal structure.

Principal's Problem The principal maximizes the expected action subject to incentive compatibility (IC), voluntary participation or individual rationality (IR), and feasibility (F) of the second-order expectation $V(\theta)$. To formulate the participation constraint, we assume that upon not participating, the agent obtains the minimal second-order expectation $\underline{\theta}$, which is the most favorable to the principal. Later, we will show that this assumption is without loss.

$$\max_{a,V:\Theta\to\mathbb{R}} \int_{\underline{\theta}}^{\overline{\theta}} a(\theta)dF(\theta) \tag{P}$$

s.t.

$$\theta V(\theta) - a(\theta) \geq \theta V(\theta') - a(\theta') \qquad \forall \theta, \theta' \tag{IC}$$

$$\theta V(\theta) - a(\theta) \geq \theta \cdot \underline{\theta} \qquad \forall \theta \tag{IR}$$

$$V \text{ is feasible.} \tag{F}$$

The principal's problem (P) looks similar to a standard screening problem where the principal allocates $V(\theta)$, usually money or the allocation probability of a good, to implement action schedule $a(\theta)$. However, the screening problem has two peculiarities. First, is has a type-dependent participation constraint (IR). Second, it has a non-standard feasibility constraint (F). We first deal with (IR) and simplify the principal's problem, before we turn to the feasibility constraint in section 3.3.

Rewriting the Principal's Problem Suppose (IR) holds for $\underline{\theta}$. By (IC) and the envelope theorem, the derivative of the left-hand side of (IR) in θ is $V(\theta)$. The derivative of the right-hand side in θ is $\underline{\theta}$. Any feasible $V(\theta)$ satisfies $V(\theta) \in [\underline{\theta}, \overline{\theta}]$, thus if (IR) holds for $\underline{\theta}$, it holds

for all θ. Thus, we can rewrite (IR) as

$$\underline{\theta} \cdot \underline{\theta} \qquad \forall \theta \tag{IR'}$$

By supermodularity of the agent's utility in θ and $V(\theta)$, (IC) implies that $V(\theta)$ is *monotone*, that is $\theta' > \theta$ implies $V(\theta') \geq V(\theta)$. Following standard arguments we can rewrite the principal's problem as a linear maximization problem P' in $V(\theta)$ subject to $V(\theta)$ being monotone and feasible.

$$\max_{V:\Theta\to\mathbb{R}} \int V(\theta) \left(\theta - \frac{1-F(\theta)}{f(\theta)}\right) dF(\theta) \tag{P'}$$

s.t.

$$V(\theta) \leq V(\theta') \qquad \forall \theta < \theta' \tag{MON}$$

$$V \text{ is feasible.} \tag{F}$$

The principal's problem reduces to a linear optimization problem over feasible, monotone V. The crucial question is what monotone $V(\theta)$ are feasible.

3.3 Risk Neutrality over Image

In this section, we consider the special case where the value of a posterior equals the expected type,

$$v(\pi) = \mathbb{E}_\pi[\boldsymbol{\theta}],$$

as assumed by virtually all of the previous literature. In this case, $V(\theta)$ equals the second-order expectation $P(\theta)$, defined as follows. A signal structure (S, σ) defines a joint distribution of $\boldsymbol{\theta}$ and $\boldsymbol{s}$: for each realized state $\boldsymbol{\theta} = \theta$, a random signal $\mathbf{s}$ is drawn from S according to the distribution $\sigma(\theta)$. We define the *second-order expectation* generated by signal structure

(S, σ) as

$$P(\theta) := \mathbb{E}[\mathbb{E}[\boldsymbol{\theta}|\boldsymbol{s}]|\boldsymbol{\theta} = \theta].$$

We say $P(\theta)\colon \Theta \to \mathbb{R}$ is *feasible* if it is generated by some signal structure (S, σ). For clarity of exposition, we use $P(\theta)$ instead of $V(\theta)$ in this section.

Saeedi and Shourideh (2023) study second-order expectations in the context of moral hazard and characterize the feasible, monotone $P(\theta)$. To state their result, we introduce the following notation.

For a subset $\Omega \subset \mathbb{R}$, let $L^1(\Omega)$ denote the real-valued, integrable, and right-continuous functions defined on Ω. For monotone functions $P \in L^1(\Theta)$, we say that P is *majorized by the identity function with respect to* F, denoted by $P \prec_F \text{id}$, if

$$\begin{aligned} \int_{\hat{\theta}}^{\overline{\theta}} P(\theta) dF(\theta) &\leq \int_{\hat{\theta}}^{\overline{\theta}} \theta dF(\theta) \quad \forall \hat{\theta} \in (\underline{\theta}, \overline{\theta}], \\ \int_{\underline{\theta}}^{\overline{\theta}} P(\theta) dF(\theta) &= \int_{\underline{\theta}}^{\overline{\theta}} \theta dF(\theta). \end{aligned} \tag{3.1}$$

Note that the usual majorization constraint is written with integrals with respect to the Lebesgue measure, while the integrals in this generalized majorization relation are with respect to the measure with CDF F. Under a continuous CDF F, one can rewrite (MAJ) as a standard majorization constraint by going into quantile space through a change of variables. This generalized majorization constraint is appropriate for the following result due to Saeedi and Shourideh (2023) because it also holds for discrete distributions (which we do not show in the present version of this draft).

Theorem 11 (Saeedi and Shourideh, 2023)**.** *Monotone $P \in L^1(\Theta)$ are feasible if and only if $P \prec_F \text{id}$, that is, P is majorized by the identity function with respect to F.*

Saeedi and Shourideh (2023) prove their results using an induction argument for finite distributions of θ and an approximation argument for continuous distributions. Here, we present a short proof of their result based on extreme point and majorization techniques,

building on a result from Kleiner et al. (2021). In addition to showing a connection between this fundamental result and results in majorization, our proof delivers additional insights relevant to our screening problem, see the following subsection 3.3.1.

Before we show our proof, we state two preliminary result about majorization and define monotone partitional signal structures. The first statement is a known result and will help us show that any feasible monotone $P(\theta)$ satisfies $P \prec_F \text{id}$. For completeness, we include a proof in the Appendix.

Lemma 32. *For monotone $P \in L^1(\Theta)$, $P \prec_F \text{id}$ if and only if the distribution of $P(\boldsymbol{\theta})$ is a mean-preserving contraction of the distribution of $\boldsymbol{\theta}$.*

For the proof of Lemma 32, note that mean-preserving contraction has a simple characterization in terms of integral inequalities similar to (3.1) (Hardy et al., 1934), although these integrals are in terms of the CDFs of the two random variables. However, by going into quantile space and using the monotonicity of $P(\theta)$, one can translate (3.1) into the integral characterization of mean-preserving contraction.

Lemma 33. *Let $g \in L^1(\Theta)$ be monotone. Then, P is an extreme point of*

$$\text{MPS}_F(g) := \{P \in L^1(\Theta) \mid P \text{ is monotone and } P \prec_F g\}$$

if and only if there exists a collection of disjoint intervals $[\underline{\theta}_i, \overline{\theta}_i)$ indexed by $i \in I$ such that for almost every $\theta \in [\underline{\theta}, \overline{\theta}]$,

$$P(\theta) = \begin{cases} g(\theta) & \text{if } \theta \notin \bigcup_{i \in I} [\underline{\theta}_i, \overline{\theta}_i), \\ \frac{\int_{\underline{\theta}_i}^{\overline{\theta}_i} g(\theta) dF(\theta)}{F(\overline{\theta}_i) - F(\underline{\theta}_i)} & \text{if } \theta \in [\underline{\theta}_i, \overline{\theta}_i). \end{cases}$$

Any element of $\text{MPS}_F(g)$ is an integral over extreme points.

This lemma is a simple generalization of the characterization of extreme points in Kleiner et al. (2021) using a coordinate change. The extreme $P(\theta)$ can be generated by monotone partitional signals, which we define as follows.

Definition 12. *A signal structure* (S, σ) *is* monotone partitional *if* $S = \mathbb{R}$ *and* $\sigma(\theta) = \delta_{m(\theta)}$ *is* degenerate *with* $m\colon [\underline{\theta}, \overline{\theta}] \to \mathbb{R}$ monotone.

An extreme point as in Lemma 33 can be implemented by a monotone partitional signal where $m(\theta) = P(\theta)$, that is the signal is revealing when $P(\theta)$ is strictly increasing and pools types that have the same second-order expectation $P(\theta)$.

Novel Proof of Theorem 1 Due to Lemma 1, necessity is easy to show. For sufficiency we use that it is enough to show that extreme points of the monotone functions $P(\theta)$ satisfying (3.1) are feasible. Then, by Choquet's theorem, all non-extreme points are also feasible by mixing over signal structures (in posterior space).

Necessity: The distribution of $P(\boldsymbol{\theta}) = \mathbb{E}[\mathbb{E}[\boldsymbol{\theta}|\boldsymbol{s}]|\boldsymbol{\theta}]$ is a mean-preserving contraction of the distribution of $\mathbb{E}[\boldsymbol{\theta}|\boldsymbol{s}]$, which is a mean-preserving contraction of the distribution of $\boldsymbol{\theta}$. By Lemma 32, this implies $P \prec_F \mathrm{id}$.

Sufficiency: By Lemma 33, any $P \in \mathrm{MPS}_F(\mathrm{id})$ can be represented as an integral over extreme points. The second-order expectation P is linear under public randomization over signal structures. If the signal spaces of two signal structures are distinct, then randomizing over the signal structures delivers the corresponding mixture over the induced second-order expectations. Thus, if we can generate the *extreme* $P \in \mathrm{MPS}_F(\mathrm{id})$, we can, after relabeling signal spaces if needed, generate *any* $P \in \mathrm{MPS}_F(\mathrm{id})$ by randomization.[4] The extreme $P\mathrm{MPS}_F(\mathrm{id})$ can be generated by monotone partitional signals, namely by pooling intervals corresponding to the quantile interval where H is constant and fully revealing intervals where H equals the majorizing function. □

[4] A similar extreme point argument is used in Yang and Zentefis (2024) to characterize the set of feasible distributions of posterior quantiles.

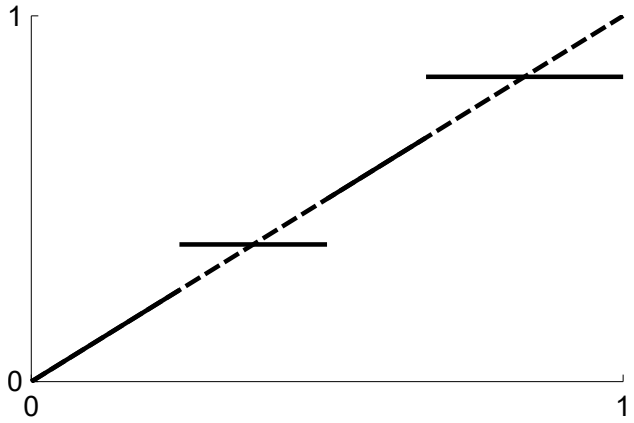


Figure 3.1: Example of an extreme $P(\theta)$ under $\theta \sim U([0,1])$.

Figure 3.1 shows an example of an extreme $P(\theta)$. This P can be generated by a monotone partitional signal that reveals each state from 0 to θ_1, pools all states in $[\theta_1, \theta_2)$ together, reveals all states in $[\theta_2, \theta_3)$, and pools all states in $[\theta_3, 1]$.

3.3.1 Implications

The first corollary follows directly from our proof of Proposition 11, but is worth stating independently.

Corollary 10. *Any feasible monotone $P(\theta) \in L^1(\Theta)$ can be generated by publicly randomizing over monotone partitional signal structures.*

The corollary states that we can generate monotone feasible second-order expectation by using very simple and widely used signal structures, namely monotone partitional signal structures. These signal structures resemble grading or rating schemes with discrete categories, which pool together intervals of the underlying variable of interest. A simple example of a monotone partitional signal structure is a grading scheme where all grades greater or equal to 90 are pooled as A's, all grades greater or equal to 80 but below 90 are pooled as B's, etc.

Second, we can use the connection to majorization and extreme points to characterize the optimal signal structures. The insight of our proof is that the extreme monotone second-

order expectations can be generated by monotone partitional signals. This is useful because our principal's problem is a linear optimization problem, in which by Bauer's maximum principal, an extreme point is optimal.

To state our result, let $h(\theta) := \theta - \frac{1-F(\theta)}{f(\theta)}$ denote the virtual valuation of type $\theta \in \Theta$. Further, define $H(\theta) := \int_{\underline{\theta}}^{\theta} h(\theta')d\theta'$ and let $\text{conv}(H)$ be the largest convex function that lies below H.

Corollary 11. *There exists a solution to the principal's problem P, that employs a monotone partitional signal structure.*

More specifically, if the virtual valuation is non-decreasing, then full revelation is optimal. Otherwise, the optimal monotone partitional signal pools exactly those intervals of types on which $\text{conv}(H)$ *is affine.*

Proof. This result follows directly from section 3.2.1 in Kleiner et al. (2021), when transforming the rewritten principal's problem P' into quantile space, combined with the insight that extreme monotone second-order expectations can be generated by monotone partitional signals, which follows from our proof of Theorem 11. □

This result generalizes and reestablishes the results from Rayo (2013) and Onuchic and Ray (2023) through majorization techniques. Both papers study versions of our principal's problem but restrict for tractability to monotone partitional signals.[5] Corollary 11 shows that this restriction is without loss. Further, it reestablishes their characterizations of the optimal signal structures under non-monotonic virtual valuation through majorization techniques.

It is important to note the limits of Corollary 2. Our principal's rewritten problem P' is linear in $P(\theta)$ because we have assumed a quasi-linear environment: the agent's cost and the principal's value are linear in the action. This is a useful benchmark and reasonable in some circumstances (see Rayo, 2013). The more complex problem with convex costs for the agent, in which case a non-extreme point can be optimal, is studied in Smith (2024).

[5]To be precise, Rayo, 2013 restricts to partitional signal structures but due to the incentive constraint it follows immediately, that we can restrict attention to monotone partitional signals.

Corollary 1 implies nevertheless that we can generate the solution through randomization over monotone partitional signals.

Finally, and beyond the screening problem studied in this paper, our proof of Theorem 11 can be used to characterize higher-order expectations. We can define higher-order expectations $P_n(\theta)$ generated by a signal structure (S, σ), recursively via

$$P_2(\theta) := P(\theta)$$
$$P_n(\theta) := \mathbb{E}[\mathbb{E}[P_{n-1}(\boldsymbol{\theta})|\boldsymbol{s}]|\boldsymbol{\theta} = \theta]$$

Intuitively, the third-order expectation can be understood as follows. Conditional on signal $\boldsymbol{s} = s$, an observer can form an expectation about the second-order expectation of the agent whose type is observed through s. The third-order expectation tells us the expectation that the agent of type θ has about the expectations other types have about her second-order expectation.

Corollary 12. *The set of feasible monotone $P_n(\theta)$ coincides with the set of feasible monotone $P(\theta)$, namely $MPS_F(id)$.*

Proof. The proof proceeds similar to the proof of Proposition 11. We need to show that any feasible monotone $P_n(\theta)$ is majorized by the identity function with respect to F.

Necessity: By construction, the distribution of $P_n(\boldsymbol{\theta})$ is a mean-preserving contraction of the distribution of $P_{n-1}(\boldsymbol{\theta})$. By induction, for any n, the distribution of $P_n(\boldsymbol{\theta})$ is a mean-preserving contraction of $\boldsymbol{\theta}$. By Lemma 32, this implies $P \prec_F \text{id}$.

Sufficiency: Because a monotone partitional signal is deterministic, all higher-order expectations are identical. More precisely, conditional on a signal realization, the second-order expectation is deterministic and thus equal to the third-order expectation,

$$\mathbb{E}[\underbrace{\mathbb{E}[P(\boldsymbol{\theta})|\boldsymbol{s}]}_{P(\boldsymbol{\theta})}|\boldsymbol{\theta} = \theta] = \mathbb{E}[P(\boldsymbol{\theta})|\boldsymbol{\theta} = \theta] = P(\theta).$$

Inductively, all higher-order expectations are the same. By our proof of Proposition 1, the extreme points of $\text{MPS}_F(\text{id})$ can be generated by publicly randomizing over monotone partitional signals. Higher-order expectations, like second-order expectations, are linear under public randomization, thus we can generate all $P_n(\theta) \in \text{MPS}_F(\text{id})$. □

3.4 Risk Aversion over Image

This section shows that under risk aversion over image, full privacy can be uniquely optimal. That is, the incentive instrument is *not used* and the agent is not screened. We contrast this with incentive provision through allocating a non-informational good, such as money. In that case, the instrument is *always used* and the principal does screen the agent, in the sense that more of the good is allocated to higher types. To make this comparison in the simplest way, we consider two versions a simple screening problem with *binary types*: first, under incentive provision through image and second, under incentive provision through allocating a non-informational good like money.

To model risk aversion, we assume that the agent values the belief $\pi \in \Delta(\Theta)$ that the population holds by $v(\pi)$ with v concave. For example, the agent's utility may be a concave function of the posterior mean, $v(\pi) = \hat{v}(\mathbb{E}_\pi[\theta])$ with $\hat{v}$ concave. Such risk aversion over image can be stem from agents intrinsic valuation of image due to psychological reasons. Butera et al. (2022) provide evidence for risk aversion over prosocial image, which can be interpreted in such a way. Alternatively, risk aversion over image can be induced from risk aversion over others goods if image is valued for instrumental reasons. For example, consider an agent who values image over her type, say ability, because her attainable wage in the labor market equals her expected ability. If she has risk-averse preferences over income, this induces risk aversion over image, because her utility will be a concave function of her expected ability.

Principal's Problem To illustrate the mechanism in the simplest example, we restrict attention to binary types, $\Theta = \{\underline{\theta}, \overline{\theta}\}$. We identify the set of beliefs $\Delta(\Theta)$ with the unit interval $[0,1]$ and let the prior $\mu \in \Delta(\Theta)$ over types be non-degenerate, that is, $\mu \in (0,1)$. Otherwise there is effectively only one type and the problem is trivial. We formulate a general version of the principal's problem (P) which allows for the incentive payoffs V to be provided through image or through a non-informational good by appropriately specifying the set $\mathcal{F}$ of feasible payoffs.

$$\max_{a, V:\, \Theta \to \mathbb{R}} (1-\mu) a(\underline{\theta}) + \mu a(\overline{\theta}) \tag{P}$$

s.t.

$$\overline{\theta} V(\overline{\theta}) - a(\overline{\theta}) \geq \overline{\theta} V(\underline{\theta}) - a(\underline{\theta}) \qquad \forall \theta \in \{\underline{\theta}, \overline{\theta}\} \tag{IC}$$

$$\underline{\theta} V(\underline{\theta}) - a(\underline{\theta}) \geq \theta \cdot \underline{\theta} \qquad \forall \theta \in \{\underline{\theta}, \overline{\theta}\} \tag{IR}$$

$$V \in \mathcal{F} \tag{F}$$

We consider two versions of the feasibility constraint (F), which correspond to incentive provision through image and to incentive provision through a non-informational good, respectively. Before we spell out these two versions, recall that, by standard arguments, the principal's problem can be written as a linear optimization problem over the incentive payoffs V subject to monotonicity and feasibility, that is,

$$\max_{V:\Theta \to \mathbb{R}} (1-\mu) h(\underline{\theta}) V(\underline{\theta}) + \mu h(\overline{\theta}) V(\overline{\theta}) \tag{P'}$$

s.t.

$$V(\underline{\theta}) \leq V(\overline{\theta}) \tag{MON}$$

$$V \in \mathcal{F} \tag{F}$$

In the Appendix 3.5, we show that, under two types and any prior μ, the virtual valuations $(h(\underline{\theta}), h(\overline{\theta}))$ are strictly increasing, $h(\underline{\theta}) = \underline{\theta} - \frac{\mu}{1-\mu}(\overline{\theta} - \underline{\theta}) < \overline{\theta} = h(\overline{\theta})$.

Incentive Provision through Image First, we consider the feasible incentive payoffs that can be induced by some revealing some information about the agent's type when there are image concerns. As introduced above, we let the value of belief $\pi \in \Delta(\Theta) = [0,1]$ to each agent be $v(\pi)$ where $v\colon [0,1] \to \mathbb{R}$ is strictly increasing and concave. Further, we assume that v is twice differentiable.

Assumption 9. *The value of image $v \in C^2([0,1],\mathbb{R})$ is strictly increasing and concave.*

We say the payoff $V : \Theta \to \mathbb{R}$ is *generated* by a distribution $\tau \in \Delta([0,1])$ over posteriors if $V(\theta)$ corresponds to the expected value $v(\pi)$ of the posterior π, conditioning on the true state being θ, that is,

$$V(\theta) = \mathbb{E}_\tau[v(\pi)|\boldsymbol{\theta} = \theta] \qquad\qquad \forall \theta \in \Theta.$$

We let the feasible set of image payoffs $\mathcal{F}_I(v)$ be the set of functions $V\colon \Theta \to \mathbb{R}$ that are generated by some Bayes-consistent distribution $\tau \in \Delta([0,1])$ over posteriors π, that is, τ which satisfy $\mathbb{E}_\tau[\pi] = \mu$. That means, we allow the principal to choose *any* signal structure without parametric restrictions.

The following theorem shows that when the level of absolute risk aversion over image is pointwise high enough, then full privacy is uniquely optimal.

Theorem 12. *Consider the principal's problem (P) with incentive provision through image, that is, $\mathcal{F} = \mathcal{F}_I(v)$. If $\forall \pi \in [0,1]$:*

$$\frac{-v''(\pi)}{v'(\pi)} > \frac{2(h(\overline{\theta}) - h(\underline{\theta}))}{(1-\pi)h(\underline{\theta}) + \pi h(\overline{\theta})}, \tag{3.2}$$

then no information revelation is uniquely optimal.

We prove this result in the Appendix, where we show that one can rewrite the principal's problem as a standard information problem. We show that the under the provided condition,

the value function is strictly concave, so no information revelation is uniquely optimal by Jensen's inequality.

Theorem 12 shows that the incentive instrument, revealing information, is sometimes not used and the agent is not screened. Next, we contrast this result with the standard setup, when incentives are provided through allocating some non-informational good. In that case, we show that the incentive instrument is always used when the agent's utility is differentiable.

Incentive Provision through Allocating a Non-Informational Good Second, we consider the feasible payoffs that can be induced by distributing among the agents some budget of a divisible good like money. We say the payoff $V : \Theta \to \mathbb{R}$ is *induced* by some allocation $x : \Theta \to \mathbb{R}_{\geq 0}$ under utility function $u : \mathbb{R}_{\geq 0} \to \mathbb{R}$ and budget $b \in \mathbb{R}_{>0}$ if

$$\begin{aligned} V(\theta) &= u(x(\theta)) && \forall \theta \in \Theta \\ b &\geq (1-\mu)x(\underline{\theta}) + \mu x(\overline{\theta}). \end{aligned}$$

The first line states that $V(\theta)$ is the utility of obtaining $x(\theta)$ units of the good. The second line states that the allocation x does not exceed the budget b. We let the feasible set of allocational payoffs $\mathcal{F}_{\neg I}(u, b)$ be the set of payoff functions $V : \Theta \to \mathbb{R}$ that are induced by some allocation x under utility u and budget b. That is, we allow the principal to choose any budget-balanced allocation. At the expense of some additional terminology, we could also allow for random allocations or allocating multiple goods; this would not change our results.

The following proposition shows that no matter the level of risk aversion over the non-informational good, as long as the utility u is differentiable at b, the agents will be screened.

Proposition 6. *Consider the principal's problem (P) with incentive provision through non-informational goods, that is, $\mathcal{F} = \mathcal{F}_{\neg I}(u, b)$. If u is differentiable at b with $u'(b) > 0$, then the optimal mechanism allocates more of the good to the high type, $x(\overline{\theta}) > x(\underline{\theta})$.*

The intuition for this result is as follows. If the utility u is differentiable at b, then u is locally approximately linear at b.[6] Thus, the utility loss from risk aversion is second-order for small deviations from the "equal allocation" $x(\underline{\theta}), x(\overline{\theta})) = (b, b)$, which allocates to both types the same. On the other hand, deviating from the equal allocation by allocating slightly more to the higher type and slightly less to the lower type, provides an *incentive* to the high type to choose a higher action. We show this gain from the generated incentive is first-order and therefore dominates the second-order loss from risk aversion.

Why does a similar logic not show that the principal should screen when providing incentives through image? The reason is that the principal has to resort to *random* allocations of image, that is, random signal structures, to implement small incentives. To explain the difference between image and non-informational goods, we first make precise the trade-off the principal has between the aggregate loss due to risk aversion and creating an incentive for the agent.

Image versus Non-Informational Goods To understand the difference between image and non-informational goods, it is useful to think of the principal as maximizing a linear combination of the *incentive* $I := V(\overline{\theta}) - V(\underline{\theta})$ and the *aggregate payoff* $A := (1-\mu)V(\underline{\theta}) + \mu V(\overline{\theta})$. In our screening problem, as usual, the principal's objective can be written as a linear combination of the allocation of incentive payoffs $V(\theta)$, see (P'). As we show in Appendix 3.5, we can, through a linear coordinate change, express the principal's objective U as a linear combination of the incentive I and the aggregate image utility A, where the weights on both I and A are necessarily positive,

$$U = (1-\mu)h(\underline{\theta})V(\underline{\theta}) + \mu h(\overline{\theta})V(\overline{\theta}) = \underbrace{\mu(1-\mu)(h(\overline{\theta}) - h(\underline{\theta}))}_{>0} I + \underbrace{((1-\mu)h(\underline{\theta}) + \mu h(\overline{\theta})))}_{>0} A. \tag{3.3}$$

We prove Proposition 6 for non-informational goods by showing that deviating from the equal allocation and allocating ε more units to the high type provides is a *first-order gain* to

[6] Footnote on Loss aversion

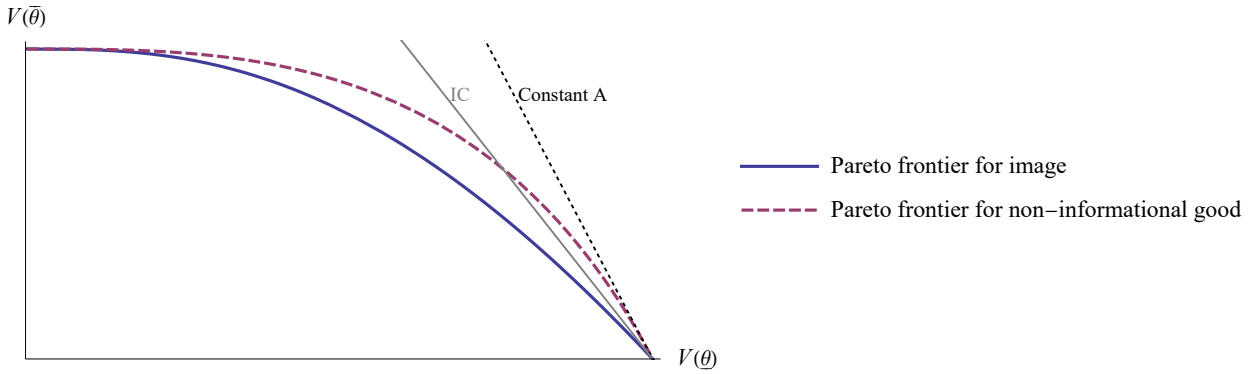


Figure 3.2: Pareto frontiers under incentive provision through image and a non-informational good.

the incentive, I, but only a *second-order loss* of aggregate payoff A from risk aversion. By the principal having a finite marginal rate of substitution between incentive I and aggregate payoff A, for small ε the principal gains overall.

By contrast, when it comes to image, deviating from no information revelation, *the incentive gain and the aggregate loss are of the same order*. Thus, if the principal cares enough about the incentive vis-à-vis the aggregate payoff, they choose to reveal no information.

This is illustrated in Figure 3.2, which plots the Pareto frontier of feasible $(V(\underline{\theta}), V(\bar{\theta}))$ under image and under a non-informational good. For comparison, we chose $v = u$ and $b = \mu$. The bottom right corner of the Pareto frontiers correspond to full privacy and the equal allocation, respectively. Moving from that point along the dotted line preserves the aggregate payoff A but increases the incentive I and thereby improves the principal's objective. Thus, the linear indifference curves of the principal is less steep than the dotted line. This rules out a corner solution for the Pareto frontier for non-informational goods, but it allows for a corner solution for the Pareto frontier for image, for example, under the gray indifference curve, labeled IC, of the principal.

Why are the incentive gain and the aggregate loss of the same order when it comes to image, by contrast to non-informational goods? The reason is that the principal has to *randomize* to achieve a small image incentive, which is not needed when creating an incentive with a divisible non-informational good. The following example illustrates. Suppose, the

principal can divide one unit of a divisible non-informational good between two agents. To create a small incentive, they can allocate $1/2+\varepsilon$ units to the high type and $1/2-\varepsilon$ units to the low type, which induces the indicated incentive payoffs $(V(\underline{\theta}), V(\overline{\theta}))$.

$$\bar{\theta} \bullet\!\!-\!\!-\!\!-

rem 12 implies that, for image, the incentive gain and aggregate loss are of the same order, regardless of which signal structure is used.

3.5 Appendix

Proof of Lemma 32

Proof. It is well known that the distribution of $P(\boldsymbol{\theta})$ is a mean-preserving contraction of the distribution of $\boldsymbol{\theta}$ if and only if the CDF G of $P(\boldsymbol{\theta})$ majorizes the CDF F of $\boldsymbol{\theta}$, that is,

$$\forall \hat{\theta} \in [\underline{\theta}, \overline{\theta}]\colon \int_{\hat{\theta}}^{\overline{\theta}} F(\theta)d\theta \leq \int_{\hat{\theta}}^{\overline{\theta}} G(\theta)d\theta,$$

with equality at $\hat{\theta} = \underline{\theta}$. The CDF G majorizes F if and only if F^{-1} majorizes G^{-1} (e.g., Shaked and Shanthikumar, 2007, Theorem 3.A.5.), that is,[7]

$$\forall \hat{q} \in [0, 1]\colon \int_{\hat{q}}^{1} G^{-1}(q)dq \leq \int_{\hat{q}}^{1} F^{-1}(q)dq,$$

with equality at $\hat{q} = 0$. Under comonotonicity (so for all θ, $P(\theta)$ and θ correspond to the same quantile) this is equivalent to

$$\forall \hat{q} \in [0, 1]\colon \int_{\hat{q}}^{1} P(F^{-1}(q))dq \leq \int_{\hat{q}}^{1} F^{-1}(q)dq, \tag{3.4}$$

with equality at $q = 0$. Recall that by assumption, F is continuously differentiable and strictly increasing. Thus, we can apply a coordinate change to $q = F(\theta)$ and rewrite the (3.4) as

$$\forall \hat{\theta} \in [\underline{\theta}, \overline{\theta}]\colon \int_{\hat{\theta}}^{\overline{\theta}} P(\theta)dF(\theta) \leq \int_{\hat{\theta}}^{\overline{\theta}} \theta dF(\theta).$$

[7]For non-continuous monotonic F, F^{-1} denotes the right-continuous inverse $F^{-1}(q) := \sup\{x|F(x) \leq q\}$ for $q \in [0, 1]$.

□

Proof of Lemma 33

Proof. The proof carries over the Kleiner et al. (2021) characterization of extreme points by going into quantile space.

Let

$$\text{MPS}_F(g) := \{P \in L^1(\mathbb{R}) | P \text{ is monotone and } P \prec_F g\}.$$

That is, $\text{MPS}_F(g)$ contains the functions $P \in L^1(\mathbb{R})$ such that

$$\begin{aligned}\int_{\hat{\theta}}^{\overline{\theta}} P(\theta)dF(\theta) &\leq \int_{\hat{\theta}}^{\overline{\theta}} g(\theta)dF(\theta) \quad \forall \hat{\theta} \in (\underline{\theta}, \overline{\theta}], \\ \int_{\underline{\theta}}^{\overline{\theta}} P(\theta)dF(\theta) &= \int_{\underline{\theta}}^{\overline{\theta}} g(\theta)dF(\theta).\end{aligned}$$

Going via the coordinate change $q = F(\theta)$ into quantile space, this is equivalent to[8]

$$\begin{aligned}\int_{\hat{q}}^{1} H(q)dq &\leq \int_{\hat{q}}^{1} g(F^{-1}(q))dq \quad \forall \hat{q} \in (0, 1] \\ \int_{0}^{1} H(q)dq &= \int_{0}^{1} g(F^{-1}(q))dq\end{aligned} \tag{3.5}$$

where $H(q) := P(F^{-1}(q))$. Let

$$\text{MPS}(g \circ F^{-1}) := \{H \in L^1([0,1]) | H \text{ is monotone and satisfies (3.5)}\}.$$

The characterization of extreme points from Kleiner et al., 2021, Theorem 1, applies to $\text{MPS}(g \circ F^{-1})$. We use the following function to show that the characterization applies

[8]Recall that F^{-1} denotes the right-continuous inverse $F^{-1}(q) := \sup\{x | F(x) \leq q\}$ for $q \in [0,1]$. Thus, we have $F^{-1}(F(\theta)) = \sup\{\theta' | F(\theta') \leq F(\theta)\} = \theta$ by right-continuity of F, so $P(\theta) = P(F^{-1}(F(\theta)))$.

analogously to $\text{MPS}_F(g)$:

$$\gamma\colon \text{MPS}_F(g) \to \text{MPS}(g(F^{-1}))$$
$$P \mapsto H = P \circ F^{-1}$$

From the above, one can easily verify that γ is well-defined, linear, and preserves the L^1-norm.

By assumption, F is strictly increasing and continuous, thus $F : [\underline{\theta}, \overline{\theta}] \to [0,1]$ is invertible. Then, γ has an inverse $\gamma'\colon H \mapsto P = H \circ F$, so γ is a bijection. Thus, γ is an isomorphism on normed vector spaces and both the Choquet representation (Proposition 1) and the characterization of extreme points (Theorem 1) from Kleiner et al. (2021) immediately carry over to $\text{MPS}_F(g)$.

□

Screening with Two Types By standard arguments, the IR constraint holds for low type and the IC constraint holds for the high type, so we get the following maximization problem.

$$\max(1-\mu)a(\underline{\theta}) + \mu a(\overline{\theta})$$
$$s.t.$$
$$\overline{\theta}V(\overline{\theta}) - a(\overline{\theta}) = \overline{\theta}V(\underline{\theta}) - a(\underline{\theta}) \qquad (\text{IC}(\overline{\theta}))$$
$$\underline{\theta}V(\underline{\theta}) - a(\underline{\theta}) = \underline{\theta}^2 \qquad (\text{IR}(\underline{\theta}))$$

We can rewrite the objective as follows:

$$a(\underline{\theta}) = \underline{\theta}V(\underline{\theta}) - \underline{\theta}^2$$
$$a(\overline{\theta}) = a(\underline{\theta}) + \overline{\theta}(V(\overline{\theta}) - V(\underline{\theta})) = \overline{\theta}V(\overline{\theta}) - (\overline{\theta} - \underline{\theta})V(\underline{\theta}) - \underline{\theta}^2$$
$$\Rightarrow (1-\mu)a(\underline{\theta}) + \mu a(\overline{\theta}) = V(\underline{\theta})((1-\mu)\underline{\theta} - \mu(\overline{\theta} - \underline{\theta})) + \mu\overline{\theta}V(\overline{\theta}) - \underline{\theta}^2$$

Modulo a constant, the objective is

$$(1-\mu)h(\underline{\theta})V(\underline{\theta}) + \mu h(\overline{\theta})V(\overline{\theta})$$

where

$$h(\underline{\theta}) := \underline{\theta} - \frac{\mu}{1-\mu}\left(\overline{\theta} - \underline{\theta}\right) < \underline{\theta}$$
$$h(\overline{\theta}) := \overline{\theta}$$

and thus $h(\underline{\theta}) < \underline{\theta} < \overline{\theta} = h(\overline{\theta})$, or $h(\overline{\theta}) - h(\underline{\theta}) > 0$.

Finally, we rewrite our objective as a linear combination of $A = \mu V(\underline{\theta}) + (1-\mu)V(\overline{\theta})$ and $I = V(\overline{\theta}) - V(\underline{\theta})$, we get

$$V = \underbrace{\left((1-\mu)h(\underline{\theta}) + \mu h(\overline{\theta})\right)}_{=\underline{\theta}>0} A + \mu(1-\mu)\underbrace{\left(h(\overline{\theta}) - h(\underline{\theta})\right)}_{>0} I.$$

That is, the principal can be understood as caring about a linear combination of aggregate image utility A and incentive I with positive weights each.

Proof of Theorem 12

Proof. We prove this result by reformulating the principal's problem as a standard information principal's problem with a strictly concave value function, making no information revelation uniquely optimal.

Let $\Delta^{\mu}([0,1]) \subset \Delta([0,1])$ denote the set of distributions τ over posteriors π that are Bayes-consistent, that is, $\mathbb{E}_{\tau}[\pi] = \mu$. We adopt the posterior approach and write the principal's objective U_{τ} as a function of the Bayes-consistent distribution $\tau \in \Delta^{\mu}([0,1])$ over posteriors $\pi \in [0,1]$,

$$U_{\tau} = (1-\mu)h(\underline{\theta})V_{\tau}(\underline{\theta}) + \mu h(\overline{\theta})V_{\tau}(\overline{\theta}), \tag{3.6}$$

where $V_\tau(\theta) = \mathbb{E}_\tau[v(\pi)|\boldsymbol{\theta} = \theta]$ is the expected value of the posterior conditional on the true state being θ. Ignoring the monotonicity constraint, $V_\tau(\underline{\theta}) \leq V_\tau(\overline{\theta})$ for now, the principal's problem is simply choosing $\tau \in \Delta^\mu([0,1])$ to maximize U_τ. Conveniently, one can reformulate this conditional expectation as an unconditional expectation with respect to τ (see also Alonso and Câmara, 2016; Doval and Smolin, 2024):

$$V_\tau(\underline{\theta}) := \mathbb{E}_\tau\Big[\frac{1-\pi}{1-\mu}v(\pi)\Big], \tag{3.7}$$

$$V_\tau(\overline{\theta}) := \mathbb{E}_\tau\Big[\frac{\pi}{\mu}v(\pi)\Big]. \tag{3.8}$$

Inserting (3.7) and (3.8) into (3.6), we rewrite the objective as an unconditional expectation as in standard information principal's problems,

$$U_\tau = \mathbb{E}_\tau[\underbrace{\big((1-\pi)h(\underline{\theta}) + \pi h(\overline{\theta})\big)v(\pi)}_{\tilde{v}(\pi):=}].$$

Taking the derivative of the value function $\tilde{v}(\pi)$ with respect to $\pi \in [0,1]$ gives

$$\tilde{v}'(\pi) = (h(\overline{\theta}) - h(\underline{\theta}))v(\pi) + ((1-\pi)h(\underline{\theta}) + \pi h(\overline{\theta}))v'(\pi)$$
$$\tilde{v}''(\pi) = 2(h(\overline{\theta}) - h(\underline{\theta}))v'(\pi) + ((1-\pi)h(\underline{\theta}) + \pi h(\overline{\theta}))v''(\pi)$$

Under $v'(\pi) \neq 0$, we have

$$\tilde{v}''(\pi) < 0 \Leftrightarrow \frac{-v''(\pi)}{v'(\pi)} > \frac{2(h(\overline{\theta}) - h(\underline{\theta}))}{(1-\pi)h(\underline{\theta}) + \pi h(\overline{\theta})}.$$

Let $\delta(\mu) \in \Delta([0,1])$ denote the degenerate distribution that reveals no information. Under a strictly concave value function $\tilde{v}(\pi)$, Jensen's inequality implies

$$U_\tau = \mathbb{E}_\tau[\tilde{v}(\pi)] \leq \tilde{v}(\mathbb{E}_\tau[\pi]) = \tilde{v}(\mu) = U_{\delta(\mu)}$$

with a strict inequality if τ is not degenerate, $\tau \neq \delta(\mu)$. Because $\tau = \delta(\mu)$ satisfies the monotonicity constraint, $V_\tau(\underline{\theta}) \leq V_\tau(\overline{\theta})$, no information revelation is uniquely optimal. $\square$

Proof of Proposition 6

Proof. Consider the allocation $(b - \frac{\mu}{1-\mu}\varepsilon, b + \varepsilon)$ for $\frac{1-\mu}{\mu}b \geq \varepsilon > 0$. This allocation is budget balanced for all ε. We show that for ε small enough, this allocation satisfies the monotonicity constraint and improves upon the equal allocation (b, b).

Recall that the principal's objective is a linear combination of the aggregate image utility $A = \mu V(\underline{\theta}) + (1-\mu)V(\overline{\theta})$ and incentive $I = V(\overline{\theta}) - V(\underline{\theta})$. By u being differentiable at b, we can apply Taylor's theorem and rewrite A and I as functions of ε up to second-order terms in ε as:

$$A(\varepsilon) = (1-\mu)u(b - \frac{\mu}{1-\mu}\varepsilon) + \mu \cdot u(b+\varepsilon) = u(b) + \underbrace{\left(\mu\varepsilon - (1-\mu)\frac{\mu}{1-\mu\varepsilon}\right)}_{=0} u'(b) + o(\varepsilon^2)$$

$$I(\varepsilon) = u(b+\varepsilon) - u(b - \frac{\mu}{1-\mu}\varepsilon) = \frac{1}{1-\mu}u'(b)\varepsilon + o(\varepsilon^2)$$

Thus, under $u'(b) > 0$ there is a first-order gain to the incentive I, whereas the change to the aggregate image utility A is second-order. Hence, even if the change to the aggregate image utility A is negative (which it is under concave u), the positive effect on the incentive dominates for small ε as long as the principal has a finite marginal rate of substitution between A and I. Let the principal's weights on A and I be h_A and h_I, respectively. By section 3.5, both weights are positive, $h_A, h_I > 0$. Taking the difference between the achieved value under $\varepsilon > 0$ and the value under $\varepsilon = 0$, which corresponds to the equal allocation, we get

$$\lim_{\varepsilon \to 0} \left(h_A A(\varepsilon) + h_I I(\varepsilon)\right) - h_A v(b) = \lim_{\varepsilon \to 0} \underbrace{h_I \frac{1}{1-\mu} v'(\mu)}_{>0} \varepsilon + o(\varepsilon^2) > 0$$

Thus, for ε small enough, the principal's value is necessarily higher than under the equal allocation. Moreover, for such ε, we have $u(b+\varepsilon) > u(b - \frac{\mu}{1-\mu}\varepsilon)$, so the monotonicity constraint is satisfied. $\square$

Bibliography

Abramowitz, A. (2010). *The disappearing center: Engaged citizens, polarization, and American democracy*. Yale University Press.

Abramowitz, A., & Saunders, K. L. (2008). Is Polarization a Myth? *The Journal of Politics*, *70*(2), 542–555.

Acemoglu, D., Chernozhukov, V., & Yildiz, M. (2016). Fragility of asymptotic agreement under Bayesian learning: Fragility of asymptotic agreement. *Theoretical Economics*, *11*(1), 187–225.

Achen, C., & Bartels, L. (2017, December). *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton University Press.

Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of economic Literature*, *54*(2), 442–92.

Albano, G. L., & Lizzeri, A. (2001). Strategic Certification and Provision of Quality. *International Economic Review*, *42*(1), 267–283.

Albrecht, B. C., & Whitmeyer, M. (2023, April). Comparison Shopping: Learning Before Buying From Duopolists.

Ali, S. N., & Bénabou, R. (2020). Image versus Information: Changing Societal Norms and Optimal Privacy. *American Economic Journal: Microeconomics*, *12*(3), 116–164.

Aliprantis, C. D., & Border, K. C. (2006). *Infinite Dimensional Analysis: A Hitchhikers Guide* (Vol. 10).

Alonso, R., & Câmara, O. (2016). Persuading voters. *American Economic Review*, *106*(11), 3590–3605.

Amari, S. (2016). *Information geometry and its applications* (Vol. 194). Springer.

Anderson, S. P., Goeree, J. K., & Ramer, R. (1997). Location, location, location. *journal of economic theory*, *77*(1), 102–127.

Andreoni, J., Harbaugh, W., & Vesterlund, L. (2003). The Carrot or the Stick: Rewards, Punishments, and Cooperation. *American Economic Review*, *93*(3), 893–902.

Ansolabehere, S., Rodden, J., & Snyder, J. M. (2008). The strength of issues: Using multiple measures to gauge preference stability, ideological constraint, and issue voting. *American Political Science Review*, *102*(2), 215–232.

Aragones, E., & Palfrey, T. R. (2002). Mixed Equilibrium in a Downsian Model with a Favored Candidate. *Journal of Economic Theory*, *103*(1), 131–161.

Armstrong, M., & Zhou, J. (2022). Consumer Information and the Limits to Competition. *American Economic Review*, *112*(2), 534–577.

Ashworth, S., & Bueno de Mesquita, E. (2009). Elections with platform and valence competition. *Games and Economic Behavior*, *67*(1), 191–216.

Aumann, R. J. (1997). *Rationality and bounded rationality*. Springer.
Bachmann, F., Sarasua, C., & Bernstein, A. (2024, April). Fast and Adaptive Questionnaires for Voting Advice Applications.
Bafumi, J., & Herron, M. C. (2010). Leapfrog Representation and Extremism: A Study of American Voters and Their Members in Congress. *American Political Science Review*, *104*(3), 519–542.
Baliga, S., Hanany, E., & Klibanoff, P. (2013). Polarization and Ambiguity. *American Economic Review*, *103*(7), 3071–3083.
Beattie, J., Baron, J., Hershey, J. C., & Spranca, M. D. (1994). Psychological determinants of decision attitude. *Journal of Behavioral Decision Making*, *7*(2), 129–144.
Bell, D. E. (1982). Regret in Decision Making under Uncertainty. *Operations Research*, *30*(5), 961–981.
Bell, D. E. (1985). Disappointment in decision making under uncertainty. *Operations research*, *33*(1), 1–27.
Bénabou, R., & Pycia, M. (2002). Dynamic inconsistency and self-control: A planner–doer interpretation. *Economics Letters*, *77*(3), 419–424.
Bénabou, R., & Tirole, J. (2004). Willpower and personal rules. *Journal of Political Economy*, *112*(4), 848–886.
Bénabou, R., & Tirole, J. (2006). Incentives and Prosocial Behavior. *American Economic Review*, *96*(5), 1652–1678.
Bénabou, R., & Tirole, J. (2011). *Laws and norms* (tech. rep.). National Bureau of Economic Research.
Bernhardt, D., Buisseret, P., & Hidir, S. (2020). The Race to the Base. *American Economic Review*, *110*(3), 922–942.
Bernheim, B. D., & Rangel, A. (2004). Addiction and cue-triggered decision processes. *American economic review*, *94*(5), 1558–1590.
Beshears, J., Choi, J. J., Clayton, C., Harris, C., Laibson, D., & Madrian, B. C. (2020, July). *Optimal Illiquidity* (Working Paper No. 27459). National Bureau of Economic Research.
Biglaiser, G., Gu, J., & Li, F. (2024). Information Acquisition and Product Differentiation Perception. *American Economic Journal: Microeconomics*.
Bischof, D., & Wagner, M. (2019). Do Voters Polarize When Radical Parties Enter Parliament? *American Journal of Political Science*, *63*(4), 888–904.
Bleichrodt, H., & Wakker, P. P. (2015). Regret theory: A bold alternative to the alternatives. *The Economic Journal*, *125*(583), 493–532.
Bloedel, A. W., & Zhong, W. (2020). *The cost of optimally-acquired information* (tech. rep.). Technical report, Working paper, Stanford University.
Blumkin, T., Margalioth, Y., & Sadka, E. (2015). Welfare stigma re-examined. *Journal of Public Economic Theory*, *17*(6), 874–886.
Bodner, R., & Prelec, D. (2003). Self-signaling and diagnostic utility in everyday decision making. *The psychology of economic decisions*, *1*(105), 26.
Bordalo, P., Gennaioli, N., & Shleifer, A. (2012). Salience theory of choice under risk. *The Quarterly journal of economics*, *127*(3), 1243–1285.
Bordalo, P., Gennaioli, N., & Shleifer, A. (2022). Salience. *Annual Review of Economics*, *14*, 521–544.

Bornschier, S. (2010a). Cleavage Politics and the Populist Right. The New Cultural Conflict in Western Europe. *The Social Logic of Politics*.

Bornschier, S. (2010b). The New Cultural Divide and the Two-Dimensional Political Space in Western Europe. *West European Politics*, *33*(3), 419–444.

Bradley, R., & Stefánsson, H. O. (2017). Counterfactual Desirability. *The British Journal for the Philosophy of Science*, *68*(2), 485–533.

Brocas, I., & Carrillo, J. D. (2008). The brain as a hierarchical organization. *American Economic Review*, *98*(4), 1312–46.

Bursztyn, L., & Jensen, R. (2017). Social image and economic behavior in the field: Identifying, understanding, and shaping social pressure. *Annual Review of Economics*, *9*, 131–153.

Butera, L., Metcalfe, R., Morrison, W., & Taubinsky, D. (2022). Measuring the welfare effects of shame and pride. *American Economic Review*, *112*(1), 122–68.

Callander, S., & Carbajal, J. C. (2022). Cause and Effect in Political Polarization: A Dynamic Analysis. *Journal of Political Economy*, *130*(4), 825–880.

Calvert, R. L. (1985). Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. *American Journal of Political Science*, 69–95.

Camille, N., Coricelli, G., Sallet, J., Pradat-Diehl, P., Duhamel, J.-R., & Sirigu, A. (2004). The Involvement of the Orbitofrontal Cortex in the Experience of Regret. *Science*, *304*(5674), 1167–1170.

Campbell, A., Converse, P. E., Miller, W. E., & Stokes, D. E. (1960). *The American Voter*. University of Chicago Press.

Caplin, A., & Dean, M. (2013). *Behavioral implications of rational inattention with shannon entropy* (tech. rep.). National Bureau of Economic Research.

Caplin, A., Dean, M., & Leahy, J. (2019). Rational inattention, optimal consideration sets, and stochastic choice. *The Review of Economic Studies*, *86*(3), 1061–1094.

Caplin, A., Dean, M., & Leahy, J. (2022). Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy. *Journal of Political Economy*, *130*(6), 1676–1715.

Caplin, A., & Nalebuff, B. (1991). Aggregation and imperfect competition: On the existence of equilibrium. *Econometrica: Journal of the Econometric Society*, 25–59.

Carpini, M. X. D., & Keeter, S. (1996). *What Americans know about politics and why it matters*. Yale University Press.

Carroll, G. (2019). Robust incentives for information acquisition. *Journal of Economic Theory*, *181*, 382–420.

Casella, A. (2005). Storable votes. *Games and Economic Behavior*, *51*(2), 391–419.

Chatterjee, K., & Krishna, R. V. (2009). A "Dual Self" Representation for Stochastic Temptation. *American Economic Journal: Microeconomics*, *1*(2), 148–67.

Choi, M., Dai, A. Y., & Kim, K. (2018). Consumer Search and Price Competition. *Econometrica*, *86*(4), 1257–1281.

Converse, P. (1964). The nature of belief systems in mass publics. *Ideology and Discontent*.

Coricelli, G., Critchley, H. D., Joffily, M., O'Doherty, J. P., Sirigu, A., & Dolan, R. J. (2005). Regret and its avoidance: A neuroimaging study of choice behavior. *Nature Neuroscience*, *8*(9), 1255–1262.

Cunha, M., Osório, A., & Ribeiro, R. M. (2022, November). Endogenous Product Design and Quality When Consumers Have Heterogeneous Limited Attention.

Curello, G., & Sinander, L. (2024, May). The comparative statics of persuasion.
d'Aspremont, C., Gabszewicz, J. J., & Thisse, J.-F. (1979). On Hotelling's" Stability in competition". *Econometrica: Journal of the Econometric Society*, 1145–1150.
Daughety, A. F., & Reinganum, J. F. (2010). Public Goods, Social Pressure, and the Choice between Privacy and Publicity. *American Economic Journal: Microeconomics*, *2*(2), 191–221.
Dekel, E., & Lipman, B. L. (2012). Costly Self-Control and Random Self-Indulgence. *Econometrica*, *80*(3), 1271–1302.
DeMarzo, P. M., Vayanos, D., & Zwiebel, J. (2003). Persuasion bias, social influence, and unidimensional opinions. *The Quarterly journal of economics*, *118*(3), 909–968.
Demski, J. S., & Sappington, D. E. (1987). Delegated expertise. *Journal of Accounting Research*, 68–89.
Dewatripont, M., Jewitt, I., & Tirole, J. (1999). The Economics of Career Concerns, Part I: Comparing Information Structures. *The Review of Economic Studies*, *66*(1), 183–198.
Doval, L., & Smolin, A. (2024). Persuasion and Welfare. *Journal of Political Economy*, *132*(7), 2451–2487.
Downs, A. (1957). An Economic Theory of Political Action in a Democracy. *Journal of Political Economy*, *65*(2), 135–150.
Dubey, P. K., & Geanakoplos, J. (2014, July). Games with Money and Status: How Best to Incentivize Work.
Duflo, E., Kremer, M., & Robinson, J. (2011). Nudging farmers to use fertilizer: Theory and experimental evidence from Kenya. *American economic review*, *101*(6), 2350–90.
Duggan, J., & Martinelli, C. (2011). A Spatial Theory of Media Slant and Voter Choice. *The Review of Economic Studies*, *78*(2), 640–666.
Duggan, J. (2017). A survey of equilibrium analysis in spatial models of elections. *Unpublished manuscript*.
Dun, L., & Jessee, S. (2020). Demographic Moderation of Spatial Voting in Presidential Elections. *American Politics Research*, *48*(6), 750–762.
Duverger, M. (1954). *Political parties, their organization and activity in the modern state*. Methuen.
Dworczak, P., & Kolotilin, A. (2023). The persuasion duality.
Edmans, A., Gabaix, X., & Jenter, D. (2017). Executive compensation: A survey of theory and evidence. *The handbook of the economics of corporate governance*, *1*, 383–539.
Eguia, J., & Hu, T.-W. (2022). *Voter polarization and extremism* (tech. rep.).
Ely, J. C. (2011). Kludged. *American Economic Journal: Microeconomics*, *3*(3), 210–231.
Enelow, J. M., & Hinich, M. J. (1989). A general probabilistic spatial theory of elections. *Public Choice*, *61*(2), 101–113.
Enke, B., Rodríguez-Padilla, R., & Zimmermann, F. (2023). Moral Universalism and the Structure of Ideology. *The Review of Economic Studies*, *90*(4), 1934–1962.
Esteban, J., & Ray, D. (1994). On the measurement of polarization. *Econometrica: Journal of the Econometric Society*, 819–851.
Esteban, J., & Ray, D. (2012, April). Comparing Polarization Measures. In M. R. Garfinkel & S. Skaperdas (Eds.), *The Oxford Handbook of the Economics of Peace and Conflict*. Oxford University Press.

Evans, R. B., Gomez, J.-P., Ma, L., & Tang, Y. (2023, April). Peer versus Pure Benchmarks in the Compensation of Mutual Fund Managers.

Feddersen, T. J., & Pesendorfer, W. (1996). The swing voter's curse. *The American economic review*, 408–424.

Feddersen, T. J., & Sandroni, A. (2006). Ethical voters and costly information acquisition. *Quarterly Journal of Political Science*, *1*(3), 287–312.

Fiorina, M. P., & Abrams, S. J. (2008). Political polarization in the American public. *Annual Review of Political Science*, *11*, 563.

Foster, D. P., & Vohra, R. (1999). Regret in the On-Line Decision Problem. *Games and Economic Behavior*, *29*(1), 7–35.

Fowler, A., Hill, S. J., Lewis, J. B., Tausanovitch, C., Vavreck, L., & Warshaw, C. (2022). Moderates. *American Political Science Review*, 1–18.

Fudenberg, D., & Levine, D. K. (2006). A dual-self model of impulse control. *American economic review*, *96*(5), 1449–1476.

Fudenberg, D., & Levine, D. K. (2012). Timing and self-control. *Econometrica*, *80*(1), 1–42.

Gabillon, E. (2020). When choosing is painful: Anticipated regret and psychological opportunity cost. *Journal of Economic Behavior & Organization*, *178*, 644–659.

Groseclose, T. (2001). A Model of Candidate Location When One Candidate Has a Valence Advantage. *American Journal of Political Science*, *45*(4), 862–886.

Gul, F. (1991). A theory of disappointment aversion. *Econometrica: Journal of the Econometric Society*, 667–686.

Halevy, Y., & Feltkamp, V. (2005). A Bayesian approach to uncertainty aversion. *The Review of Economic Studies*, *72*(2), 449–466.

Halff, H. M. (1976). Choice theories for differentially comparable alternatives. *Journal of Mathematical Psychology*, *14*(3), 244–246.

Hansson, I., & Stuart, C. (1984). Voting competitions with interested politicians: Platforms do not converge to the preferences of the median voter. *Public Choice*, *44*(3), 431–441.

Hardy, G. H., Littlewood, J. E., & Pólya, G. (1934). *Inequalities*. Cambridge University Press.

Hare, C. (2022). Constrained Citizens? Ideological Structure and Conflict Extension in the US Electorate, 1980–2016. *British Journal of Political Science*, *52*(4), 1602–1621.

Hare, C., Armstrong, D. A., Bakker, R., Carroll, R., & Poole, K. T. (2015). Using Bayesian Aldrich-McKelvey Scaling to Study Citizens' Ideological Preferences and Perceptions. *American Journal of Political Science*, *59*(3), 759–774.

Hare, C., Highton, B., & Jones, B. (2023). Assessing the Structure of Policy Preferences: A Hard Test of the Low Dimensionality Hypothesis. *The Journal of Politics*.

He, J., & Natenzon, P. (2024). Moderate utility. *American Economic Review: Insights*, *6*(2), 176–195.

Hébert, B., & Woodford, M. (2021). Neighborhood-Based Information Costs. *American Economic Review*, *111*(10), 3225–3255.

Hébert, B., & Woodford, M. (2023). Rational inattention when decisions take time. *Journal of Economic Theory*, *208*, 105612.

Heiner, R. A. (1983). The origin of predictable behavior. *The American economic review*, *73*(4), 560–595.

Hill, S. J., & Tausanovitch, C. (2015). A Disconnect in Representation? Comparison of Trends in Congressional and Public Polarization. *The Journal of Politics*, *77*(4), 1058–1075.

Hinich, M. J. (1977). Equilibrium in spatial voting: The median voter result is an artifact. *Journal of Economic Theory*, *16*(2), 208–219.

Holmström, B. (1979). Moral Hazard and Observability. *The Bell Journal of Economics*, *10*(1), 74–91.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *The review of Economic studies*, *66*(1), 169–182.

Howard, R. A. (1992). In Praise of the Old Time Religion. In W. Edwards (Ed.), *Utility Theories: Measurements and Applications* (pp. 27–55). Springer Netherlands.

Hsiaw, A. (2013). Goal-setting and self-control. *Journal of Economic Theory*, *148*(2), 601–626.

Hu, L., Li, A., & Segal, I. (2023). The Politics of Personalized News Aggregation. *Journal of Political Economy Microeconomics*, *1*(3), 463–505.

Humphrey, S. J. (2004). Feedback-conditional regret theory and testing regret-aversion in risky choice. *Journal of Economic Psychology*, *25*(6), 839–857.

Inglehart, R., & Klingemann, H.-D. (1976). Party identification, ideological preference and the left-right dimension among Western mass publics. *Party identification and beyond*, 243–273.

Iyengar, S. S., Wells, R. E., & Schwartz, B. (2006). Doing Better but Feeling Worse: Looking for the "Best" Job Undermines Satisfaction. *Psychological Science*, *17*(2), 143–150.

Janis, I. L., & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice, and commitment.* Free press.

Jessee, S. A. (2009). Spatial voting in the 2004 presidential election. *American Political Science Review*, *103*(1), 59–81.

Jessee, S. A. (2012). *Ideology and spatial voting in American elections*. Cambridge University Press.

Jewitt, I. (2004). Notes on the 'Shape' of Distributions.

Jung, J., Kim, J. H., Matějka, F., & Sims, C. A. (2019). Discrete Actions in Information-Constrained Decision Problems. *The Review of Economic Studies*, *86*(6), 2643–2667.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, *47*(2), 263–292.

Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, *101*(6), 2590–2615.

Kaplow, L., & Shavell, S. (2007). Moral rules, the moral sentiments, and behavior: Toward a theory of an optimal moral system. *Journal of Political Economy*, *115*(3), 494–514.

Kartik, N., Lee, S., Liu, T., & Rappoport, D. (2022). *Beyond Unbounded Beliefs: How Preferences and Information Interplay in Social Learning* (tech. rep.).

Kaur, S., Kremer, M., & Mullainathan, S. (2015). Self-Control at Work. *Journal of Political Economy*, *123*(6), 1227–1277.

Keltner, D., & Gross, J. J. (1999). Functional accounts of emotions. *Cognition & Emotion*, *13*(5), 467–480.

Kleiner, A., Moldovanu, B., & Strack, P. (2021). Extreme points and majorization: Economic applications. *Econometrica*, *89*(4), 1557–1593.

Koch, A. K., & Nafziger, J. (2011). Self-regulation through goal setting. *Scandinavian Journal of Economics*, *113*(1), 212–227.

Kőszegi, B. (2003). Health anxiety and patient behavior. *Journal of health economics*, *22*(6), 1073–1084.

Kőszegi, B. (2014). Behavioral contract theory. *Journal of Economic Literature*, *52*(4), 1075–1118.

Kőszegi, B., & Matějka, F. (2020). Choice simplification: A theory of mental budgeting and naive diversification. *The Quarterly Journal of Economics*, *135*(2), 1153–1207.

Kőszegi, B., & Rabin, M. (2006). A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics*, *121*(4), 1133–1165.

Kőszegi, B., & Szeidl, A. (2013). A model of focusing in economic choice. *The Quarterly Journal of Economics*, *128*(1), 53–104.

Kriesi, H., Grande, E., Lachat, R., Dolezal, M., Bornschier, S., & Frey, T. (2006). Globalization and the transformation of the national political space: Six European countries compared. *European Journal of Political Research*, *45*(6), 921–956.

Kriesi, H., Grande, E., Lachat, R., Dolezal, M., Bornschier, S., & Frey, T. (2008). *West European Politics in the Age of Globalization*. Cambridge University Press.

Ladha, K. (1991). A spatial model of legislative voting with perceptual error. *Public Choice*, *68*(1-3).

Larrick, R. P., & Boles, T. L. (1995). Avoiding Regret in Decisions with Feedback: A Negotiation Example. *Organizational Behavior and Human Decision Processes*, *63*(1), 87–97.

Lauderdale, B. E. (2013). Does Inattention to Political Debate Explain the Polarization Gap between the U.S. Congress and Public? *Public Opinion Quarterly*, *77*(S1), 2–23.

Ledyard, J. O. (1984). The pure theory of large two-candidate elections. *Public Choice*, *44*(1), 7–41.

Leeper, T. J., & Slothuus, R. (2014). Political Parties, Motivated Reasoning, and Public Opinion Formation. *Political Psychology*, *35*(S1), 129–156.

Lelkes, Y. (2016). Mass polarization: Manifestations and measurements. *Public Opinion Quarterly*, *80*(S1), 392–410.

Lenz, G. S. (2012, October). *Follow the Leader?: How Voters Respond to Politicians' Policies and Performance*. University of Chicago Press.

Levy, G., & Razin, R. (2015). Correlation neglect, voting behavior, and information aggregation. *American Economic Review*, *105*(4), 1634–1645.

Li, A., & Hu, L. (2023). Electoral accountability and selection with personalized information aggregation. *Games and Economic Behavior*, *140*, 296–315.

Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological review*, *125*(1), 1.

Lindbeck, A., & Weibull, J. (2020). Delegation of investment decisions, and optimal remuneration of agents. *European Economic Review*, *129*, 103559.

Lindbeck, A., & Weibull, J. W. (1987). Balanced-budget redistribution as the outcome of political competition. *Public choice*, *52*(3), 273–297.

Linden, D. J. (2007). *The accidental mind*. Harvard University Press.

Lipnowski, E., & Ravid, D. (2023, April). Predicting Choice from Information Costs.

Lockwood, B. B. (2020). Optimal Income Taxation with Present Bias. *American Economic Journal: Economic Policy*, *12*(4), 298–327.

Loewenstein, G., & Lerner, J. S. (2003). The role of affect in decision making. *Handbook of affective sciences*, 619–642.

Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The economic journal*, *92*(368), 805–824.

Loomes, G., & Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *The Review of Economic Studies*, *53*(2), 271–282.

Loomes, G., & Sugden, R. (1987). Some implications of a more general form of regret theory. *Journal of Economic Theory*, *41*(2), 270–287.

Ma, L., Tang, Y., & Gómez, J.-P. (2019). Portfolio Manager Compensation in the U.S. Mutual Fund Industry. *The Journal of Finance*, *74*(2), 587–638.

Maćkowiak, B., Matějka, F., & Wiederholt, M. (2023). Rational inattention: A review. *Journal of Economic Literature*, *61*(1), 226–273.

Malamud, S., & Schrimpf, A. (2022, October). Persuasion by Dimension Reduction.

Malka, A., Lelkes, Y., & Soto, C. J. (2019). Are Cultural and Economic Conservatism Positively Correlated? A Large-Scale Cross-National Test. *British Journal of Political Science*, *49*(3), 1045–1069.

Marcus, G. (2008). *Kluge: The haphazard evolution of the human mind*. Houghton Mifflin Harcourt.

Martinelli, C. (2001). Elections with Privately Informed Parties and Voters. *Public Choice*, *108*(1), 147–167.

Martinelli, C. (2006). Would rational voters acquire costly information? *Journal of Economic Theory*, *129*(1), 225–251.

Matějka, F., & McKay, A. (2015). Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. *American Economic Review*, *105*(1), 272–298.

Matějka, F., & Tabellini, G. (2021). Electoral Competition with Rationally Inattentive Voters. *Journal of the European Economic Association*, *19*(3), 1899–1935.

McCarty, N. (2011, March). *Measuring Legislative Preferences*. Oxford University Press.

McCarty, N. (2019). *Polarization: What everyone needs to know®*. Oxford University Press.

McCarty, N., Poole, K. T., & Rosenthal, H. (2016). *Polarized America: The dance of ideology and unequal riches*. mit Press.

McCarty, N., Rodden, J., Shor, B., Tausanovitch, C., & Warshaw, C. (2019). Geography, uncertainty, and polarization. *Political Science Research and Methods*, *7*(4), 775–794.

McMurray, J. C. (2023). Why the Political World is Flat: An Endogenous "Left" and "Right" in Multidimensional Elections.

Mellers, B., Schwartz, A., & Ritov, I. (1999). Emotion-based choice. *Journal of Experimental Psychology: General*, *128*(3), 332.

Mensch, J., & Malik, K. (2024, November). Posterior-Mean Separable Costs of Information Acquisition.

Milgrom, P., & Shannon, C. (1994). Monotone Comparative Statics. *Econometrica*, *62*(1), 157–180.

Miller, E. K., & Cohen, J. D. (2001). An Integrative Theory of Prefrontal Cortex Function. *Annual Review of Neuroscience*, *24*(1), 167–202.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The review of economic studies*, *38*(2), 175–208.

Moldovanu, B., Sela, A., & Shi, X. (2007). Contests for Status. *Journal of Political Economy*, *115*(2), 338–363.

Morris, S., & Strack, P. (2019). The wald problem and the relation of sequential sampling and ex-ante information costs. *Available at SSRN 2991567*.

Moscarini, G., & Ottaviani, M. (2001). Price Competition for an Informed Buyer. *Journal of Economic Theory*, *101*(2), 457–493.

Mu, X., Pomatto, L., Strack, P., & Tamuz, O. (2021). From Blackwell Dominance in Large Samples to Rényi Divergences and Back Again. *Econometrica*, *89*(1), 475–506.

Nimark, K. P., & Sundaresan, S. (2019). Inattention and belief polarization. *Journal of Economic Theory*, *180*, 203–228.

Novák, V., Matveenko, A., & Ravaioli, S. (2024). The Status Quo and Belief Polarization of Inattentive Agents: Theory and Experiment. *American Economic Journal: Microeconomics*.

O'Connor, E., McCormack, T., & Feeney, A. (2012). The development of regret. *Journal of Experimental Child Psychology*, *111*(1), 120–127.

O'Donoghue, T., & Rabin, M. (2006). Optimal sin taxes. *Journal of Public Economics*, *90*(10-11), 1825–1849.

Onuchic, P., & Ray, D. (2023). Conveying value via categories. *Theoretical Economics*, *18*(4), 1407–1439.

Ortoleva, P., & Snowberg, E. (2015). Overconfidence in Political Behavior. *American Economic Review*, *105*(2), 504–535.

Osband, K. (1989). Optimal forecasting incentives. *Journal of Political Economy*, *97*(5), 1091–1112.

Palfrey, T. R., & Poole, K. T. (1987). The relationship between information, ideology, and voting behavior. *American journal of political science*, 511–530.

Patty, J. W. (2002). Equivalence of Objectives in Two Candidate Elections. *Public Choice*, *112*(1/2), 151–166.

Patty, J. W. (2005). Local equilibrium equivalence in probabilistic voting models. *Games and Economic Behavior*, *51*(2), 523–536.

Perego, J., & Yuksel, S. (2022). Media Competition and Social Disagreement. *Econometrica*, *90*(1), 223–265.

Persson, T., & Tabellini, G. (2002). *Political economics: Explaining economic policy*. MIT press.

Pew Reseach Center, W. D. (2014). Political polarization in the american public.

Pomatto, L., Strack, P., & Tamuz, O. (2023). The Cost of Information: The Case of Constant Marginal Costs. *American Economic Review*, *113*(5), 1360–1393.

Poole, K. T., & Rosenthal, H. (1985). A Spatial Model for Legislative Roll Call Analysis. *American Journal of Political Science*, *29*(2), 357–384.

Posner, E. (1975). Random coding strategies for minimum entropy. *IEEE Transactions on Information Theory*, *21*(4), 388–391.

Posner, R. A. (1978). The Right of Privacy. *Georgie Law Review*, *12*, 393–422.

Qin, J. (2020). Regret-based capital asset pricing model. *Journal of Banking & Finance*, *114*, 105784.

Quiggin, J. (1994). Regret theory with general choice sets. *Journal of Risk and Uncertainty*, *8*(2), 153–165.

Rappoport, D., & Somma, V. (2017, August). *Incentivizing Information Design* (SSRN Scholarly Paper No. ID 3001416). Social Science Research Network. Rochester, NY.

Ravid, D., Roesler, A.-K., & Szentes, B. (2022). Learning before trading: On the inefficiency of ignoring free information. *Journal of Political Economy*, *130*(2), 346–387.

Rayo, L. (2013). Monopolistic Signal Provision. *The B.E. Journal of Theoretical Economics*, *13*(1), 27–58.

Rayo, L., & Becker, G. S. (2007). Evolutionary efficiency and happiness. *Journal of Political Economy*, *115*(2), 302–337.

Rayo, L., & Segal, I. (2010). Optimal information disclosure. *Journal of political Economy*, *118*(5), 949–987.

Robson, A. J. (2001). The biological basis of economic behavior. *Journal of economic literature*, *39*(1), 11–33.

Roemer, J. E. (1997). Political–economic equilibrium when parties represent constituents: The unidimensional case. *Social Choice and Welfare*, *14*(4), 479–502.

Roemer, J. (1994). A theory of policy differentiation in single issue electoral politics. *Social Choice and Welfare*, *11*(4).

Saeedi, M., & Shourideh, A. (2023). Optimal Rating Design under Moral Hazard.

Samuelson, L., & Swinkels, J. M. (2006). Information, evolution and utility. *Theoretical Economics*, *1*(1), 119–142.

Sarver, T. (2008). Anticipating regret: Why fewer options may be better. *Econometrica*, *76*(2), 263–305.

Schelling, T. C. (1980). *The Strategy of Conflict: With a new Preface by the Author*. Harvard university press.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, *275*(5306), 1593–1599.

Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., & Lehman, D. R. (2002). Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology*, *83*(5), 1178–1197.

Shaked, M., & Shanthikumar, J. G. (Eds.). (2007). *Stochastic Orders*. Springer New York.

Shor, B., & Rogowski, J. C. (2018). Ideology and the US Congressional Vote. *Political Science Research and Methods*, *6*(2), 323–341.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, *50*(3), 665–690.

Smith, A. (2024). Optimal Signalling Incentives. *Working Paper*.

Smithies, A. (1941). Optimum Location in Spatial Competition. *Journal of Political Economy*, *49*(3), 423–439.

Spector, D. (2000). Rational debate and one-dimensional conflict. *The Quarterly Journal of Economics*, *115*(1), 181–200.

Sugden, R. (1986). Regret, recrimination and rationality. In *Recent developments in the foundations of utility and risk theory* (pp. 67–80). Springer.

Sugden, R. (1993). An Axiomatic Foundation for Regret Theory. *Journal of Economic Theory*, *60*(1), 159–180.

Sunstein, C. (2018, December). *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.

Szalay, D. (2005). The economics of clear advice and extreme options. *The Review of Economic Studies*, *72*(4), 1173–1198.

Tausanovitch, C., & Warshaw, C. (2012). *How Should We Choose Survey Questions to Measure Citizens' Policy Preferences?* (Tech. rep.). Citeseer.

Thaler, R. H., & Shefrin, H. M. (1981). An Economic Theory of Self-Control. *Journal of Political Economy*, 15.

Thereze, J. (2022a). Screening Costly Information.

Thereze, J. (2022b). Adverse Selection and Endogenous Information.

Tian, G., & Zhou, J. (1992). The maximum theorem and the existence of Nash equilibrium of (generalized) games without lower semicontinuities. *Journal of Mathematical Analysis and Applications*, *166*(2), 351–364.

Tirole, J. (2021). Digital dystopia. *American Economic Review*, *111*(6), 2007–2048.

Topkis, D. M. (1978). Minimizing a submodular function on a lattice. *Operations research*, *26*(2), 305–321.

Uhlig, H. (1996). A law of large numbers for large economies. *Economic Theory*, *8*, 41–50.

Vaeth, M. (2022). Self-Control through Emotions: Goal-Setting, Mental Accounting, and Demand for Commitment. *Working Paper*.

Van Der Brug, W., & Van Spanje, J. (2009). Immigration, Europe and the 'new' cultural dimension. *European Journal of Political Research*, *48*(3), 309–334.

Villani, C. (2009). *Optimal Transport* (M. Berger, B. Eckmann, P. De La Harpe, F. Hirzebruch, N. Hitchin, L. Hörmander, A. Kupiainen, G. Lebeau, M. Ratner, D. Serre, Y. G. Sinai, N. J. A. Sloane, A. M. Vershik, & M. Waldschmidt, Eds.; Vol. 338). Springer Berlin Heidelberg.

Whitmeyer, M., & Zhang, K. (2022). Buying Opinions. *arXiv preprint arXiv:2202.05249*.

Wittman, D. (1973). Parties as utility maximizers. *American Political Science Review*, *67*(2), 490–498.

Wittman, D. (1983). Candidate Motivation: A Synthesis of Alternative Theories. *The American Political Science Review*, *77*(1), 142–157.

Wittman, D. (1990). Spatial strategies when candidates have policy preferences. *Advances in the spatial theory of voting*, 66–98.

Yoder, N. (2022). Designing incentives for heterogeneous researchers. *Journal of Political Economy*, *130*(8), 2018–2054.

Yuksel, S. (2022). Specialized Learning and Political Polarization. *International Economic Review*, *63*(1), 457–474.

Zaller, J. (1992). *The nature and origins of mass opinion*. Cambridge University Press.

Zapechelnyuk, A. (2020). Optimal Quality Certification. *American Economic Review: Insights*, *2*(2), 161–176.

Zeelenberg, M. (1999a). Anticipated regret, expected feedback and behavioral decision making. *Journal of behavioral decision making*, *12*(2), 93–106.

Zeelenberg, M. (1999b). The use of crying over spilled milk: A note on the rationality and functionality of regret. *Philosophical Psychology*, *12*(3), 325–340.

Zeelenberg, M., Van Dijk, W. W., Manstead, A. S., & vanr de Pligt, J. (2000). On bad decisions and disconfirmed expectancies: The psychology of regret and disappointment. *Cognition & Emotion*, *14*(4), 521–541.

Zubrickas, R. (2015). Optimal Grading. *International Economic Review*, *56*(3), 751–776.