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

2016

Essays on Information Economics

Tolvanen, Juha Kasperi

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp010v838305d

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 on Information Economics

Juha Tolvanen

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: Professor Bo Honoré

September 2016

© Copyright by Juha Tolvanen, 2016.

All rights reserved.

Abstract

This thesis is both an empirical and a theoretical study of how imperfect information affects economic decision making and how to measure those effects. The first chapter describes a new method for testing and measuring moral hazard in insurance contracts for which panel data is available. The method is able to measure moral hazard separately from adverse selection. It uses the variation in risk over different objects with different levels of insurance owned by the same person to isolate the average effect that insurance coverage has on precautionary effort. The method requires relatively weak symmetry between the items and applies to a large variety of standard contracts.

The second chapter builds on the first chapter by utilizing the time dimension of panel data to measure moral hazard. I discuss how many every-day contracts imply different incentives to reduce risks at different times over the span of the contract. These differences can be used to measure how sophisticated agents are in allocating precautionary effort over time. I also discuss how the method presented in this chapter can be combined with the method from first chapter to obtain an identification strategy that is highly robust to alternative explanations. I apply each of my methods to a detailed insurance data set of virtual space ship insurance from a large virtual world called EVE Online and find strong evidence of moral hazard.

The third chapter (with Nemanja Antić) is a theoretical study of an alternating offers bargaining model where the players have private information about their own deadline after which they have to accept any non-negative offer. Such situations arise, for example, when a firm is bargaining with its employees and both sides are facing fixed costs and finite budgets. We show that the game has an intuitive screening equilibrium with delay. The equilibrium is wasteful in the strict sense that both players would prefer to commit to reveal their private information if such commitment device was available. The equilibrium highlights the potential benefits from third-party arbitration assuming that it helps players to commit to revealing private information.

Acknowledgements

First and foremost I am indebted to my advisor Bo Honoré. His advice and encouragement were invaluable throughout the research process. I thank Jan De Loecker and Jakub Kastl for their extremely helpful comments on earlier drafts of these chapters.

This dissertation has profoundly benefited from discussions (as well as classes) with Dilip Abreu, Matias Iaryczower, Yuliy Sannikov, Myrto Kalouptsidi, Roland Bénabou, Juuso Toikka and Klaus Kultti. I would also like to thank all of the seminar participants at both the theory and applied theory seminars at Princeton as well as at the University of Vienna and Einaudi Institute for Economics and Finance for their helpful comments.

I am indebted to my colleagues at the Helsinki Institute for Information Technology, especially Vili Lehdonvirta, Kai Huotari, Juho Hamari and Antti Ukkonen, for their invaluable input in the early stages of this project. I would like to thank Kristófer Hannesson, Kjartan Halldórsson and Eyjólfur Gudmundsson from CCP for providing me the data and helping me understand it. Thank you to Laura Hedden for all of the administrative help and guidance.

I am grateful beyond words to my peers Nemanja Antić, Olivier Darmouni, Pedro Olea, Sharon Traiberman and Dan Zeltzer, who made my years at Princeton constantly enjoyable and insightful and who helped me tirelessly at all stages of this project.

Parts of this dissertation have been presented at the 17th CEPR-JIE School on Applied Industrial Organization. I have benefited from financial support from the Princeton University Fellowship, the Department of Economics at Princeton University, OP-Pohjola Research Fund and Marcus Wallenberg Fund for Economic Research.

To my father.

Contents

Abstract iii

Acknowledgements iv

List of Tables ix

List of Figures x

1 Multiple Contracts and Measurement of Moral Hazard 1

1.1 Introduction 1

1.2 Theoretical Model and Hypotheses 7

1.3 Data 15

1.3.1 EVE Online and Its Spaceship Insurance 15

1.3.2 Some Descriptive Statistics 20

1.4 Econometric Model and Results 26

1.5 Experience and Moral Hazard 33

1.6 Conclusions and Future Work 36

1.A Omitted Proofs 37

1.A.1 Proofs for Section 1.2 38

1.B Supporting Tables 48

1.C Proportional Hazards Assumption and Goodness of Fit 49

2 Measuring Dynamic Moral Hazard with Panel Data 54

2.1 Introduction 54

2.2 Theoretical Model and Hypotheses 58
2.2.1 A Dynamic Model with a Single Insurable Unit 58
2.2.2 Dynamic Model with Two Insurable Items 67
2.3 Econometric Model and Results 71
2.3.1 The Evolution of Risk over the Insurance Contract in the Cross Section 72
2.3.2 Test for Dynamic Moral Hazard Using Multiple Ships Owned by the Same Agent 75
2.3.3 Experience and Timing of Moral Hazard 81
2.4 Conclusions 84
2.A Omitted Proofs 86
2.A.1 Proofs for Section 2.2.1 87
2.A.2 Proofs for Section 2.2.2 97

3 Bargaining with Private Deadlines **118**

3.1 Introduction 118
3.2 Related Literature 120
3.3 The Symmetric Model and Its Screening Equilibrium 122
3.3.1 A Partial Screening Equilibrium for the Subgames Starting from Period L and $L-1$ 123
3.3.2 The Equilibrium Strategies for Periods $L-k$, $k \geq 2$ 126
3.4 Asymmetric Deadlines 135
3.4.1 A Recursive Formulation for a Partial Screening Equilibrium for Periods after L_1+1 139
3.4.2 Period L_1+1 143
3.4.3 Equilibrium Strategies for Periods $t \leq L_1$ 146
3.5 Discussion of the Properties of the Equilibria 154
3.5.1 Efficiency Properties of the Equilibria 154

3.5.2 Trembling Hand Perfection and Conjecture on Uniqueness . . 155
3.5.3 The Effect of Deadlines on Bargaining Power 156
3.6 Conclusions and Future Research 157

Bibliography **158**

List of Tables

1.1 Descriptive statistics about players and their ship ownership 21
1.2 Descriptive statistics about characters and their ships 22
1.3 Distribution of assembled ships by insurance level 25
1.4 Estimates of the insurance parameter from stratified Cox regressions . 30
1.5 Estimates of the insurance parameter from stratified Cox regressions for the most and least experienced players 35
1.6 Ownership of uninsured and insured ships by ship type 48

List of Figures

1.1 Value of the maximization problem as a function of risk type 13

1.2 Accident probabilities as a function of risk type 13

1.3 The number of destruction events over the observation period. 28

1.4 Kaplan-Meier and Cox proportional hazards survival curves for insured and uninsured ships. 51

1.5 Kaplan-Meier and Cox proportional hazards survival curves for insured and uninsured Battlecruisers. 52

1.6 Kaplan-Meier and Cox proportional hazards survival curves for insured and uninsured Drakes. 53

2.1 Hazard ratio as a function of time left in the insurance contract. 95% confidence intervals calculated using standard errors clustered at the player level. 73

2.2 Hazard ratio as a function of time left in the insurance contract excluding first 8 days of the contract. 95% confidence intervals calculated using standard errors clustered at the player level. 74

2.3 Hazard ratio as a function of time left in the insurance contract. Model includes a user specific baseline hazard rate. 95% confidence intervals calculated using standard errors clustered at the player level. 76

2.4 Hazard ratio as a function of time left in the insurance contract excluding first 8 days of the contract. Model includes a user specific baseline hazard rate. 95% confidence intervals calculated using standard errors clustered at the player level. 77

2.5 Hazard ratio as a function of time left in the insurance contract. Model includes a user specific baseline hazard rate and ship age dummies. 95% confidence intervals calculated using standard errors clustered at the player level. 79

2.6 Hazard ratio as a function of time left in the insurance contract excluding first 8 days of the contract. Model includes a user specific baseline hazard rate and ship age dummies. 95% confidence intervals calculated using standard errors clustered at the player level. 80

2.7 Drakes' hazard ratios as a function of time left in the insurance contract. Model includes a user specific baseline hazard rate and ship age dummies. 95% confidence intervals calculated using standard errors clustered at the player level. 82

2.8 Drake's hazard ratios as a function of time left in the insurance contract excluding first 8 days of the contract. Model includes a user specific baseline hazard rate and ship age dummies. 95% confidence intervals calculated using standard errors clustered at the player level. 83

2.9 Hazard ratio as a function of time left in the insurance contract for the *least* experienced quartile of players. Model includes a user specific baseline hazard rate. 95% confidence intervals calculated using standard errors clustered at the player level. 84

2.10 Hazard ratio as a function of time left in the insurance contract for the *most* experienced quartile. Model includes a user specific baseline hazard rate. 95% confidence intervals calculated using standard errors clustered at the player level. 85

Chapter 1

Multiple Contracts and Measurement of Moral Hazard

1.1 Introduction

Moral hazard and adverse selection can cause severe market failures in insurance markets.[1] Both phenomena can lead to inefficient levels of coverage and prices or even market breakdown. Being able to measure the two phenomena separately would help to direct resources more optimally to either alleviate moral hazard by investing in monitoring technologies or to reduce adverse selection by more accurately screening the insured population. As was already noted by Chiappori and Salanié (2000), separately estimating the two is difficult, if not impossible, using only cross-sectional data on accidents and coverage decisions. In this chapter I introduce a new method for separately measuring moral hazard with insurance panel data by using variation in risk between differently insured items owned by the same agent. The second chapter of this dissertation introduces another method that focuses on the dynamic variation in risk over the contract period.

[1] See, for example, Pauly, 1968 and Holmström, 1979 on moral hazard and Rothschild and Stiglitz, 1976 on adverse selection.

The method from this chapter can be applied to situations where an agent purchases different levels of insurance for a number of similar items. Because all these items will share the agent's risk type, any variation in risk between them should be attributed to moral hazard. For example, consider a person who buys two televisions of the same make and buys product insurance or an additional warranty for one but not the other. Since both TVs are identical and owned by the same agent, differences in their accident risks cannot be due to adverse selection. If the insured television is still more likely to break down or get stolen, this difference in risk must be a result of differences in precautionary effort. Hence, estimating a model akin to a fixed effects model with agent fixed effects will allow the estimation of the average effect an increase in insurance coverage has on the decision maker's precautionary effort.

The key identification assumption guaranteeing the separation of moral hazard from adverse selection is that, from the agent's perspective, the insurable items are fairly similar. As I show in the theory section of the paper, the items need not be completely identical. For the empirical strategy to work, it is enough that the marginal cost and benefit from precautionary effort are the same for the two items when equal amounts of effort are spent on protecting them. The method works especially well with new items for which first insurance decisions are made at the time of purchase, because prior to owning the items the agent is unlikely to have any private information about their riskiness.

Furthermore, there is a wealth of situations that can be mapped to the theoretical model even if the insurable items seem initially different. For example, consider a situation where a person owns two otherwise identical cars but knows that one of them has some problems with its brakes and insures only that car. If, in absence of insurance, the agent repaired the brakes, then failing to repair them when the car is insured should still be classified as moral hazard. On the other hand, if the person did not fix the brakes even when the car is uninsured, then the broken brakes

of the insured car should be classified as car-specific adverse selection. In the latter case my method will still provide useful information by separating the decision maker-specific risk type from the item-specific combined effect of adverse selection and moral hazard. The separation is valuable, because it helps in evaluating the relative benefit of screening agents versus screening the property they own.

On the econometric side, I use versions of the Cox proportional hazards model (Cox, 1972) to construct my tests for moral hazard. The model has multiple benefits over commonly used probit and logit models used in earlier cross-sectional studies. First, it allows for highly flexible agent-specific base line hazard rates that let me capture dynamically varying heterogeneity in both tastes and risk types. Second, the model helps me to paint a more detailed picture of the agent's effort choices over time and flexibly estimate dynamic changes in moral hazard. For example, the model easily allows for time-varying covariates such as the time left in the insurance contract or the age of an insured item.

To illustrate the application of my methods I use data from a large multiplayer online game called EVE Online, where the players have the option to insure their virtual space ship against destruction. The game offers an ideal, almost laboratory-like environment for applying the methods for three reasons: First, the incentives in the game are clearly defined and, for example, the value of a ship or an insurance contract can be measured in real monetary terms. Second, the insurance contracts in the game are very similar to many real world insurance contracts and they satisfy all of the identifying assumptions required by my methods. Last, the server log data allows me to accurately measure all of the key variables without error and with a high dynamic frequency.

When I estimate the average difference in risk between insured and uninsured items owned by the same person over the whole contract period, I find that moral hazard leads to an insured ship being about 6 times more likely to get destroyed in any

time period compared to an uninsured ship. I also show that experience or learning is a key factor in the magnitude of moral hazard with the daily effect of moral hazard on risk being almost five times as big for experienced players than for new-comers.

It is noteworthy that, even though my model and econometric methods are framed in terms of insurance, their logic applies potentially to other contracts such as employment contracts. For example, when workers or companies are faced with similar projects with different incentive schemes, one could use my method to estimate the moral hazard associated with each scheme. Potential applications in employment markets range from crowdsourcing marketplaces such as Amazon Mechanical Turk or TaskRabbit to more ordinary project work such as consulting. Similarly, one could compare the spent resources and the quality of the output of building contractors when they are contracted with a fixed-price contract versus the outcomes under a cost-plus contract where the contractors costs are reimbursed.

The literature on measuring moral hazard separately from adverse selection is fairly new.[2] The papers that come thematically closest to mine are the ones by Abbring, Chiappori and Pinquet (2003) and Dionne, Michaud and Dahchour (2013).[3] Both papers use the bonus malus feature of French car insurance that increases the premium for all future years after an accident and reduces the premium after a year without an accident. The authors argue that moral hazard would imply that there should be a discontinuous increase in effort after an accident. The authors find little evidence of moral hazard in their data. Dionne, Michaud and Dachour (2013) look at a panel data over a longer period and are able to control for learning about one's risk type that could confound the results if bad drivers start exerting more effort after learning about their risk type from previous accidents. Their parametric model finds statistically significant evidence for moral hazard among drivers with less than 15

[2]There is a large body of empirical literature measuring the combined effect of moral hazard and adverse selection. For recent reviews, see Einav et al. (2010) and Cohen and Siegelman (2010).

[3]See also Abbring, Heckman, Chiappori and Pinquet (2003).

years of driving experience but does not find any evidence of moral hazard for drivers with more than 15 years of driving experience. It is noteworthy that a method that bases its identification on previous accidents may yield biased results. If an accident leads to the car being damaged or the driver being injured or traumatized, the agent is likely to drive her car much less than normally. Consequently, the likelihood of future accidents is lower during the months following an accident for reasons that have little to do with the agent's effort choices. The empirical pattern is much like what is generated by the incentives from the bonus malus system.

Another strain of literature estimating moral hazard uses discontinuities or kinks in the incentive structure of offered contracts. Dobbie and Skiba (2013) consider the payday lending market where moral hazard would imply that borrowers with larger loans have a smaller opportunity cost of defaulting, while adverse selection implies that borrowers with an ex-ante higher risk of defaulting know they are likely to default and hence select a loan with a higher value. Dobbie and Skiba notice that offered loan sizes change discontinuously with respect to a borrower's pay. Assuming that borrowers around the discontinuity are on average fairly similar, the authors are able to measure how much increased loan size affects the likelihood of a given borrower paying back the loan and thus get a measure of moral hazard. Einav, Finkelstein and Schrimpf (2015) is conceptually fairly similar.[4] It uses discontinuities in the co-insurance rate of Medicare Part D prescription drug coverage to estimate the effect insurance has on people's drug purchases.

My paper is unique in the way it uses variation in risk within a person, either between separate items owned by the same agent or over time for a given item, to factor out adverse selection and identify moral hazard by keeping the agent's risk type constant. Adams, Einav and Levin (2009) share a similarity with this idea, as they also use a special feature of their data to hold the agent's risk type constant while

[4]See also Einav, Finkelstein, Kluender and Schrimpf (2015).

varying the incentives for effort.[5] It uses data from a large auto sales company that lends money to its customers who have a hard time obtaining credit elsewhere. The forms of moral hazard and adverse selection are conceptually the same as in Dobbie and Skiba (2013). Their data contains two sources of variation in loan size which allow them to separate moral hazard from adverse selection: the choice of the down payment and the variation in the pricing of the purchased car which they claim to be orthogonal to the borrower's risk type conditional on his/her down payment choice. The first source of variation allows the authors to pin down the contract choice and hence control for any private information that the customer has about their credit worthiness. As long as the remaining variation in loan size that comes from the price of the offered car is not related to the customer's risk type in a way not captured by the down payment choice, then this variation will not be influenced by adverse selection and can be used to measure the effect that loan size has on moral hazard.

I also illustrate how virtual worlds can be used as living laboratories where economists can test their theories under fairly idealized incentive structures while still retaining extremely large sample sizes. The possibility of using virtual worlds as a testing ground for theories from social sciences is a fairly novel idea (see Bainbridge, 2007). These environments usually include many of the institutions found in normal societies, while the total complexity of social interaction is considerably reduced. This creates laboratory-like environments with study populations that are a thousand fold to what researchers usually can afford to hire for their experiments. Furthermore, if the researcher has access to the log data of the virtual world, he or she can accurately observe practically everything that happens in the world. Last, observing log data allows the researcher to run studies where the population under study is unaware of the fact that they are being observed. Such studies can avoid

[5]See also Einav, Jenkins and Levin (2012).

the usual threats to internal validity due to reactivity of the subjects to being in an experiment such as the Hawthorne effect.[6]

Economists have slowly become aware of the research potential of data gathered from the virtual worlds. Initial studies were mostly explorative trying to document what was happening in these online hang-outs (see, for example, Castronova 2001). The more recent papers have used virtual worlds to study social interaction in trust games (Fiedler et al., 2011 and Fiedler and Haruvy, 2009), online labor markets (Horton et al., 2011), the market for avatars (Castronova, 2004) and the general potential of virtual worlds for experiments (Chesney et al. 2009). All of these papers are experiments and not observational studies that would directly use the server log data the way I do here. For a previous study that uses log data, see Golde (2008) on the efficiency of thick markets.

The rest of the paper is organized as follows. The next section presents my theoretical models and formalizes my two identification strategies. Section 3 describes my virtual world data. Section 4 presents the results from applying my methods to the data. Section 5 investigates the effect of experience on moral hazard and Section 6 concludes the study.

1.2 Theoretical Model and Hypotheses

This section starts by introducing a simple static model wherein each agent owns two insurable objects and which provides my identificaton method with the agent "fixed effects". In short, the model is as follows: A population of heterogeneous, risk averse agents each owns two insurable items. The heterogeneity can be in terms of risk types, preferences or effort costs. They first choose which of the items to insure (if any) and then select a separate level of precautionary effort for both items. Finally,

[6]See, for example, Adair (1984) for a more detailed description of the Hawthorne effect and a review of empirical literature documenting it.

some items are randomly involved in an accident that destroys them resulting in a monetary loss for the agent. The probability of an accident is inversely related to how much precautionary effort was taken to protect it. If the item was insured, some of the loss is covered by the insurance. Conditional on the agent's type and effort choices the accidents happen independently of each other. I show that in a model like this, for a given person owning both an insured and an uninsured item, the person will always exercise more precautionary effort to protect the uninsured item. This difference can then be used to test moral hazard. All of the proofs for this section can be found in the Appendix.

Formally, I assume that every agent $j \in \mathbb{N}$ owns two items indexed by $i = 1, 2$, both of which can be insured against a monetary loss of L for a premium C to get a pay-out P in case there is an accident. I assume that $L - P > 0$, i.e. that there is a deductible in the insurance contract. Let $I_j \in \{0, 1\}^2$ be the agent's choice of insurance contracts where $I_{ij} = 1$ is interpreted to mean that item i is insured.

Suppose that agent j's monetary endowment is equal to ω_j. I assume that ω_j is independently and identically distributed according to some probability distribution F_ω for all $j \in \mathbb{N}$. Furthermore, I allow heterogeneity in terms of a random variable $\varepsilon_j \in \mathbb{R}$ that is independently and identically distributed according to some F_ε. This ε_j captures arbitrary variation in tastes, risk types and risk-aversion in ways that will become clear below.

The consumer's von Neumann-Morgenstern utility from consuming wealth c when his/her heterogeneity term is ε is given by $u(c, \varepsilon)$, where

$$u \colon \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}_+$$

I assume that u is continuous everywhere and twice continuously differentiable with respect to c. Furthermore, I require that $\partial_c u(c,\varepsilon) > 0,\ \partial_c^2 u(c,\varepsilon) < 0$ for all $c \in \mathbb{R}_+$ and $\varepsilon \in \mathbb{R}$. Notice that u can depend arbitrarily on ε.

Each agent chooses precautionary effort levels (x_{1j}, x_{2j}) that determine the probability of an accident independently for each item. Let $D_j \in \{0,1\}^2$ be a binary random vector where $D_{ij} = 1$ is interpreted as an accident that happened to item $i = 1, 2$ owned by agent j. I assume that

$$\mathbb{P}(D_{ij} = 1 \mid \varepsilon_j, x_{1j}, x_{2j}) = p(x_{ij}, \varepsilon_j),$$

where $p\colon \mathbb{R}_{++} \times \mathbb{R} \to [0,1]$ is a mapping that translates effort choices and risk/taste types to accident probabilities. Notice that I am assuming that given the same precautionary effort for both items owned by the same person, their accident probabilities will also be given by the same mapping. I require that conditional on ε_j and (x_{1j}, x_{2j}), D_{1j} is independent of D_{2j}. The assumption of conditional independence combined with the assumption that the accident probability for item 1 does not depend on the effort choice for item 2 precludes settings where investing in precautionary effort may cause spillover effects to the accident probability of the other item. For example, investing in fire alarms for your garage will not only protect the building but also the car parked in it and any other buildings near-by. This type of interdependence is not allowed by the model.

Furthermore, D_j is assumed to be independent of D_k for $j \neq k$. I also assume that $p(\cdot, \varepsilon_j)$ is twice continuously differentiable, decreasing and convex for every ε_j and that the mapping $(x_{1j}, x_{2j}) \mapsto p(x_{1j}, \varepsilon_j)p(x_{2j}, \varepsilon_j)$ is strictly convex. The latter assumption guarantees the concavity of the agent's optimization problem.

The utility cost of effort is measured by a function $E\colon \mathbb{R}_+^2 \times \mathbb{R} \to \mathbb{R}_+$. Here, $E(x_{1j}, x_{2j}, \varepsilon_j)$ is interpreted as the utility cost from effort levels (x_{1j}, x_{2j}) for person

j. This function is also allowed to depend arbitrarily on ε. The key identifying assumption for this section is that the marginal cost of effort for the two items is equal at equal effort levels:

Assumption 1 (Diagonal symmmetry)**.** *The cost function satisfies:* $\partial_1 E(x, x, \varepsilon_j) = \partial_2 E(x, x, \varepsilon_j)$ *for all* ε_j *and all* $x \in \mathbb{R}_+$.

This is clearly more general than assuming that E is symmetric with respect to x_{1j} and x_{2j}. Hence, the two items need not be completely identical, as long as the marginal return to effort is equal for the two items at equal effort levels. I also require that E is twice continuously differentiable, increasing and convex in its first two arguments for every ε_j and that $\partial_{12} E(x_1, x_2, \varepsilon_j) \geq 0$. To guarantee that first-order conditions are satisfied at some (x_{1j}, x_{2j}) I assume that for all x_{-i}, $\lim_{x_i \to 0} \partial_i E(x_1, x_2) = 0$ and $\lim_{x_i \to \infty} \partial_i E(x_1, x_2) = \infty$.

The utility maximization problem for agent j with a risk/taste type ε_j is given by

$$\begin{aligned}
&\max_{I_j, x_{1j}, x_{2j}} p(x_{1j}, \varepsilon) p(x_{2j}, \varepsilon_j) u(\omega_j - 2L + (I_{1j} + I_{2j})(P - C), \varepsilon_j) \\
&+ p(x_{1j}, \varepsilon_j)(1 - p(x_{2j}, \varepsilon_j)) u(\omega_j - L - (I_{1j} + I_{2j})C + I_{1j}P, \varepsilon_j) \\
&+ p(x_{2j}, \varepsilon_j)(1 - p(x_{1j}, \varepsilon_j)) u(\omega_j - L - (I_{1j} + I_{2j})C + I_{2j}P, \varepsilon_j) \\
&+ (1 - p(x_{1j}, \varepsilon_j))(1 - p(x_{2j}, \varepsilon_j)) u(\omega_j - (I_{1j} + I_{2j})C, \varepsilon_j) \\
&- E(x_{1j}, x_{2j}, \varepsilon_j). \qquad (1.2.1)
\end{aligned}$$

The first result of this section shows that, if an agent finds herself with both or none of her items insured, then the accident probabilities for both of the items will be the same. The point of this result is to highlight that even extreme asymmetries in the cost function that are off the diagonal will not translate into differences in effort choices when both items have the same insurance status.[7]

[7]Notice that I allow for selection into insurance based on differing marginal costs of effort between the items off the diagonal. However, I classify any resulting differences in risk between the insured

Proposition 1. *The following equalities hold:*

1. $\mathbb{E}[p(x_{1j}, \varepsilon_j) \mid I_{1j} = 0, I_{2j} = 0] = \mathbb{E}[p(x_{2j}, \varepsilon_j) \mid I_{1j} = 0, I_{2j} = 0]$

2. $\mathbb{E}[p(x_{1j}, \varepsilon_j) \mid I_{1j} = 1, I_{2j} = 1] = \mathbb{E}[p(x_{2j}, \varepsilon_j) \mid I_{1j} = 1, I_{2j} = 1]$.

By potentially re-indexing the items, it is without loss of generality to assume that if the agent insures only one item, she insures item number 1. The key identification result from this model is given in the proposition below:

Proposition 2. *Optimal effort choices imply that*

$$\mathbb{E}[p(x_{1j}, \varepsilon_j) \mid I_{1j} = 1, I_{2j} = 0] - \mathbb{E}[p(x_{2j}, \varepsilon_j) \mid I_{1j} = 1, I_{2j} = 0] > 0. \tag{1.2.2}$$

That is, if it is known that a person has insured one of her items but not the other, then one would expect a higher accident probability for the insured item. The result is highly intuitive: ceteris paribus, an agent should be taking better care of her uninsured items. It is easy to also see that given the assumptions on identical items, if I shut down the choice of effort and consider a model with only heterogeneity in risk type, tastes and wealth, then the accident probabilities for each insurable object owned by each person will be the same and hence the inequality above does not hold. Consequently, the moment condition in Proposition 2 can be used to identify moral hazard separately from adverse selection.

The key assumptions are that at least some agents own multiple *nearly identical* objects and purchase varying degrees of coverage for them. When this is satisfied the test accounts for a large variety of possible sources of heterogeneity ranging from differences in wealth through differences in risk aversion all the way to differences in risk types.

---

and uninsured items as moral hazard, because when both items are uninsured the items would have the same accident probability.

As argued by Chiappori and Salanié (2000), testing whether the difference

$$E[p(x_{ij}, \varepsilon_j) \mid I_{ij} = 1] - \mathbb{E}[p(x_{ik}, \varepsilon_k) \mid I_{ik} = 0] \tag{1.2.3}$$

is positive yields a test of the combined effect of adverse selection and moral hazard. Insured items are more risky, because more risky agents buy more insurance and insurance reduces incentives for precautionary effort. It is tempting to argue that one gets a measure for adverse selection by taking the difference between the combined effect of moral hazard and adverse selection from 1.2.3 and the pure effect of moral hazard from 1.2.2. However, this intuition is not strictly correct. The problem is that it may be optimal for agents who insure one of their items to shift effort from the insured item to the uninsured item. If this "substitution" effect is strong enough it is possible to generate situations where the agents who buy insurance have a higher risk type but their uninsured item ends up being less risky than the items of agents who do not buy insurance. If this happens, it is possible that the combined effect of moral hazard and adverse selection from 1.2.3 is in fact lower than the effect of moral hazard from 1.2.2.

The idea is illustrated by the following numerical example. Consider agents for whom $u(x) = x^{\frac{9}{10}}$, $E(x, y) = (x + y)^2$ and $p(x, \varepsilon) = \frac{1}{1+x} + \varepsilon$. Let $\omega_j = 15$ for all agents and let $L = 7$, $P = 1.2$ and $C = 1$. In Figure 1.1 I plot the agents' expected utility given optimal effort choices as a function of their risk type ε and depending on their choice of insurance. One line shows the value of the maximization problem, if they insure one unit and the other shows what they get, if they leave both items uninsured. The dashed lines show two risk types of whom the left one never wants to insure her items while the one on the right chooses to insure one item.[8]

In Figure 1.2 I plot the optimal accident probabilities for items under different insurance regimes as a function of the agent's risk type. The highest line shows the

[8]It can be shown that neither of the types wants to insure both items.

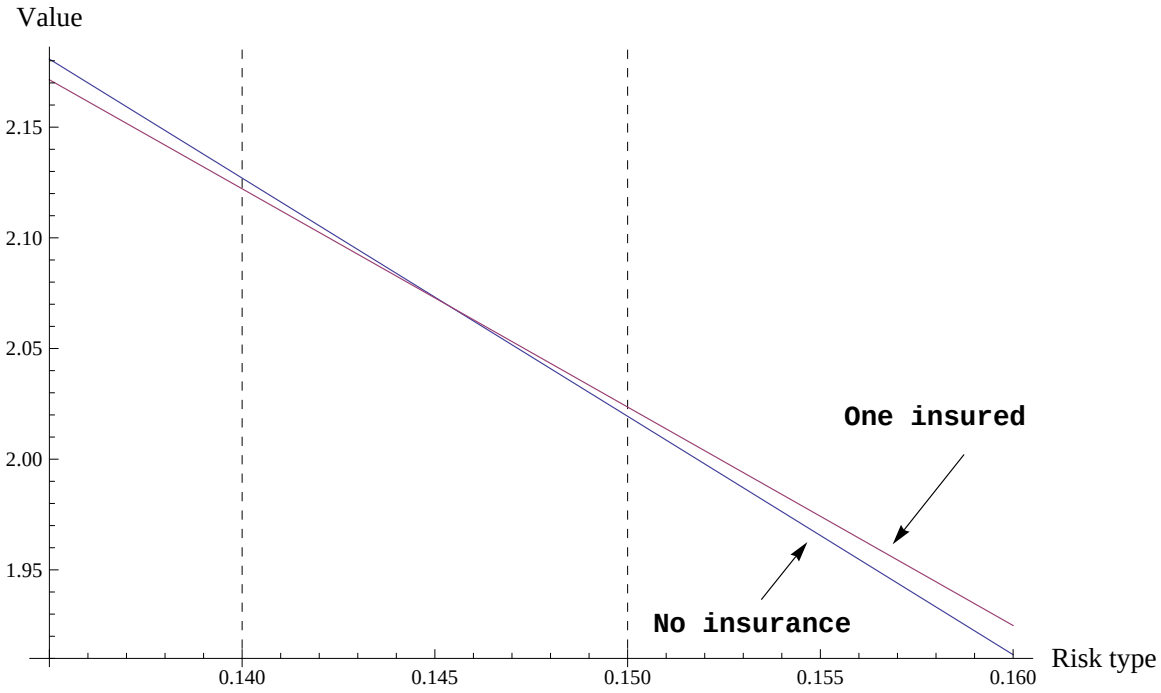


Figure 1.1: Value of the maximization problem as a function of risk type

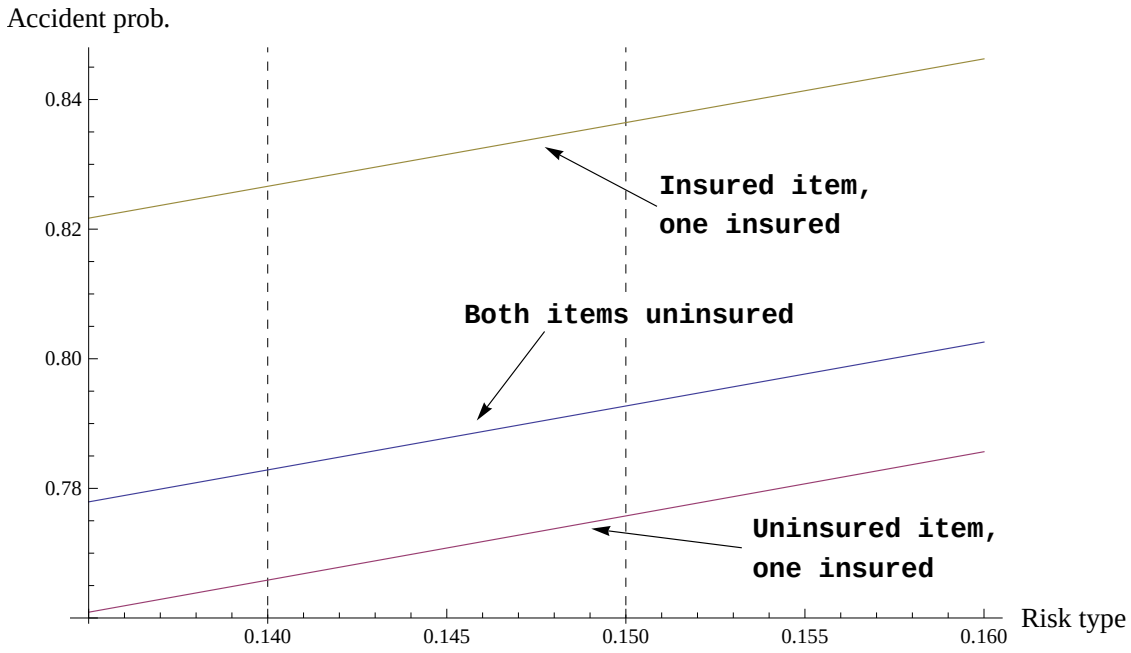


Figure 1.2: Accident probabilities as a function of risk type

accident probability of an insured item when the other unit is uninsured. The middle line shows the accident probabilities of the items of an agent who left both uninsured

and finally, the lowest line depicts the accident probability of the uninsured unit of a person who insured one of her items. As can be seen from the picture, with this parametric specification the agent always chooses lower risk for the uninsured item when she owns one insured item than what she would choose if both items were uninsured. In other words, she "transfers" risk from the uninsured item to the insured one.

Consider then a scenario where most of the population is of the lower risk type with $\varepsilon = 0.14$ (the first dashed line) and a very small fraction has risk type $\varepsilon = 0.15$. It is easy to check that neither of these types wants to insure both items. Now, the first picture shows that the lower risk type never insures any of her items while the higher risk type prefers to insure one unit. If I estimate the ratio of accident probabilities between insured and uninsured items owned by the same person, I will be estimating the ratio between the highest and the lowest line at $\varepsilon = 0.15$, because these are the only types in the population who own both an insured and an uninsured item. The ratio of these accident probabilities is about $0.836/0.776 \approx 1.08$. However, when estimating 1.2.3 most of the uninsured observations come from the population for whom $\varepsilon = 0.14$ as they were an overwhelming majority in the population. Consequently, the compound estimate will end up comparing the middle line at $\varepsilon = 0.14$ to the highest line at $\varepsilon = 1.50$. This results in a hazard ratio of about $0.836/0.783 \approx 1.07 < 1.08$. Consequently, a too hasty application of this technique would lead me to claim that there is advantageous selection in the population even though in reality the converse is true.

The problem with looking at the residual positive correlation in the cross section that is not explained by my fixed effects method for moral hazard is, that it always contains some element of comparing apples with oranges: To measure adverse selection, I would like to compare the risks of either insured or uninsured items of two risk types conditional on them facing the same incentives for precautionary effort. How-

ever, when one of the types is insuring one item and the other is insuring none, not only the insured items, but also the uninsured items have different returns to protective effort. As I never observe the two risk types making the same insurance choices, the incentives for precautionary effort will also vary between both items owned by the two types. Hence, without fixing functional forms for the utility function and the way effort is transformed into accident risk, it is hard to clean the estimate of adverse selection of the effect that insurance has on the effort choice for the uninsured item when the person's other item is insured.

1.3 Data

I will start this section by introducing my data source and describing its insurance market. After that I will present some descriptive statistics of my data.

1.3.1 EVE Online and Its Spaceship Insurance

My data source, EVE Online, is a science fiction multiplayer game played by over 300,000 players worldwide in a single virtual world. The game is run by an Iceland based game company called CCP and is part of the massively multiplayer online game industry that was estimated to generate about $10 billion in revenue in 2014, which was about 15% of the total video game market value. During the data collection period, to play EVE a player had to pay a subscription fee between $14 and $20 per month depending on the method of payment and how many months were purchased at once.

A player in EVE is a spaceship pilot in a complex society consisting of all of the game's subscribers. The game itself does not give players any clear goals that would lead to somehow completing the game. Hence, players are left to their own devices at deciding what activities they find worth doing. For many players, important measures

of how well they are doing in the game are such things as economic wealth, their status in the game's social community, the fraction of the virtual universe that their group of players, called a corporation, controls or how well they are doing in spaceship combat against other players.

All of these activities require some economic resources and hence some involvement in the game's player-run economy. The economy consists of over 5,000 spatially distinct markets where players trade thousands of intermediate and final products manufactured by players who are often highly specialized in certain manufacturing activities. Common tasks in the economy involve mining asteroids for raw materials, manufacturing goods for markets, freighting products from where they are cheap to where they are expensive, running errands for richer players or non-player characters, or highway robbery, i.e. threatening a mining vessel or a freighter with a war-ship in the hope that the owner will rather pay some protection money than see hours of his or her work blown to bits. The game has even seen multiple player-run banks as well as some bank runs that resembled their real-life counterparts.

Entering many of the areas in the game runs a high risk of getting one's ship attacked by space pirates or members of a hostile corporation who have a stake on the area and its economic wealth. Because large fractions of players' wealth consists often of the ships they own, the threat of violence creates a demand for ship insurance. As much of the time in the game goes into gathering means to purchase space-ships and to tuning them with parts and weapons sold in the market, losing the ship in combat may mean losing many days of work. The value of a given ship can even be measured in real currency using the unofficial floating exchange rate between the in-game currency ISK and euros. At the time of the data collection a player could buy 6 months of play time from the game company for 105 euros. The purchased play time could then be turned into six single month play-time extensions that could be sold in the in-game markets for in-game currency. At the time of data collection, a play-time

extension cost around 480,000,000 ISK. Hence, a euro bought about 27,400,000 ISK. A Drake, one of the most popular combat ships that will appear multiple times later in this article, cost approximately 55,000,000 ISK. Using the implied exchange rate above, this is equivalent to approximately 2 euros of real money. The most expensive ships, Titans, cost around 4,950,000,000 ISK or around 180 euros.

EVE's spaceship insurance is provided by the game company and it offers seven levels of insurance for every ship type. The contracts differ by the percentage of the full payout covered by the insurance ranging from the free 40% insurance to the highest 100% insurance. To determine the full payout, the developers of the game periodically calculate the market prices of the raw materials that go into producing each ship. This cost is usually considerably less than the market value of the ship, because the formula does not take into account the profits of the ship builder and all the market intermediaries nor the materials wasted in the production process. The different insurance contracts are then based solely on this cost and what is called the ship type group. Ship type groups are categories of ships that broadly share the same purpose (mining barge, freighter, different combat roles, etc.) and size. Every ship type group has a different percentage of the material costs covered by the full payout. For instance, for the Battleship group all of the material costs are covered while for the Heavy Interdictors only 70% of the material costs get covered. Since the payout of the insurance is based on material costs and not on the market price, even the 100% coverage contract for Battleships entails a de facto deductible. Also the premiums of different insurance contracts are based on the ship type group level full payout. Adding 10 percentage points more coverage increases the premium by 5 percentage points of the full payout. Thus starting from the free 40% insurance, the contract that yields the maximum payout costs 30% of the full payout. All of the contracts cover the ship for 12 weeks. If a ship is destroyed before 12 weeks have passed, the insured agent is immediately awarded the predetermined payout and the

contract ends. Hence the structure of the contracts corresponds to the structure in my model.

Although the goal of the insurance system is primarily to make the game more entertaining and not to maximize profits, it still adheres to both of the supply side predictions described by Chiappori and Salanié (2000, p. 58) which should prevail in an insurance market with adverse selection: "(1) observationally equivalent agents are faced with menus of contracts, among which they are free to choose and (2) within the menu, contracts with more comprehensive coverage are sold at a higher premium". Furthermore, Chiappori and Salanié point out that complex non-linear pricing of insurance contracts may cause a spurious correlation between coverage and accident probabilities, if the characteristics on which the pricing is based are not properly controlled for. In EVE Online, the pricing categories are extremely simple (ship types) and hence easily controlled in an econometric model. In this sense the setup has the transparency, clarity and simplicity of a laboratory experiment. More complex contracts may also be hard for the prospective insurees to comprehend and, as was demonstrated by Fang, Keane and Silverman (2008), cognitive ability may become a more important factor than risk type in deciding how much insurance people will buy. This is unlikely to happen with EVE's very simple contracts.

Another benefit from using log data from an online game is that it is high frequency data that is extremely detailed, extremely reliable and the potential number of observations is very large.[9] Furthermore, because the payout is automatic and no player action is required for reporting accidents, I observe all accidents and not only the ones that get reported. As is commonly pointed out, people usually do not report claims that are less than or barely exceed their deductible. Thus tests based on claims as a proxy for risk and the deductible as a measure of coverage may often

[9]In particular, the only things not collected automatically by the servers are the player demographics, age and gender, which the players fill in when they subscribe to the game. These are not visible to other players and hence there should be little incentives to misreport.

yield spurious correlation between coverage and the probability of an accident, as people with higher deductibles simply report fewer claims even though their number of accidents may be higher (see e.g. Cohen and Siegelman 2010). Since all accidents are automatically reported, I am able to circumvent this problem. I also observe all of the objects that could have been insured, not only those that end up in an insurance company's books. Consequently, my data is able to capture the selection into the insurance pool, not only selection effects between different levels of coverage offered by a given company. Finally, as I noted in the introduction, compared to experimental studies, my empirical results are considerably less susceptible to psychological reactivity effects, such as the Hawthorne effect, as the players do not know that they are being studied.

Some people have voiced a concern over the generalizability of even qualitative empirical results from a computer game. Especially, people have worried about whether the players' actions are driven by the same incentives as in the real world, because the players in a game might be just "playing" and hence might act irrationally even when risking the loss of property that has real monetary value. This may be a potential explanation for when empirical results show players acting irrationally or not adhering to the proposed theory. However, whenever players appear to be acting rationally, the results should be generalizable to any populations acting at least as rationally as the player population when faced with similar incentives.

The dog-eat-dog environment of EVE strongly rewards rational actions and harshly punishes blatant irrationality, as those who are the most able to utilize all available information are much more likely to gather most wealth, build the best space-ships and stations and hence get the upper hand in the game's fierce competition. For instance, the traders in the EVE markets often use special automated Excel spreadsheets with macros that take in large amounts of market data and use it to suggest optimal trade routes and goods that yield the largest profits. There are

also multiple sites on the internet where speculators can follow the development of market prices of different commodities in all spatial markets.[10] Hence, at least on the surface, it seems that EVE players respond to the incentives of the virtual economy very much like people in the real economy.

1.3.2 Some Descriptive Statistics

My data consists of a random sample of 60,022 EVE players. For these players I have the basic, self-reported real life demographics, gender and age. One player may own multiple characters or avatars which are their embodiments in the game. I have a comprehensive set of in-game demographic variables for each of the 99,730 characters owned by my 60,022 players, including attributes such as the character gender, monetary wealth, number of ships owned, the market value of these ships, the distribution of skill points[11] and the total number of times the character lost a space ship. Last, I have a record of all the ships that were owned by my players and in working condition some time between May 22, 2011 and July 18, 2011 (meaning 1,873,343 ships). However, due to server capacity constraints the game company does not save event logs for the cheapest of ships. Consequently, my sample contains full records that detail how, where and when a ship got destroyed for 759,865 ships excluding all of the ships meant primarily for the new players. The total monetary value of the observed ships is about 75 times as large as the total monetary value of the excluded ships.[12] Hence, in terms of monetary value, my sample contains arguably the more interesting half of the stock of ships. Furthermore, the cheaper ships are mostly used by inexperienced, new players and hence excluding them is likely to yield a sample of players that better understand both the insurance system and the risks

[10] Good examples are `http://www.eve-markets.net` and `http://eve-central.com`.

[11] Players use skill points to develop skills which make them more proficient in different areas of the game.

[12] The value of the stock of cheap, excluded ships in my sample, using the implied exchange rate above, was about €32,000. The total value of the included ships was about €2,400,000.

that come with their actions. I observe these ships for nearly three months or until they get destroyed starting from May 22, 2011 with the last observations being on July 18, 2011. My data on the insurance spells for these ships goes even further. The insurance contracts are sold in two month spells and the first contract start date in my data is February 25, 2006.

Table 1.1: Descriptive statistics about players and their ship ownership

	Variable	Mean	Median	Std. Dev.
Player characteristics	Male	0.96	1	0.20
	Age	31.52	30	9.31
	Hrs. played	1737	880	2294
	Hrs. July 1-18	30.60	10.46	45.63
	Characters	1.66	2	0.45
Ship ownership	Active ships	31.21	20	40.27
	Insured ships	1.54	0	3.16

Number of players=60,022

Tables 1.1 and 1.2 report some descriptive statistics of the players in the sample. Table 1.1 contains normal player level statistics. However, as players may have multiple characters, some statistics such as the fraction of male characters in the population are more naturally reported at the character level. These variables are collected in Table 1.2.

Some of the variables may require a small explanation: The Hrs. played and the Hrs. played in July 1-18 variables both report the total number of hours a player stayed logged on to the game in total and between July 1-18, 2011, respectively.

The Active ships and Insured ships show the mean, median and standard deviation for the number of ships owned by a player or a character. I call a ship active, if it

Table 1.2: Descriptive statistics about characters and their ships

	Variable	Mean	Median	Std. Dev.
Character attributes	Male	0.68	1	0.47
	Money (M ISK)	490	13	3300
	Life inc. (bn ISK)	6.8	0.41	60
	Skill points (M)	17	5.0	24
	Security status	1.3	0.29	2.3
	Low sec. visit	0.09	0.35	0.58
Ship variables	Owned ships	36.79	11	390.37
	Ships (bn ISK)	1.6	0.11	5.1
	Active ships	19	9	33
	Ships lost	1.2	0	10
	Ships destr.	0.82	0	7.4

Number of characters=90,730

is not in a storage container. For any market-related purposes such as freighting or selling the ship, it needs to be packaged into a container before it can be put on sale or shipped on a freighter. Packaged ships cannot be used before unpacking. Therefore, an unpacked ship is a good indicator of the ship being currently in use and not just stored to be sold in the market. In Table 1.2 Ships lost and Others' ships destroyed show the same statistics for how many ships a player lost and how many other players' ships she destroyed in combat during the data collection period.

Skill points are something that players accumulate automatically over time in the game and what the players can then invest to develop the efficiency of their characters. Skill points can be used for a wide variety of purposes ranging from reducing waste in manufacturing and the costs incurred from selling objects to higher combat efficiency. Skill points accumulate automatically whenever the player has paid

her monthly subscription fee. Hence, the total number of skill points can be viewed as the age of the character. Furthermore, it reflects the level of proficiency of the character.

Last, there are the Security status and Lowest sec. visit variables. The first one is lower the more crimes the player has committed in the space policed by the in-game police, Concord. The lower this rating, the more there are star systems where the galactic police will automatically attack the player. Furthermore, players with security status -5 or less are considered outlaws and can be attacked by other players without the Concord intervening or having the attackers security status lowered. The Lowest sec. system visited is the security rating of the star system with the lowest security rating visited by the player during the data collection period. The lower this number, the more lawless and perilous the system is and the easier it is to get away with crime in that system. Hence, visits to systems with a low security rating are indicative of high risk taking.

A striking feature of these tables is the sex distribution: only 4% of the whole player population are female. This is possibly even more extreme than in many other online role-playing games. As my aim is to showcase my methods and test the existence of moral hazard, not so much to give accurate point estimates with external validity, the skewed gender distribution is not a big problem. Another noteworthy fact is that only 68% of the characters are male. This implies that a large fraction of the male players chose to play with a female character. At the time of the data collection, the character gender in the game affected only the small picture of the avatar visible to other players.

The mean and median player ages are slightly lower than the ones found in many real world countries.[13] However, they are higher than what one might expect from a

[13]For instance, the median age in the USA in 2010 was 35.3 years (Howden and Meyer 2011)

computer game. This reflects the fact that EVE Online is a relatively complex game directed to more mature audiences.

Many of the remaining variables in both of the tables show a common feature where the mean is much higher than the median. As most of these variables are measures of wealth, success and time invested in the game, this is compatible with the commonly shared preconception that EVE's laissez-faire, dog-eat-dog world is a very hard place for newcomers and success requires a lot of playing. At the same time, a small elite controls large quantities of capital with which they are able to generate large income flows. These players may daily sell dozens of ships that might cost over a hundred times a newcomer's monthly in-game income. In general, taking part in the more economically lucrative activities in EVE demands considerable amounts of wealth as these activities often happen in the lawless 0.0 space where operations usually require well equipped, expensive and high quality ships. Furthermore, the risk of losing one's ship in combat is relatively high in 0.0 and consequently players operating there must have enough liquid property to be able to overcome occasional losses.[14] This picture is reflected by the fact that a median player had not lost any ships or destroyed any opposing players' ships, and on average both of these numbers were close to one. However, a player one standard deviation away from the mean lost over 10 ships and destroyed over 7 other players' ships.

The descriptive statistics of both of the security status variables reinforce this interpretation. Security status is accumulated by destroying pirates in low security space and running errands for computerized agents which often require destroying pirates or traveling to low security space. Wealthier players in their more expensive ships have a much better chance to successfully pursue both of these activities. Similarly, due to their superior equipment wealthier players have a much higher chance to

[14] Although there are some player-run banks the financial markets in EVE are arguably far from perfect. Hence, most players are heavily credit constrained and cannot borrow the assets required to operate in 0.0.

survive trips to low security space, and replacing any ships lost there does not imply personal bankruptcy. Thus it is not surprising to find that the mean security status is higher than the median and that the median character stayed in a much safer space than the mean.

The insurance status of the ships in my data is fairly polarized between the two extremes. This can be seen in Table 1.3. About 91% of the ships are insured only with the free 0.4 insurance level and out of the people who paid for their insurance (categories 0.5-1.0) over 85% have the highest insurance status. Consequently, I will reclassify all ships in categories 0.5-1.0 as insured and any ship in the free 0.4 category as uninsured. All results presented in this paper are robust to different potential classifications.

Table 1.3: Distribution of assembled ships by insurance level

Insurance level	Frequency	Percent
0.4	693,459	91.26
0.5	4,128	0.54
0.6	1,375	0.18
0.7	1,256	0.17
0.8	1,552	0.20
0.9	1,632	0.21
1.0	56,463	7.43
Total	759,865	100

Since two out of the three of my identification strategies use variation in risk between ships of the same type, owned by the same person but have different amounts of insurance coverage, it is worthwhile to document that such individuals exist. In total, there are 11,272 players who own at least 2 ships of the same type out of which

at least one is uninsured and at least one is insured. Table 1.4 in the appendix breaks this down by ship type over the ten most popular ship types.

1.4 Econometric Model and Results

My identification strategy measures the average difference in risk between differently insured items owned by the same agent. As was formally argued in Section 1.2, controlling for the agent's identity controls for agent-specific adverse selection. Any remaining positive correlation between the insurance coverage and accident risk between items owned by the same agent is caused by different choices of precautionary effort. I will measure this remaining correlation using a highly flexible semiparametric econometric model.

I assume that the probability that ship j of type k owned by person i survives without an accident for more than t units of time is given by

$$S_{ji}(t;k) := \exp\left[-\int_0^t \lambda_{ik}(s)\exp(\beta Z_{ji}(s))\,\mathrm{d}s\right],$$

where $Z_{ji}(s)$ equals one if ship j is insured at time s, and zero otherwise. The function $\lambda_{ik}(\cdot)$ is a person-ship-type specific baseline hazard rate which can vary arbitrarily over time. The parameter β measures the effect that insurance status has on accident risk because of the agent's effort choices. This is a standard Cox proportional hazards model (Cox 1972 and 1975) with time varying covariates stratified at the person-ship type level. The parameter β can be consistently estimated using partial maximum likelihood without estimating the potentially infinite dimensional λ_{ik} (see Ridder and Tunali, 1999 or Kalbfleisch and Prentice, 2011).

Proposition 2 implies that if there is moral hazard, β will be positive. If both insured and uninsured ships of the same type, owned by the same person have the same hazard rate, it will still be captured by $\lambda_{ik}(t)$. This is true even if this common

hazard rate varies over time. Consequently, if there is no moral hazard, β will be zero.

My identification strategy requires that I compare multiple *simultaneous* effort choices made by the same agent. In other words, I need to compare items owned by the same person with different insurance statuses at the same *calendar* time. Hence the origin of the time axis in the econometric model should not vary between the items owned by the same person. My empirical models will therefore use the same calendar date as their origin.

The model allows me to control for multiple sources of possible dynamic variation in the agent's risk. First, it allows for times when the person was logged out of the game and hence had a zero baseline hazard rate. The same holds true for episodes of high intensity gaming such as weekends or private vacation days. Figure 1.3 presents the frequency of ship destruction events as a function of time. One bar of the histogram represents a day during my observation period. As can be seen from the figure, there is a strong weekly cycle in destruction events. Furthermore, there are three days with abnormally few destruction events that are likely due to server down time. My model accommodates not only these features but any heterogeneity in gaming habits. It also allows the player's baseline hazard rate to change when an agent becomes more experienced with a specific ship type or just generally learns more about the environment.[15] All of these sources of variation have clear analogues in most real-life insurance settings. The flexibility of the model is therefore valuable also outside the game.

The model does not account for variation in risk over the life cycle of an item. The most prominent reason why this type of variation might matter in most settings is depreciation. If old or depreciated items are both riskier and insured more often, the model above would confound moral hazard with item-specific adverse selection. Many

[15] I discuss the effect of experience on moral hazard in Section 1.5.

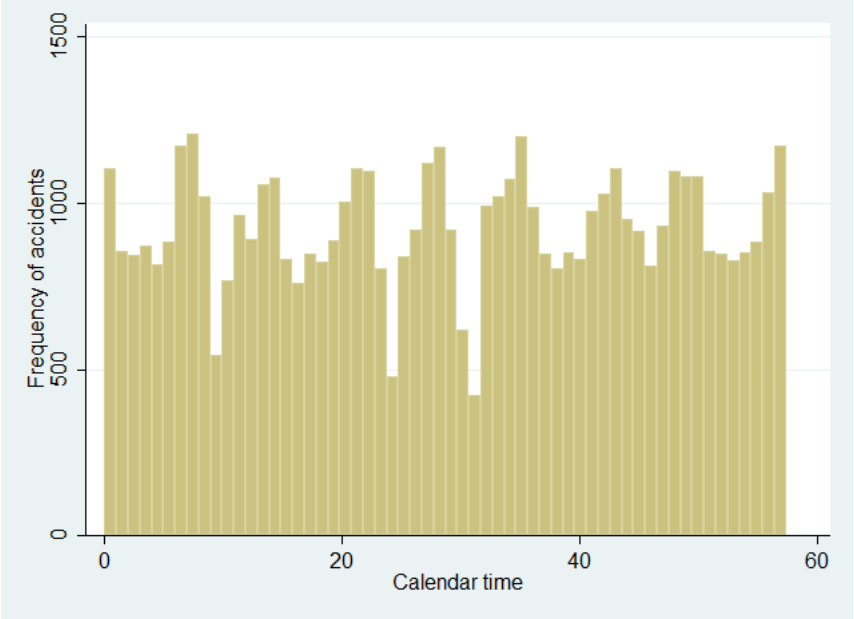


Figure 1.3: The number of destruction events over the observation period.

real world insurance contracts account for depreciation. In these settings stratifying the model at the contract and agent level will alleviate the problem.

In EVE, a ship can get partially damaged in combat. The more damaged the ship, the easier it will get destroyed. Damage can be repaired for a low cost but the insurance does not cover this cost. It is also likely that older ships are on average more damaged than newer ships. This depreciation may generate a potential source of a bias, if damaged old ships are insured more than new ships. However, in EVE depreciation is unlikely to be a big factor, since players can fit their ships with very cheap repair units that repair any damage for free.

In some settings, depreciation can be approximated by the item's age. As long as the effect of the age of an item on the hazard rate is uniform across individuals I can control for it by appending the model with a set of age fixed effects. Formally, the survival probability is then given by

$$S_{ji}(t;k) := \exp\left[-\int_0^t \lambda_{ik}(s)\exp(\beta Z_{ji}(s) + \sum_{h=1}^{H} \delta_{A_h}\mathbb{1}\{a_j(s) \in A_h\})\,\mathrm{d}s\right],$$

where $a_j(s)$ is the age of ship j at time s, the set $A_h \in \{A_1, \ldots, A_H\}$ is an arbitrary partition of possible ship ages and δ_{A_h} is the effect that belonging to age group A_h has on the ship's hazard rate.

Table 1.4 presents estimation results from different specifications of the model. Rather than reporting the estimate for β which is harder to interpret, it gives the estimates for $\exp(\beta)$. This corresponds to the ratio of the hazard rates between an insured and an uninsured observation from the same strata with otherwise identical covariates. Formally, for example in the model with age dummies, for any t,

$$\exp(\beta) = \frac{\lambda_{ik}(t)\exp(\beta + \sum_{h=1}^{H} \delta_{A_h} \mathbb{1}\{a_j(s) \in A_h\})}{\lambda_{ik}(t)\exp(\sum_{h=1}^{H} \delta_{A_h} \mathbb{1}\{a_j(s) \in A_h\})}.$$

Consequently, $\exp(\beta)$ measures the relative change in the hazard rate due to the item having insurance. Below its estimates I also report the 95% confidence interval for this hazard ratio. In the model with ship age controls, I add dummies for the first 7 days of a given ship, as well as a dummy for ships that are older than 180 days.[16]

The non-stratified estimates measure the compound effect of moral hazard and adverse selection given in (1.2.3). These estimates allow for seasonality in the common baseline hazard rate shared by the population. The first row of the table shows that this combined effect implies that insured ships are on average eight times as likely to get destroyed each day compared to their uninsured counterparts. Estimates from different markets are not highly comparable due to differences in the fraction of the risk carried by the agent under different contracts, opportunity cost of effort and losses from poor outcomes. Nevertheless, eightfold hazard rate of insured items compared to uninsured items is much higher than the estimates found in the earlier literature. For example, Dobbie and Skiba (2013) find that agents who choose a two-week payday loan that is \$50 larger have a 28-44% higher default risk. For car

[16]This is approximately the median age of a ship in my sample both at the end of the data collection period and at the time of destruction for destroyed ships. I experimented with other partitions of the age distribution without considerable effect on the results.

Table 1.4: Estimates of the insurance parameter from stratified Cox regressions

	Strata	Hazard ratio No age dummies	Hazard ratio Age dummies
Pooled Sample	Ship type	8.02	4.38
		(7.68, 8.37)	(4.23, 4.54)
	Ship type,	5.83	6.16
	User ID	(5.32, 6.40)	(5.64, 6.73)
	Ship type,	5.72	6.09
	Character ID	(5.22, 6.27)	(5.58, 6.66)
Drakes	None	6.56	4.04
		(6.02, 7.14)	(3.73, 4.37)
	User ID	6.44	6.19
		(5.14, 8.08)	(5.06, 7.56)
	Character ID	6.38	6.13
		(5.09, 7.80)	(5.02, 7.49)
Badger IIs	None	7.04	3.39
		(5.15, 9.63)	(2.51, 4.58)
	User ID	4.28	9.25
		(1.42, 12.89)	(1.58, 54.14)
	Character ID	4.28	10.00
		(1.42, 12.89)	(1.55, 64.43)

Note: 95% confidence intervals in parenthesis. Implicit standard errors clustered at the user level.

insurance, Cohen (2005) shows that having a low deductible is associated with about a 4% higher probability of having at least one accident during the year in which the policyholder took out the insurance. The high combined effect of moral hazard

and adverse selection in EVE's spaceship insurance is not that surprising given the nature of the game and its insurance contracts. One of the design goals of the game's insurance contracts is to make highly risky modes of gameplay, such as spaceship combat and exploring the high risk space, more appealing to the players.

When baseline hazard rates are allowed to vary from player to player, I get the effect of moral hazard on hazard rates. These estimates are given on the second row for the pooled sample. An insured ship is almost 6 times more likely to get destroyed each day compared to an uninsured, identical ship owned by the same individual. The result indicates a high level of moral hazard in EVE's insurance market. The third row of the table allows the baseline hazard rate to vary also between different characters owned by the same user. In principle, it could be possible that the player played differently with different characters. Some of this effect could be incorrectly captured in my agent-level estimate of moral hazard. However, the effect of moral hazard from this regression is commensurate with the one that stratifies only at the user level. This indicates that players' risk taking behavior does not vary considerably between characters.

The level of moral hazard is likely to vary by the purpose of the insurable item. When an agent is driving a sports car she may behave differently than when driving a truck. The rows 4-9 of the table run the same model but restrict the pool of observations only to Drakes or Badger Mark IIs. Drake is the most popular ship in my sample and it is mostly used for fighting other players or hostile ships controlled by the game server. Badger Mark II, on the other hand, is the most popular mainly non-combat ship.[17] Looking at these separate estimates I can get some indication of how the purpose of the ship affects the effect of moral hazard. From rows 5 and 6 one can see that the effect of moral hazard on the risk of Drakes is somewhat higher than the average effect in the population of all ships. However, the Badger's

[17]Badger Mark II is one of the best freight carriers in the game.

point estimate is somewhat smaller (rows 8-9). This seems natural, since the peaceful purpose of the Badger implies that the benefits from taking risks with it are likely to be smaller. However, the standard errors for these estimates are so high that the population average is inside the 95% confidence interval for both ships.

The last column shows what happens to the estimates if I add the ship age dummies into the model. In general, if there is persistent heterogeneity in the baseline hazard rates of the items owned by the same person, then an average item with a lower risk type has survived longer and is consequently older than an average item of a higher risk type. Therefore, the item's age would be negatively correlated with its risk type. If there is item-specific adverse selection, adding item age as a control should therefore reduce the estimate of the effect of insurance on hazard rates in the model stratified at the agent level.[18] In this sense, the last column serves as a test of my identifying assumption that ships of the same type owned by the same person have the same risk type. As can be seen from Table 1.4, adding these dummies has a negligible effect on the estimates stratifying at the user or character level. This finding supports my claim that ships of the same type are fairly identical in EVE. Adding the time dummies for Badgers doubles the estimate on the effect of moral hazard on insured ships' hazard rates. However, the increase in the estimate can be completely due to the increased inaccuracy in the estimates. Furthermore, as argued above, in case of ship-specific adverse selection, controlling for ship age should decrease, not increase, the estimate for the effect of insurance. The reason for the decreased accuracy is the relatively small number of players (704) that own both an insured and an uninsured Badger Mark II.

Comparing the first and the second row of Table 1.4, I find that the hazard ratio between insured and uninsured items is at least 40% higher than what would be implied by only moral hazard in a fully homogeneous population. As was pointed out

[18]If there can be multiple accidents per insurable object, time since the previous accident serves the same purpose.

in the end of Section 1.2 this way of measuring adverse selection is not completely accurate. When I add age dummies to the model stratified only at the ship type level, the hazard ratio falls from 8.02 to 4.38. This is even lower than what is obtained in the model stratified also at the user level. Since adding age dummies controls for at least some of the adverse selection, the relatively low hazard ratio could be indicative of the "substitution effect" outlined in Section 1.2, dominating any remaining adverse selection.

I assess the goodness of fit of the model and especially the appropriateness of the proportional hazards assumption in Section 1.C of the appendix. The model seems to fit the data exceptionally well.

1.5 Experience and Moral Hazard

There are at least three reasons why agents' level of experience can be an important factor in insurance markets. First, more experienced agents may better understand what type of behavior is likely to lead to an accident and choose safer actions. Second, learning one's risk type potentially increases the information asymmetry between the insurance provider and the agent. Hence, more experienced policy holders may be more adversely selected. Finally, better understanding the relationship between available actions and risk helps agents to choose more optimal precautionary effort when faced with different levels of insurance coverage. More experienced agents may hence exhibit more moral hazard if they are better able to reduce their effort costs by choosing risky actions when insured and exercise more relevant precautionary effort when uninsured. Consequently, offering separate contracts for experienced and inexperienced agents can be an important way of increasing efficiency in the insurance market. Furthermore, insights on the channels through which experience affects risk can help target monitoring and screening effort to the right subpopulation.

I will measure player's experience by the number of days she has spent logged into the game. This is analogous to many real world studies that measure experience, for example, by using the time since a driver got her driver's license. My measure is potentially more accurate, because being logged into the game is a strong indicator of actively using a ship. Most activities in EVE involve the use of ships and hence actively playing the game is almost synonymous with using a ship. In comparison, two persons who got their licenses at the same time are likely to have considerably different amounts of driving experience.

To establish the fact that players become better at protecting their property over time, I compared the ship destruction rates of the bottom and top quartiles of players sorted by their time spent in the game. As expected, new players seem to get their spacecraft destroyed much more frequently than the old: 11.6% of the ships owned by the least experienced 25% of the sample got destroyed, while only 5.4% of the ships owned by the most experienced 25% faced the same fate. In other words, the ship of a new player was almost twice as likely to get destroyed than that of a seasoned veteran.

Next, I estimated the model from Section 1.4 separately for these two quartiles. The results are presented in Table 1.5. The first thing to notice is that the combined effect of moral hazard and adverse selection is over 3 times as high for the more experienced players than for the new-comers. Even though the new players are very likely to get their ships destroyed, the experienced players' ships are much more likely to be insured *when* they get destroyed. The direction of the effect is in line with what one would expect if it takes time for the players to learn their risk type or the mapping from effort to accident probabilities. The magnitude of the effect, however, is striking.

The results on the row where the estimates are stratified at the user level show that the effect of *moral hazard* on the ship's risk is almost five times as high for the

Table 1.5: Estimates of the insurance parameter from stratified Cox regressions for the most and least experienced players

Strata	Included observations	Hazard ratio Bottom 25%	Hazard ratio Top 25%
Ship type	All	3.53	10.82
		(3.30, 3.78)	(10.04, 11.66)
Ship type,	All	2.19	9.94
User ID		(1.76, 2.73)	(8.40, 11.76)
None	Drakes	3.00	11.09
		(2.34, 3.83)	(9.32, 13.20)
User ID	Drakes	4.08	12.26
		(1.68, 9.96)	(7.54, 19.95)

Note: 95% confidence intervals in parenthesis. Implicit standard errors clustered at the user level.

veterans compared to the new players. The impact of experience on moral hazard is slightly smaller when I look only at Drakes. However, due to the relatively small number of Drakes in the sample, the estimates become inaccurate when stratified at the user level. For the players in EVE, learning and experience seem to play a key role in figuring out the optimal amount of precautionary effort and how insurance factors into it. If a similar finding is true in other insurance contexts, having higher deductibles or coinsurance rates for young people might be neither profit maximizing nor socially optimal. If information asymmetries are less of a concern with young agents, a better practice would have them paying higher premiums but face lower coinsurance rates than what is offered to more experienced agents.

1.6 Conclusions and Future Work

This chapter has discussed how multiple contracts taken by the same agent can be used to separate moral hazard from adverse selection in an intuitive way by "factoring out" the agent's risk type. An application of the method to a laboratory-like setup shows how moral hazard can induce very high increases in risk-taking, especially among experienced agents.

In the second chapter of this dissertation I discuss how the time dimension can often be used to measure dynamic changes in risk-taking and moral hazard. Both dimensions of panel data yield useful information on how agents respond to different incentive structures. This information is likely to be valuable in calibrating and designing new contracts. It allows us to better understand the level of sophistication of the agents, what the relevant incentive constraints are and how to reach less inefficient contracts.

In future, putting more structure on the model and applying it to a dataset of property insurance or a pool of labor contracts could potentially allow measuring the size of the dead-weight loss, as well as calculating counterfactuals for how much efficiency can be gained with more optimal contracts. As was discussed earlier, the method seems especially suitable for measuring moral hazard and its effect on fail rates in project work and between costs-plus and fixed fee contracts taken by the same agent.

Another interesting avenue of future research would be to assume more on the contracting setting and agent's utilities to see how much the symmetry assumption can be relaxed or how its violations could be corrected for in the estimation procedure. This would be especially useful for many types insurance contracts, such as car insurance, where agents are likely to own multiple differently covered cars but where those cars may not satisfy the symmetry assumption.

From a purely theoretical point of view, this chapter discusses optimal responses of agents when facing multiple discrete insurance decisions. I discuss how insurance purchased for one item affects the riskiness of the uninsured item when risks are independent. Optimal design of contracts for agents with multiple insurable items could potentially yield welfare improving insights on how the agents' riskiness as a whole changes when they hold multiple insurance contracts with different levels of coverage. Furthermore, it could be beneficial to understand what kind of a menu an oligopolist insurer would offer if it can insure only one of the agent's items and the other one is likely to be insured by its competitor. This question seems especially interesting when the effort chosen for one insurable item affects the riskiness of the other. The prevalence of moral hazard and adverse selection make competition in the insurance sector quite different from ordinary goods oligopolies. Nevertheless, there is very little theoretical work that discusses how competition work agent's effort affects the offered contracts when agents can purchase multiple contracts from multiple providers for multiple insurable items.

1.A Omitted Proofs

The following well-known lemma on decreasing differences of a concave function is useful throughout the proofs:

Lemma 1.A.1. *Let $u\colon \mathbb{R} \to \mathbb{R}$ be differentiable, increasing and strictly concave. If $z > 0$, then*

$$\frac{\mathrm{d}\left(u(x) - u(x-z)\right)}{\mathrm{d}x} < 0$$

for all $x \in \mathbb{R}$.

Proof. By direct calculation,

$$\frac{\mathrm{d}\left(u(x) - u(x-z)\right)}{\mathrm{d}x} = u'(x) - u'(x-z).$$

Since u is strictly concave, u' is decreasing. This in turn implies that $u'(x-z) > u'(x)$ proving the lemma. □

1.A.1 Proofs for Section 1.2

For the proofs in this section, it is useful to define the utility value from any given effort levels and insurance choices as:

$$\begin{aligned} V(x_{1j}, x_{2j}; I_j, \varepsilon_j) \;:=\; & p(x_{1j}, \varepsilon_j)p(x_{2j}, \varepsilon_j)u(\omega_j - 2L + (I_{1j} + I_{2j})(P - C), \varepsilon_j) \\ & + p(x_{1j}, \varepsilon_j)(1 - p(x_{2j}, \varepsilon_j))u(\omega_j - L - (I_{1j} + I_{2j})C + I_{1j}P, \varepsilon_j) \\ & + p(x_{2j}, \varepsilon_j)(1 - p(x_{1j}, \varepsilon_j))u(\omega_j - L - (I_{1j} + I_{2j})C + I_{2j}P, \varepsilon_j) \\ & + (1 - p(x_{1j}, \varepsilon_j))(1 - p(x_{2j}, \varepsilon_j))u(\omega_j - (I_{1j} + I_{2j})C, \varepsilon_j) \\ & - E(x_{1j}, x_{2j}, \varepsilon_j), \end{aligned}$$

I will first show that for any I_j there exists levels of effort that satisfy the first-order conditions of the agent's optimization problem:

Lemma 1.A.2. *For any I_j and ε_j there exists (x_{1j}, x_{2j}) such that $\nabla V(x_{1j}, x_{2j}; I_j, \varepsilon_j) = 0$.*

Proof. Fix any I_j and ε_j. For this proof I will drop the dependence on both ε_j and j since the result clearly does not depend on them. The two first-order conditions are now given by:

$$\begin{aligned} & p'(x_1)p(x_2)u(\omega - 2L + (I_1 + I_2)(P - C)) \\ & + p'(x_1)(1 - p(x_2))u(\omega - L - (I_1 + I_2)C + I_1P) \\ & - p'(x_1)p(x_2)u(\omega - L - (I_1 + I_2)C + I_2P) \\ & - p'(x_1)(1 - p(x_2))u(\omega - (I_1 + I_2)C) \\ & - \partial_1 E(x_1, x_2) = 0 \end{aligned} \tag{1.A.1}$$

and

$$
\begin{aligned}
&p'(x_2)p(x_1)u(\omega - 2L + (I_1 + I_2)(P - C)) \\
&-p'(x_2)p(x_1))u(\omega - L - (I_1 + I_2)C + I_1P) \\
&+p'(x_2)(1 - p(x_1))u(\omega - L - (I_1 + I_2)C + I_2P) \\
&-p'(x_2)(1 - p(x_1))u(\omega - (I_1 + I_2)C) \\
&-\partial_2 E(x_1, x_2) = 0. \qquad (1.A.2)
\end{aligned}
$$

Fix any x_2 and write equation (1.A.1) as

$$
p'(x_1)\left(p(x_2)(a - c) + (1 - p(x_2))(b - d)\right) - \partial_1 E(x_1, x_2) = 0 \qquad (1.A.3)
$$

where

$$
\begin{aligned}
a &:= u(\omega - 2L + (I_1 + I_2)(P - C)) \\
b &:= u(\omega - L - (I_1 + I_2)C + I_1P) \\
c &:= u(\omega - L - (I_1 + I_2)C + I_2P) \\
d &:= u(\omega - (I_1 + I_2)C)
\end{aligned}
$$

Now since u is increasing and $L > P$ both $a-c$ and $b-d$ are strictly negative. Furthermore, since p is decreasing and strictly convex, both $\lim_{x\to 0} p'(x)$ and $\lim_{x\to\infty} p'(x)$ exist in $[-\infty, 0]$ and $\lim_{x\to 0} p'(x) < 0$ and $\lim_{x\to\infty} p'(x) > -\infty$. Consequently, $\lim_{x_1\to 0} E(x_1, x_2) = 0$ implies then that

$$
\lim_{x_1\to 0} p'(x_1)\left(p(x_2)(a - c) + (1 - p(x_2))(b - d)\right) - \partial_1 E(x_1, x_2) > 0.
$$

Similarly, $\lim_{x_1\to\infty} \partial_1 E(x_1, x_2) = \infty$ implies that

$$\lim_{x_1\to\infty} p'(x_1)\left(p(x_2)(a-c) + (1-p(x_2))(b-d)\right) - \partial_1 E(x_1, x_2) = -\infty.$$

But then the intermediate value theorem implies that for each x_2 there exists $x_1(x_2)$ such that

$$p'(x_1(x_2))\left(p(x_2)(a-c) + (1-p(x_2))(b-d)\right) - \partial_1 E(x_1(x_2), x_2) = 0.$$

Plugging this back in to (1.A.2), it will be enought to show that

$$p'(x_2)\left(p(x_1(x_2))(a-b) + (1-p(x_1(x_2)))(c-d)\right) - \partial_2 E(x_1(x_2), x_2) = 0$$

has a solution. Just as above, $a-b$ and $c-d$ are both strictly negative. This together with the assumption that $\partial_{12} E(x_1, x_2) > 0$ implies that

$$\begin{aligned} & \lim_{x_2\to\infty} p'(x_2)\left(p(x_1(x_2))(a-b) + (1-p(x_1(x_2)))(c-d)\right) - \partial_2 E(x_1(x_2), x_2) \\ \leq\ & \lim_{x_2\to\infty} p'(x_2)\left(p(x_1(x_2))(a-b) + (1-p(x_1(x_2)))(c-d)\right) - \partial_2 E(0, x_2) = -\infty \end{aligned}$$

Now if I consider $x_2 = 0$, then (1.A.2) becomes

$$p'(0)\left(p(x_1(0))(a-b) + (1-p(x_1(0)))(c-d)\right) - \partial_2 E(x_1(0), 0) > 0.$$

Applying the intermediate value theorem again implies that there exists x_2 such that

$$p'(x_2)\left(p(x_1(x_2))(a-b) + (1-p(x_1(x_2)))(c-d)\right) - \partial_2 E(x_1(x_2), x_2) = 0$$

proving the lemma. □

The next lemma implies that for any choice of insurance, a choice of effort levels where the gradient of V vanishes yields maximal utility for that level of insurance:

Lemma 1.A.3. *V is strictly concave in (x_{1j}, x_{2j}) .*

Proof. Fix any I_j. I will again drop the dependence on j and ε_j in this proof. Furthermore, I will make the notation somewhat clearer by writting:

$$\begin{aligned} W(x_1, x_2; I) &= p(x_1)p(x_2)a + p(x_1)(1-p(x_2))b + p(x_2)(1-p(x_1))c \\ &\quad +(1-p(x_1))(1-p(x_2))d, \end{aligned}$$

where again

$$\begin{aligned} a &:= u(\omega - 2L + (I_1 + I_2)(P - C)) \\ b &:= u(\omega - L - (I_1 + I_2)C + I_1 P)) \\ c &:= u(\omega - L - (I_1 + I_2)C + I_2 P) \\ d &:= u(\omega - (I_1 + I_2)C) \end{aligned}$$

Since E is convex and the sum of two concave functions is concave, it is enough to show that W is concave. Now,

$$\frac{\partial^2}{\partial x_1^2} W(x_1, x_2) = p''(x_1)\left((a-c)\,p(x_2) + (b-d)\,(1-p(x_2))\right) < 0,$$

where the negativity follows, because u is increasing and p is convex. Similarly,

$$\frac{\partial^2}{\partial x_2^2} W(x_1, x_2) = p''(x_2)\left((a-b)\,p(x_1) + (c-d)\,(1-p(x_1))\right) < 0.$$

Thus it is enough to show that the determinant of the Hessian of W, $|H_W|$, is strictly positive. Calculating this Hessian requires the cross-derivative:

$$\frac{\partial^2}{\partial x_1 \partial x_2} W(x_1, x_2) = (a - b - c + d) p'(x_1) p'(x_2).$$

Now a straightforward calculation yields,

$$\begin{aligned} |H_W| &= \frac{\partial^2}{\partial x_1^2} W(x_1, x_2) \frac{\partial^2}{\partial x_2^2} W(x_1, x_2) - \left(\frac{\partial^2}{\partial x_1 \partial x_2} W(x_1, x_2) \right)^2 \\ &= p''(x_1) p''(x_2) \left((p(x_2)(a - b - c + d) + b - d) \right) \left(p(x_1)(a - b - c + d) + c - d \right) \\ &\quad - \left((a - b - c + d) \, p'(x_1) p'(x_2) \right)^2 \end{aligned} \tag{1.A.4}$$

Now, since u is concave, $a - b - c + d$ is easily verified negative. Furthermore, since $b - d < 0$ and $c - d < 0$, it follows that

$$\begin{aligned} & \left((p(x_2)(a - b - c + d) + b - d) \right) \left(p(x_1)(a - b - c + d) + c - d \right) \\ = \; & \left| (p(x_2)(a - b - c + d) + b - d \right| \left| p(x_1)(a - b - c + d) + c - d \right| \\ > \; & p(x_1) p(x_2) (a - b - c + d)^2 \end{aligned}$$

Now, I can plugg this back into (1.A.4) and use the fact that $p'' > 0$ to obtain

$$|H_W| > \left(p''(x_1) p''(x_2) p(x_1) p(x_2) - p'(x_1)^2 p'(x_2)^2 \right) (a - b - c + d)^2 \tag{1.A.5}$$

Now the first term is simply the determinant of the Hessian of the mapping $(x_1, x_2) \mapsto p(x_1)p(x_2)$ which was assumed to be strictly convex. Hence, that term is strictly positive implying that $|H_W|$ must be strictly positive. □

Together Lemmas 1.A.2 and 1.*A*.3 and the fact that I_j is selected from a finite set imply that the agent's maximization problem has a unique solution. Next I will

show that an insured item always has a higher accident risk than an uninsured item owned by the same person.

Lemma 1.A.4. *Take any person j with $(\varepsilon_j, \omega_j)$ such that $I_j = (0,0)$ or $I_j = (1,1)$. Then for this person the optimal choice of effort (x_{1j}^*, x_{2j}^*) satisfies: $p(x_{1j}^*, \varepsilon_j) = p(x_{2j}^*, \varepsilon_j)$.*

Proof. For the proof of this lemma I may again drop the dependence on j and ε_j. The proofs for the two cases are identical and I will only consider the slightly more complex case where $I_j = (1,1)$. The case where $I_j = (0,0)$ follows by dropping both C and P from the proof below. I will abuse the notation slightly and write

$$\begin{aligned} V_2(x_1, x_2) \;=\; & p(x_1)p(x_2)u(\omega - 2L + 2P - 2C) \\ & + p(x_1)(1 - p(x_2))u(\omega - L - 2C + P) \\ & + p(x_2)(1 - p(x_1))u(\omega - L - 2C + P) \\ & + (1 - p(x_1))(1 - p(x_2))u(\omega - 2C) \\ & - E(x_1, x_2), \end{aligned}$$

where the subscript 2 stands for the two insured units. With this notation, it is enough to prove that if (x_1^*, x_2^*) maximizes V_2, then $x_1^* = x_2^*$. The first-order conditions for this maximization problem are given by

$$\begin{aligned} x_1\colon \quad & \partial_1 V_2(x_1^*, x_2^*) = p'(x_1^*)p(x_2^*)A - p'(x_1^*)\left(u(\omega - 2C) - u(\omega - L - 2C + P)\right) \\ & - \partial_1 E(x_1^*, x_2^*) = 0 \qquad (1.A.6) \\ x_2\colon \quad & \partial_2 V_2(x_1^*, x_2^*) = p'(x_2^*)p(x_1^*)A - p'(x_2^*)\left(u(\omega - 2C) - u(\omega - L - 2C + P)\right) \\ & - \partial_2 E(x_1^*, x_2^*) = 0 \qquad (1.A.7) \end{aligned}$$

where

$$A\colon = u(\omega - 2C) - 2u(\omega - L - 2C + P) + u(\omega - 2C - 2L + 2P) < 0,$$

because u is concave.

Notice that $\lim_{x\to 0} \partial_1 V_2(x,x) > 0$, because $\lim_{x\to 0} \partial_1 E(x,x) = 0$, $p'(x) < 0$, $A < 0$, and $u(\omega - 2C) - u(\omega - L - 2C + P > 0)$, by the fact that u is increasing. On the other hand, $\lim_{x\to\infty} \partial_1 V_2(x,x) = -\infty$, because $\lim_{x\to\infty} E(x,x) = \infty$. Consequently, there exists an $\hat{x} \in \mathbb{R}_+$ such that $\partial_1 V_2(\hat{x},\hat{x}) = 0$. Moreover,

$$\partial_2 V_2(\hat{x},\hat{x}) = \partial_1 V_2(\hat{x},\hat{x}) + \partial_1 E(\hat{x},\hat{x}) - \partial_2 E(\hat{x},\hat{x}) = 0,$$

as, by assumption $\partial_1 E(\hat{x},\hat{x}) - \partial_2 E(\hat{x},\hat{x})$. Therefore, $\hat{x}$ satisfies both of the first-order conditions. As, by the previous lemma, V_2 is strictly concave, $(\hat{x},\hat{x})$ must be the unique maximizer of V_2. □

Then the proof of Proposition 1 follows trivially from the law of iterated expectations:

Proof of Proposition 1. Let $k \in \{0,1\}$. The law of iterated expectations together with the previous lemma imply:

$$\begin{aligned}
& \mathbb{E}[p(x_{1j},\varepsilon_j) \mid I_{1j} = k, I_{2j} = k] \\
= \;& \mathbb{E}\left[\mathbb{E}\left[p(x_{1j},\varepsilon_j) \mid I_{1j} = k, I_{2j} = k, \varepsilon_j, \omega_j\right] \mid I_{1j} = k, I_{2j} = k\right] \\
= \;& \mathbb{E}\left[\mathbb{E}\left[p(x_{2j},\varepsilon_j) \mid I_{1j} = k, I_{2j} = k, \varepsilon_j, \omega_j\right] \mid I_{1j} = k, I_{2j} = k\right] \\
= \;& \mathbb{E}[p(x_{2j},\varepsilon_j) \mid I_{1j} = k, I_{2j} = k].
\end{aligned}$$

□

Lemma 1.A.5. *Take any person j with $(\varepsilon_j, \omega_j)$ such that $I_j = (1, 0)$. Then for this person the optimal choice of effort (x_{1j}^*, x_{2j}^*) satisfies: $p(x_{1j}^*, \varepsilon_j) > p(x_{2j}^*, \varepsilon_j)$.*

Proof. I continue suppressing the dependence on j and ε_j. I will abuse the notation slightly and write

$$\begin{aligned} V_1(x_1, x_2) &= p(x_1)p(x_2)u(\omega - 2L + P - C) \\ &\quad + p(x_1)(1 - p(x_2))u(\omega - L - C + P) \\ &\quad + p(x_2)(1 - p(x_1))u(\omega - L - C) \\ &\quad + (1 - p(x_1))(1 - p(x_2))u(\omega - C) \\ &\quad - E(x_1, x_2). \end{aligned}$$

Now, it is enough to prove that if (x_1^*, x_2^*) maximizes V, then $x_1^* < x_2^*$. The first-order conditions for this maximization problem are given by

$$\begin{aligned} x_1: \quad & \partial_1 V_1(x_1^*, x_2^*) = p'(x_1^*)p(x_2^*)B - p'(x_1^*)\left(u(\omega - C) - u(\omega - L - C + P)\right) \\ & -\partial_1 E(x_1^*, x_2^*) = 0 \qquad (1.A.8) \\ x_2: \quad & \partial_2 V_1(x_1^*, x_2^*) = p'(x_2^*)p(x_1^*)B - p'(x_2^*)\left(u(\omega - C) - u(\omega - L - C)\right) \\ & -\partial_2 E(x_1^*, x_2^*) = 0 \qquad (1.A.9) \end{aligned}$$

where

$$B: = u(\omega - C) - u(\omega - L - C + P) - u(\omega - C - L) + u(\omega - C - 2L + P),$$

which is negative, since u is concave. For any x_2 define

$$\hat{x}_1(x_2) := \arg\max_{x_1} V_1(x_1, x_2).$$

Now, $\hat{x}_1(x_2)$ solves $\partial_1 V_1(x_1, x_2) = 0$ for any given x_2. Thus, by the implicit function theorem,

$$\begin{aligned}\frac{d\hat{x}_1(x_2)}{dx_2} &= -\frac{\partial_{12} V_1(\hat{x}_1(x_2), x_2)}{\partial_{11} V_1(\hat{x}_1(x_2), x_2)} \\ &= -\frac{p'(\hat{x}_1)p'(x_2)B - \partial_{12}E(\hat{x}_1, x_2)}{p''(\hat{x}_1)p(x_2)B - p''(\hat{x}_1)\left(u(\omega - C) - u(\omega - L - C + P)\right) - \partial_{11}E(\hat{x}_1, x_2)},\end{aligned}$$

where $\hat{x}_1$ stands for $\hat{x}_1(x_2)$. Because p' and B are negative and $\partial_{12}E$ is positive, the numerator is clearly negative. Furthermore, since the convexity of E and p imply that $\partial_{11}E$ and p'' are positive, the denominator must be negative. Thus I have that $\frac{d\hat{x}_1(x_2)}{dx_2} < 0$. This in turn implies that the graph of $\hat{x}_1(x_2)$ has one and only one intersection with the graph of $f(x_2) = x_2$. Hence, $\hat{x}_1(x_2)$ has a unique fixed point, $\hat{x}_2$.

Define next an auxiliary function

$$F(x_1, x_2) := \partial_1 V_1(x_1, x_2) - \partial_2 V_1(x_1, x_2).$$

The first-order conditions imply that

$$F(\hat{x}_1(x_2^*), x_2^*) = F(x_1^*, x_2^*) = 0. \tag{1.A.10}$$

On the other hand, $\partial_1 E(\hat{x}_2, \hat{x}_2) = \partial_2 E(\hat{x}_2, \hat{x}_2)$ implies that

$$F(\hat{x}_1(\hat{x}_2), \hat{x}_2) = F(\hat{x}_2, \hat{x}_2) = p'(\hat{x}_2)\left(u(\omega - L - C + P) - u(\omega - L - C)\right) < 0. \tag{1.A.11}$$

This implies that the lemma holds, if I can show that $F(\hat{x}_1(x_2), x_2)$ is strictly increasing. To complete the proof, notice that

$$\begin{aligned}
\frac{dF(\hat{x}_1(x_2), x_2)}{dx_2} &= \left(\partial_{11}V_1(\hat{x}_1(x_2), x_2) - \partial_{12}V_1(\hat{x}_1(x_2), x_2)\right) \frac{d\hat{x}_1(x_2)}{dx_2} \\
&\quad + \partial_{12}V_1(\hat{x}_1(x_2), x_2) - \partial_{22}V_1(\hat{x}_1(x_2), x_2) \\
&= -\left(\partial_{11}V_1(\hat{x}_1(x_2), x_2) - \partial_{12}V_1(\hat{x}_1(x_2), x_2)\right) \frac{\partial_{12}V_1(\hat{x}_1(x_2), x_2)}{\partial_{11}V_1(\hat{x}_1(x_2), x_2)} \\
&\quad + \partial_{12}V_1(\hat{x}_1(x_2), x_2) - \partial_{22}V_1(\hat{x}_1(x_2), x_2) \\
&= \frac{\partial_{12}V_1(\hat{x}_1(x_2), x_2)^2}{\partial_{11}V_1(\hat{x}_1(x_2), x_2)} - \partial_{22}V_1(\hat{x}_1(x_2), x_2) \qquad (1.A.12)
\end{aligned}$$

Since V_1 is striclty concave, $\partial_{11}V_1 < 0$. Thus (1.A.12) is strictly positive, if and only if

$$\partial_{11}V_1(\hat{x}_1(x_2), x_2)\partial_{22}V_1(\hat{x}_1(x_2), x_2) - \partial_{12}V_1(\hat{x}_1(x_2), x_2)^2 > 0.$$

But this holds always, since V is strictly concave and hence its Hessian must have a strictly positive determinant. □

Proof of Proposition 2. After Lemma 1.A.5, the proof for this proposition follows easily from the law of total probability:

$$\begin{aligned}
& \mathbb{E}[p(x_{1j}, \varepsilon_j) \mid I_{1j} = 1, I_{2j} = 0] \\
= \; & \mathbb{E}\left[\mathbb{E}\left[p(x_{1j}, \varepsilon_j) \mid I_{1j} = 1, I_{2j} = 0, \varepsilon_j, \omega_j\right] \mid I_{1j} = 1, I_{2j} = 0\right] \\
> \; & \mathbb{E}\left[\mathbb{E}\left[p(x_{2j}, \varepsilon_j) \mid I_{1j} = 1, I_{2j} = 0, \varepsilon_j, \omega_j\right] \mid I_{1j} = 1, I_{2j} = 0\right] \\
= \; & \mathbb{E}[p(x_{2j}, \varepsilon_j) \mid I_{1j} = 1, I_{2j} = 0].
\end{aligned}$$

□

1.B Supporting Tables

Table 1.6: Ownership of uninsured and insured ships by ship type

Ship type	Ships	Fraction insured	Owners	Own insured & uninsured
Drake	34,039	0.20	16,970	2,363
Badger Mk. II	22,248	0.03	11,396	705
Hurricane	21,634	0.23	10,358	1,608
Catalyst	16,909	0.12	10,448	577
Bestower	16,638	0.04	9,008	150
Hulk	16,084	0.06	12,266	179
Noctis	15,946	0.07	13,239	157
Thrasher	15,781	0.12	8,683	518
Iteron Mk. V	15,731	0.02	7,993	92
Cormorant	15,464	0.10	10,041	406

The first column of the table documents the total number of ships, the second column shows the fraction of those ships that are insured, the third column contains the

number of players who own a ship of that type and the last column is the most important one, containing the number of players who own both an uninsured and an insured ship of the given type. As can be seen from the table, there is considerable variation in how large a fraction of ships of a given type is insured. These differences are likely to be driven by the ship's use. All of the ship types of which more than 10% are insured are combat ships. Since this is the riskiest activity in the game, the result is not surprising. The rest of the ships in the table are ship types with non-combative intended purpose such as mining, salvaging of destroyed ships or freighting. A peaceful main purpose translates also to a low number of individuals that own both insured and uninsured ships of that type. Hence the moral hazard estimates that use variation in risk between ships of a given type owned by the same individual are mostly driven by the dedicated combat ships.

1.C Proportional Hazards Assumption and Goodness of Fit

Let T_i be the destruction time of ship i and let $Z_i(t) \in \{0, 1\}$ denote the insurance status of ship i. To assess the fit of the model I can first non-parametrically estimate the probability:

$$S_{0i}(t) := \mathbb{P}\left[T_i \geq t \mid Z_i(s) = 0 \; \forall s \leq t\right].$$

A standard way of estimating this is using the so called Kaplan-Meier estimator for the subsample where the insurance status is zero. Call this estimator $\hat{S}_0(t)$. The Kaplan-Meier estimator is a straightforward maximum likelihood estimator of the cumulative distribution function of accidents. For more details see Kalbfleisch and Prentice (2011). Now take any Borel measurable insurance process $z_i : [0, \infty) \to \{0, 1\}$. If I assume the accident risk in the population to be given by the Cox proportional

hazards model

$$\mathbb{P}\left[T \geq t \mid Z_i(s) = z_i(s)\ \forall s \leq t\right] = \exp\left[-\int_0^t \lambda(s)\exp(\beta z_i(s))\ \mathrm{d}s\right] =: S_{cox,i}(t),$$

then for $z_i \equiv 1$, I get that the survival probability is given by

$$S_{cox,1}(t) := \exp\left[-\int_0^t \lambda(s)\ \mathrm{d}s\right]^{\exp(\beta)} \equiv S_{0i}(t)^{\exp(\beta)}.$$

This suggests an estimator of $S_{cox,1}(t)$ by setting $\hat{S}_{cox,1}(t) = \hat{S}_{0i}(t)^{\exp(\hat{\beta})}$, where $\hat{\beta}$ is the standard partial maximum likelihood estimator of β. On the other hand, I can apply the Kaplan-Meyer estimator to the the population of insured ships to directly obtain an estimator of

$$\mathbb{P}\left[T_i \geq t \mid Z_i(s) = 1\ \forall s \leq t\right] =: S_{1i}(t).$$

If I denote this other estimator by $\hat{S}_1(t)$, then the appropriateness of the Cox model can be assessed by evaluating how well the graphs of the estimated $\hat{S}_1$ and $\hat{S}_{cox,1}$ agree.

This method does not allow for many strata, as one would need to draw a separate line for each of them. Arguably, adding strata to the model can only improve the fit and hence the pictures presented below can be thought of as lower bounds for the fit of the models used, especially, in Section 1.4.

Figure 1.4 plots the survival curves treating all ship types equal. Even with this very crude choice the fit of the model seems extremely good. The model slightly overestimates the likelihood of survival of insured ships. The figure also illustrates the magnitude of the effect that insurance has on survival probabilities either because of adverse selection or moral hazard with insured ships having much lower survival probabilities throughout the sample.

I next apply the technique to the subpopulation of Battlecruisers, a group of combat ship types that includes also the Drake, as well as only Drakes. The results are shown in Figures 1.5 and 1.6. As can be seen from the pictures, the fit becomes the better the fewer ship types are included. Notice also that the model in Figure 1.6 is exactly the same model as on the fourth row of Table 1.4. The fit of that model is almost perfect in comparison to the completely non-parametric Kaplan-Meier estimates with only slight overestimation of survival probabilities in the middle of the analysis period.

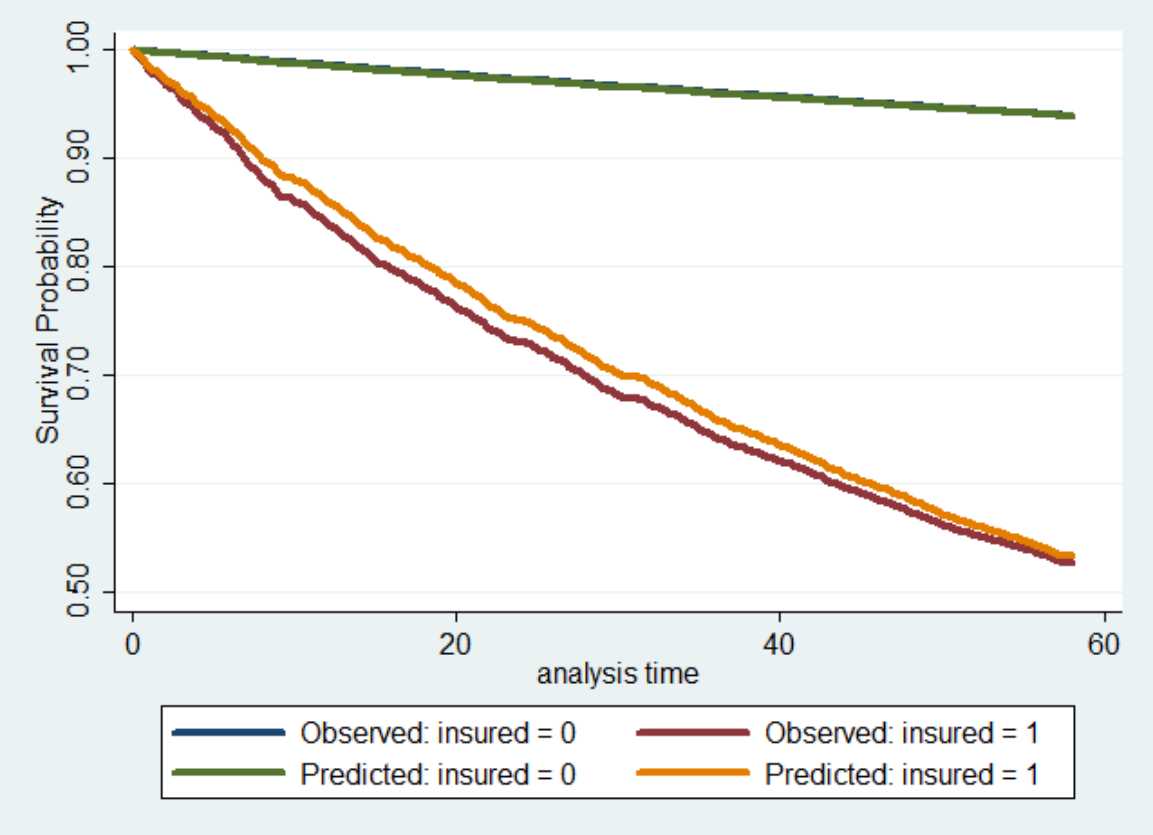


Figure 1.4: Kaplan-Meier and Cox proportional hazards survival curves for insured and uninsured ships.

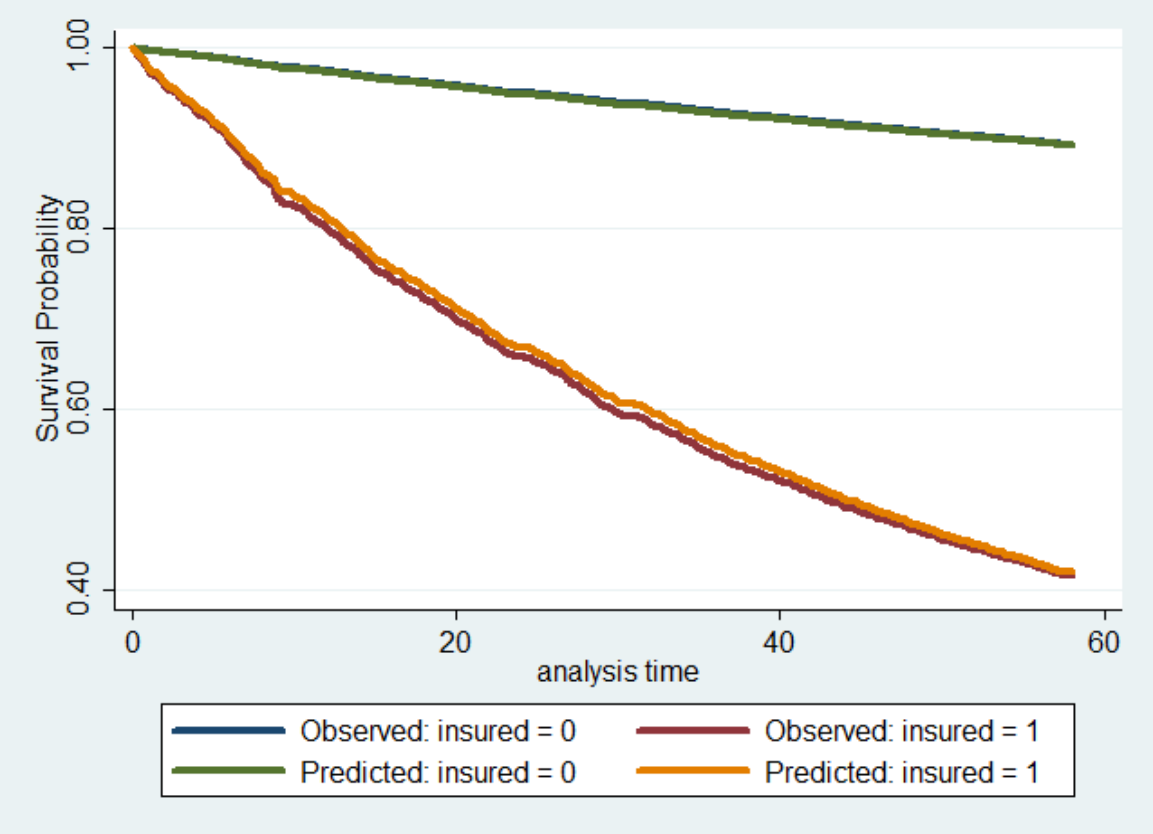


Figure 1.5: Kaplan-Meier and Cox proportional hazards survival curves for insured and uninsured Battlecruisers.

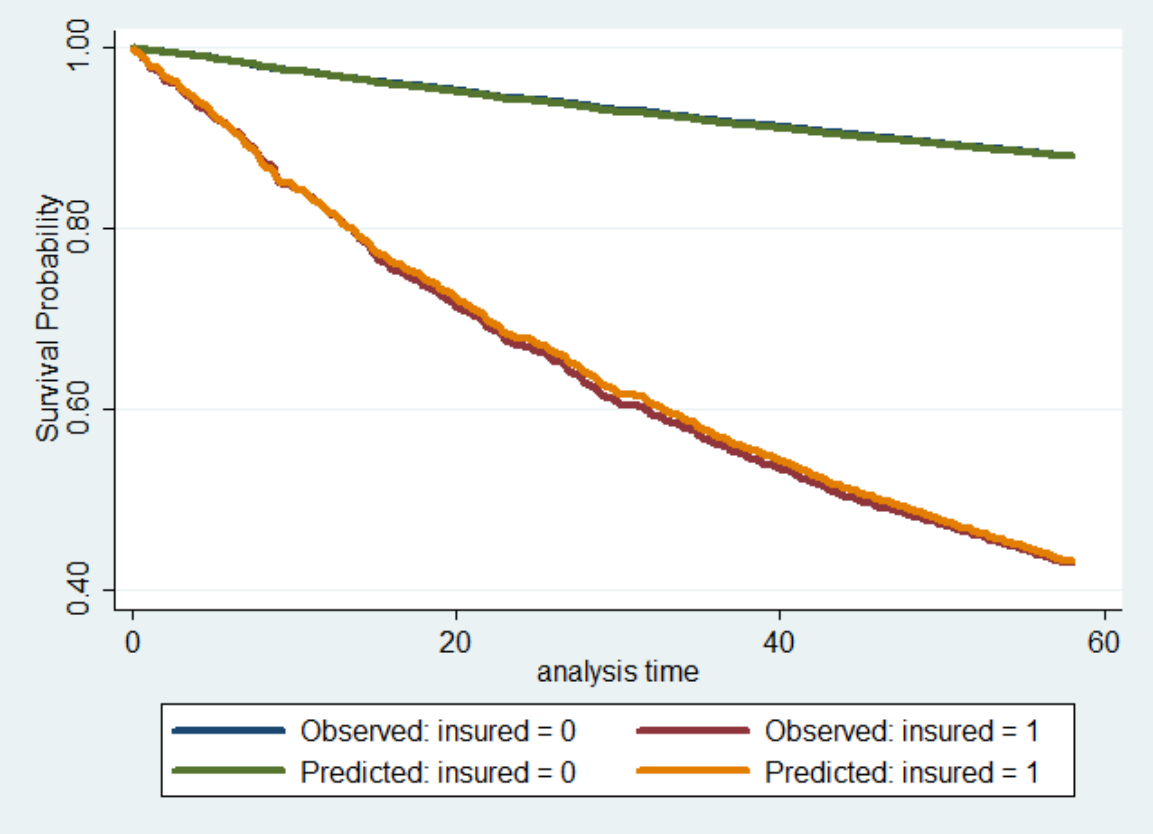


Figure 1.6: Kaplan-Meier and Cox proportional hazards survival curves for insured and uninsured Drakes.

Chapter 2

Measuring Dynamic Moral Hazard with Panel Data

2.1 Introduction

The general dynamic principal-agent problem is known to be highly complex and rarely tractable especially in discrete time (see, for example, Phelan and Townsend, 1991 and Edmans and Gabaix, 2011). In continuous time with Brownian noise, the problem becomes more tractable (Sannikov, 2008). However, many every-day contracts do not adjust their incentives continuously and can superficially have fairly obviously sub-optimal dynamic incentive structures. If agents are insensitive to the dynamics or unable to optimize under highly complex dynamic contracts, the static-looking contracts can prove to be optimal. This chapter develops a flexible econometric method for assessing how sensitive agents are to the dynamic provision of incentives over the span of the contract. The method will hopefully help us design more efficient contracts that take into account the way agent's in real contracting situations respond to dynamically changing incentives.

I apply my econometric method to measure dynamic moral hazard in insurance contracts with the following two commonly observed features: (1) protection is offered for a fixed period time and (2) accidents that result in the loss of the insured property also terminate the insurance contract and only the payout of the contract is awarded to the insured agent. Specifically, many common insurance contracts do not refund any fraction of the price of the contract if it is terminated before the contract period ends. One example of contracts that have this structure are warranties and additional product insurance policies. If a TV with a warranty breaks down beyond repair, the owner will get a refund and the warranty will then terminate, as there is no TV to be covered by the warranty. Similarly, car insurance policies in the United Kingdom often have a structure where no premiums are returned upon termination of the contract, if the insured person has made a claim during the period of insurance.

This contract structure implies that, in case of an accident, the insured agent will not only lose her deductible but also the intrinsic value of the remaining insurance coverage. Consequently, the agent faces different incentives for precautionary effort depending on how much time is left in the insurance contract. The value of the contract will be higher, the longer the remaining coverage period. Therefore, the incentives for precautionary effort should be higher in the beginning of the insurance contract and decrease towards its end. As long as agents' risk types are static over time, adverse selection will not produce a similar dynamic pattern.

I also show how this method can be combined with the method from the first chapter if there is data that simultaneously fulfills the prerequisites of both of the methods. The resulting method is highly robust to alternative explanations such as adverse selection or risk preferences. It boils down to estimating the variation in risk over items owned by the same person with different amounts of time left in their insurance contracts. If the contracts satisfy the two properties above, then at any given point in time the item with more time left in its contract should be less risky

than any item owned by the same agent that is close to the end of its contract period. This combined test is robust to seasonality and heterogeneity in risk types, wealth and tastes, because comparisons are made separately within each time period and between different objects owned by the same agent.

Just like in the first chapter, I will use versions of the Cox proportional hazards model (Cox, 1972) to construct my test for moral hazard. The general benefits of the proportional hazards model were outlined in the first chapter. The ability to incorporate time-varying covariates is of critical importance, as the time since the beginning of the insurance contract is obviously time-varying over the observation period. Much like the first chapter, I apply my methods to a data set on virtual space ship insurance from a large multiplayer online game called EVE Online. The data set is described in more detail in the first chapter. The key feature of the data is that the insurance contracts in the game satisfy both of the 2 assumptions outlined above. Furthermore, agents in the game own multiple perfectly identical ships which allows me to demonstrate the method combining the ideas from this and the first chapter.

When I measure the evolution of risk over the span of the game's insurance contracts, an interesting pattern emerges: I show that a newly insured ship is about 13 times more likely to get destroyed each day compared to an identical ship with only 15 days left in its contract when both ships are owned simultaneously by the same agent. Although moral hazard causes insured ships to be more risky over the whole contract period, this empirical pattern of decreasing dynamic risk contradicts the theoretical prediction outlined above for how optimally behaving agents should be reducing their precautionary effort towards the end of the contract period. I also discuss how this effect is somewhat mitigated by experience or learning.

It is noteworthy that many other contracts outside insurance imply identical dynamic incentives to the ones in the game. For example, many temporary labor contracts with hard deadlines have highly similar dynamic incentives. If poor output can

lead to an early termination of the contract and if at least some wage payments are made over the course of the contract, then the incentive structure for the employee is identical to the one given in my insurance model: When the project is close to its end, by shirking the employee loses only the value of the remaining few wage payments. However, early in the project the value of the potential stream of payments is much higher. Consequently, effort should be decreasing towards the end of the contract. The result is robust to even reputation effects, as long as the reputational cost of a premature termination is not lower in the beginning of the contract. These ideas are somewhat similar to what was presented in Gibbons and Murphy (1992).

On a more general level, this paper highlights that despite our imperfect understanding of optimal dynamic contracting with both adverse selection and moral hazard, theoretical modeling of existing contract structures can still be used to derive predictions of optimal responses to these contracts that are testable using panel data. The econometric methods used in this paper to measure dynamic variation in risk can be flexibly applied to other contracting situations with completely different incentive structures.

Much of the related literature was discussed in the first chapter of this dissertation. There are very few papers that discuss the empirics of timing of moral hazard.[1] Most of the papers that discuss the dynamics of the contract are concerned with adverse selection (for examle, Cabral, 2016 and Finkelstein, McGarry and Sufi, 2005).

The rest of the paper is organized as follows. The next section presents my theoretical model and formalizes my identification strategy. Section 3 discusses the econometric method and presents the results from applying it to the data. Section 3 investigates also the dynamic effect of experience on moral hazard. Section 4 concludes

[1] Abbring, Chiappori and Pinquet (2003) being one of the only exceptions. They use the bonus malus feature of French car insurance to estimate the effect of past claims on the propensity of accumulating future claims.

the study. For a detailed description of EVE Online and its insurance contracts, see the first chapter, especially Section 3.

2.2 Theoretical Model and Hypotheses

This section starts by introducing a simple dynamic model of effort choice and shows how the changing value of the contract over the contract period affects effort choices and accident risk. I then add the choice of effort for two items to the model. This produces a method that requires the most from the data but that controls for the richest variety of alternative explanations. By a logic similar to the one presented in the first chapter of this dissertation, similar items with different amounts of time left in their contracts all share the owners risk type and differences in riskiness can be explained only by different effort choices. All of the proofs for this section can be found in the Appendix.

2.2.1 A Dynamic Model with a Single Insurable Unit

In this section each agent $j \in N$ owns a single item and faces a dynamic problem of insurance and precautionary effort choice. I assume that the insurance contracts have the form described in the introduction: they are taken for l units of time or until the next accident happens. The offered length l is set exogenously by the insurance provider. I show that in the beginning of the contract the net utility loss from an accident is higher than in the end because the longer the remaining insurance coverage period the more valuable it is and hence the more it hurts to lose it. The section gives a set of assumptions under which items with more time left in their contracts are also cross-sectionally less risky than items whose contract is about to end. My goal in this section is to first formally illustrate the dynamic moral hazard outlined above and then discuss some of the restrictions of trying to measure this effect across different

agents. The assumptions of this section are fairly demanding but most of them are relaxed in the following section where I discuss my model with dynamic effort choice for two insurable items.

I assume that time is discrete and the timing of events in each period, $t = 1, 2, \ldots$, is as follows: In the beginning of the period the agent observes how much time $I'_{jt} \in \{0, 1, \ldots, l-1\}$ she has left in her insurance contract. Here $l \in \{1, 2, \ldots\}$ is the length of the offered contract and it is set by the insurance company. After observing I'_{jt} the agent decides how much precautionary effort x_{jt} she would like to exercise and whether she wants to renew her contract by setting $I_{jt} = l$ by paying C or just keep whatever time she has in her contract (if any) by setting $I_{jt} = I'_{jt}$ for free. After the agent's choice the accident risk is realized. If I let D_{jt} be an indicator of an accident for person j in period t, then person j's monetary endowment at the end of period t is given by

$$\omega - D_{jt}(L - \mathbb{1}\{I_{jt} > 0\}P) - \mathbb{1}\{I_{jt} = l\}C,$$

where $\mathbb{1}A$ is an indicator of event A. Just as before, ω is the person's per period monetary income, L is the size of the loss, P is the insurance payout and C is the cost of renewing the insurance contract. The agent then enjoys her utility for period t given by

$$u(\omega - D_{jt}(L - \mathbb{1}\{I_{jt} > 0\}P) - \mathbb{1}\{I_{jt} = l\}C) - E(x_{jt}, \varepsilon_j)$$

I assume that the agent's von Neumann-Morgenstern utility from consumption

$$u \colon \mathbb{R}_+ \to \mathbb{R}_+$$

is continuous everywhere and twice continuously differentiable. Furthermore, I require that $u'(c) > 0$, $u''(c) < 0$ for all $c \in \mathbb{R}_+$. The cost of effort $E(\cdot, \varepsilon_j)$ is a twice

continuously differentiable, increasing and convex. For this subsection, I assume that there is only heterogeneity in risk types and effort costs, and this heterogeneity is captured by $\varepsilon_j \in \mathbb{R}$ with higher values interpreted as higher risk types. I require that

$$\partial_x \partial_\varepsilon E(x, \varepsilon) \geq 0.$$

This single-crossing condition guarantees that, true to my interpretation of ε, increasing risk type does not imply that effort becomes relatively cheaper. Otherwise, agents with higher ε might optimally choose lower accident probabilities and hence, calling them higher risk types would be misleading. I also assume that $\lim_{x\to 0} \partial_x E(x, \varepsilon_j) = 0$ and $\lim_{x\to\infty} \partial_x E(x, \varepsilon_j) = \infty$. Given an effort level x, agent j's accident probability is $p(x, \varepsilon_j)$ where $p(\cdot, \varepsilon_j)$ is twice continuously differentiable, decreasing and convex for every ε_j. Conditional on the risk type and effort choices, accidents are assumed to be independent over time and across individuals. Furthermore, I assume that $\partial_\varepsilon p(x, \varepsilon) > 0$, and $\partial_x \partial_\varepsilon p(x, \varepsilon) \geq 0$. The interpretation for the first is that higher risk types have higher accident probabilities at every effort level and the second implies that increases in effort do not buy larger decreases in risk for higher risk types. The latter assumption complements the assumption that $\partial_x \partial_\varepsilon E(x, \varepsilon) \geq 0$ in making sure that, ceteris paribus, higher risk types will not have lower realized risk. Last, I assume that every agent discounts the future by a common discount factor $\delta \in (0, 1)$.

A central identification assumption in this section is that an accident resets the time left in the insurance contract to zero. Formally, I assume that the time left in the insurance contract at the beginning of period $t + 1$ is given by $I'_{t+1j} = (1 - D_{jt})(\max\{0, I_{tj} - 1\})$. In other words, either an accident resets the time left to 0 or, in the absence of an accident, the person starts her next period with one unit of time less left in her contract. Of course, if her contract has already run out of time, she will start the period with zero units of time in her contract.

I assume for simplicity that there are no endogenous saving decisions or heterogeneity in wealth and that the agent's per period monetary income is fixed.[2] Without assuming constant absolute risk aversion, heterogeneity or fluctuations in unobserved wealth translate to potential heterogeneity in risk preferences that considerably complicate the problem. The biggest complication is that, without strong assumptions, risk preferences are not monotonously related risk choices (see Dionne and Eeckhoudt, 1985 and Briys and Schlesinger, 1990)

Let $h_t = \{x_{js}, I_{js}, D_{js}\}_{s=0}^{t-1}$ be the t length history of accidents and insurance and effort choices before period t and let $\mathcal{H}$ be the set of all of these histories. The initial levels (x_{j0}, I_{j0}, D_{j0}) are assumed exogenous.[3] If the agent starts at $t = 0$ with I' units of insurance, then the agent's problem is to choose functions $x_j : \mathcal{H} \to \mathbb{R}_+$ and $I_j : \mathcal{H} \to \{0, 1, 2, \ldots, l\}$ to solve

$$V_j(I') := \max_{x_j, I_j} \mathbb{E}\Bigg[\sum_{t=0}^{\infty} \Big\{u(\omega - D_{jt+1}(L - \mathbb{1}\{I_j(h_t) > 0\}P) - \mathbb{1}\{I_j(h_t) = l\}C) \\ -E(x_j(h_t), \varepsilon_j)\Big\}\delta^t \mid x_j, I_j, \varepsilon_j\Bigg], \tag{2.2.1}$$

where the expectation is over the realization of histories and where for a history h_t, $I_j(h_t)$ is constrained to the set $\{\max\{0, (1 - D_{jt-1})I_j(h_{t-1}) - 1\}, l\}$. This is a Markov decision process where the time left in the insurance contract when entering the period, I'_t, can be taken as the state. Arguments given in, for example Bertsekas (1995), can be used to show that a maximizing policy can be found from the set of policies that depend only on I'_t. A solution to (2.2.1) is then obtained by solving the

[2] I allow for saving in my most general model below.

[3] Of these only the level of I_{j0} will affect the solution to the agent's maximization problem and none of them change my results.

following Bellman equation:

$$
\begin{aligned}
&V(I', \varepsilon_j) = \\
&\max_{x_j(I'), I_j(I')} \Big\{ p(x_j(I'), \varepsilon_j \left[u(\omega - L + \mathbb{1}\{I_j(I') > 0\}P - \mathbb{1}\{I_j(I') = l\}C) + \delta V_j(0) \right] \\
&+ (1 - p(x_j(I'), \varepsilon_j)) \left[u(\omega - \mathbb{1}\{I_j(I') = l\}C) + \delta_j V_j(\max\{0, I_j(I') - 1\}) \right] \\
&- E(x_j(I'), \varepsilon_j) \Big\}, \qquad (2.2.2)
\end{aligned}
$$

subject to $I_j(I') \in \{l, I'\}$.

The first key observation is that for any individual the probability of an accident increases as the end of the insurance contract becomes closer. The formal version of this idea is given in the following lemma:

Lemma 2.2.1. *If $I(I') > I(I'') > 0$, then $p(x_j(I'), \varepsilon_j) < p(x_j(I''), \varepsilon_j)$ for all ε_j.*

The full proof of this lemma is in the Appendix. Here I highlight the key step behind the proof. I consider first the harder case, where $I(I') < l$. For $\hat{I} \in \{I', I''\}$ the first-order condition with respect to x is:

$$
p'(x, \varepsilon_j) \Big(u(\omega - L + P) - u(\omega) + \delta_j V(0, \varepsilon_j) - \delta_j V(I(\hat{I}) - 1, \varepsilon_j) \Big) = E'(x, \varepsilon_j).
$$

This can be interpreted as a standard balancing of marginal cost and marginal benefit of effort. On the left-hand side the benefit comes from two sources: First, higher effort (x) decreases the probability of an accident (p). Increasing x moves probability mass at rate p' from the event where the flow utility is $u(\omega - L + P)$ to the event where it is $u(\omega) > u(\omega - L + P)$. It also increases the likelihood of the continuation utility being $V(I(\hat{I}) - 1, \varepsilon_j)$ rather than $V(0, \varepsilon_j)$. Because the agent is risk averse, she finds insurance valuable and the more periods that remain in the contract, the better. In other words, V is increasing in its first argument. I show this intuitive result formally in the Appendix. But this then implies that $V(I(I') - 1, \varepsilon_j) > V(I(I'') - 1, \varepsilon_j) \geq$

$V(0, \varepsilon_j)$. Thus, the marginal benefit for effort is higher when the insurance choice is $I(I')$ than when it is $I(I'')$ implying higher effort choices and hence lower accident probabilities. When $I(I') = l$ the effect is even stronger, since the contemporaneous utility loss from an accident is $u(\omega - L + P - C) - u(\omega - C) < u(\omega - L + P) - u(\omega)$ where the inequality follows from concavity of u. The income effect of the momentary wealth loss due to the insurance payment makes the agent more sensitive to the absolute monetary loss from the accident. Of course, allowing saving and borrowing would smooth some of this effect over several periods. However, this observation is otherwise very general and holds with much richer assumptions about the individuals.

Now, consider contracts that are fully separating in the sense that all of the agents who buy the insurance contract share the same risk type ε. Then all of the agents buying insurance will be identical and consequently will also choose the same level of precautionary effort when faced with the same amount of time left in their contracts. Lemma 2.2.1 then implies that on expectation an agent with less time left in her contract will be more risky than an agent who is in the beginning of her contract. Formally, this idea is given by the following corollary:

Corollary 2.2.2. *Assume that the contract is such that if $I_j > 0$ and $I_k > 0$ then $\varepsilon_j = \varepsilon_k =: \hat{\varepsilon}$. In that case $I(I', \varepsilon) > I(I'', \varepsilon) > 0$ implies that*

$$p(x(I'), \hat{\varepsilon}) = \mathbb{E}[p(x, \varepsilon) \mid I(I', \varepsilon)] < \mathbb{E}[p(x, \varepsilon) \mid I(I'', \varepsilon)] = p(x(I''), \hat{\varepsilon}).$$

This is a testable implication with my data.

Consider then a model without effort choice and hence without moral hazard where each agent solves:

$$V_j(I') = \max_{I_j(I')} \Big\{ p(\varepsilon_j)\left[u(\omega - L + \mathbb{1}\{I_j(I') > 0\}P - \mathbb{1}\{I_j(I') = l\}C) + \delta_j V_j(0)\right] + (1 - p(\varepsilon_j))\left[u(\omega - \mathbb{1}\{I_j(I') = l\}C) + \delta_j V_j(\max\{0, I_j(I') - 1\})\right] \Big\} \quad (2.2.3)$$

Here, observing $I_{jt} > I_{kt} > 0$ implies that person j has on average been a shorter time without an accident than person k. In other words, observing person k with less time in her contract than what j has implies that she is on average *less risky* than j because she is known to have seen at least $l - I_{kt}$ periods without an accident, while for person j the same lower bound is only $l - I_{jt}$. This insight can be used to prove the following result:

Proposition 3. *In the model without moral hazard given in (2.2.3), if* $I_{jt} > I_{kt}$*, then* $\mathbb{E}[D_{jt} \mid I_{jt}] \geq \mathbb{E}[D_{kt} \mid I_{kt}]$.

This inequality holds as an equality only if there is almost surely no heterogeneity of risk types in the population buying the contract – in other words, if the contract is separating in the sense described above. This is my second testable result. To summarize, whenever the contract is non-separating and there is both moral hazard and adverse selection, the econometrician may observe new contracts being riskier than old contracts or vice versa. If old contracts are riskier than new ones, that can be taken as proof of moral hazard while new contracts being riskier than old ones implies that there is a possibility of adverse selection. If the hazard rate is constant over the span of the insurance contract, the test is inconclusive, either the contract is separating in the above sense[4] and there is no moral hazard, or the heterogeneity in risk types completely offsets the effect of moral hazard. I would like to emphasize that

[4]Notice that this contains the case where there is no adverse selection.

heterogeneity of risk types in the pool of insuring agents does not strictly speaking imply adverse selection, as the selection to the insurance pool may be based on factors such as tastes that may be completely orthogonal to the person's risk type. This could yield the same distribution of risk types in both insuring and non-insuring populations with any differences in risk between the two populations being completely explained by moral hazard. This ambiguity is fully resolved by the richer model in the next section.

A notable restriction of the model is that it does not allow for risk to vary over time. Even very uniform seasonal variation in risk may cause problems, if it results in different risk types buying insurance at different times. Patterns where riskier types buy insurance when seasonal risk is lower can potentially lead to observing periods where agents with more time left in their contracts are less risky even if the agents cannot affect their riskiness. Controlling for seasonality with something like time fixed effects can mitigate the problem as long as the effect of the seasonality is relatively homogeneous for different risk types. However, this problem is not present in my third identification strategy presented in the next subsection.

Second, the model above does not allow for heterogeneity in risk preferences or wealth. Allowing for heterogeneity in wealth is clearly almost synonymous to allowing for variations in risk preferences unless I assume constant absolute risk aversion. Variation in risk preferences, on the other hand, is problematic because it does not yield clear predictions of how it will translate into variations in accident probabilities. This is the old dichotomy between self-insurance and self-protection (see Ehrlich and Becker, 1972, Dionne and Eeckhoudt, 1984 and Briys and Schleisinger, 1990). While more risk averse agents are always willing to accept higher reductions to their welfare in the good state in exchange for an increase in welfare in the bad state (self-insurance), the same is not true for paying a cost in both states for reducing the probability of the bad state (self-protection). In fact, the authors above present

examples where more risk averse agents exhibit more risky behavior. The intuition behind the result is that self-protection, i.e. increasing effort to reduce the accident probability, makes the agent worse off in both the good and the bad state. In contrast, the mean preserving contractions that are always preferred by a risk averse agent have the feature that they always increase the payoff in the bad state.

The result from Proposition 3 clearly still holds if I allow for variation in risk preferences, because it assumes no variation in effort. If I assume enough on the utilities to guarantee that more risk averse agents choose higher levels of effort, and if there are multiple risk preference types in the pool of insured agents, then it is more likely to observe average risk decreasing over the span of the contract, as the agents with higher risk aversion choose more effort in each period and hence are more likely to survive longer. This effect still has to offset each individual's decreasing effort choices over the span of the contract. The model in the following section goes around this ambiguity by comparing multiple items owned by the same person with different amounts of time left in their contracts. This allows me to factor out the effects that risk type and preferences have on survivability.

As I alluded to earlier in the introduction, this section's incentive structure is not unique to insurance contracts. Many labor market contracts, especially in service and construction work, have a structure where a contract lasts for a set amount of time, payments happen periodically within that time frame and poor outcomes may lead to an early termination of the contract. For instance, payments for construction work occur over the span of the project and the contractor or an individual worker can be fired any time during the contract because of shoddy workmanship. It is easy to modify the model above to show that these contracts will potentially lead to similarly decreasing incentives for effort over the span of the contract. As long as the agent is relatively patient compared to the value of the periodic payments, terminating the contract early will lead to a higher loss in expected value to the agent

than terminating it late, because more installments are lost with an early termination. This idea is fairly similar with the one given in Gibbbons and Murphy (1992).

2.2.2 Dynamic Model with Two Insurable Items

Next, I will show how many of the assumptions can be weakened if agents own two similar items. Intuitively the identification strategy is a combination of the idea presented above and the one from the first chapter of this dissertation. It is the most demanding on the data but also the most robust to alternative explanations such as seasonality. Each agent owns two similar insurable objects. Either of these can be insured at any point in time with a separate contract for l units of time. In the case of an accident, the insurance offers a fixed payout of P after which the contract on that item ends, no matter how much time was left in the contract. Just as in the model above, I assume that new contracts purchased while an old one is still running only renew the length of the contract but do not affect the coverage in any other way. The main result of this section is that if a person owns two items that are both insured but one of them has less time left in its contract than the other, then the item with less time in its contract is on average more risky. Here, the fact that I observe multiple items in different states from the same person allows me to factor out all heterogeneity and seasonality - even when these affect the agent's effort choices. Furthermore, since both of the compared items are insured under the same contract, the model is robust to item-specific selection that leads to different items being insured with different contracts.

Let $x_j^s = (x_{j1}^s, x_{j2}^s) \in \mathbb{R}_+^2$ be the vector of effort choices of the agent in period s. Similarly, $I_j^s = (I_{1j}^s, I_{2j}^s) \in \{0, 1, \ldots, l\}^2$ is the time left in the insurance contracts in period s after new insurance purchases have been made. Denote by $D_j^s = (D_{1j}^s, D_{2j}^s) \in \{0, 1\}^2$ the indicators of accidents for both items in period s. Now, in the beginning of period t the agent recalls the history $h_j^t = \{x_j^s, I_j^s, D_j^s\}_{s=0}^{t-1}$. The timing of events in

each period is the same as in the previous section with the modification that whenever the agent is making a choice, she is simultaneously choosing effort or insurance levels for both of the two items. Similarly, she learns whether each item had an accident simultaneously. Define $\mathcal{H}$ as the set of all histories. I assume that conditional on the history of effort choices, D_{ij}^s is independent of D_{-ij}^s and of D_{kj}^r for all $r < s$ and all k. I also assume that $\mathbb{P}(D_{ij}^t = 1 \mid x_{1j}^t, x_{2j}^t, \varepsilon_j^t, t) = p(x_{ij}^t, \varepsilon_j^t, t)$ and that the agent is fully aware of this mapping. In particular, she is able to perfectly forecast all possible future accident probabilities. I continue assuming that $p(\cdot, \varepsilon_j^t, t)$ is continuously twice differentiable, decreasing and strictly convex and that $(x_1, x_2) \mapsto p(x_1, \varepsilon_j^t, t)p(x_2, \varepsilon_j^t, t)$ is also strictly convex for all ε_j^t and t.

I allow the agent's risk/taste type ε_j^t and her periodic monetary endowment ω_j^t to vary arbitrarily over time. For simplicity, I assume that the agent knows the full paths of these processes. Notice that this assumption includes the case where the agent earns a deterministic stream of income and is potentially committed to a predetermined plan for saving and borrowing decisions. However, I am not allowing the wealth distribution to depend on the history of accidents, exercised effort or insurance choices. Relaxing this assumption is unlikely to completely change the results in this section. Nevertheless, introducing, for example, rational saving decisions into the model can dampen some of the dynamic results provided below through consumption smoothing.

Similarly to the previous section, the consumer's von Neumann-Morgenstern utility from consuming wealth c when his/her heterogeneity term is ε is given by $u(c, \varepsilon)$, where

$$u \colon \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}_+$$

Just as above u is continuous everywhere and twice continuously differentiable with respect to c. Furthermore, I require that $\partial_c u(c, \varepsilon) > 0$, $\partial_c^2 u(c, \varepsilon) < 0$ for all $c \in \mathbb{R}_+$ and $\varepsilon \in \mathbb{R}$. Notice, however, that here u can depend arbitrarily on ε.

The utility cost of effort is measured by a function $E\colon \mathbb{R}^2_+ \times \mathbb{R} \to \mathbb{R}_+$. Here, $E(x_{1j}, x_{2j}, \varepsilon_j)$ is interpreted as the utility cost from effort levels (x_{1j}, x_{2j}) for person j. This function is also allowed to depend arbitrarily on ε. I require that E is twice continuously differentiable, increasing and convex in its first two arguments for every ε_j and that $\partial_{12}E(x_1, x_2, \varepsilon_j) \geq 0$. To guarantee that first-order conditions are satisfied at some (x_{1j}, x_{2j}) I assume that for all x_{-i}, $\lim_{x_i \to 0} \partial_i E(x_1, x_2, \varepsilon_j) = 0$ and $\lim_{x_i \to \infty} \partial_i E(x_1, x_2, \varepsilon_j) = \infty$. Compared to the first chapter, I also require the stronger assumption of full symmetry:

Assumption 2 (Symmetric effort cost)**.** *The cost function satisfies $E(x, y) = E(y, x)$ for all $x, y \in \mathbb{R}_+$.*

This is likely to be stronger than what is needed and could potentially be weakened using a similar technique to the one presented in the first chapter.

The agent's problem is then to choose policy functions $x_j : \mathcal{H} \to \mathbb{R}^2_+$ and $I_j : \mathcal{H} \to \{0, 1, 2, \ldots, l\}^2$ to maximize the expectation

$$\mathbb{E}\Bigg[\sum_{t=0}^{\infty} u\left(\omega_j^t + \sum_{i=1}^{2} D_{ij}^{t+1}(\mathbb{1}\{I_{ij}(h_j^t) > 0\}P - L) - \mathbb{1}\{I_{ij}(h_j^t) = l\}C, \varepsilon_j^t \right) \delta_j^t \\ - E(x_{1j}(h_j^t), x_{2j}(h_j^t), \varepsilon_j)\delta_j^t \mid (\varepsilon_j^t)_{t\in\mathbb{N}}, (\omega_j^t)_{t\in\mathbb{N}}, x_j, I_j \Bigg], \tag{2.2.4}$$

where $\delta_j \in (0, 1)$ is the discount factor and the maximization is subject to

$$I_{ij}(h_j^t) \in \left\{ l, \max\left\{0, (1 - D_{ij}^{t-1})(I_{ij}(h_j^{t-1}) - 1)\right\}\right\}.$$

In other words, the choice of insurance for each period must be made from the set $\{0, l\}$, if the object was in an accident in the previous period. Otherwise, one can either just keep the existing contract (with one unit less time left) or alternatively buy a new contract with l units of time.

Results from Bertsekas (1995) can again be applied to show that the problem has a recursive structure with respect to the state (t, I'_{j1}, I'_{j2}) where I'_{jk} is the time left in the insurance contract $k \in \{1, 2\}$ owned by agent j when entering period t before any choices are made. I let the accident probabilities, monetary endowments and risk/taste type depend on the time period. I therefore have to add it as a state.

The next proposition argues that the idea of the changing option value embedded in the contract carries over from the previous model to choices made by a single person on how to distribute effort over two insured items with different amounts of time left in their contracts. Specifically, a person who owns two insured items is going to take better care of the item that has more time left in its contract.

Proposition 4. *With dynamic moral hazard, the following inequalities hold for all* $t \in \mathbb{N}$*:*

1. *If* $0 < I'_1 < I'_2$*, then*

$$\begin{aligned} & \mathbb{E}[p(x_{1j}(t, I'_1, I'_2), \varepsilon_j^t, t) \mid I'_1, I'_2, I_{ij}(t, I'_1, I'_2) = I'_i, i = 1, 2] \\ \geq\ & \mathbb{E}[p(x_{2j}(t, I'_1, I'_2), \varepsilon_j^t, t) \mid I'_1, I'_2, I_{ij}(t, I'_1, I'_2) = I'_i, i = 1, 2]. \end{aligned} \tag{2.2.5}$$

2. *If* $0 < I'_1$*, then*

$$\begin{aligned} & \mathbb{E}[p(x_{1j}(t, I'_1, I'_2), \varepsilon_j^t, t) \mid I'_1, I_{1j}(t, I'_1, I'_2) = I'_1, I_{2j} = l] \\ \geq\ & \mathbb{E}[p(x_{2j}(t, I'_1, I'_2), \varepsilon_j^t, t) \mid I'_1, I'_2, I_{1j}(t, I'_1, I'_2) = I'_1, I_{2j} = l]. \end{aligned} \tag{2.2.6}$$

3. *Both inequalities are strict, if an agent never re-insures an item that still has time left in its contract.*

The third point does not necessarily always hold, because buying insurance "too early" can serve as a way to smooth out consumption. For example, a very high

accident risk tomorrow implies lower expected utility. Buying insurance today lets the agent shift some of the higher welfare from today forward in time and smooth expected consumption over periods. However, if the agent is allowed to make rational borrowing and saving decisions, smoothing with front loading insurance decisions is unlikely to be optimal, and hence the third point is likely to hold in most applications. Its validity can also be easily checked from the data.

The proof of Proposition 4 is somewhat more involved than the proof of the result in the previous section but the intuition behind it is the same: More time left in the contract of one of the items implies that the marginal return for effort is higher for that item, as an accident will result not only in losing the item but also in the termination of the longer remaining insurance coverage.

In the absence of moral hazard the two sides of the inequality should clearly be equal, because the two items share the agent's risk type. Since I am comparing the risk of two items owned by the same person in the same time period, heterogeneity of risk types combined with survivorship is not dampening the effect like it was in the previous section. Hence, even when both phenomena are present, rationality of the agents would imply that the inequalities above should hold in the population. Notice also that the model allows for arbitrary seasonal variation in risk across individuals, as well as variation in wealth and risk preferences. This flexibility is a result of the ability to compare the risk across two items owned by the same person at any given time.

2.3 Econometric Model and Results

In this section I will briefly discuss the econometric model that I will use to estimate the dynamic effect of effort on accident risk and then apply it to the space ship insurance data from EVE Online described in the first Chapter. A detailed discussion

of the data and descriptive statistics of the data can be found in the first chapter. The three key features of the data are that the contracts are taken for a fixed period of 12 weeks, accidents terminate the contracts and players own multiple simultaneously insured, identical ships. Hence, the assumptions of the model above are perfectly satisfied.

2.3.1 The Evolution of Risk over the Insurance Contract in the Cross Section

The goal here is to estimate how the time left in the insurance contract affects the item's hazard rate. If the assumptions from Section 2.2.1 hold, moral hazard should increase the hazard rate towards the end of the contract. Conversely, heterogeneity of risk types in the insurance pool will have the opposite effect of decreasing the average hazard rate over the span of the contract.

Just as in the first chapter, I will use a Cox proportional hazards model (Cox, 1972 and 1975) to estimate the effect of effort choices on accident risk. I assume that the probability that ship j of type k owned by person i survives without an accident for more than t units of time is given by

$$S_j(t;k) := \exp\left[-\int_0^t \lambda_k(s)\exp\left(\sum_{m=1}^{M}\gamma_m \mathbb{1}\left\{I_j(s)\in\left(\frac{l(m-1)}{M},\frac{lm}{M}\right]\right\}\right)\,\mathrm{d}s\right].$$

Here, $\lambda_k(s)$ is a ship type specific, possibly time varying baseline hazard rate, and γ_m is the effect of having more than $\frac{l(m-1)}{M}$ but no more than $\frac{lm}{M}$ days left in the insurance contract. In other words, I split each insurance contract into M parts of equal length and add dummies for each of the parts.

In contrast to the model presented in the first chapter, I do not allow the baseline hazard rate to depend on the identity of the person owning the ship. This restriction is applied simply to demonstrate the method from section 2.2.1 when only a single

observation per agent is available. I will relax it in the next section. Just as in the first chapter, the parameters γ_m can be consistently estimated using partial maximum likelihood without estimating the potentially infinite dimensional λ_k (see Ridder and Tunali, 1999 or Kalbfleisch and Prentice, 2011). The model allows for arbitrary seasonality in risk that is common across agents but may vary across ships of different type.

Notice that the econometric model is not specific to the EVE's insurance setting. It can be used to flexibly estimate the evolution of risk over the span of a contract in a multitude of other insurance settings. More specifically, the model is able measure any changes in the accident risk over the course of a contract and is not constrained to estimating the effects discussed in the previous section.

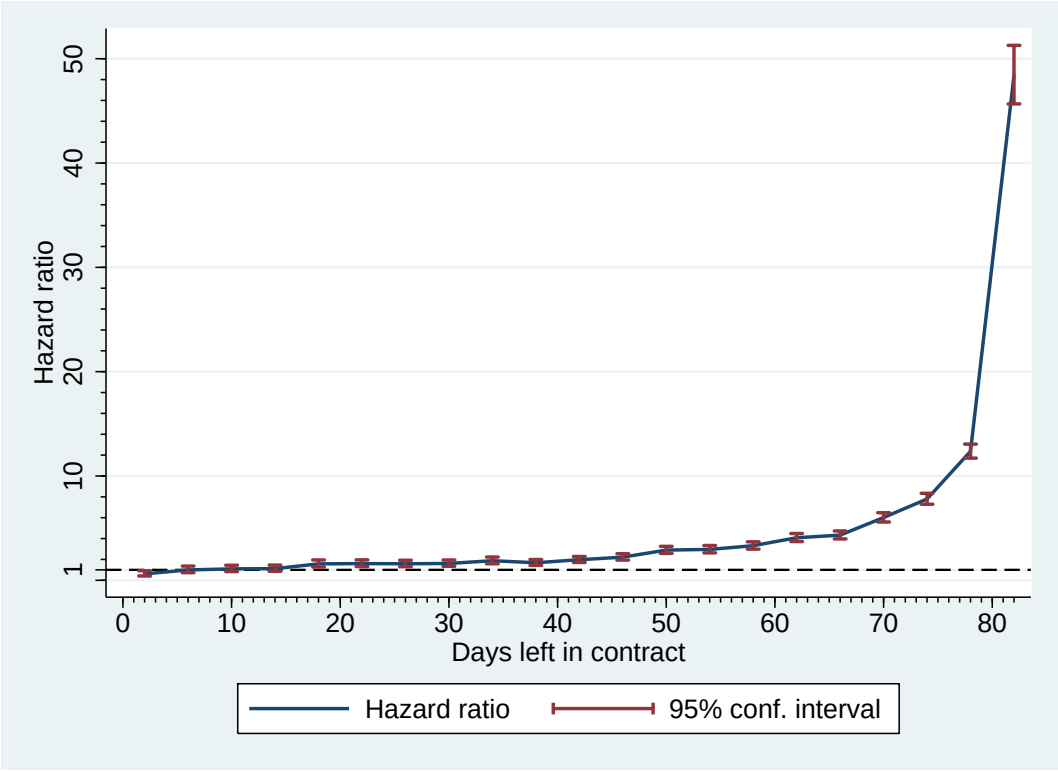


Figure 2.1: Hazard ratio as a function of time left in the insurance contract. 95% confidence intervals calculated using standard errors clustered at the player level.

Figure 2.1 plots the results from estimating the model when the insurance contract is split into episodes of 4 days. To ease interpretation I plot the estimates for the

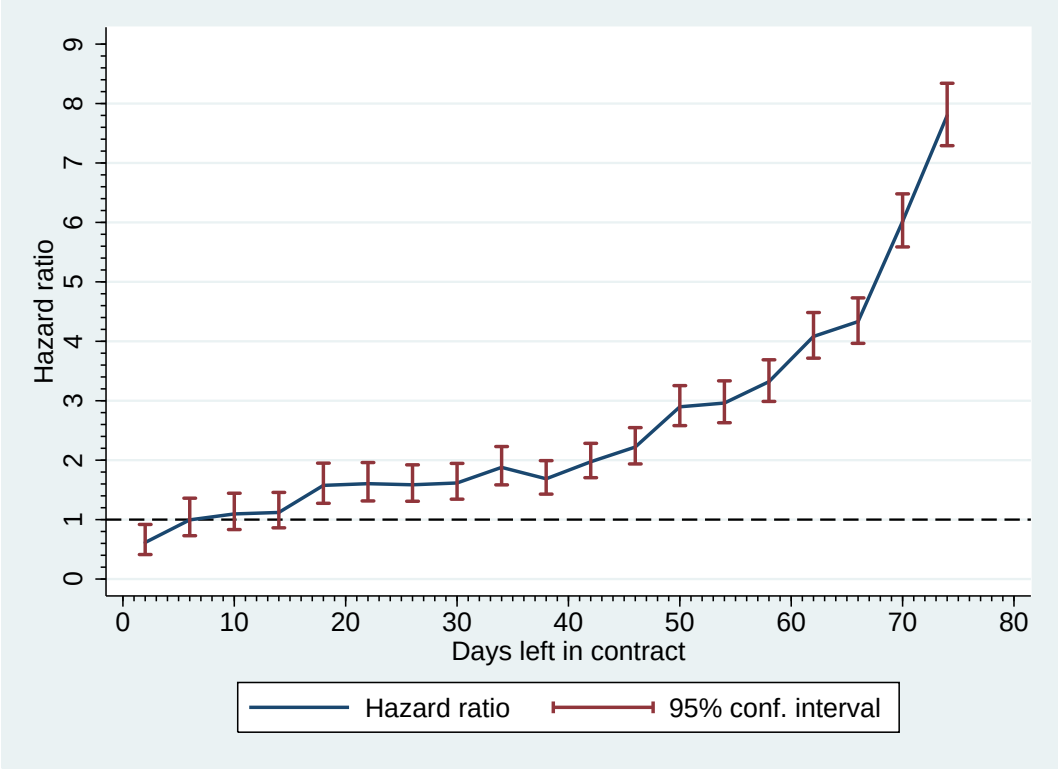


Figure 2.2: Hazard ratio as a function of time left in the insurance contract excluding first 8 days of the contract. 95% confidence intervals calculated using standard errors clustered at the player level.

hazard ratio e_m^{γ}.[5] Each estimate is plotted with its 95% confidence interval. Notice that, for a given ship, time in the figure runs from right to left. As can be seen from the figure, the first days of the insurance contract are a period of extremely high risk with hazard rates nearly 50 times as high as what they are for uninsured ships. The hazard ratio falls rapidly during the first few days. Once about 10 days have passed from the purchase of the contract, the hazard rate of an average ship is only about 5 times higher than the hazard rate of an uninsured ship. To get a clearer idea of what happens after 10 days from the purchase of the contract, I plot the same graph excluding the first 8 days of the contract. The result can be seen in Figure 2.2. The risk is fairly monotonously decreasing also over the remaining 76 days of insurance.

[5] A discussion on how to interpret hazard ratios can be found in the first chapter.

If one is willing to accept the theoretical model from Section 2.2.1, the results above imply that heterogeneity of risk types and potentially adverse selection is dominating moral hazard in EVE's spaceship insurance market. The hazard ratio is not increasing even at the end of the contract. Hence, the dynamic effect that the contract structure has on players' incentives seems relatively weak. However, the next section illustrates that interpreting the decreasing risk over the contract period as adverse selection could be erroneous. It is likely caused mostly by behavioral patterns or transaction costs not present in the theoretical model. I discuss these alternative explanations for the observed empirical pattern in the following section.

The results are robust to splitting the contract into even shorter time intervals as well as looking only at agents who own only one ship of a given kind. A specification where a person owns only a single ship corresponds better, strictly speaking, to my theoretical model in Section 2.2.1.

2.3.2 Test for Dynamic Moral Hazard Using Multiple Ships Owned by the Same Agent

As was outlined in Section 2.2.2, here I will estimate the difference in hazard rates for two items, conditional on them being owned by the same person and them having different amounts of time left in their insurance contracts. To this end, I need to only slightly alter the model from the previous section by allowing the baseline hazard rate to be person specific. In other words, the survival probability for ship j of type k, owned by person i is given by:

$$S_{ij}(t;k) := \exp\left[-\int_0^t \lambda_{ik}(s)\exp\left(\sum_{m=1}^{M}\gamma_m \mathbb{1}\left\{I_{ij}(s)\in\left(\frac{l(m-1)}{M},\frac{lm}{M}\right]\right\}\right)\,\mathrm{d}s\right].$$

The only difference compared to the previous section is that the baseline hazard rate λ_{ik} is allowed to depend arbitrarily on the person in addition to the ship's type.

The coefficients γ_m are identified using the variation in risk between ships of the same type owned by the same person that have different amount of time left in their contracts at a given point in time. As long as ships of the same type and owned by the same person share a risk type at each point in time,[6] the same amount of effort from the user should result in equal hazard rates for the two ships. Consequently, any observed differences in the hazard rates should be attributed to differences in precautionary effort. If the players find longer remaining contracts more valuable, I should observe ships with less time left in their contracts being more hazardous than ships closer to the beginning of their contract. In other words, I expect the pictures from the previous section to be reversed.

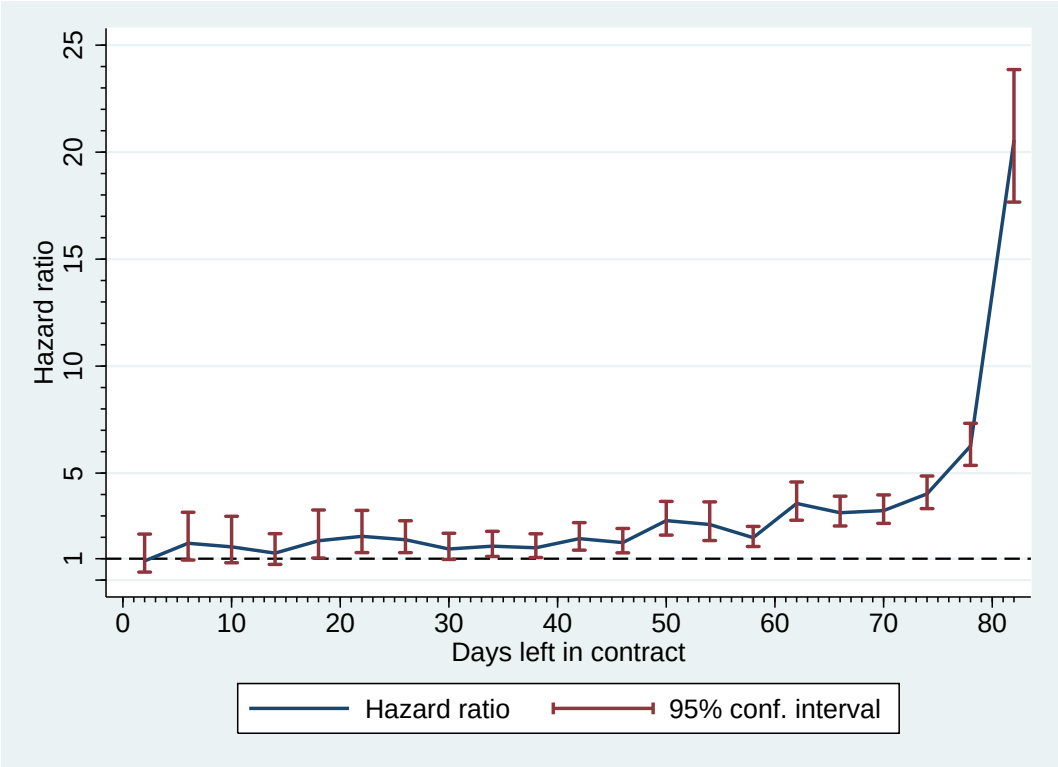


Figure 2.3: Hazard ratio as a function of time left in the insurance contract. Model includes a user specific baseline hazard rate. 95% confidence intervals calculated using standard errors clustered at the player level.

[6]This risk type can be arbitrarily time varying.

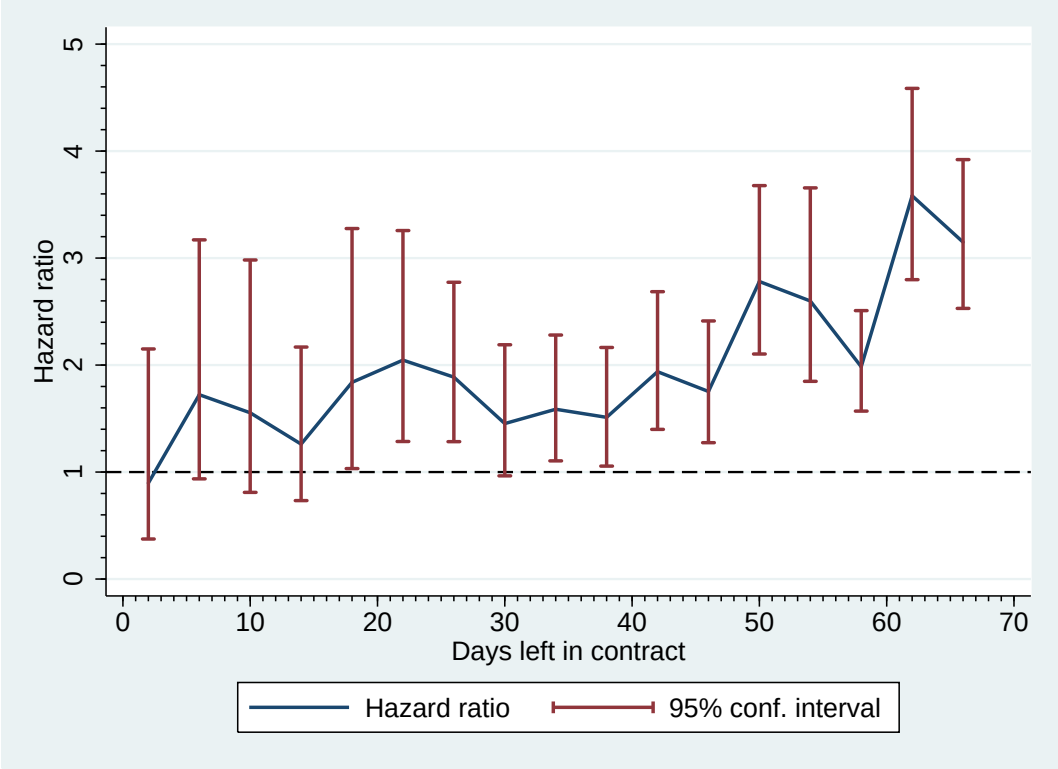


Figure 2.4: Hazard ratio as a function of time left in the insurance contract excluding first 8 days of the contract. Model includes a user specific baseline hazard rate. 95% confidence intervals calculated using standard errors clustered at the player level.

The results from estimating the model are presented in Figures 2.3 and 2.4. The striking first finding is that the decreasing overall hazard ratio from the previous section is mostly preserved. Computing the estimates using only insured ships does not change the results.

Rationalizing the decreasing hazard ratio requires an explanation for why an agent might use the recently insured item more than the item that has only a single day left in its contract. The benefit from allocating effort this way has to be larger than the expected value of the remaining long period of future insurance protection enjoyed by the newly insured ship. One possible explanation is that the items are being used in different locations and the environmental risk or the benefit from risk taking in these locations varies over time. For example, consider an agent who owns two cars, one in Chicago and the other in New York. Assume that the car in Chicago has only

liability insurance while the vehicle in New York has more comprehensive coverage. If Chicago is hit by a snow storm, using the car there becomes more risky for a short while. The agent has two options: she can either drive the car with more coverage from New York to Chicago or buy more coverage for the car in Chicago. The second choice might be more appealing if the opportunity cost of the long drive is higher than the price of additional coverage. If this is true, we might observe a similar pattern to that presented in Figure 2.3. Hence, even if the insurable items are otherwise identical, the combination of highly localized use, dynamically and spatially varying risk and high cost of moving items between locations can pose a threat to my dynamic identification strategy.

Ships in EVE are also used in different locations and the reward for taking risky actions is likely to vary between different locations and over time. The explanation still requires that most players find driving the existing insured ship to the new location a dreadful task whose cost overweighs the expected cost of lost insurance protection. The shocks to the riskiness of different locations need to be relatively persistent to be able to explain why a ship that has been insured for over 30 days still has a hazard rate that is over twofold compared to a ship at the end of its insurance contract.

Another plausible explanation is more behavioral: the ship that got recently insured has been the target of the user's attention. The longer it has passed from the purchase of the previous insurance contract the more likely it is that the user has forgotten that the ship is available for use. In other words, the insurance purchase acts as a proxy for attention paid to the ship. This story could also generate the monotonous decreasing relationship between hazard ratios and time since the previous insurance purchase.

For reasons explained in the first chapter, adding dummies for ship age should partially control for ship-specific heterogeneity in the insurance pool. Heterogeneity

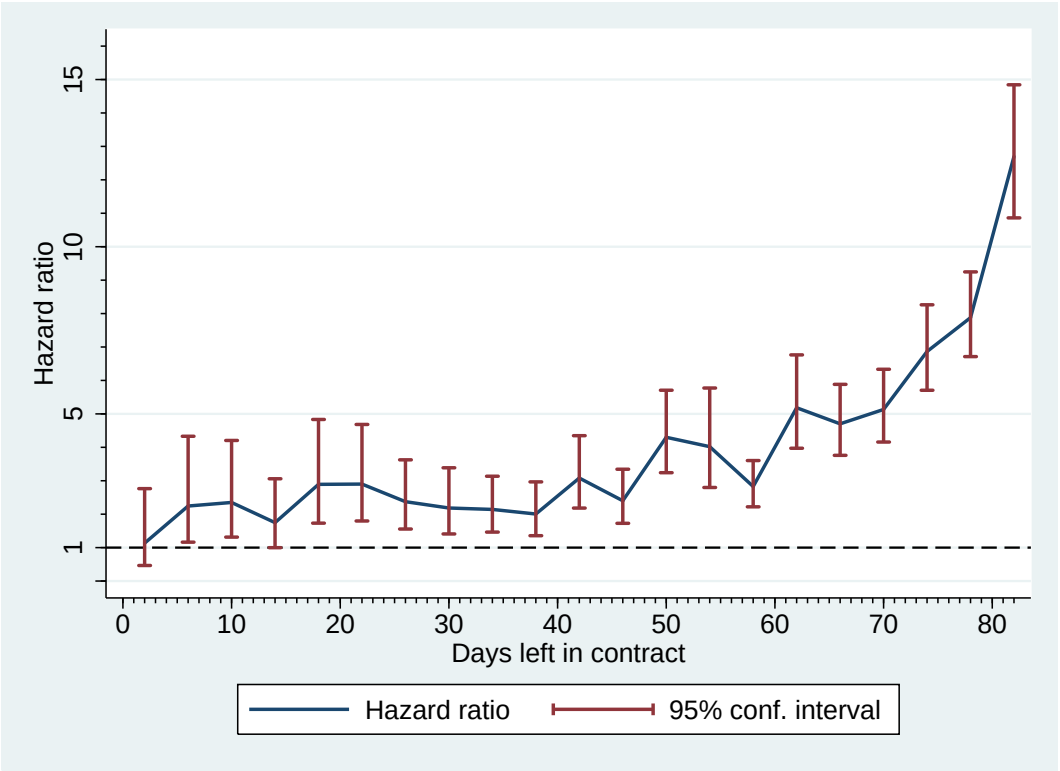


Figure 2.5: Hazard ratio as a function of time left in the insurance contract. Model includes a user specific baseline hazard rate and ship age dummies. 95% confidence intervals calculated using standard errors clustered at the player level.

would imply reductions in hazard ratios over the whole span of the insurance contract. In Figures 2.5 and 2.6 I present the results from the estimation with age dummies for the first 7 days of the ship and a separate dummy for ships older than the cross-sectional median of 180 days. The effect of the first days of insurance on the hazard ratio drops from about 20 to 13. However, compared to the model without the age dummies, the remainder of the graph stays at a considerably higher level. For example, for the last 40 days of the contract the hazard ratio with age dummies is consistently above 2 compared to the earlier range between 1.5 and 2 in the model without age dummies.[7] Because adding the dummies reduces the hazard ratio only in the very beginning of the contract, the presence of ship-specific adverse selection is unlikely.

[7] Alternative specifications with differently paced age dummies produced qualitatively very similar results.

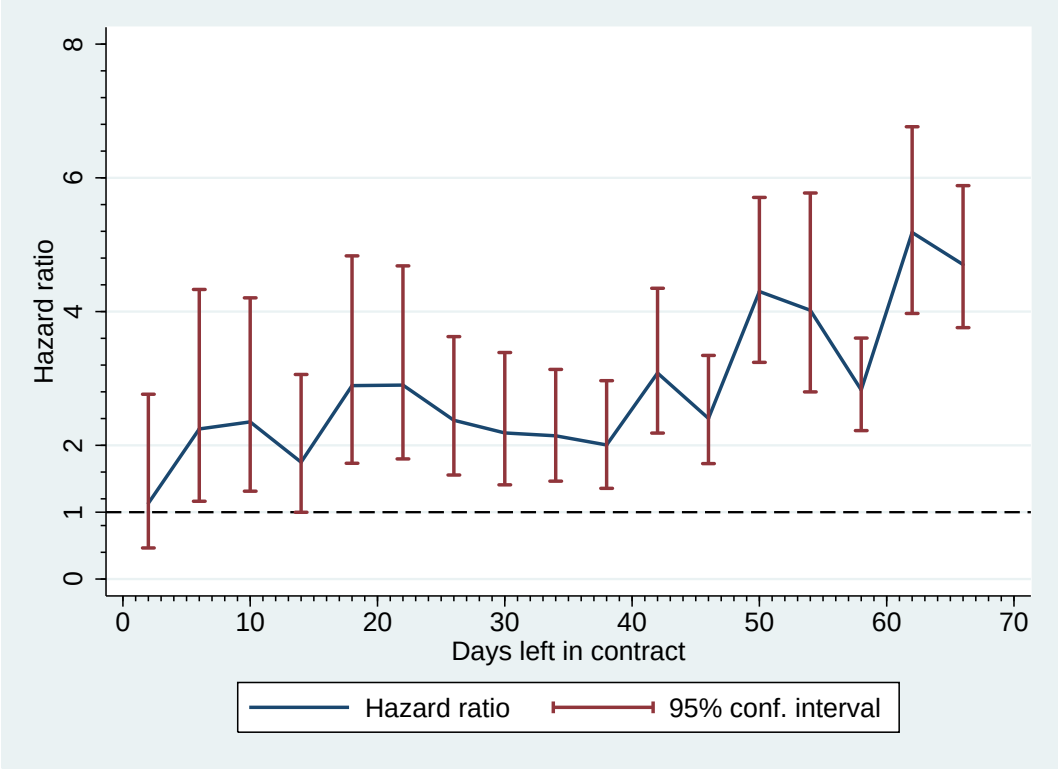


Figure 2.6: Hazard ratio as a function of time left in the insurance contract excluding first 8 days of the contract. Model includes a user specific baseline hazard rate and ship age dummies. 95% confidence intervals calculated using standard errors clustered at the player level.

The age dummies themselves are statistically significant. The first day after a ship is being unpacked is especially risky with a hazard ratio of 15 while ships older than 180 days are fairly safe with a hazard ratio of about 0.38. This pattern is consistent with the attention based story, because the act of unpacking a ship is a likely indicator of an intention to use it. Attention would also explain why adding age dummies had a negligible effect on the stratified estimates in the first chapter, while the age dummies themselves are statistically significant: Old ships are less risky only because they are more likely to be forgotten. Once an agent insures an old ship, she is fully aware of its existence. Consequently, she will treat it just like a recently unpacked, insured ship.

Both, Figure 2.4 and 2.6, show a considerable increase in risk when about 20 days of the contract is left. According to the point estimate from the model with the ship age dummies, when a ship has 24 days left in its contract its hazard rate is about 45% higher than the hazard rate of a ship that has 40 days left in its contract. In the model without the ship age dummies this gap is about 36%. When the model is estimated separately for Drakes, the most popular ship in the data, the spike in risk for ships with 20-24 days left in their contracts is even more pronounced. Looking at Figures 2.7 and 2.8, one can see that the hazard ratio for these days is over 5 times as high as the hazard ratio for the ships with 44-48 days left in their contract. This effect could be taken as partial evidence for dynamic moral hazard increasing as the contract comes close to its end. However, the accuracy of these estimates is so poor that I cannot completely rule out the gap being purely a result of random variation. Moreover, the fall in risk over the last 20 days of the contract is difficult to reconcile with the theory.

2.3.3 Experience and Timing of Moral Hazard

As was seen in the first chapter, experience and learning can strongly influence the level of moral hazard. If it also affects timing of effort choices over the contract period, experienced and inexperienced agents should be offered contracts with different dynamic incentives.

I next apply the method from above separately to the bottom and top 25% of players ranked by the hours they have spent logged into the game. As was discussed in the first chapter this is likely to be good indicator of how much experience the player has with the game. To gain enough within person variation in both insurance contracts and accidents to accurately estimate the model, I split the insurance contracts into 12-day periods instead of the previous 4-day periods. The results are presented in Figures 2.9 and 2.10.

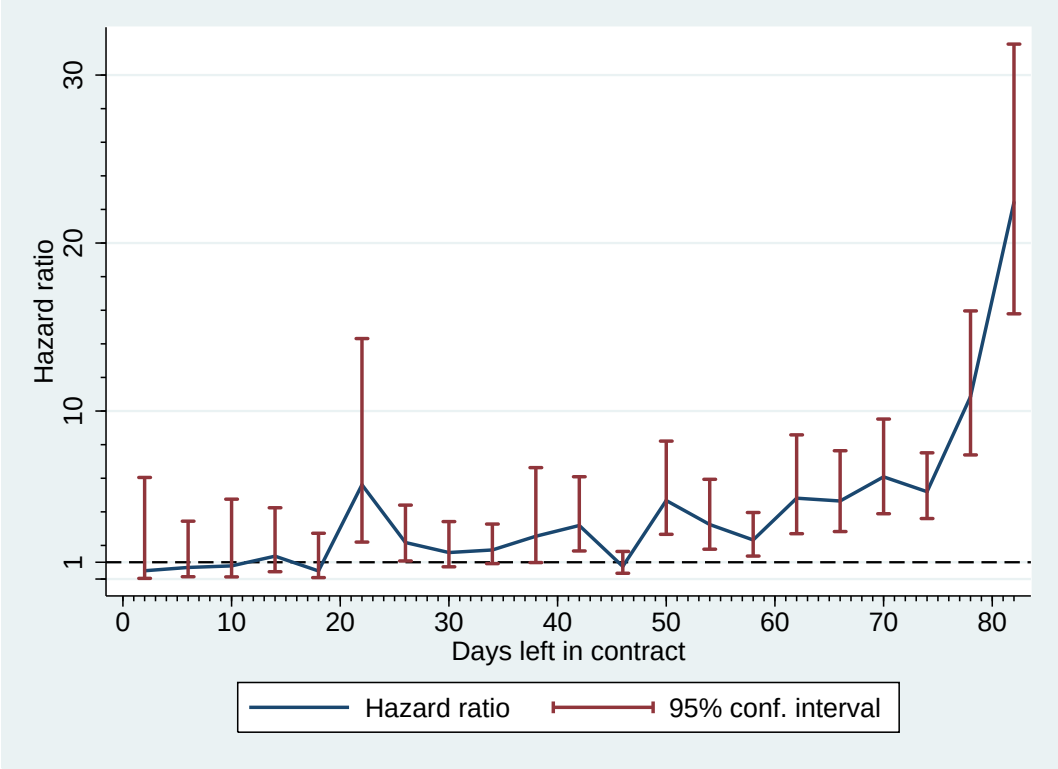


Figure 2.7: Drakes' hazard ratios as a function of time left in the insurance contract. Model includes a user specific baseline hazard rate and ship age dummies. 95% confidence intervals calculated using standard errors clustered at the player level.

The first interesting pattern from Figure 2.9 is that the insured ships of the least experienced players are statistically significantly more risky than uninsured ships only in the beginning of the insurance contract. After 12 days of the contract has passed, insured ships have hazard rates that are indistinguishable from uninsured ships. The same is not true for experienced players for whom the hazard ratio between insured and uninsured ships is clearly above one for the whole contract period. Hence, the inexperienced players seem to be planning their risky activities only for a very short horizon.

As can be seen in Figure 2.10, the spike in risk in the beginning of the contract is even higher for the experienced players. The daily hazard rate of a ship that was insured less than 12 days ago is almost 25 times as high as the hazard rate of an uninsured ship and about 8 times as high as the hazard rate of a ship that is close to

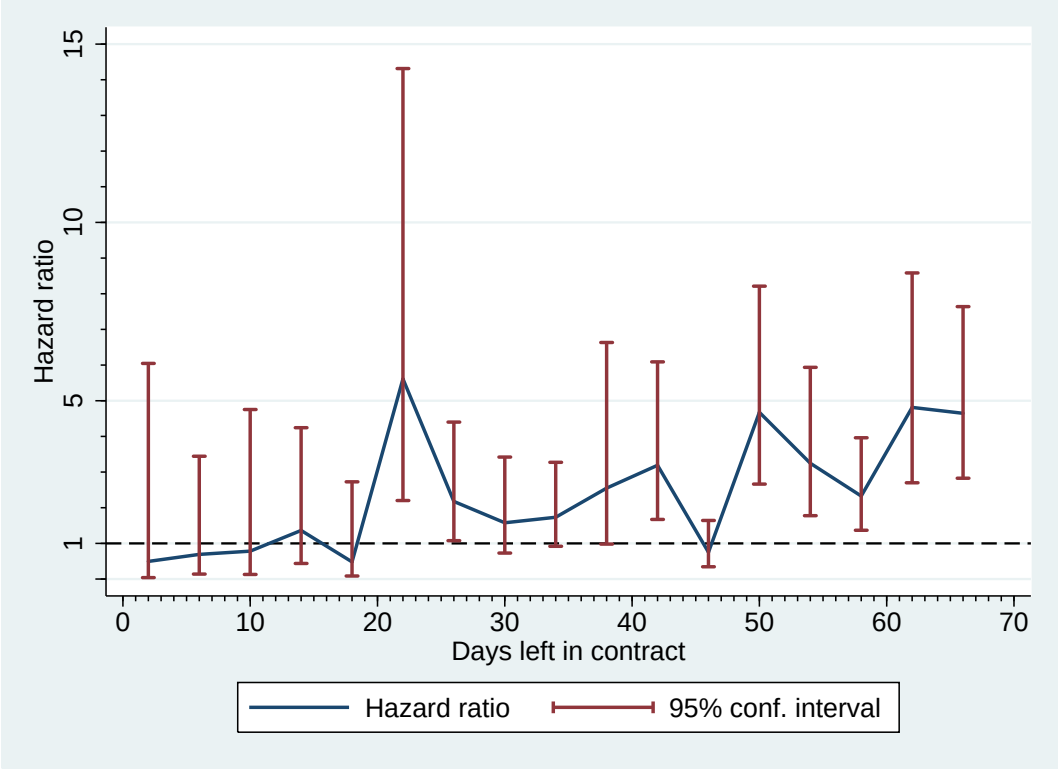


Figure 2.8: Drake's hazard ratios as a function of time left in the insurance contract excluding first 8 days of the contract. Model includes a user specific baseline hazard rate and ship age dummies. 95% confidence intervals calculated using standard errors clustered at the player level.

the end of its insurance contract. This is consistent with experienced players better knowing when they need insurance. Notice also that after the first 12 days, the hazard ratio settles to a relatively stable level between 2 and 3. Hence, it seems that these players do not find dynamic optimization of their precautionary effort worthwhile but just choose to use their insured ships equally more over the remaining 72 days of the contract. In contrast to the inexperienced agents, these players exhibit strong moral hazard over the whole span of the insurance contract but do not seem to adjust their effort choices after the first 12-24 days.

The results have interesting implications for contract design, if similar patterns hold even in real world insurance. If purchase of insurance is a strong predictor of increased risk for the short term but not for the long horizon, and if this effect is

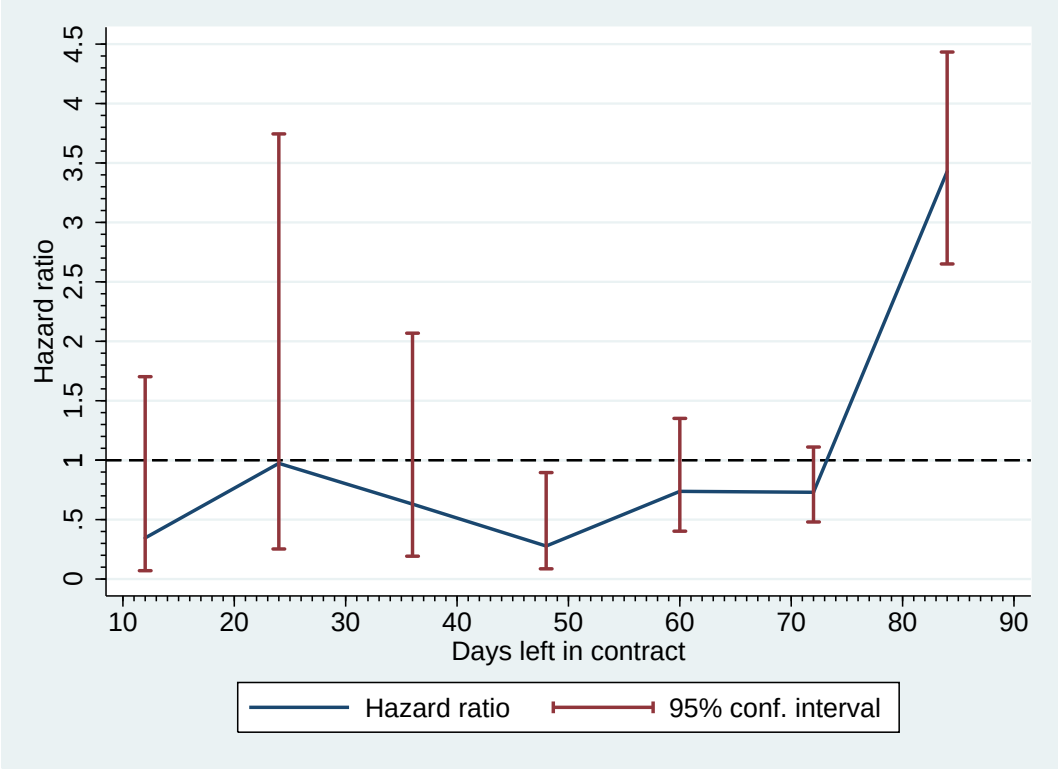


Figure 2.9: Hazard ratio as a function of time left in the insurance contract for the *least* experienced quartile of players. Model includes a user specific baseline hazard rate. 95% confidence intervals calculated using standard errors clustered at the player level.

due to effort choices, insurance contracts should have higher deductibles for accidents that happen early in the contract. Furthermore, if there are inexperienced agents who exhibit moral hazard only in the beginning of the contract, then these agents could be potentially nearly fully insured for accidents that happen later in the contract.

2.4 Conclusions

Better measurement of moral hazard and how people respond to incentives that change over time may translate into more efficient contracts. Interesting dynamic patterns, such as the high initial risk in the beginning of the contract, as presented in this paper, could potentially be used to generate slack in the more binding constraints of the contract design problem. Hence, a better understanding of these dynamic pat-

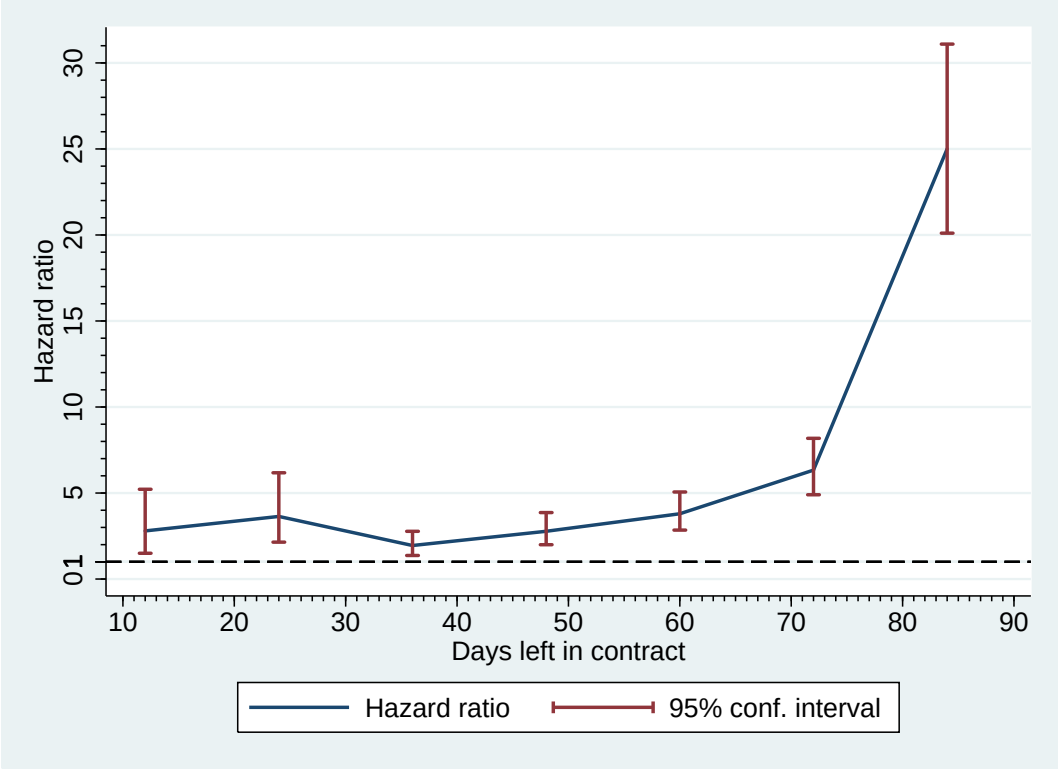


Figure 2.10: Hazard ratio as a function of time left in the insurance contract for the *most* experienced quartile. Model includes a user specific baseline hazard rate. 95% confidence intervals calculated using standard errors clustered at the player level.

terns could lead to contracts that are able to reduce information rents and yield socially more desirable outcomes. This paper has introduced methods that hopefully help us to achieve these goals.

More specifically, it would interesting to learn whether the front-loading of risk to the beginning of the contract period is a more ubiquitous phenomenon or whether it is special to the insurance contracts in EVE. Similarly, documenting the effect of experience on moral hazard in other insurance contracts can have high social returns. If inexperienced contracting agents exhibit less moral hazard also in other institutional environments, it could be beneficial for both society and the insurance companies to sell young or otherwise inexperienced agents contracts that cover a higher fraction of all losses.

Applying my methods to short-term labor contracting is another interesting avenue of future research. Measuring how output and the risk of premature termination of the contract varies over the span of a temporary work contract and how the dynamic variation in output is affected by the timing of wage payments can be especially important if the role of temporary and flexible work continues to grow in modern societies.

Last, my econometric method for measuring the changes in hazard rates over the course of the contract is not specific to EVE's institutional setup and could be applied to measure dynamic variation in risk even when the intertemporal incentives for risk taking are completely different from the ones induced by EVE's insurance contracts. Tolvanen (in progress) uses a similar method to identify dynamic incentives for moral hazard in Nordic car insurance. In some of the insurance company's contracts the owner of a new car is guaranteed an identical replacement vehicle during the first 3 years of the car in case the original car is totaled. Because the car still depreciates during the first 3 years of the car, the contract structure should imply decreasing effort over the first 3 years. Furthermore, since after the first 3 years the depreciation is taken discontinuously into account when valuing the totaled car, there should be a discontinuous increase in effort after the third year. This identification strategy, even though being different on the surface shares a strong similarity with the dynamic identification strategy in this paper and similar econometric modeling can be used to measure the dynamic changes in risk induced by the structure of these contracts.

2.A Omitted Proofs

The following well-known lemma on decreasing differences of a concave function is useful throughout the proofs:

Lemma 2.A.1. *Let $u\colon \mathbb{R} \to \mathbb{R}$ be differentiable, increasing and strictly concave. If $z > 0$, then*

$$\frac{\mathrm{d}\left(u(x) - u(x-z)\right)}{\mathrm{d}x} < 0$$

for all $x \in \mathbb{R}$.

Proof. By direct calculation,

$$\frac{\mathrm{d}\left(u(x) - u(x-z)\right)}{\mathrm{d}x} = u'(x) - u'(x-z).$$

Since u is strictly concave, u' is decreasing. This in turn implies that $u'(x-z) > u'(x)$ proving the lemma. □

2.A.1 Proofs for Section 2.2.1

I will start by showing that the agent's maximization problem has a recursive structure with respect to $I'_t = \max\{0, (1 - D_t)I_t - 1\}$. For this end I will need some auxiliary definitions. Let

$$U_j(x, I, D) = u(\omega - D(L - \mathbb{1}\{I\}P) - \mathbb{1}\{I = l\}C) - E(x, \varepsilon_j).$$

Then the agent's maximization problem in this notation is

$$\sup_{(x_j, I_j) \in C} \mathbb{E}\left[\sum_{t=0}^{\infty} \delta^t U(x_j(h_t), I_j(h_t), D_{t+1}) \mid x_j, I_j, h_0\right],$$

where

$$\begin{aligned} C \;:=\; & \{(x_j, I_j)\colon \\ & \mathcal{H} \mapsto \mathbb{R}_+ \times \{1, \ldots, l\}, I_j(h_t) \in \{l, \max\{0, (1 - D_{jt-1})I_{jt-1} - 1\}\}\, \forall h_t \in \mathcal{H}\} \end{aligned}$$

Notice that since

$$\lim_{x\to\infty} U_j(x, I, D) = -\infty$$

and $U_j(0, 0, 1) > -\infty$, there exists $M > 0$ such that choosing the controls from

$$\begin{aligned}\hat{C} \;:=\; & \{(x_j, I_j):\\ & \mathcal{H} \mapsto [0, M] \times \{1, \ldots, l\}, I_j(h_t) \in \{l, \max\{0, (1 - D_{jt-1})I_{jt-1} - 1\}\}\, \forall h_t \in \mathcal{H}\}\end{aligned}$$

will yield the same supremum as choosing them from C. If one endows $\{0, 1, \ldots, l\}$ with its discreet topology, then $\hat{C}$ is clearly compact in the product topology by the Tychonoff's Theorem.

For a given policy x_j, I_j, let $W(x_j, I_j, h_t)$ be the continuation utility from history h_s onwards:

$$W(x_j, I_j, h_s) := \mathbb{E}\left[\sum_{t=s}^{\infty} \delta^{t-s} U(x_j(h_t), I_j(h_t), D_{t+1}) \mid x_j, I_j, h_s\right]$$

Since both p and E are continuous in x, one can also show that $W(\cdot, \cdot, h_0)$ is continuous for each h_0. Consequently, $\max_{(x_j, I_j)\in C} W(x_j, I_j, h_0)$ is well defined.

Denote an arbitrary outcome in period t by $\theta_{jt} = (x_{jt}, I_{jt}, D_{jt})$. Then I can write an arbitrary history $h_t = (\theta_{jt})_{s=0}^{t-1}$. For any $s < u$ I will denote the partial history $(\theta_{jr})_{r=s}^{u-1}$ of the full history $h_t = (\theta_r)_{r=0}^{t-1}$ by h_s^u. Furthermore, for any two histories $h_t = (\theta_{jr})_{r=0}^{t-1}$ and $h'_s = (\theta'_{jr})_{r=0}^{s-1}$ and for any $0 \geq a < b \leq t$ and $0 < c < d \leq s$ I will define the combined history

$$h_a^b \otimes h_c^{'d} := (h_a^b, h_c^{'d}) = (\theta_{ja}, \theta_{ja+1}, \ldots, \theta_{jb-1}, \theta'_{jc}, \theta'_{jc+1}, \ldots, \theta'_{jd-1}).$$

The following lemma shows that for any two histories that share the same amount of insurance in the beginning of the period, the supremum of the continuation payoffs must be the same.

Lemma 2.A.2. *Let h_t and h'_s be such that* $\max\{0,(1-D_{jt-1})I_{t-1}-1\} = \max\{0,(1-D'_{js-1})I'_{s-1}-1\}$. *Then*

$$\max_{(x_j,I_j)\in C} W(x_j,I_j,h_t) = \max_{(x_j,I_j)\in C} W(x_j,I_j,h'_s).$$

Proof. Let (x_j,I_j) be an arbitrary policy that satisfies

$$I_j(h_t) \in \{l, \max\{0,(1-D_{jt-1})I_{t-1}\}\}$$

for every h_t. Partition $\mathcal{H}$ by setting

$$\mathcal{H} \mid h'_s = \left\{h_r \colon \exists h_q \in \mathcal{H} \text{ such that } \; h_r = h_0^{'s} \otimes h_0^q\right\}$$

Then define (x'_j, I'_j) by setting $(x'_j(h_r), I'_j(h_r)) = (x_j(h_r), I_j(h_r))$ if $h_r \notin \mathcal{H} \mid h'_s$. If $h_r = h_0^{'s} \otimes h_0^q$, then set

$$(x'_j(h_r), I'_j(h_r)) = (x_j(h_0^t \otimes h_0^q), I_j(h_0^t \otimes h_0^q)).$$

Notice that this policy is trivially an element of C, because (x_j,I_j) is an element in it *and* the same choices are available in C after histories h'_s and h_t by assumption. Notice then that for any $(x,I,D) \in \mathbb{R}_+ \times \{0,\ldots,l\} \times \{0,1\}$

$$\mathbb{P}\left[(x_{js},I_{js},D_{js}) = (x,I,D) \mid x'_j, I'_j, h'_s\right] = \mathbb{P}\left[(x_{jt},I_{jt},D_{jt}) = (x,I,D) \mid x_j, I_j, h_t\right],$$

since both are equal to $Dp(x) + (1-D)(1-p(x))$, if $x = x_j(h_t)$ and $I = I_j(h_t)$, and zero otherwise. Then an easy induction step can be used to establish that more generally, for any $h_q \in \mathcal{H}$

$$\mathbb{P}\left[h_0^{'s} \otimes h_0^q \mid x'_j, I'_j, h'_s\right] = \mathbb{P}\left[h_0^t \otimes h_0^q \mid x_j, I_j, h_t\right].$$

Since the utility stream U depends in any arbitrary period r only on θ_r, we must have that

$$\begin{aligned} W(x_j, I_j, h_t) &= \mathbb{E}\left[\sum_{r=t}^{\infty} \delta^{r-t} U(x_j(h_r), I_j(h_r), D_{r+1}) \mid x_j, I_j, h_t\right] \\ &= \mathbb{E}\left[\sum_{r=s}^{\infty} \delta^{r-s} U(x'_j(h_r), I'_j(h_r), D_{r+1}) \mid x'_j, I'_j, h'_s\right] \\ &= W(x'_j, I'_j, h'_s) \leq \max_{(\hat{x}_j, \hat{I}_j) \in C} W(\hat{x}_j, \hat{I}_j, h'_s), \end{aligned} \tag{2.A.1}$$

Because (x_j, I_j) was arbitrary element of C, it must be the case that

$$\max_{(\hat{x}_j, \hat{I}_j) \in C} W(\hat{x}_j, \hat{I}_j, h_t) \leq \max_{(\hat{x}_j, \hat{I}_j) \in C} W(\hat{x}_j, \hat{I}_j, h'_s).$$

As h'_s and h_t were arbitrary histories sharing the same amount of time left in the contract, one can switch the roles of h'_s and h_t in the argument above to obtain

$$\max_{(\hat{x}_j, \hat{I}_j) \in C} W(\hat{x}_j, \hat{I}_j, h_t) \geq \max_{(\hat{x}_j, \hat{I}_j) \in C} W(\hat{x}_j, \hat{I}_j, h'_s),$$

which proves the lemma. □

Applying this lemma to the case of initial history yields the following corollary:

Corollary 2.A.3. *If h_0 and h'_t are such that*

$$\max\{0, (1 - D_{j0})I_{j0} - 1\} = \max\{0, (1 - D'_{jt})I'_{jt} - 1\},$$

then $V_j(h_0) = \max_{(x_j, I_j) \in C} W(x_j, I_j, h'_t)$.

Furthermore, with slight abuse of notation, I can now write $V(I'_0)$ instead of $V(h_0)$ where $I'_0 = \max\{0, (1 - D_0)I_0 - 1\}$.

Then one can apply standard verification arguments such as those given for Propositions 2.2 and 2.3 in Bertsekas (1995) to argue that the optimal value function is given by 2.2.2 and that optimal policies can be chosen to depend only on $I'_t = \max\{0, (1 - D_t)I_t - 1\}$, if such policies exist that solve 2.2.2. It is easy to see that the maximization problem 2.2.2 is concave and maximizing functions $x_j(I')$ and $I_j(I')$ exist. I will prove these details for the more general case with 2 items in the following section. The next lemma shows that more time left in the insurance contract is valuable:

Lemma 2.A.4. *$V(I', \varepsilon) \geq V(I'', \varepsilon)$ for any $I' > I''$ and $\varepsilon \in \mathbb{R}$.*

Proof. This proof is a simpler version of the proof of Lemma 2.A.9. Since the result does not depend on ε, I will opt for clearer notation and write all functions as if they did not depend on ε. Assume first that the optimal choice of insurance level in state I'' is l. Then,

$$\begin{aligned} V(I') &= p(x(I'))\left[u\left(\omega - L + \mathbb{1}\left\{I(I') > 0\right\} P - \mathbb{1}\left\{I(I') = l\right\} C\right) + \delta V(0)\right] \\ &\quad + (1 - p(x(I')))\left[u(\omega - \mathbb{1}\left\{I(I') = l\right\} C) + \delta V(\max\{0, I(I') - 1\})\right] - E(x(I')) \\ &\geq p(x(I''))\left[u\left(\omega - L + P - C\right) + \delta V(0)\right] \\ &\quad + (1 - p(x(I'')))\left[u(\omega - C) + \delta V(l - 1)\right] - E(x(I'')), \\ &= V(I'') \end{aligned} \tag{2.A.2}$$

where the inequality follows since $x(I'')$ and l are available also when the state is I'. I will show the case where $I(I'') = I''$ using induction. Assume first that $I'' = 0$. Then

$$\begin{aligned} V(0) &= p(x(0))\left[u(\omega - L) + \delta V(0)\right] \\ &\quad +(1 - p(x(0)))\left[u(\omega) + \delta V(0)\right] - E(x(0)) \\ &< p(x(0))\left[u(\omega - L + P) + \delta V(0)\right] \\ &\quad +(1 - p(x(0)))\left[u(\omega) + \delta V(0)\right] - E(x(0)) \\ &\leq V(1), \end{aligned} \tag{2.A.3}$$

where the first inequality follows, since u is strictly increasing and $p(x(0)) > 0$. The second follows, because both $x(0)$ and $I(1) = 1$ are available policies when the state is 1. Now, for the induction step, assume that $V(k) \geq V(k-1)$ for $1 \leq k < l - 1$. It is enough to show that if $I(k) = k$, then $V(k+1) \geq V(k)$. Using estimates like those above I get that

$$\begin{aligned} V(k) &= p(x(k))\left[u(\omega - L + P) + \delta V(0)\right] \\ &\quad +(1 - p(x(k)))\left[u(\omega) + \delta V(k-1)\right] - E(x(k)) \\ &\leq p(x(k))\left[u(\omega - L + P) + \delta V(0)\right] \\ &\quad +(1 - p(x(k)))\left[u(\omega) + \delta V(k)\right] - E(x(k)) \\ &\leq V(k+1), \end{aligned}$$

where the first inequality follows from the induction assumption and the second from the fact that $x(k)$ and $I(k+1) = k+1$ are available in the state $k+1$. This completes the induction step.

Notice also that $V(I', \varepsilon) > V(I'', \varepsilon)$, if I can show that it is never optimal to buy a new contract when there is still time left in the old one. The result follows, because in that case (2.A.2) will be a strict inequality. Since (2.A.3) is strict, the whole induction

argument can be argued using strict inequalities. The next lemma shows that it is indeed true that renewing contracts when there is still time left is never optimal. □

Lemma 2.A.5. *If $I' > 0$, then $I(I') = I'$.*

Proof. Define the following set of numbers:

$$\begin{aligned} W_l &= \max_x \big\{ p(x) \left(u(\omega - L + P - C) + \delta V(0) \right) \\ &\quad + (1 - p(x)) \left(u(\omega - C) + \delta V(l-1) \right) - E(x) \big\} \end{aligned} \tag{2.A.4}$$

and

$$\begin{aligned} W_I &= \max_x \big\{ p(x) \left(u(\omega - L + P) + \delta V(0) \right) \\ &\quad + (1 - p(x)) \left(u(\omega) + \delta V(I-1) \right) - E(x) \big\} \end{aligned} \tag{2.A.5}$$

where $I \in \{1, \ldots, l-1\}$. It is enough to show that $W_I > W_l$ for all $I \in \{1, \ldots, l-1\}$. Let $x(l)$ and $x(I)$ be maximizers for the problems (2.A.4) and (2.A.5), respectively. The proof is by induction over k, for W_{l-k}. Assume first that $k = 1$, $l - 1 > 0$ and assume by way of contradiction that $W_l \geq W_{l-1}$. This directly implies that $V(l-1) = W_l$. On the other hand,

$$\begin{aligned} V(l-1) &\geq W_{l-1} \\ &\geq p(x(l))u(\omega - L + P) + (1 - p(x(l)))u(\omega) \\ &\quad + p(x(l))\delta V(0) + (1 - p(x(l)))\delta W_l - E(x(l)) \\ &> p(x(l))u(\omega - L + P - C) + (1 - p(x(l)))u(\omega - C) \\ &\quad + p(x(l))\delta V(0) + (1 - p(x(l)))\delta V(l-1) - E(x(l)) \\ &= W_l, \end{aligned}$$

which clearly contradicts $V(l-1) = W_l$. The second inequality follows by noticing that $x(l)$ is available also when the state is $l-1$ and the agent is always free to re-insure in the state $l-2$ which will yield her W_l. The third inequality results from u being strictly increasing.

Assume next that $W_{l-k} > W_l$ for some $k \in \{1, \ldots, l-1\}$ and $l-k-1 > 0$. Assume again by way of contradiction that $W_l \geq W_{l-k-1}$. In other words, $V(l-k-1) = W_l$. Now, arguments identical to those above yield that

$$\begin{aligned} V(l-k-1) &\geq W_{l-k-1} \\ &\geq p(x(l-k))u(\omega - L + P) + (1 - p(x(l-k)))u(\omega) \\ &\quad + p(x(l-k))\delta V(0) + (1 - p(x(l-k)))\delta W_l - E(x(l-k)) \\ &= p(x(l-k))u(\omega - L + P) + (1 - p(x(l-k)))u(\omega) \\ &\quad + p(x(l-k))\delta V(0) + (1 - p(x(l-k)))\delta V(l-k-1) - E(x(l-k)) \\ &= W_{l-k} > W_l, \end{aligned}$$

where the second inequality follows, since $x(l-k)$ is available when the state is $l-k-1$ and the agent can re-insure in state $l-k-2$ if she gets that far without an accident. The last inequality follows from the induction assumption. The result contradicts $V(l-k-1) = W_l$ and hence completes the proof. □

By the arguments given at the end of the proof for Lemma 2.A.4, I now get the following corollary:

Corollary 2.A.6. *$V(I', \varepsilon) > V(I'', \varepsilon)$ for any $I' > I''$ and $\varepsilon \in \mathbb{R}$.*

With these results I have the ingredients for the key result of the section:

Proof of Lemma 2.2.1. Let first $I(I') = I' < l$. Then the optimal $x(I')$ solves the first-order condition:

$$\begin{aligned}\frac{E'(x(I'))}{p'(x(I'))} &= (u(\omega - L + P) - u(\omega) + \delta V(0) - \delta V(I' - 1)) \\ &< (u(\omega - L + P) - u(\omega) + \delta V(0) - \delta V(I'' - 1)) \\ &= \frac{E'(x(I''))}{p'(x(I''))},\end{aligned}$$

where the inequality follows from the corollary above and the last equality from the first-order condition for $x(I'')$. Notice then that

$$\frac{\mathrm{d}}{\mathrm{d}x}\frac{E'(x)}{p'(x)} = \frac{E''(x)p'(x) - p''(x)E'(x)}{p'(x)^2} < 0$$

which implies that $x(I'') < x(I')$ proving the first part.

Assume then that $I(I') = l$. The first-order condition for $x(I')$ now solves:

$$\frac{E'(x(I'))}{p'(x(I'))} = (u(\omega - L + P - C) - u(\omega - C) + \delta V(0) - \delta V(I' - 1)).$$

Hence, the same proof as above will work as long as I can show that

$$u(\omega - L + P - C) - u(\omega - C) < u(\omega - L + P) - u(\omega).$$

Notice then that this inequality can be written as

$$u(\omega' + P - L) - u(\omega') < u(\omega + P - L) - u(\omega),$$

where

$$\omega' := \omega - C < \omega.$$

But then the inequality follows the concavity of u by Lemma 2.A.1, as $P - L < 0$. □

The previous lemma implies the trivial result that if only one risk type buys insurance, then more time left in the contract implies lower expected risk:

Corollary 2.A.7. *Assume that the contract is such that if $I_j > 0$ and $I_k > 0$ then $\varepsilon_j = \varepsilon_k =: \hat{\varepsilon}$. In that case $I(I', \varepsilon) > I(I'', \varepsilon) > 0$ implies that*

$$p(x(I'), \hat{\varepsilon}) = \mathbb{E}[p(x, \varepsilon) \mid I(I', \varepsilon)] < \mathbb{E}[p(x, \varepsilon) \mid I(I'', \varepsilon)] = p(x(I''), \hat{\varepsilon}).$$

The next result shows that heterogeneity of risk types in the pool of insured has the opposite effect, because less risky agents are more likely to survive longer into the contract causing a selection effect.

Lemma 2.A.8. *Consider the model without moral hazard. In other words, assume that $p(x, \varepsilon) = p(x', \varepsilon) \equiv p(\varepsilon)$ for all x. Then, if $I > I' > 0$,*

$$\mathbb{E}[p(\varepsilon) \mid I] \geq \mathbb{E}[p(\varepsilon) \mid I'].$$

Proof. Let $F_\varepsilon(e \mid I)$ be the cumulative distribution function of ε conditional on I and let $I > 0$. I will show that $F_\varepsilon(e \mid I + 1)$ first-order stochastically dominates $F_\varepsilon(e \mid I)$. The result then follows from the fact that $p(\varepsilon)$ is increasing.

For notational convenience I will reintroduce the time period into the notation. In what follows $I = I_t$ is the time left in the insurance contract at the beginning of time period t and $D_t \in \{0, 1\}$, where $D_t = 1$ is interpreted as the item getting destroyed in the end of period t.

Now,

$$\begin{aligned}
F_\varepsilon(e \mid I_t = I) &= \mathbb{P}(\varepsilon \leq e \mid I_{t-1} = I+1, D_{t-1} = 0) \\
&= \frac{\mathbb{P}(\varepsilon \leq e, D_{t-1} = 0 \mid I_{t-1} = I+1)}{\mathbb{P}(D_{t-1} = 0 \mid I_{t-1} = I+1)} \\
&= \frac{\int_{-\infty}^{e} 1 - p(z)\, \mathrm{d}F_\varepsilon(z \mid I_{t-1} = I+1)}{\int_{-\infty}^{\infty} 1 - p(z)\, \mathrm{d}F_\varepsilon(z \mid I_{t-1} = I+1)} \\
&= F_\varepsilon(e \mid I_{t-1} = I+1) \frac{1 - \frac{1}{F_\varepsilon(e|I_{t-1}=I+1)} \int_{-\infty}^{e} p(z)\, \mathrm{d}F_\varepsilon(z \mid I_{t-1} = I+1)}{1 - \mathbb{E}\left[p(\varepsilon) \mid I_{t-1} = I+1\right]} \\
&= F_\varepsilon(e \mid I_{t-1} = I+1) \frac{1 - \mathbb{E}\left[p(\varepsilon) \mid \varepsilon \leq e, I_{t-1} = I+1\right]}{1 - \mathbb{E}\left[p(\varepsilon) \mid I_{t-1} = I+1\right]} \\
&\geq F_\varepsilon(e \mid I_{t-1} = I+1), \qquad (2.A.6)
\end{aligned}$$

where the inequality follows since, p is increasing and hence

$$\mathbb{E}\left[p(\varepsilon) \mid \varepsilon \leq e, I_{t-1} = I+1\right] \leq \mathbb{E}\left[p(\varepsilon) \mid I_{t-1} = I+1\right].$$

□

2.A.2 Proofs for Section 2.2.2

It is straightforward to adapt the proofs from the section above to show that given a person with some e_j and ω_j, then policies that depend only on (t, I_1', I_2'), where $I_i' = \max\{0, (1 - D_i^t) I_i^t - 1\}$, the time left in insurance policy i when entering period t, can always do at least as well as fully history dependent policies. This follows from two facts: First, the available controls after any two histories that share (t, I_1', I_2') is the same. Second, given a triplet (t, I_1', I_2'), the probability measure over the tomorrow's history h^{t+1} depends only on today's choice of controls. Hence, for any two histories that share the same (t, I_1', I_2'), one can use an argument identical to the one above to show that the continuation value from those two histories must be the same. Just as above one can then apply the results from Bertsekas (1995) to show that the optimal

value of the problem satisfies the Bellman equation:

$$
\begin{aligned}
& V(t, I_1', I_2') \\
= & \sup_{x_1, x_2, \hat{I}_1, \hat{I}_2} \Big\{ p(x_1,t)p(x_2,t) \\
& \times \Big(u(\omega - 2L + (\mathbb{1}\{\hat{I}_1 > 0\} + \mathbb{1}\{\hat{I}_2 > 0\})P - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, 0, 0) \Big) \\
+ & \; p(x_1,t)(1 - p(x_2,t)) \Big(u(\omega - L + \mathbb{1}\{\hat{I}_1 > 0\}P - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, 0, \max\{\hat{I}_2 - 1, 0\}) \Big) \\
+ & \; (1 - p(x_1,t))p(x_2,t) \Big(u(\omega - L + \mathbb{1}\{\hat{I}_2 > 0\}P - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, \max\{\hat{I}_1 - 1, 0\}, 0) \Big) \\
+ & \; (1 - p(x_1,t))(1 - p(x_2,t)) \Big(u(\omega - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, \max\{\hat{I}_1 - 1, 0\}, \max\{\hat{I}_2 - 1, 0\}) \Big) - E(x_1, x_2) \Big\} \qquad \text{(2.A.7)}
\end{aligned}
$$

It is also straightforward to apply standard results, like those in Bertsekas (1995), to show that this functional equation has a unique solution in the space of bounded functions and that solution is the actual maximal value function. I will next show that there exists a policy that achieves the supremum.

I will start by showing the following result which shows that the agent always prefers more insurance to less insurance:

Lemma 2.A.9. *The following inequalities hold:*

1. $V(t, I_1', I_2') \geq V(t, I_1'', I_2')$ *for all* $I_1' > I_1''$

2. $V(t, I_1', I_2') \geq V(t, I_1', I_2'')$ *for all* $I_2' > I_2''$.

Proof. By the symmetry of the problem, it is clearly enough to prove just the first inequality. Since the set from which $(\hat{I}_1, \hat{I}_2)$ is chosen is finite, I can find $(\hat{I}_1, \hat{I}_2)$ and

a sequence $(x''_{1n}, x''_{2n})_{n\in\mathbb{N}}$ such that

$$
\begin{aligned}
& p(x''_{1n},t)p(x''_{2n},t) \\
& \times\Big(u(\omega-2L+(\mathbb{1}\{\hat{I}_1>0\}+\mathbb{1}\{\hat{I}_2>0\})P-(\mathbb{1}\{\hat{I}_1=l\}+\mathbb{1}\{\hat{I}_2=l\})C) \\
& +\delta V(t+1,0,0)\Big) \\
+\quad & p(x''_{1n},t)(1-p(x''_{2n},t))\Big(u(\omega-L+\mathbb{1}\{\hat{I}_1>0\}P-(\mathbb{1}\{\hat{I}_1=l\}+\mathbb{1}\{\hat{I}_2=l\})C) \\
& +\delta V(t+1,0,\max\{\hat{I}_2-1,0\})\Big) \\
+\quad & (1-p(x''_{1n},t))p(x''_{2n},t)\Big(u(\omega-L+\mathbb{1}\{\hat{I}_2>0\}P-(\mathbb{1}\{\hat{I}_1=l\}+\mathbb{1}\{\hat{I}_2=l\})C) \\
& +\delta V(t+1,\max\{\hat{I}_1-1,0\},0)\Big) \\
+\quad & (1-p(x''_{1n},t))(1-p(x''_{2n},t))\Big(u(\omega-(\mathbb{1}\{\hat{I}_1=l\}+\mathbb{1}\{\hat{I}_2=l\})C) \\
& +\delta V(t+1,\max\{\hat{I}_1-1,0\},\max\{\hat{I}_2-1,0\})\Big)-E(x''_{1n},x''_{2n})\underset{n\to\infty}{\longrightarrow}V(t,I''_1,I'_2).
\end{aligned}
$$

Consider first the case when $\hat{I}_1 = l$ achieves the supremum when the state is (t, I_1'', I_2'). Then I have:

$$
\begin{aligned}
& V(t, I_1', I_2') \\
= & \sup_{x_1, x_2, \hat{I}_1, \hat{I}_2} \Big\{ p(x_1, t) p(x_2, t) \\
& \times \Big(u(\omega - 2L + (\mathbb{1}\{\hat{I}_1 > 0\} + \mathbb{1}\{\hat{I}_2 > 0\})P - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, 0, 0) \Big) \\
+ & \; p(x_1, t)(1 - p(x_2, t)) \Big(u(\omega - L + \mathbb{1}\{\hat{I}_1 > 0\}P - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, 0, \max\{\hat{I}_2 - 1, 0\}) \Big) \\
+ & \; (1 - p(x_1, t)) p(x_2, t) \Big(u(\omega - L + \mathbb{1}\{\hat{I}_2 > 0\}P - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, \max\{\hat{I}_1 - 1, 0\}, 0) \Big) \\
+ & \; (1 - p(x_1, t))(1 - p(x_2, t)) \Big(u(\omega - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, \max\{\hat{I}_1 - 1, 0\}, \max\{\hat{I}_2 - 1, 0\}) \Big) - E(x_1, x_2) \Big\} \\
\geq & \; p(x_{1n}'', t) p(x_{2n}'', t) \Big(u(\omega - 2L + (1 + \mathbb{1}\{\hat{I}_2 > 0\})P - (1 + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, 0, 0) \Big) \\
+ & \; p(x_{1n}'', t)(1 - p(x_{2n}'', t)) \Big(u(\omega - L + P - (1 + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, 0, \max\{\hat{I}_2 - 1, 0\}) \Big) \\
+ & \; (1 - p(x_{1n}'', t)) p(x_{2n}'', t) \Big(u(\omega - L + \mathbb{1}\{\hat{I}_2 > 0\}P - (1 + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, l-1, 0) \Big) \\
+ & \; (1 - p(x_{1n}'', t))(1 - p(x_{2n}'', t)) \Big(u(\omega - (1 + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, l-1, \max\{\hat{I}_2 - 1, 0\}) \Big) - E(x_{1n}'', x_{2n}'')
\end{aligned}
\tag{2.A.8}
$$

for every $n \in \mathbb{N}$. The inequality follows, since $(x_{1n}'', x_{2n}'', l, \hat{I}_2)$ is also available when the state is (t, I_1', I_2'). Now the right-hand side converges to $V(t, I_1'', I_2')$ as $n \to \infty$

and hence the claim must hold when in the state (t, I_1'', I_2') the supremum is achieved with $\hat{I}_1 = l$.

To cover the case when the supremum is achieved at $\hat{I}_1 = I_1''$ when the state is (t, I_2'', I_2') I proceed by induction to show that $V(t, I_1'' + 1, I_2') \geq V(t, I_1'', I_1')$. This together with the result above clearly implies the whole lemma. Assume first that $I_1'' = 0$ and that the supremum is attained at $\hat{I}_1 = 0$ when the state is (t, I_1'', I_2'), i.e.

$$
\begin{aligned}
V(t,0,I_2') \underset{n\to\infty}{\longleftarrow} \; & p(x_{1n}'',t)p(x_{2n}'',t)\Big(u(\omega - 2L + \mathbb{1}\{\hat{I}_2 > 0\}P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,0)\Big) \\
+ \; & p(x_{1n}'',t)(1-p(x_{2n}'',t))\Big(u(\omega - L - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,\max\{\hat{I}_2 - 1, 0\})\Big) \\
+ \; & (1-p(x_{1n}'',t))p(x_{2n}'',t)\Big(u(\omega - L + P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,0)\Big) \\
+ \; & (1-p(x_{1n}'',t))(1-p(x_{2n}'',t))\Big(u(\omega - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,\max\{\hat{I}_2 - 1, 0\})\Big) - E(x_{1n}'', x_{2n}'') \\
\leq \; & p(x_{1n}'',t)p(x_{2n}'',t)\Big(u(\omega - 2L + (1 + \mathbb{1}\{\hat{I}_2 > 0\})P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,0)\Big) \\
+ \; & p(x_{1n}'',t)(1-p(x_{2n}'',t))\Big(u(\omega - L + P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,\max\{\hat{I}_2 - 1, 0\})\Big) \\
+ \; & (1-p(x_{1n}'',t))p(x_{2n}'',t)\Big(u(\omega - L + P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,0)\Big) \\
+ \; & (1-p(x_{1n}'',t))(1-p(x_{2n}'',t))\Big(u(\omega - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,\max\{\hat{I}_2 - 1, 0\})\Big) - E(x_{1n}'', x_{2n}''),
\end{aligned}
$$

where the inequality follows, because u is increasing. The right-hand side is exactly the value from choosing policy $(x_{1n}'', x_{2n}'', I_1' - 1, \hat{I}_2)$ when the state is $(t, I_1' = I_1'' + 1, I_2') =$

$(t, 1, I_2')$. Since this policy is a valid choice in that state, the value of that policy must be at most $V(t, 1, I_2')$. Now by the convergence on the first row, I know that for every $\varepsilon > 0$ there exists $n \in \mathbb{N}$ such that

$$\begin{aligned}
V(t,0,I_2') \leq\ & p(x_{1n}'',t)p(x_{2n}'',t)\Big(u(\omega - 2L + (1 + \mathbb{1}\{\hat{I}_2 > 0\})P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& + \delta V(t+1,0,0)\Big) \\
+\ & p(x_{1n}'',t)(1 - p(x_{2n}'',t))\Big(u(\omega - L + P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& + \delta V(t+1,0,\max\{\hat{I}_2 - 1, 0\})\Big) \\
+\ & (1 - p(x_{1n}'',t))p(x_{2n}'',t)\Big(u(\omega - L + P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& + \delta V(t+1,0,0)\Big) \\
+\ & (1 - p(x_{1n}'',t))(1 - p(x_{2n}'',t))\Big(u(\omega - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& + \delta V(t+1,0,\max\{\hat{I}_2 - 1, 0\})\Big) - E(x_{1n}'', x_{2n}'') + \varepsilon \\
\leq\ & V(t,1,I_2') + \varepsilon.
\end{aligned}$$

Taking $\varepsilon \to 0$ yields the first step of the induction.

Next, suppose that for any $t \in \mathbb{N}$ and $k \in \{1, \ldots, l-1\}$, $V(t,k,I_2') \geq V(t,k-1,I_2')$. I want to show that $V(t,k+1,I_2') \geq V(t,k,I_2')$. If $V(t,k,I_2')$ is a limit of a sequence of effort levels when, $\hat{I}_1 = l$ then our claim is proven by the first part of the lemma. Hence assume that optimal insurance choice for state (t,k,I_2') is k. Then for every

$\varepsilon > 0$ there exists $x_{1\varepsilon}$ and $x_{2\varepsilon}$ such that

$$\begin{aligned}
V(t,k,I_2') \;\leq\;& p(x_{1\varepsilon},t)p(x_{2\varepsilon},t)\Big(u(\omega - 2L + (1+\mathbb{1}\{\hat{I}_2>0\})P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,0)\Big) \\
+\;& p(x_{1\varepsilon},t)(1-p(x_{2\varepsilon},t))\Big(u(\omega - L + P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,\max\{\hat{I}_2 - 1,0\})\Big) \\
+\;& (1-p(x_{1\varepsilon},t))p(x_{2\varepsilon},t)\Big(u(\omega - L + \mathbb{1}\{\hat{I}_2>0\}P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,k-1,0)\Big) \\
+\;& (1-p(x_{1\varepsilon},t))(1-p(x_{2\varepsilon},t))\Big(u(\omega - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,k-1,\max\{\hat{I}_2 - 1,0\})\Big) - E(x_{1\varepsilon},x_{2\varepsilon}) + \varepsilon \\
\leq\;& p(x_{1\varepsilon},t)p(x_{2\varepsilon},t)\Big(u(\omega - 2L + (1+\mathbb{1}\{\hat{I}_2>0\})P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,0)\Big) \\
+\;& p(x_{1\varepsilon},t)(1-p(x_{2\varepsilon},t))\Big(u(\omega - L + P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,0,\max\{\hat{I}_2 - 1,0\})\Big) \\
+\;& (1-p(x_{1\varepsilon},t))p(x_{2\varepsilon},t)\Big(u(\omega - L + \mathbb{1}\{\hat{I}_2>0\}P - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,k,0)\Big) \\
+\;& (1-p(x_{1\varepsilon},t))(1-p(x_{2\varepsilon},t))\Big(u(\omega - \mathbb{1}\{\hat{I}_2 = l\}C) \\
& +\delta V(t+1,k,\max\{\hat{I}_2 - 1,0\})\Big) - E(x_{1\varepsilon},x_{2\varepsilon}) + \varepsilon \qquad (2.A.9)
\end{aligned}$$

where the inequality follows from the induction assumption. Now the right-hand side is ε plus the value from choosing $(x_{1\varepsilon}, x_{2\varepsilon}, k+1, \hat{I}_2)$ when the state is $(t, k+1, I_2')$. Since this option is available in that state, I get that $V(t,k,I_2') \leq V(t,k+1,I_2') + \varepsilon$ for all $\varepsilon > 0$. Taking the limit as $\varepsilon \to 0$ completes the induction. □

Given the value function V, define the value of choosing policy $(x_1, x_2, \hat{I}_1, \hat{I}_2)$ in period t when the insurance status is (I_1', I_2') as

$$
\begin{aligned}
& W(x_1, x_2, \hat{I}_1, \hat{I}_2; t, I_1', I_2') \\
:= \; & p(x_1, t)p(x_2, t) \\
& \times \Big(u(\omega - 2L + (\mathbb{1}\{\hat{I}_1 > 0\} + \mathbb{1}\{\hat{I}_2 > 0\})P - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, 0, 0) \Big) \\
+ \; & p(x_1, t)(1 - p(x_2, t)) \Big(u(\omega - L + \mathbb{1}\{\hat{I}_1 > 0\}P - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, 0, \max\{\hat{I}_2 - 1, 0\}) \Big) \\
+ \; & (1 - p(x_1, t))p(x_2, t) \Big(u(\omega - L + \mathbb{1}\{\hat{I}_2 > 0\}P - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, \max\{\hat{I}_1 - 1, 0\}, 0) \Big) \\
+ \; & (1 - p(x_1, t))(1 - p(x_2, t)) \Big(u(\omega - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\
& + \delta V(t+1, \max\{\hat{I}_1 - 1, 0\}, \max\{\hat{I}_2 - 1, 0\}) \Big) - E(x_1, x_2) \qquad (2.A.10)
\end{aligned}
$$

In what follows I will suppress the dependence of W on (t, I_1', I_2'). The next two lemmas show that there exists an optimal effort choice vector (x_1, x_2) and that effort choice is given by the first-order condition over (x_1, x_2) at some insurance choice $(\hat{I}_1, \hat{I}_2)$.

Lemma 2.A.10. *Given any choice of insurance levels $(\hat{I}_1, \hat{I}_2)$ (and any state (t, I_1', I_2')), the mapping $(x_1, x_2) \mapsto W(x_1, x_2, \hat{I}_1, \hat{I}_2)$ is strictly concave.*

Proof. Define $\hat{W}\colon (x_1, x_2) \mapsto W(x_1, x_2, \hat{I}_1, \hat{I}_2) + E(x_1, x_2)$. Because E is strictly convex, it is enough to show that $\hat{W}$ is concave. I will do this by proving that $\partial_{11}\hat{W}(x_1, x_2) < 0$ and

$$\partial_{11}\hat{W}(x_1, x_2)\partial_{22}\hat{W}(x_1, x_2) - \partial_{12}\hat{W}(x_1, x_2)^2 > 0.$$

Since the notation for the general case is nearly illegible, I will prove only the case where $\hat{I}_1 = l$ and $0 < \hat{I}_2 < l$. The argument for the other cases is identical. Notice first that

$$\partial_{11}\hat{W}(x_1, x_2) = p''(x_1, t)\Big[p(x_2, t)A + B + p(x_2, t)C_1 + (1 - p(x_2, t))D_1\Big],$$

where I write $p'(x, t)$ instead of $\partial_x p(x, t)$ and where

$$A := u(\omega - 2L + 2P - C) - 2u(\omega - L + P - C) + u(\omega - C),$$

$$B := u(\omega - L + P - C) - u(\omega - C),$$

$$C_1 := \delta\left[V(t + 1, 0, 0) - V(t + 1, l - 1, 0)\right] \tag{2.A.11}$$

and

$$D_1 := \delta\left[V(t + 1, 0, \hat{I}_2 - 1) - V(t + 1, l - 1, \hat{I}_2 - 1)\right]. \tag{2.A.12}$$

All of these terms are negative: the negativity of A follows from concavity of u, B is negative because u is strictly increasing and C_1 and D_1 are negative by the previous lemma. Hence, $\partial_{11}\hat{W}$ is negative.

Next, notice that

$$\partial_{22}\hat{W}(x_1, x_2) = p''(x_2, t)\Big[p(x_1, t)A + B + p(x_1, t)C_2 + (1 - p(x_1, t))D_2\Big],$$

where A and B were defined above and

$$C_2 := \delta \left[V(t+1,0,0) - V(t+1,0,\hat{I}_2 - 1) \right] \tag{2.A.13}$$

and

$$D_2 := \delta \left[V(t+1,l-1,0) - V(t+1,l-1,\hat{I}_2 - 1) \right] \tag{2.A.14}$$

which are also all negative. Hence, $\partial_{22}\hat{W}(x_1,x_2) < 0$. It is also easy to verify that $E := C_1 - D_1 = C_2 - D_2$. Hence, I have:

$$\begin{aligned} & \partial_{11}\hat{W}(x_1,x_2)\partial_{22}\hat{W}(x_1,x_2) \\ = \; & p''(x_1,t)p''(x_2,t)\left[p(x_2,t)(A+E)+B+D_1\right] \\ & \times \left[p(x_1,t)(A+E)+B+D_2\right] \\ = \; & p''(x_1,t)p''(x_2,t)\left|p(x_2,t)(A+E)+B+D_1\right| \\ & \times \left|p(x_1,t)(A+E)+B+D_2\right|. \end{aligned} \tag{2.A.15}$$

I already argued that $p(x_i,t)(A+E)+B+D_i < 0$ and that $B, D_i < 0$ for $i = 1,2$. This then implies that

$$(2.A.15) > p''(x_1,t)p''(x_2,t)p(x_1,t)p(x_2,t)\left(A+E\right)^2. \tag{2.A.16}$$

Next, consider the cross partial derivative:

$$\begin{aligned}
& \partial_{12} W(x_1, x_2) \\
= \; & p'(x_1, t) p'(x_2, t) \Bigg[u(\omega - 2L + 2P - C) - 2u(\omega - L + P - C) + u(\omega - C) \\
& + \delta \Big(V(t+1, 0, 0) - V(t+1, 0, \hat{I}_2 - 1) \\
& - V(t+1, l-1, 0) + V(t+1, l-1, \hat{I}_2 - 1) \Big) \Bigg] \\
= \; & p'(x_1, t) p'(x_2, t) \left(A + E \right) \qquad (2.A.17)
\end{aligned}$$

Now, if I combine (2.A.16) and (2.A.17), I get that

$$\begin{aligned}
& \partial_{11} W(x_1, x_2) \partial_{22} W(x_1, x_2) - \partial_{12} W(x_1, x_2)^2 \\
> \; & (A + E)^2 \left(p''(x_1, t) p''(x_2, t) p(x_1, t) p(x_2, t) - \left(p'(x_1, t) p'(x_2, t) \right)^2 \right) > 0,
\end{aligned}$$

where the inequality follows from the maintained assumption that $(x_1, x_2) \mapsto p(x_1, t)p(x_2, t)$ is convex. □

Lemma 2.A.11. *For each $(\hat{I}_1, \hat{I}_2)$ (and (t, I_1', I_2')) there exists a unique (x_1, x_2) that maximizes $W(x_1, x_2, \hat{I}_1, \hat{I}_2)$. Furthermore, this maximizer is the unique solution to*

$$\nabla_{(x_1, x_2)} W(x_1, x_2, \hat{I}_1, \hat{I}_2) = 0 \qquad (2.A.18)$$

Proof. For any $M > 0$, let $S_M = \{(x_1, x_2) \in \mathbb{R}^2_+ : x_1 \leq M, x_2 \leq M\}$. I start by showing that there exists $M > 0$ such that

$$\sup_{(x_1, x_2) \in \mathbb{R}^2_+} W(x_1, x_2, \hat{I}_1, \hat{I}_2) = \sup_{(x_1, x_2) \in S_M} W(x_1, x_2, \hat{I}_1, \hat{I}_2). \qquad (2.A.19)$$

This together with the compactness of S_M and continuity of W guarantees the existence of a maximum. To prove the equality, notice that

$$\begin{aligned} & W(x_1, x_2, \hat{I}_1, \hat{I}_2) \leq u(\omega - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\ & +\delta V(t+1, \max\{\hat{I}_1 - 1, 0\}, \max\{\hat{I}_2 - 1, 0\}) - E(x_1, x_2) \\ \leq \; & u(\omega - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\ & +\delta V(t+1, \max\{\hat{I}_1 - 1, 0\}, \max\{\hat{I}_2 - 1, 0\}) - E(0, x_2) \underset{x_2 \to \infty}{\to} -\infty, \end{aligned}$$

where the first inequality follows by Lemma 2.A.9 and because u is increasing, the second inequality follows because E is increasing in both of its arguments and the limit uses the fact that V is bounded and E is increasing and strictly convex and hence $\lim_{x_2 \to \infty} E(x_1, x_2) = \infty$ for all $x_1 \in \mathbb{R}_+$. Then there exists $M_2 > 0$ such that $W(x_1, x_2, \hat{I}_1, \hat{I}_2) < 0$ for any $x_1 \in \mathbb{R}_+$ and $x_2 \geq M_2$. Choosing any such x_2 is clearly suboptimal, since choosing $(0, 0)$ instead yields:

$$\begin{aligned} & W(0, 0, \hat{I}_1, \hat{I}_2) \\ \geq \; & u(\omega - 2L + (\mathbb{1}\{\hat{I}_1 > 0\} + \mathbb{1}\{\hat{I}_2 > 0\})P - (\mathbb{1}\{\hat{I}_1 = l\} + \mathbb{1}\{\hat{I}_2 = l\})C) \\ & +\delta V(t+1, 0, 0) - E(0, 0) > -\infty. \end{aligned}$$

A symmetrical argument shows that choosing a policy pair with $x_1 \geq M_1$ for some large enough M_1 is also suboptimal. Setting $M = \max\{M_1, M_2\}$ yields the equality (2.A.19).

Potential maximizers in the square S_M must be either on the boundary of S_M or in the interior. If I can rule out any maximizers on the boundary, then the strict concavity of W from Lemma 2.A.10 guarantees that there can be only one maximizer in the interior. Furthermore, the system of equations (2.A.18) is a necessary condition

for an interior optimum and hence ruling out maximizers on the boundary of S_M will prove the lemma.

Notice first that by the arguments above, there exists M large enough such that, $W(x_1, x_2) < -E(0,0)$ on the set

$$\{(x_1, x_2) \in \mathbb{R}^2_+ : x_1 = M, x_2 \leq M\} \cup \{(x_1, x_2) \in \mathbb{R}^2_+ : x_1 \leq M, x_2 = M\}.$$

Just like above, choosing $(0,0)$ yields a better result and hence I am left with only the lower part of the boundary given by

$$\{(x_1, x_2) \in \mathbb{R}^2_+ : x_1 = 0, x_2 < M\} \cup \{(x_1, x_2) \in \mathbb{R}^2_+ : x_2 = 0, x_1 < M\}.$$

By symmetry, it is enough to show that it is never optimal to set $x_1 = 0$. Because $\partial_1 E(0, x_2) = 0$, it is enough to show that $\partial_1 \hat{W}(0, x_2) > 0$, where $\hat{W}$ was defined in the proof of the previous lemma, since this would imply that $\partial_1 W(0, x_2, \hat{I}_1, \hat{I}_2) > 0$ which in turn is enough to show that setting $x_1 = 0$ is not optimal for any level of x_2. Again, the proof follows by checking all of the possible configurations of $(\hat{I}_1, \hat{I}_2)$. The steps behind all of the cases are the same and I again consider only the one where $(\hat{I}_1, \hat{I}_2) = (l, I_2')$ for some $I_2' > 0$, then it is easy to see that,

$$\partial_1 \hat{W}(x_1, x_2) = p'(x_1, t)\Big[p(x_2, t)A + B + p(x_2, t)C_1 + (1 - p(x_2, t))D_1\Big],$$

where A, B, C_1, D_1 are just as in the proof of Lemma 2.A.10. There I also argued that the term in square brackets is negative. Now the result simply follows from the assumption that p is strictly decreasing, so $\partial_1 \hat{W}(0, x_2) > 0$. □

The following result shows that at individual level effort levels are ordered the same way as is required by Proposition 4

Lemma 2.A.12. *If $0 < I_1' < I_2'$, and t, ε_j and ω_j are such that $I_{ij}(t, I_1', I_2') = I_i'$ for $i = 1, 2$ then*

$$p(x_1(t, I_1', I_2'), \varepsilon_j, t) \geq p(x_2(t, I_1', I_2'), \varepsilon_j, t)$$

The same inequality holds also, if $0 < I_1'$, t, ε_j and ω_j are such that $I_{1j}(t, I_1', I_2') = I_1'$ and $I_{2j}(t, I_1', I_2') = l$ are optimal without further restrictions on I_2'.

Proof. Again, after fixing j or ε_j so that the assumptions of the lemma are met, the result does not depend on these variables and hence I can drop this dependence in favor of notational clarity. I will also write $p'(x,t)$ instead of $\partial_x p(x,t)$. Consider first an agent for whom the assumptions of the first part hold. Now since $I_{1j}(t, I_1', I_2') = I_1'$ and $I_{2j}(t, I_1', I_2') = I_2'$ are optimal it is enough to look at the first-order conditions on effort choices given these levels of insurance. Define the value from choosing effort level (x_1, x_2) as

$$\begin{aligned} W(x_1,x_2) &:= p(x_1,t)p(x_2,t)\left(u\left(\omega - 2L + P\right) + \delta V\left(t+1,0,0\right)\right) \\ &+ p(x_1,t)(1-p(x_2,t))\left(u\left(\omega - L + P\right) + \delta V\left(t+1,0,I_2'-1\right)\right) \\ &+ (1-p(x_1,t))p(x_2,t)\left(u\left(\omega - L + P\right) + \delta V\left(t+1,I_1'-1,0\right)\right) \\ &+ (1-p(x_1,t))(1-p(x_2,t))\left(u\left(\omega\right) + \delta V\left(t+1,I_1'-1,I_2'-1\right)\right) \\ &- E(x_1,x_2) \end{aligned} \tag{2.A.20}$$

The first-order conditions of the agent's maximization problem are then given by:

$$\begin{aligned} \partial_1 W(x_1,x_2) &= p'(x_1,t)p(x_2,t)\left(u\left(\omega - 2L + 2P\right) + \delta V\left(t+1,0,0\right)\right) \\ &+ p'(x_1,t)(1-p(x_2,t))\left(u\left(\omega - L + P\right) + \delta V\left(t+1,0,I_2'-1\right)\right) \\ &- p'(x_1,t)p(x_2,t)\left(u\left(\omega - L + P\right) + \delta V\left(t+1,I_1'-1,0\right)\right) \\ &- p'(x_1,t)(1-p(x_2,t))\left(u\left(\omega\right) + \delta V\left(t+1,I_1'-1,I_2'-1\right)\right) \\ &- \partial_1 E(x_1,x_2) = 0 \end{aligned} \tag{2.A.21}$$

and

$$\begin{aligned}
\partial_2 W(x_1,x_2) &= p'(x_2,t)p(x_1,t)\left(u\left(\omega-2L+2P\right)+\delta V\left(t+1,0,0\right)\right)\\
&- p'(x_2,t)p(x_1,t)\left(u\left(\omega-L+P\right)+\delta V\left(t+1,0,I_2'-1\right)\right)\\
&+ p'(x_2,t)(1-p(x_1,t))\left(u\left(\omega-L+P\right)+\delta V\left(t+1,I_1'-1,0\right)\right)\\
&- p'(x_2,t)(1-p(x_1,t))\left(u\left(\omega\right)+\delta V\left(t+1,I_1'-1,I_2'-1\right)\right)\\
&- \partial_2 E(x_1,x_2)=0 \qquad (2.A.22)
\end{aligned}$$

Now, define

$$\begin{aligned}
A' &:= u(\omega-2L+P)-2u(\omega-L+P)+u(\omega)\\
&\quad +\delta V(t+1,0,0)-\delta V(t+1,0,I_2'-1)\\
&\quad -\delta V(t+1,I_1'-1,0)+\delta V(t+1,I_1'-1,I_2'-1).
\end{aligned}$$

I will prove the lemma in two parts: first when $A' \leq 0$ and then when $A' > 0$.

Assume that $A' \leq 0$. Equation (2.A.21) implicitly defines a function $\hat{x}_1(x_2)$, a level of effort for taking care of item 1 that maximizes the agents utility given that item 2 is protected at effort level x_2. I next want to show that there exists $\hat{x} \in \mathbb{R}$ such that $\partial_1 W(\hat{x},\hat{x}) = 0$. To see that this holds, notice first that

$$\begin{aligned}
\partial_1 W(x,x) &= p'(x,t)\Bigg\{p(x,t)\Big[u\left(\omega-2L+P\right)-2u(\omega-L+P)+u(\omega)\\
&+ \delta V\left(t+1,0,0\right)-\delta V\left(t+1,0,I_2'-1\right)\\
&- \delta V\left(t+1,I_1'-1,0\right)+\delta V\left(t+1,I_1'-1,I_2'-1\right)\Big]\\
&+ u(\omega-L+P)-u(\omega)+\delta V(t+1,0,I_2'-1)-\delta V(t+1,I_1'-1,I_2'-1)\Bigg\}\\
&- \partial_1 E(x,x). \qquad (2.A.23)
\end{aligned}$$

Denote the term in the curly brackets by T. Then,

$$\begin{aligned} T &\leq p(x,t)\Big[u\left(\omega - 2L + P\right) - 2u(\omega - L + P) + u(\omega) \\ &+ \delta V\left(t+1,0,0\right) - \delta V\left(t+1,0,I_2' - 1\right)\Big] \\ &+ u(\omega - L + P) - u(\omega) < 0, \end{aligned} \tag{2.A.24}$$

where the first inequality holds, since $V(t+1,0,I_2'-1) - V(t+1,I_1',I_2') \leq 0$ by Lemma 2.A.9 and because $p(x,t) \leq 1$. The second follows because the sum of the first three terms is negative by concavity of u and the sum of the V terms is non-positive by Lemma 2.A.9. Inequality (2.A.24) then implies that

$$\partial_1 W(x,x) > \partial_1 E(x,x),$$

because $p'(x) < 0$. Taking a limit as $x \to 0$ I get that

$$\lim_{x\to 0} \partial_1 W(x,x) > \lim_{x\to 0} \partial_1 E(x,x) = 0.$$

Consider then the limit as $x \to \infty$. By assumption,

$$\lim_{x\to\infty} \partial_1 E(x,x) = \infty \text{ and } \lim_{x\to\infty} p'(x) = 0.$$

Because T is always finite, I get that

$$\lim_{x\to\infty} \partial_1 W(x,x) = 0 - \lim_{x\to\infty} \partial_1 E(x,x) = -\infty.$$

Since, $\partial_1 W(x,x)$ is continuous, there must exist $\hat{x} \in \mathbb{R}_+$ such that $\partial_1 W(\hat{x},\hat{x}) = 0$. In other words, $\hat{x}_1(\hat{x}) = \hat{x}$.

It also turns out that $\hat{x}$ is unique. This clearly follows, if I can show that $g(x) := \partial_1 W(x,x)$ is strictly monotonous. Differentiate g to get

$$g'(x) = \partial_{11} W(x,x) + \partial_{12} W(x,x). \tag{2.A.25}$$

Since W is concave, I also have:

$$\begin{aligned} & \partial_{11} W(x,x)\partial_{22} W(x,x) - \partial_{12} W(x,x)^2 > 0 \\ \Leftrightarrow \quad & \partial_{11} W(x,x)^2 > \partial_{12} W(x,x)^2 \\ \Rightarrow \quad & |\partial_{11} W(x,x)| > |\partial_{12} W(x,x)| \geq \partial_{12} W(x,x), \end{aligned} \tag{2.A.26}$$

where the second line follows, since E is symmetric. Note further that

$$\partial_{11} W(x,x) = p''(x,t)T - \partial_{11} E(x,x) < 0,$$

because $p'' > 0$, $T < 0$ and $\partial_{11} E > 0$. Hence,

$$|\partial_{11} W(x,x)| = -\partial_{11} W(x,x).$$

Plugging this back into (2.A.26) and reorganizing terms yields:

$$g'(x) = \partial_{11} W(x,x) + \partial_{12} W(x,x) < 0.$$

Hence g is strictly decreasing implying the uniqueness of $\hat{x}$.

Define then

$$\begin{aligned} F(x_1,x_2) &:= \partial_1 W(x_1,x_2) - \partial_2 W(x_1,x_2) \\ &= A'\left(p'(x_1,t)p(x_2,t) - p'(x_2,t)p(x_1,t)\right) \\ &+ \left(p'(x_1,t) - p'(x_2,t)\right)\left(u(\omega - L + P) - u(\omega) - \delta V(t+1, I_1' - 1, I_2' - 1)\right) \\ &+ p'(x_1,t)\delta V(t+1, 0, I_2' - 1) - p'(x_2,t)\delta V(t+1, I_1' - 1, 0) \\ &- \partial_1 E(x_1,x_2) + \partial_2 E(x_1,x_2) \end{aligned} \tag{2.A.27}$$

If (x_1^*, x_2^*) is the optimal effort level (which is guaranteed to be unique by the strict concavity of W), then the first-order conditions imply that $F(x_1^*, x_2^*) = 0$. Notice then that

$$\begin{aligned} F(\hat{x},\hat{x}) &= p'(\hat{x},t)\delta V(t+1, 0, I_2' - 1) \\ &\quad -p'(\hat{x},t)\delta V(t+1, I_1' - 1, 0) \leq 0 \end{aligned} \tag{2.A.28}$$

because $p'(\hat{x}, t) < 0$.

It is straight-forward to show that the concavity of W implies that $F(\hat{x}_1(x_2), x_2)$ is strictly increasing in x_2.[8] This together with (2.A.28) implies that $x_2^* \geq \hat{x}$ for any $\hat{x}$ that solves $\hat{x}_1(\hat{x}) = \hat{x}$. Now, suppose I was able to show that $\hat{x}_1'(x_2) < 1$. Then it would follow that

$$\begin{aligned} x_1^* = \hat{x}_1(x_2^*) &= \int_{\hat{x}}^{x_2^*} \hat{x}_1'(z)\,\mathrm{d}z + \hat{x}_1(\hat{x}) \\ \leq \int_{\hat{x}}^{x_2^*} 1\,\mathrm{d}z + \hat{x}_1(\hat{x}) &= x_2^* - \hat{x} + \hat{x}_1(\hat{x}) = x_2^*, \end{aligned} \tag{2.A.29}$$

where the second equality follows from the fundamental theorem of calculus and the last one is implied by $\hat{x}$ being a fixed point of $\hat{x}_1$. The inequality $x_1^* \leq x_2^*$ together

[8] For a proof apply the same argument as what is given in the proof of Lemma 5 in the first chapter.

with $p' < 0$ would clearly imply the claim of the lemma. Hence, the proof of the first part will be complete, if I can show that $\hat{x}_1'(x_2) < 1$.

Now, by the implicit function theorem,

$$\begin{aligned}\frac{d\hat{x}_1(x_2)}{dx_2} &= -\frac{\partial_{12}W(\hat{x}_1(x_2),x_2)}{\partial_{11}W(\hat{x}_1(x_2),x_2)} \\ &= -\frac{p'(x_1,t)p'(x_2,t)A' - \partial_{12}E(x_1,x_2)}{\partial_{11}W(\hat{x}_1(x_2),x_2)}.\end{aligned}$$

The numerator is clearly negative, because $p' < 0$, $A' \leq 0$ and $\partial_{12}E \geq 0$. Furthermore, I already showed in Lemma 2.A.10 that $\partial_{11}W < 0$. Consequently, $\frac{d\hat{x}_1(x_2)}{dx_2} \leq 0 < 1$ proving the first part.

Assume then that $A' > 0$ and assume by way of contradiction that $x_1^* > x_2^*$. By the first-order conditions, I must still have $F(x_1^*, x_2^*) = 0$. Notice first that, since E is convex and symmetric, it is also Schur convex and hence[9]

$$\partial_2 E(x_1^*, x_2^*) \leq \partial_1 E(x_1^*, x_2^*).$$

This together with Lemma 2.A.9 implies that

$$\begin{aligned}F(x_1^*,x_2^*) \leq\ & A'\left(p'(x_1^*,t)p(x_2^*,t) - p'(x_2^*,t)p(x_1^*,t)\right) \\ +\ & \left(p'(x_1^*,t) - p'(x_2^*,t)\right)\left(u(\omega - L + P) - u(\omega)\right) \\ +\ & \left(p'(x_1^*,t) - p'(x_2^*,t)\right)\left(\delta V(t+1, I_1' - 1, 0) - \delta V(t+1, I_1' - 1, I_2' - 1)\right).\end{aligned}$$

[9]See, for example, Theorem C on p. 259 in Roberts and Varberg (1973).

Now, $p'(x_1^*,t)p(x_2^*,t) \leq p'(x_1^*,t)p(x_1^*,t)$, because p is decreasing. Now, since $A' > 0$, I also get that

$$\begin{aligned} & A'\left(p'(x_1^*,t)p(x_2^*,t) - p'(x_1^*,t)p(x_1^*,t)\right) \\ \leq\ & A'p(x_1^*,t)(p'(x_1^*,t) - p'(x_2^*,t)) \\ \leq\ & A'(p'(x_1^*,t) - p'(x_2^*,t)), \end{aligned}$$

where the last inequality follows, since the the whole term is positive and the left-hand side is the right-hand side multiplied by $p(x_1,t) \leq 1$. Plugging this back into the inequality I had for $F(x_1^*, x_2^*)$ yields:

$$\begin{aligned} F(x_1^*,x_2^*) \ \leq\ & (p'(x_1^*,t) - p'(x_2^*,t))\Big(A' + u(\omega - L + P) - u(\omega) \\ & +\delta V(t+1, I_1' - 1, 0) - \delta V(t+1, I_1' - 1, I_2' - 1)\Big) \\ =\ & (p'(x_1^*,t) - p'(x_2^*,t))\Big(u(\omega - 2L + P) - u(\omega - L + P) \\ & +\delta V(t+1,0,0) - \delta V(t+1,0,I_2' - 1)\Big) < 0, \end{aligned} \tag{2.A.30}$$

by the definition of A' and since $p'(x_1^*,t) > p'(x_2,t)$ as p' is increasing. This clearly contradicts $F(x_1^*, x_2^*) = 0$.

Identical arguments hold for the case when $\hat{I}_{1j} = I_1'$ and $\hat{I}_{2j} = l$ by simply replacing $u(x)$ with $u(x-C)$ for all values of x and I_2' by l everywhere in the argument above. □

Corollary 2.A.13. *If it is never optimal to buy a new contract for an item that still has time left in its contract and if* $0 < I_1' < I_2'$, *and* t, ε_j *and* ω_j *are such that* $\hat{I}_{i,j} = I_i'$ *for* $i = 1, 2$ *then*

$$p(x_1(t, I_1', I_2'), \varepsilon_j, t) > p(x_2(t, I_1', I_2'), \varepsilon_j, t)$$

The same inequality holds also, if $0 < I_1'$, t, ε_j *and* ω_j *are such that* $\hat{I}_{1j} = I_1'$ *and* $\hat{I}_{2j} = l$ *are optimal.*

Proof. It is easy to check first that the uniqueness of the optimal policy and the additional assumption of renewing contracts before the previous contract ends being suboptimal imply that Lemma 2.A.9 holds with a strict inequality. This in turn allows replacing the relevant inequalities in the proof above by strict inequalities. □

Chapter 3

Bargaining with Private Deadlines

3.1 Introduction

This paper studies a model of two-sided private information in a variation of Ståhl (1972) and Rubinstein (1982) alternating-offer bargaining. Players have private information about their deadline: a time after which the player is willing to accept any feasible offer. One natural interpretation of this is that players are endowed with a stock of resources which are consumed during the bargaining process and once a player's stock has run out, rejecting offers becomes extremely costly or the bargaining surplus loses all of its value to the player. For example, when an upstream firm is bargaining with a downstream firm over the value of an intermediate product, both might be unable to continue production before the bargaining impasse is settled. When not producing, the firms may still need to pay fixed costs that run down their savings. As financial institutions are unable verify the expected bargaining result, borrowing with favorable terms during the impasse can be difficult. To avoid bankruptcy, a firm that has run out of resources may be willing to accept any offer that leaves it with a non-negative surplus. Similar situations arise when a labor union is striking. Both the employer and the employees have to pay fixed costs during the strike. Once

a party runs out of savings they have to accept any offer that leaves them with a non-negative surplus. Our paper shows how imperfect information about the timing of the deadline may generate both delay and waste in the bargaining process.

We derive a relatively intuitive equilibrium for a version of the model where both players have two possible deadlines. In this equilibrium both players make offers as if they were the type with the highest deadline and their opponent had the lowest deadline. The offers correspond to the offers that would be made in a standard Rubinstein bargaining game with complete information. If the player has a low deadline, she will be accepting these offers with a positive probability, while if the player has a high deadline she always wants to reject these offers. Hence, the equilibrium has partial screening in the sense that at least one player's type will be revealed: if a player has the low deadline, the game will always end with one player's type fully revealed. On the other hand, if both players have a high deadline, neither will accept any equilibrium path offers before both players' low deadlines have passed and thus both players will be known to be high types. The structure of the equilibrium resembles equilibria from standard reputational games.

An interesting feature of the equilibrium is that it provides no information rents for the type with the low deadline for partially revealing her type. Her expected utility in every period corresponds directly to what she would get in a complete information version of the game where she was known to have the low deadline and her opponent was known to have a high deadline. The equilibrium is wasteful in a strict Pareto sense: both types of both players would prefer being able to commit to perfectly revealing their type in the beginning of the bargaining process. This highlights the potential efficiency gains from an arbitration process if an impartial arbitrator can be trusted with the players' private information.

Another interesting feature of our equilibrium is that the player's bargaining power, measured by how close their low type's deadline is, affects only the offers

that their opponent makes. The equilibrium acceptance probabilities of both players' low types are highly symmetrical even if one player's low type's deadline is much after the other.

We also conjecture that for patient players our equilibrium path outcome is the unique trembling hand perfect equilibrium (Selten, 1975). The intuition is that, if deviations leave probability mass on both types, a patient high type has a large incentive to just wait past her own lower deadline to perfectly signal her type. Our equilibrium offers are consistent with this type of signaling but no other equilibria seem to be able to simultaneously guarantee both player's high types this signaling payoff without inducing incentives for the low type to mimic. Trembling hand equilibrium is a relatively weak refinement and is usually not enough in bargaining with asymmetric information to guarantee unique predictions. This is an important extension since typically in bargaining games with two-sided asymmetric information, it is hard to obtain uniqueness of screening equilibria and authors have to resort to somewhat specialized refinements.

The rest of the paper is organized as follows. The next section gives a brief review of related literature, Section 3 derives a partial screening equilibrium for a model where both players' low deadlines are the same. Section 4 relaxes this assumption by allowing players to have different low types. Section 5 discusses the properties of the derived equilibria and Section 6 concludes.

3.2 Related Literature

Our paper contributes to the large literature on bargaining with asymmetric information. The previous papers that assume a deadline, such as such as Fanning (2013), Simsek and Yildiz (2014), Fuchs and Skrzypacz (2012) and Hart (1989), assume either there is no uncertainty about the timing of the deadline, or that the deadline is

common for the two players, or both. We introduce the new dimension of uncertainty in deadlines and show how it affects the players' bargaining power. Our equilibrium does not exhibit waiting until the eleventh hour before the deadline, often known as the deadline effect, but rather there is a two-sided screening equilibrium. The key reason for this is the asymmetric information about the deadlines. For example, in labor market bargaining, papers such as Fanning (forthcoming) and Simsek and Yildiz (2014) could be interpreted as explaining the behavior prior to a strike when the start date of the strike is common knowledge. Our equilibrium, on the other hand, describes what happens during the strike when the ability to continue the strike is mostly private information. It also explains why the parties often reach an agreement much before either goes bankrupt while simultaneously allowing for very prolonged bargaining impasses. The only work dealing with the case where deadlines are private information which we are aware of is Sandholm and Vulkan (1999). However, their results which allow for discounting (Proposition 4.1 and Theorem 4.1) have straightforward counter-examples, e.g., take the familiar Rubinstein (1982) construction, while the results not involving discounting are of less interest (as there is no loss to higher types in waiting-out for deadlines of lower types to pass).

The structure of the equilibrium has some strong similarity with the war of attrition (see Hendricks, Weiss and Wilson, 1988) and how it comes up in bargaining with behavioral types as in Abreu and Gul (2000). In particular, in our equilibrium, in all but potentially the first period, both types of a given player always demand the same share and concessions happen only when players decide whether to accept the opponents offer. There is also a characteristic indifference for the low type on the equilibrium path.

A large part of the earlier literature without behavioral types assumes one-sided uncertainty, only one player making offers or both of these restrictions. Cramton (1992) studies a model that is essentially identical to ours except that the asymmetric

information is about the size of the bargaining surplus. Similar to our equilibrium, he suggests intuitively an equilibrium where the agreement is based on offers derived from a full information Rubinstein bargaining model. However, in his equilibrium bargaining continues for at most two periods. In our equilibrium agreement might not be reached before both players' low deadlines have passed and hence all hidden information have been revealed.

3.3 The Symmetric Model and Its Screening Equilibrium

Consider two players bargaining over the division of a surplus. The baragaining protocol involves alternating offers (with player 1 making offers and player 2 responding in odd periods, and vice versa in even periods) and an infinite horizon. Assume for notational simplicity that both players have the same discount factor $\delta \in (0,1)$. Furthermore, suppose that each player has a type (low or high), measuring the number of periods of bargaining this agent can undertake. In the example of a strike, think of this as the amount of time the worker's union or the employer has before they run out of resources to continue the strike.

In this section we assume that the Players are ex-ante symmetric and have type $H \in \mathbb{N} \cup \{\infty\}$ with probability ρ_1 and type $L \in \mathbb{N}$, $L < H$ with probability $1 - \rho_1$. If an agent is of type $\tau \in \{L, H\}$, then an agreement on an offer that allocates share x of the surplus to the agent in period $t \leq \tau$ will yield the agent a present value payoff equal to $\delta^t x$ while any agreement in period $t > \tau$ yields the player 0.

3.3.1 A Partial Screening Equilibrium for the Subgames Starting from Period L and $L-1$

Suppose that L is odd.[1] This means that Player 1 is offering in period L and Player 2 is offering in period $L+1$. In period $L+1$ Player 1 of type L is indifferent between accepting and rejecting any positive offer. Furthermore, Player 2 of type L is indifferent between all available offers. Consider the following strategy profile starting from period $L+1$: Consider first any odd period $t > L+1$ where Player 2 is choosing to accept or reject, type H accepts any offer that allocates her x where $x \geq \frac{\delta(1-\delta^{H-t})}{1+\delta}$. Low type Player 2 accepts any x such that $x \geq 0$. Let $\rho_2(h_t)$ be the belief that player 1 holds in period t after history h_t that player 2 is type H. Then Player 1 of type H demands 1 for herself if:

$$\rho_2(h_t)\delta\frac{\delta(1+\delta^{H-t-1})}{1+\delta} + (1-\rho_2(h_t)) \geq \frac{1+\delta^{H-t+1}}{1+\delta}.$$

Otherwise, she demands $\frac{1+\delta^{H-t+1}}{1+\delta}$ for herself. This is the standard demand from a Rubinstein bargaining game where both players are known to be type H. Type L demands 0 for herself.

In period $t \geq L+1$ where t is even, type H Player 1 accepts any offer that allocates her $x \geq \frac{\delta(1+\delta^{H-t})}{1+\delta}$. The low type accepts any non-negative offer. Similar to above, if $\rho_1(h_t)$ is the probability with which Player 1 is believed to be the high type, then Player 2 demands the whole surplus if:

$$\rho_1(h_t)\delta\frac{\delta(1-\delta^{H-t-1})}{1+\delta} + (1-\rho_1(h_t)) \geq \frac{1-\delta^{H-t+1}}{1+\delta}.$$

[1] The case when it is even can be solved similarly.

Otherwise, she demands $\frac{1-\delta^{H-t+1}}{1+\delta}$. Similiarly to above, the low type always demands 0. This profile clearly constitutes a perfect Bayesian equilibrium for the subgame starting from period $L+1$ and the equilibrium does not depend on the off-equilibrium beliefs.

Suppose that players play this equilibrium after period L. Consider then the following strategies for period L: The high type of Player 1 offers $(1,0)$ and the low type offers $(1-\delta,\delta)$. The low type of Player 2 accepts any offer that guarantees her at least zero.[2] The high type Player 2 accepts anything that yields her at least δ. Clearly, the low type of Player 2 does not have an incentive to deviate. As long as the weight on the low type after a deviation from Player 1 is close to 1, also the high type of Player 2 does not have an incentive to deviate.

Since, given the high weight on the low type off the equilibrium path, the high type of Player 2 does not want to accept any offer that yields Player 1 more than $1-\delta$, the best deviation for the high type is to offer $1-\delta$ (other offers either yield her less in this period or get still rejected by the high type in this period *and* lead to the low $(0,1)$ offer from the high type in the following period). This deviation does not benefit Player 1 as long as

$$1-\delta < (1-\rho_{L,2}) + \rho_{L,2}\delta\frac{\delta(1+\delta^{H-L-1})}{1+\delta},$$

where by $\rho_{L,2}$ we denote the probability with which Player 2 is believed to be the high type if period L is reached on the equilibrum path. This inequality clearly holds as long as δ is sufficiently high.

[2]Notice that mixing here is not possible in trembling hand perfect equilibria for reasonable values of δ. If the low type mixes, the high type Player 1 can simply deviate to offering a little more than 0 to Player 2. The low type must accept this. Even if the high type Player 2 puts a very high off-equilibrium belief on Player 1 being the low type, as long as this belief is not equal to 1, Player 1 can reject the offer $(0,1)$ in the following period and convince Player 2 on the fact that her type is high.

The more interesting question is whether the low type of Player 1 wants to deviate to offering $(1,0)$.[3] This deviation does not benefit the low type as long as

$$(1-\rho_{L,2}) \leq 1-\delta \Leftrightarrow \rho_{L,2} \geq \delta.$$

We will show below how these beliefs can be reached using mixing in the first period and hence we will concentrate on this case. When this does not hold, we have an equilibrium where both types of Player 1 offer $(1,0)$. This equilibrium is important for off-equilibrium path considerations. However, it implies a potentially higher continuation value for the low type Player 1 in the previous periods which in turn puts more constraints on the beliefs when constructing separating equilibria in those periods. The equilibrium constructed for $\rho_{L,2} \geq \delta$ is also the one that is the analogue of the equilibrium for the asymmetric model below.

Let us then move to period $L-1$. Notice that if Player 2's offer in this period gets rejected the continuation value for the low type Player 1 from the equilibrium in period L is $1-\delta = \frac{1-\delta^{L-L+2}}{1+\delta}$. We will look for an equilibrium where both types of Player 2 offer $(\delta(1-\delta), 1-\delta(1-\delta))$. The high type of Player 1 will always reject this and the low type accepts it with probability $\alpha_{L-1,1}$. Any deviation in offers from Player 2 will lead to a weight close to 1 being put on the low type. This implies that there is no incentive to try to offer the low type slightly more to get her to accept with certainty. The large change in her beliefs will lead her to reject and offer $(1,0)$ in the following period which is clearly not beneficial for neither type of Player 2. Both types of Player 1 must however accept any offer that yields them at least δ. This is the present value of asking the whole surplus in the following period from the low type. Now, the low type of Player 2 does not want to deviate to offering $(\delta, 1-\delta)$ as

[3]Clearly, other deviations cannot be more beneficial.

long as

$$(1-\rho_{L-1,1})\alpha_{L-1,1}(1-\delta(1-\delta))+(1-\rho_{L-1,1})(1-\alpha_{L-1,1})\delta(1-\delta)\geq 1-\delta$$

or equivalently when

$$\alpha_{L-1,1}\geq \frac{(1-\delta)(1-\delta(1-\rho_{L-1,1}))}{(1-2(1-\delta)\delta)(1-\rho_{L-1,1}))}. \tag{3.3.1}$$

The right-hand side is less than one as long as $\rho_{L-1,1}\leq \frac{\delta}{1-\delta+\delta^2}$. It turns out that the low type of Player 2 has the lowest incentive to deviate in previous periods when (3.3.1) holds as an equality. This implies also that Player 2 is indifferent between signaling her type or submitting the pooling offer $(\delta(1-\delta), 1-\delta(1-\delta))$. Notice tha if period $L-1$ is the first period of the game and if Player 1 is mixing so that (3.3.1) holds as an equality, then the low type Player 2 can play $(\delta, 1-\delta)$ with probability $q_{L-1,2}$ such that in period L after seeing the pooling offer, Player 1's beliefs satisfy:

$$\begin{aligned}\rho_{L,2} &= \frac{\rho_{L-1,2}}{\rho_{L-1,2}+(1-\rho_{L-1,2})(1-q_{L-1,2})}\geq\delta \\ &\Leftrightarrow\ q_{L-1,2}\geq\frac{\delta-\rho_{L-1,2}}{\delta(1-\rho_{L-1,2})}\end{aligned} \tag{3.3.2}$$

which is clearly a probability as long as $\rho_{L-1,2}\leq\delta$. Similar indifference will hold through-out the game and can be used to get enough weight on the high type for period L. Having this mixing happen in the first period makes sure that it does not incentivize the low type of Player 1 to wait for the period when the mixing happens. Consequently, it puts the least restrictions on the prior in the first period.

3.3.2 The Equilibrium Strategies for Periods $L-k$, $k\geq 2$

Let k be even. We will construct the equilibrium for periods $L-k$ and $L-k-1$ inductively using the equilibrium from above for the last two periods. Since k is even,

Player 1 will be offering in period k. We assume that in period $L-k+1$ both types of Player 2 will offer

$$\left(\frac{\delta(1-\delta^k)}{1+\delta}, \frac{1+\delta^{k+1}}{1+\delta}\right).$$

Just as before, this corresponds to the sub-game perfect equilibrium offer in a game where Player 2 is known to be the low type and Player 1 is known to be the high type. We assume that the low type of Player 1 is indifferent between accepting and rejecting this. The high type will always reject it. By the construction above, both of these claims are true for period $L-1$. We also assume that, conditional on reaching period $L-k+1$, the low type Player 2's expected utility on the equilibrium path is equal to $\frac{1-\delta^k}{1+\delta}$. This is also true for period $L-1$.

In period $L-k$ both types of Player 1 will offer

$$\left(\frac{1+\delta^{k+1}}{1+\delta}, \frac{\delta(1-\delta^k)}{1+\delta}\right) \tag{3.3.3}$$

Similarly to above, this is the sub-game perfect equilibrium offer for the game where Player 1 is known to be the high type and Player 2 is known to be the low type. We look for an equilibrium where the high type of Player 2 always rejects anything that yields her less than $\frac{\delta(1+\delta^{k+1})}{1+\delta}$ and the low type mixes between accepting and rejecting Player 1's equilibrium offer and rejects any off-equilibrium offer that yields her less than $\frac{\delta(1+\delta^{k+1})}{1+\delta}$. Notice that by the induction assumption, rejection yields $\frac{1-\delta^k}{1+\delta}$ in expectation to the low type Player 2. This has exactly the same present value as accepting Player 1's offer. Hence Player 2 is free to mix in Period $L-k$.

We look for an equilibrium where the low type Player 1 is indifferent between this equilibrium offer and signaling her type by offering

$$\left(\frac{1-\delta^{k+2}}{1+\delta}, \frac{\delta(1+\delta^{k+1})}{1+\delta}\right).$$

This is the sub-game perfect equilibrium offer from a game where Player 1 is known to be the low type and Player 2 is known to be the high type. In any equilibrium this offer must be accepted by both types of Player 2. The indifference holds, if and only if

$$\begin{aligned}\frac{1-\delta^{k+2}}{1+\delta} &= (1-\rho_{L-k,2})\alpha_{L-k,2}\frac{1+\delta^{k+1}}{1+\delta} \\ &\quad + (1-(1-\rho_{L-k,2})\alpha_{L-k,2})\,\delta\frac{\delta(1-\delta^{k})}{1+\delta}\end{aligned} \tag{3.3.4}$$

Solving this for $\alpha_{L-k,2}$ yields

$$\alpha_{L-k,2} = \frac{1-\delta}{(1-\delta+\delta^{k+1})(1-\rho_{L-k,2})}, \tag{3.3.5}$$

which is at most 1 as long as

$$\rho_{L-k,2} \leq \frac{\delta^{k+1}}{1-\delta-\delta^{k+1}}. \tag{3.3.6}$$

The beliefs on the equilibrium path evolve according to the Bayes' Rule:

$$\rho_{L-k+1,2} = \rho_{L-k+2,2} = \frac{\rho_{L-k,2}}{\rho_{L-k,2}+(1-\alpha_{L-k,2})(1-\rho_{L-k,2})}.$$

Combining this with (3.3.5) yields the following a recursive formula for $\rho_{L-k,2}$

$$\rho_{L-k,2} = \begin{cases} \rho_{L-k+1,2} & k \text{ odd} \\ \frac{\delta^{k+1}}{1-\delta+\delta^{k+1}}\rho_{L-k+1,2} & k \text{ even.} \end{cases} \tag{3.3.7}$$

Solving this backwards gives us:

$$\rho_{L-t,2} = \begin{cases} \rho_{L,2}\prod_{k=1}^{t/2}\frac{\delta^{2k+1}}{1-\delta+\delta^{2k+1}} & t \text{ even} \\ \rho_{L,2}\prod_{k=1}^{(t-1)/2}\frac{\delta^{2k+1}}{1-\delta+\delta^{2k+1}} & t \text{ odd.} \end{cases} \tag{3.3.8}$$

The reason for solving this backwards is that we know that we need $\rho_{L,2} \geq \delta$. This implies that without any changes to the strategy we need:

$$\rho_{1,2} \geq \delta \prod_{k=1}^{(L-2)/2} \frac{\delta^{2k+1}}{1-\delta+\delta^{2k+1}}. \tag{3.3.9}$$

If the prior does not satisfy this, it can be easily fixed, as the low type of Player 1 need not be indifferent between her signaling and pooling offers in period 1. It turns out that on the equilibrium path, the correct beliefs will be achieved as long as

$$\rho_{2,2} \geq \delta \prod_{k=1}^{(L-2)/2} \frac{\delta^{2k+1}}{1-\delta+\delta^{2k+1}}. \tag{3.3.10}$$

and if the prior in period 1 is below this the posterior after period 1 actions can be always adjusted to satisfy this inequality with the correct mixing in period 1. We will discuss this after describing Player 1's strategy below.

If the restriction in (3.3.9) is naturally satisfied by the prior, then the strategy described above will work without any adjustments. Solving (3.3.7) forward yields

$$\rho_{L-t,2} = \begin{cases} \rho_{1,2} \prod_{k=1}^{(L-t-1)/2} \frac{1-\delta+\delta^{L-2k}}{\delta^{L-2k}} & t \text{ even} \\ \rho_{1,2} \prod_{k=1}^{(L-t)/2} \frac{1-\delta+\delta^{L-2k}}{\delta^{L-2k}} & t \text{ odd.} \end{cases} \tag{3.3.11}$$

Substituting this back into (3.3.4) yields the mixing probabilities in terms of the prior. Furthermore, combining with (3.3.6) yields:

$$\begin{aligned} \rho_{1,2} \prod_{t=1}^{(L-k-1)/2} \frac{1-\delta+\delta^{L-2t}}{\delta^{L-2t}} &\leq \frac{\delta^{k+1}}{1-\delta+\delta^{k+1}} \\ \Leftrightarrow \rho_{1,2} &\leq \frac{\delta^{k+1}}{1-\delta+\delta^{k+1}} \prod_{t=1}^{(L-k-1)/2} \frac{\delta^{L-2t}}{1-\delta+\delta^{L-2t}}. \end{aligned}$$

The right-hand side is increasing in k and hence the inequality is the most binding when $k = 2$. Hence, all of these constraints are satisfied when

$$\rho_{1,2} \leq \frac{\delta^3}{1-\delta+\delta^3} \prod_{t=1}^{(L-3)/2} \frac{\delta^{L-2t}}{1-\delta+\delta^{L-2t}}$$

The right-hand side above is increasing in δ and it equals 1 when $\delta = 1$. Furthermore, it is decreasing in L.

We assume that any deviations from Player 1 lead to beliefs that put almost all probability mass on Player 1 being the low type. The continuation equilibrium starting with these beliefs conditional on reaching period $L-k+1$ will be as if the players were playing an analogous equilibrium to the one constructed here assuming that the game starts with the off-equilibrium belief $(\hat{\rho}_{L-k+1}, \rho_{L-k+1})$ with $\hat{\rho}_{L-k+1} \approx 0$ in period $L-k+1$. In particular, we can have the low type of Player 1 accepting in period $L-k+1$ with a high enough a probability so that both types of Player 2 will reject any off-equilibrium offer that yields her less than $\frac{\delta(1+\delta^{k+1})}{1+\delta}$. We will construct the precise off-equilibrium offers and acceptance behavior after constructing the equilibrium strategies for period $L-k-1$.

In period $L-k-1$ Player 2 is offering. We will look for an equilibrium where both types of Player 2 offer

$$\left(\frac{\delta(1-\delta^{k+2})}{1+\delta}, \frac{1+\delta^{k+3}}{1+\delta}\right).$$

The high type of Player 1 rejects any offer that yields her less than $\frac{\delta(1+\delta^{k+1})}{1+\delta}$. Just like above the low type will be indifferent between accepting and rejecting the equilibrium offer while rejecting any off-equilibrium offer that yields her less than $\frac{\delta(1+\delta^{k+1})}{1+\delta}$. To close the induction on the equilibrium path, we need the low type of Player 2 to be

indifferent between the equilibrium offer and signaling her type by offering

$$\left(\frac{\delta(1+\delta^{k+1})}{1+\delta}, \frac{1-\delta^{k+2}}{1+\delta}\right).$$

Notice first that, if one multiplies (3.3.4) by δ, the left-hand side becomes what Player 1 gets if she accepts Player 2's equilibrium offer in period $L-k-1$. The right-hand side is what the low type expects to get if she rejects the offer. Hence, the low type of Player 1 is, by construction indifferent between accepting and rejecting the equilibrium offer.

To achieve indifference between signaling and pooling for the low type of Player 2, the low type of Player 1 must accept with probability $\alpha_{L-k-1,1}$ such that

$$\begin{aligned}\frac{1-\delta^{k+2}}{1+\delta} &= (1-\rho_{L-k-1,1})\alpha_{L-k-1,1}\frac{1+\delta^{k+3}}{1+\delta} \\ &\quad + (1-(1-\rho_{L-k-1,1})\alpha_{L-k-1,1})\,\delta\frac{\delta(1-\delta^{k})}{1+\delta}. \qquad (3.3.12)\end{aligned}$$

Solving this for $\alpha_{L-k-1,1}$ yields

$$\alpha_{L-k-1,1} = \frac{1-\delta}{(1-\delta+\delta^{k+2})(1-\rho_{L-k-1,1})}. \qquad (3.3.13)$$

This is at most 1 as long as

$$\rho_{L-k-1,1} \leq \frac{\delta^{k+2}}{1-\delta+\delta^{k+2}} \qquad (3.3.14)$$

Combining Bayes' Rule with (3.3.13) yields the way beliefs are updated on the equilibrium path:

$$\rho_{L-k,1} = \begin{cases} \rho_{L-k-1,1} & k \text{ odd} \\ \frac{1-\delta+\delta^{k+2}}{\delta^{k+2}}\rho_{L-k-1,1} & k \text{ even.} \end{cases}$$

Since there are no additional restrictions on these beliefs in the last period, we can simply solve this forward starting from the prior:

$$\rho_{L-k,1} = \begin{cases} \rho_{1,1} \prod_{t=1}^{(L-k-1)/2} \frac{1-\delta+\delta^{L-2t+1}}{\delta^{L-2t+1}} & k \text{ even} \\ \rho_{1,1} \prod_{t=1}^{(L-k-2)/2} \frac{1-\delta+\delta^{L-2t+1}}{\delta^{L-2t+1}} & k \text{ odd,} \end{cases} \tag{3.3.15}$$

which in turn can be used to solve the mixing probabilities in terms of the prior.

Note that inequality 3.3.14 has to be satisfied for all of the above expressions. That inequality, once decremented and when expression (3.3.15) is subsituted in becomes:

$$\begin{aligned} \rho_{1,1} \prod_{t=1}^{(L-k-2)/2} \frac{1-\delta+\delta^{L-2t+1}}{\delta^{L-2t+1}} &\leq \frac{\delta^{k+1}}{1-\delta+\delta^{k+1}} \\ \Leftrightarrow \rho_{1,1} &\leq \frac{\delta^{k+1}}{1-\delta+\delta^{k+1}} \prod_{t=1}^{(L-k-2)/2} \frac{\delta^{L-2t+1}}{1-\delta+\delta^{L-2t+1}}. \end{aligned}$$

Now, the RHS is increasing in k, so that the above inequality is most binding for $k = 1$, thus we require that:

$$\rho_{1,1} \leq \frac{\delta^2}{1-\delta+\delta^2} \prod_{t=1}^{(L-3)/2} \frac{\delta^{L-2t+1}}{1-\delta+\delta^{L-2t+1}}.$$

We note that the RHS above is increasing in δ, and it equals 1 when $\delta = 1$. It is decreasing in L.

Notice now that that the low type of Player 1 need not be indifferent in the first period and the equilibrium still works as long as the low type of Player 2 mixes in the first period in a way that guarantees that (3.3.6) holds in every subsequent period. Hence, if (3.3.9) does not hold, then Player 2 can accept in the first period with probability $\alpha_{1,2}$ such that

$$\delta \prod_{k=1}^{(L-2)/2} \frac{\delta^{2k+1}}{1-\delta+\delta^{2k+1}} = \frac{\rho_{1,2}}{\rho_{1,2} + (1-\rho_{1,2})(1-\alpha_{1,2})},$$

which yields

$$\alpha_{1,2} = \frac{\rho_{1,2} - \delta \prod_{k=1}^{(L-2)/2} \frac{\delta^{2k+1}}{1-\delta+\delta^{2k+1}}}{(1-\rho_{1,2})\delta \prod_{k=1}^{(L-2)/2} \frac{\delta^{2k+1}}{1-\delta+\delta^{2k+1}}}$$

as the first period mixing probability.

Last we need to construct the strategies off the equilibrium path. Since the model is highly symmetric, we will only consider a deviation by Player 2 in period $L-k-1$ to a non-equilibrium offer. Deviations by Player 1 can be handled symmetrically. Suppose the off-equilibrium belief after Player 2 deviates to offering $(x, 1-x)$ in period $L-k-1$ is given by $\hat{\rho}^x_{L-k,2}$. Since we are looking for a trembling hand perfect equilibrium, we assume that $\hat{\rho}^x_{L-k,2} \in (0,1)$. Notice first that we can treat the subgame starting from period $L-k$ as a separate game where there are k periods until the low types' deadline is up and where the initial beliefs are given by $(\rho_{L-k,1}, \hat{\rho}^x_{L-k,2})$. Just as above in the first period of the game, the low type of Player 1 can strictly prefer offering

$$\left(\frac{1+\delta^{k+1}}{1+\delta}, \frac{\delta(1-\delta^k)}{1+\delta}\right)$$

in the first period following the deviation from Player 2. Assuming that the offers and acceptance behavior is otherwise the same as on the equilibrium path following the deviation, Player 2 can still mix in period $L-k$. However, her acceptance probability can be higher than the equilibrium path acceptance probability as there is no indifference that needs to be maintained for Player 1 in period $L-k$. The key is to look for an acceptance probability and off-equilibrium beliefs for period $L-k$ that induce both types of Player 1 to reject any off-equilibrium offer in period $L-k-1$ that yields her less than $\frac{\delta(1+\delta^{k+1})}{1+\delta}$. This would rule out any profitable deviations because the low type of Player 2 is exactly indifferent between offering Player 1 this quantity and the high type strictly prefers the equilibrium path to this offer.

Now, the low type of Player 1 will reject an off-equilibrium offer $(x, 1-x)$ as long as

$$x < (1-\hat{\rho}^x_{L-k,2})\hat{\alpha}^x_{L-k,2}\frac{\delta(1+\delta^{k+1})}{1+\delta} + (1-(1-\hat{\rho}^x_{L-k,2})\hat{\alpha}^x_{L-k,2})\delta^2\frac{\delta(1-\delta^k)}{1+\delta}, \quad (3.3.16)$$

where $\hat{\alpha}^x_{L-k,2}$ is the off-equilibrium probability with which the low type accepts Player 1's offer in period $L-k$. The mixing probability is constrained by the requirement that

$$\frac{\hat{\rho}^x_{L-k,2}}{\hat{\rho}^x_{L-k,2} + (1-\hat{\rho}^x_{L-k,2})(1-\hat{\alpha}^x_{L-k,2})} = \hat{\rho}_{L-k+2,2} \leq \frac{\delta^{k-1}}{(1-\delta+\delta^{k-1})}.$$

Solving for $\hat{\alpha}^x_{L-k,2}$ yields

$$\hat{\alpha}^x_{L-k,2} \leq 1 - \frac{(1-\delta)\hat{\rho}^x_{L-k,2}}{\delta^{k-1}(1-\hat{\rho}^x_{L-k,2})}.$$

Notice that by choosing $\hat{\rho}^x_{L-k,2}$ close to 0, also $\hat{\alpha}^x_{L-k,2}$ can be chosen arbitrarily close to 1. This in turn implies that for any $x < \frac{\delta(1+\delta^{k+1})}{1+\delta}$ the inequality (3.3.16) can always be satisfied by choosing $\hat{\rho}^x_{L-k,2}$ close enough to 0 and $\hat{\alpha}^x_{L-k,2}$ close to 1. With any such beliefs and a continuation strategy from the low type Player 2, Player 1 will find it optimal to reject any $x < \frac{\delta(1+\delta^{k+1})}{1+\delta}$ and hence deviating to such x will not be beneficial to either type of Player 2. Fixing beliefs and acceptance probabilities similarly after Player 1 deviates to an off-equilibrium path offer shows that the profile constructed here constitutes a perfect Bayesian equilibrium. Furthermore, the equilibrium is clearly trembling hand perfect, since the off-equilibrium beliefs put mass on both types after each deviation.

Summarizing the results from above, we have the following Theorem:

Theorem 3.3.1. *If*

$$\rho_{1,1} \leq \frac{\delta^2}{1-\delta+\delta^2} \prod_{t=1}^{(L-3)/2} \frac{\delta^{L-2t+1}}{1-\delta+\delta^{L-2t+1}}$$

and

$$\rho_{1,2} \leq \frac{\delta^3}{1-\delta+\delta^3} \prod_{t=1}^{(L-3)/2} \frac{\delta^{L-2t}}{1-\delta+\delta^{L-2t}},$$

then there exists a trembling hand perfect equilibrium where in each odd period $L-k,\ k>0$, both types of Player 1 offer $\left(\frac{1+\delta^{k+1}}{1+\delta}, \frac{\delta(1-\delta^k)}{1+\delta}\right)$. The high type of Player 2 always rejects this and the low type accepts it with a non-zero probability. Similarly, in every even period $L-k,\ k>0$, Player 2 offer $\left(\frac{\delta(1-\delta^{k+1})}{1+\delta}, \frac{1+\delta^{k+2}}{1+\delta}\right)$, which the high type of Player 1 always rejects and the low type accepts with a non-zero probability. In period L, high type Player 1 offers $(1,0)$ and the low type offers $(1-\delta,\delta)$. The low type of Player 2 accepts any non-negative offer and the high type accepts any offer that yields her at least δ.

3.4 Asymmetric Deadlines

Consider next a version of the model where the two players have different early deadlines. In other words, let the deadline for player i's low type be L_i and assume that $L_1 < L_2$.[4] Suppose for now that L_1 is odd, L_2 is even and H is odd or infinite. Assume also that $L_2+1<H$. Suppose we are in a subgame where $\rho_{L_2,1}=1$. Then, just as before, in period L_2+1 Player 1 knows that both types of Player 2 will accept $(\frac{1+\delta^{H-L_2}}{1+\delta}, \frac{\delta(1-\delta^{H-L_2-1})}{1+\delta})$ and the low type will always accept $(1,0)$ while the high type rejects it. If Player 1 offers $(1,0)$ and it gets rejected, the game reduces to a normal Rubinstein bargaining with two high types. Hence, offering $(1,0)$ in period L_2+1 is

[4]The case where the opposite inequality holds can be solved similarly.

the *unique* equilibrium for the subgame, if and only if

$$(1-\rho_{L_2,2}(h_{L_2}))+\rho_{L_2,2}(h_{L_2})\frac{\delta^2(1+\delta^{H-L_2-2})}{1+\delta}>\frac{1+\delta^{H-L_2}}{1+\delta}.$$

Solving this for $\rho_{L_2,2}$ yields all of the possible equilibria for period L_2+1.

Lemma 3.4.1. *In all equilibria where $\rho_{1,L_2}=0$ in period L_2+1 Player 1 offers:*

- $(1,0)$ *if* $\rho_{L_2,2}(h_{L_2})<\frac{\delta-\delta^{H-L_2}}{1+\delta-\delta^2-\delta^{H-L_2}}$,
- $(\frac{1+\delta^{H-L_2}}{1+\delta},\frac{\delta(1-\delta^{H-L_2-1})}{1+\delta})$ *if* $\rho_{L_2,2}(h_{L_2})>\frac{\delta-\delta^{H-L_2}}{1+\delta-\delta^2-\delta^{H-L_2}}$,
- *Can mix with any probability between the two offers above if* $\rho_{L_2,2}(h_{L_2})=\frac{\delta-\delta^{H-L_2}}{1+\delta-\delta^2-\delta^{H-L_2}}$.

Consider then period L_2. Notice that both types of Player 2 pooling on offering $(\frac{\delta(1+\delta^{H-L_2})}{1+\delta},\frac{1-\delta^{H-L_2+1}}{1+\delta})$ can be an equilibrium as long as Player 1 accepts it with a high enough a probability. For Player 1 to be willing to accept, it must be the case that

$$\frac{\delta(1+\delta^{H-L_2})}{1+\delta}\geq\rho_{L_2,2}\frac{\delta^3(1+\delta^{H-L_2-2})}{1+\delta}+(1-\rho_{L_2,2})\delta.$$

The key case for constructing equilibria in other periods is the one where this holds as an equality. Being indifferent between accepting and rejecting the offer will allow Player 1 to mix.[5]

Solving for the threshold belief yields:

$$\rho_{L_2,2}=\frac{\delta^{L_2+1}-\delta^H}{\delta^{L_2}(1+(1-\delta)\delta)-\delta^H}. \tag{3.4.1}$$

[5]It can be always reached from previous periods with the right mixing probabilities. We should consider uniqueness of equilibrium offers given these beliefs. There is a strong intuitive sense why the high type should offer this given strong beliefs like this. For high levels of patience, indifference is needed in the last period. Otherwise, the low type will want to deviate accepting earlier or wait until the last period.

The low type of Player 2 can always deviate to $(\delta, 1-\delta)$ which is the highest demand that Player 1 will agree to knowing that Player 2 is low. If $\alpha_{L_2,1}$ denotes the probability with which Player 1 accepts the pooling offer, then deviating to $(\delta, 1-\delta)$ is not beneficial as long as $\alpha_{L_2,1}\frac{1-\delta^{H-L_2+1}}{1+\delta} \geq 1-\delta$ which is equivalent to

$$\alpha_{L_2,1} \geq \frac{1-\delta^2}{1-\delta^{H-L_2+1}}$$

The right-hand side is less than 1 as long as $H \geq L_2 + 2$. Any other deviations by Player 2 can be discouraged by off-equilibrium beliefs that put full weight on the low type.

Let us then consider how this "knife-edge" equilibrium can be used to construct a signaling equilibrium for period $L_2 - 1$. Now, suppose that in equilibrium Player 1 is offering $(1-\delta(1-\delta), \delta(1-\delta))$.We will look for an equilibrium where the low type of Player 2 will mix between accepting and rejecting this offer and the high type will always reject it.

The low type is willing to mix exactly when

$$\alpha_{L_2,1}\delta\frac{1-\delta^{H-L_2+1}}{1+\delta} = \delta(1-\delta).$$

Solving this for $\alpha_{L_2,1}$ yields exactly the limiting probability from above:

$$\alpha_{L_2,1} = \frac{1-\delta^2}{1-\delta^{H-L_2+1}}$$

Furthermore, since the high type gets a positive utility in the next period even when Player 1 rejects the pooling offer, it is clear that the high type will reject this offer. Notice that if we denote the mixing probability with which the low type Player 2 accepts in period $L_2 - 1$ by $\alpha_{L_2-1,2}$ then Bayes' law implies that if period L_2 is

reached through a rejection,

$$\rho_{L_2,2} = \frac{\rho_{L_2-1,2}}{\rho_{L_2-1,2} + (1-\rho_{L_2-1,2})(1-\alpha_{L_2-1,2})}.$$

Combining this with (3.4.1) and solving for $\alpha_{L_2-1,2}$ yields:

$$\begin{aligned}\alpha_{L_2-1,2} &= \frac{\delta^{L_2}(\delta - (1+\delta-\delta^2)\,\rho_{L_2-1,2}) - \delta^H\,(1-\rho_{L_2-1,2})}{(\delta^{L_2+1}-\delta^H)(1-\rho_{L_2-1,2})} \\ &= 1 - \frac{\delta^{L_2}\,(1-\delta^2)\,\rho_{L_2-1,2}}{(\delta^{L_2+1}-\delta^H)(1-\rho_{L_2-1,2})}\end{aligned}$$

To discourage deviations from Player 1 to offering Player 2 slightly more to get low Player 2 accept with a higher probability we need to construct the equilibrium response from both types of Player 2 and the continuation of the game after such offer. Notice that since Player 2 is not deviating, all beliefs will still need to be consistent with Bayesian updating.

Suppose that Player 1 deviates to offering $(1-\delta(1-\delta)-\delta\varepsilon, \delta(1-\delta)+\delta\varepsilon)$, $\varepsilon > 0$. Let $\alpha_{L_2,1}(\varepsilon)$ be the probability with which Player 1 accepts $(\frac{\delta(1+\delta^{H-L_2})}{1+\delta}, \frac{1-\delta^{H-L_2+1}}{1+\delta})$ in period L_2. Now, low Player 2 is indifferent between accepting and rejecting $\delta(1-\delta)+\varepsilon$ as long as

$$\delta(1-\delta) + \delta\varepsilon = \alpha_{L_2,1}(\varepsilon)\delta\frac{1-\delta^{H-L_2+1}}{1+\delta}.$$

Solving for $\alpha_{L_2,1}(\varepsilon)$ yields

$$\alpha_{L_2,1}(\varepsilon) = \frac{1-\delta^2}{1-\delta^{H-L_2+1}} + \varepsilon\frac{1+\delta}{1-\delta^{H-L_2+1}} = \left(1+\frac{\varepsilon}{1-\delta}\right)\frac{1-\delta^2}{1-\delta^{H-L_2+1}} \qquad (3.4.2)$$

For Player 1 to be willing to mix, (3.4.1) must still hold. This in turn implies that low Player 2 must be accepting $\delta(1-\delta)+\varepsilon\delta$ with the same probability as with which she is accepting $\delta(1-\delta)$. Hence, as long as $\alpha_{L_2,1}(\varepsilon)$ is a probability, deviating to offer $\delta\varepsilon$ more to Player 2 in period L_2-1 will not benefit Player 1. This will get accepted

with exactly the same probability as $\delta(1-\delta)$, leads to the same expected utility in the following period if rejected but yields $\delta\varepsilon$ less if accepted in this period.

Notice then that Player 2 will accept $\delta(1-\delta)+\delta\varepsilon$ with certainty when $\alpha_{L_2,1}(\varepsilon)=1$ only if

$$\delta(1-\delta)+\delta\varepsilon > \delta\frac{1-\delta^{H-L_2+1}}{1+\delta}.$$

But this implies that the high type will also accept this offer. Deviating to pooling on this high offer is not beneficial for Player 1 if

$$\begin{aligned}\frac{1+\delta^{H-L_2+2}}{1+\delta} \quad &\leq \quad (1-\rho_{L_2-1,2})\alpha_{L_2-1,2}(1-\delta(1-\delta)) \\ &\quad + ((1-\rho_{L_2-1,2})(1-\alpha_{L_2-1,2})+\rho_{L_2-1,2})\,\frac{\delta^2(1+\delta^{H-L_2})}{1+\delta}.\end{aligned}$$

This together with the formula for $\alpha_{L_2-1,2}$ yields a upper bound on $\rho_{L_2-1,2}$ that tells us when pooling to the offer where both players are claimed to be high is not beneficial for Player 1:

$$\rho_{L_2-1,2} \leq \frac{\delta^2(\delta^{L_2+1}-\delta^H)^2}{(\delta^{L_2}+\delta^{L_2+1}-\delta^{L_2+2}-\delta^H)(\delta^{L_2}-\delta^{L_2+2}+\delta^{L_2+3}-\delta^{H+2})}. \tag{3.4.3}$$

As $\delta \to 1$, the right-hand side can be shown to have the limit $\left(\frac{H-L_2-1}{H-L_2+1}\right)^2$, which in turn is arbitrarily close to 1, if H is large.

3.4.1 A Recursive Formulation for a Partial Screening Equilibrium for Periods after L_1+1

Consider now period L_2-k where k is even and less than L_2-L_1-1. We will look for an equilibrium where both types of Player 2 offer $(\frac{\delta(1+\delta^{H-L_2+k})}{1+\delta}, \frac{1-\delta^{H-L_2+k+1}}{1+\delta})$ and Player 1 accepts this with probability $\alpha_{L_2-k,1}$. This is the offer that corresponds to the unique equilibrium offer in a Rubinstein bargaining game where both players

are known to be high. In periods where Player 1 makes the offer, she will offer as if Player 2 were known to be low. For example, in period $L_2 - k - 1$ she will offer $(\frac{1+\delta^{k+3}}{1+\delta}, \frac{\delta(1-\delta^{k+2})}{1+\delta})$. The mixing probability $\alpha_{L_2-k,1}$ in period k will make the low type indifferent between accepting and rejecting this offer in period $L_2 - k - 1$. She will mix accepting the offer with probability $\alpha_{L_2-k-1,2}$ that generates the correct beliefs for Player 1 that make her indifferent between accepting and rejecting in period $L_2 - k$.

Let $V_{1,L_2-k+1}(\rho_{L_2-k,2})$ be the expected utilitity for Player 1 from the equilibrium starting from period $L_2 - k + 1$. For player 1 to be willing to mix, her belief about the probability of a high type Player 2, $\rho_{L_2-k,2}$ must satisfy:

$$\frac{\delta(1+\delta^{H-L_2+k})}{1+\delta} = \delta V_{L_2-k+1,1}(\rho_{L_2-k,2}). \tag{3.4.4}$$

If we assume inductively that the above structure generates equilibria for all of the subgames starting between periods $L_2 - k + 1$ and L_2, then

$$\begin{aligned} & V_{L_2-k+1,}(\rho_{L_2-k+1,2}) \\ = \; & \alpha_{L_2-k+1,2}(1-\rho_{L_2-k+1,2})\frac{1+\delta^{k+1}}{1+\delta} \\ & + (1-\alpha_{L_2-k+1,2}(1-\rho_{L_2-k+1,2}))\frac{\delta^2(1+\delta^{H-L_2+k-2})}{1+\delta} \end{aligned} \tag{3.4.5}$$

Player 1's beliefs after non-terminal equilibrium histories follow Bayes' law:

$$\rho_{L_2-k+1,2} = \rho_{L_2-k,2} = \frac{\rho_{L_2-k-1,2}}{\rho_{L_2-k-1,2} + (1-\rho_{L_2-k-1,2})(1-\alpha_{L_2-k-1,2})} \tag{3.4.6}$$

Combining equations (3.4.4) and (3.4.5) and using (3.4.6) to eliminate $\alpha_{L_2-k-1,2}$ from this equation yields a recursive formula for Player 1's equilibrium beliefs. Equation (3.4.1) yields a terminal condition that allows then to solve for the equilibrium beliefs and mixing probability $\alpha_{L_2-k+1,2}$ using equation (3.4.6).

Let us then consider the low type Player 2. In period $L_2 - k - 1$ she is supposed to mix between accepting and rejecting $(\frac{1+\delta^{k+3}}{1+\delta}, \frac{\delta(1-\delta^{k+2})}{1+\delta})$. The same is true for the offer $(\frac{1+\delta^{k+1}}{1+\delta}, \frac{\delta(1-\delta^{k})}{1+\delta})$ in period $L_2 - k + 1$. Hence, using the indifference in period $L_2 - k + 1$, we get that the low type Player 2 is indifferent in period $L_2 - k - 1$ if and only if

$$\frac{\delta(1-\delta^{k+2})}{1+\delta} = \delta\left(\alpha_{L_2-k,1}\frac{1-\delta^{H-L_2+k+1}}{1+\delta} + (1-\alpha_{L_2-k,1})\delta\frac{\delta(1-\delta^{k})}{1+\delta}\right). \tag{3.4.7}$$

From here we can solve for the unique mixing probability for Player 1:

$$\alpha_{L_2-k,1} = \frac{1-\delta^2}{1-\delta^2+\delta^{2+k}-\delta^{k+1+H-L_2}},$$

which is clearly always in $(0,1)$ as long as $H - L_2 \geq 2$. The low type Player 2 has always the option to reveal her type. If Player 1 knows that Player 2 is low, the least that Player 1 is willing to accept is

$$\left(\frac{\delta(1+\delta^{k+1})}{1+\delta}, \frac{1-\delta^{k+2}}{1+\delta}\right).$$

Hence, Player 1's mixing probability has to satisfy

$$\alpha_{L_2-k,1}\frac{1-\delta^{H-L_2+k+1}}{1+\delta} + (1-\alpha_{L_2-k,1})\delta\frac{\delta(1-\delta^{k})}{1+\delta} \geq \frac{1-\delta^{k+2}}{1+\delta}.$$

But this is trivially implied by (3.4.7). Hence, neither type of Player 2 will want to deviate as long as in periods where they are offering, any deviation leads to beliefs that put full mass on the low type.

We still need to characterize the equilibrium strategies following a deviation from Player 1 in a turn where she is offering. Notice first that offering $(\frac{1+\delta^{H-L_2+k+2}}{1+\delta}, \frac{\delta(1-\delta^{H-L_2-k+1})}{1+\delta})$ in period $L_2 - k - 1$ must lead to both types of Player 2 accepting, as this is the unique equilibrium offer from a game where both players are

known to be high. Player 1 does not want to deviate to offering this as long as

$$\begin{aligned} & \frac{1+\delta^{H-L_2+k+2}}{1+\delta} \\ \leq \; & \alpha_{L_2-k-1,2}(1-\rho_{L_2-k-1,2})\frac{1+\delta^{k+3}}{1+\delta} \\ & + (1-\alpha_{L_2-k-1,2}(1-\rho_{L_2-k-1,2}))\frac{\delta^2(1+\delta^{H-L_2+k})}{1+\delta}. \end{aligned} \tag{3.4.8}$$

Given the suggested equlibrium mixing probability above, this yields an upper bound on $\rho_{L_2-k-1,2}$.

For deviating offers in period L_2-k-1 from Player 1 that give

$$x \in \left(\frac{\delta(1-\delta^{k+2})}{1+\delta}, \frac{\delta(1-\delta^{H-L_2-k+1})}{1+\delta}\right)$$

to Player 2, we apply a similar logic to the one presented already for period L_2-1. We want Player 2 to mix after this offer using exactly the same mixing probability as when faced with the equilibrium path offer. To incentivise this behavior, we must adjust the probability with which Player 1 accepts $\frac{\delta(1+\delta^{H-L_2+k})}{1+\delta}$ in the following period. Clearly, this will leave the incentives in all of the following periods untouched and the Players' utilities change only in these two periods. Since the low type Player 2 is accepting with the equilibrium path probability, Player 1's beliefs evolve just as before. Thus she will be losing compared to the equilibrium in period L_2-k-1 due to the increased offer to Player 2 while her expected utility in all of the following periods remains the same. Hence, as long as we can choose the mixing probabilities in a way that make the low type Player 2 indifferent, Player 1 will not have an incentive to deviate to any of these offers.

Consider Player 1 offering

$$\left(\frac{1+\delta^{k+3}}{1+\delta}-\delta\varepsilon, \frac{\delta(1-\delta^{k+2})}{1+\delta}+\delta\varepsilon\right),$$

where $\varepsilon > 0$. The low type of Player 2 is indifferent between accepting and rejecting as long as

$$\frac{\delta(1-\delta^{k+2})}{1+\delta} + \delta\varepsilon = \delta\left(\alpha_{L_2-k,1}(\varepsilon)\frac{1-\delta^{H-L_2+k+1}}{1+\delta} + (1-\alpha_{L_2-k,1}(\varepsilon))\delta\frac{\delta(1-\delta^{k})}{1+\delta}\right) \quad (3.4.9)$$

Just as above, this always has a solution, since setting $\alpha_{L_2-k,1}(\varepsilon) = 1$ implies indifference between accepting the pooling offer $\frac{\delta(1-\delta^{H-L_2-k+1})}{1+\delta}$ today or waiting for the respective pooling counter-offer $\frac{1-\delta^{H-L_2+k+1}}{1+\delta}$ getting accepted tomorrow. Hence, the strategies outlined above will yield a sub-game perfect equilibrium as long as (3.4.8) is satisfied.

3.4.2 Period $L_1 + 1$

There are two reasons to treat this period separately. First, since potentially $\rho_{L_1+1,1} \leq 1$ and a low type Player 1 is willing to accept any non-negative offer we need to make sure that Player 2 does not want to deviate to offering $(0, 1)$ even when off-equilibrium beliefs "punish" her for doing so. The second reason is that we use mixing in this period from Player 2 to make sure that the high type Player 1 has the "correct" on-equilibrium path beliefs.

Since L_1 is odd, Player 2 will be offering in period $L_1 + 1$. Consider first when Player 2 will want to offer

$$\left(\frac{\delta(1+\delta^{H-L_1-1})}{1+\delta}, \frac{1-\delta^{H-L_1}}{1+\delta}\right).$$

This is the natural extension of the strategy that she will be playing in all of the future periods where when offering she offers as if both players were high types. Let us assume first that Player 1's beliefs are consistent with the equilibrium outlined above for periods following $L_1 + 1$.

Player 2 has the lowest incentives to deviate from offering

$$\left(\frac{\delta(1+\delta^{H-L_1-1})}{1+\delta}, \frac{1-\delta^{H-L_1}}{1+\delta}\right)$$

if any deviation leads to off-equilibrium beliefs that put full mass on Player 2 being the low type. Given these off-equilibrium beliefs, the low type Player 2 has always the higher incentives to deviate, since her continuation value on the equilibrium path is lower and hence the opportunity cost of deviation is lower. Since deviating benefits Player 2 only in the case where Player 1's type is low, the best possible deviation is to offer $(0,1)$, as the low type is obliged to accept this in equilibrium. Using the construction from the previous section, we get that the low type Player 2 does not want to deviate as long as:

$$\begin{aligned} & (1-\rho_{L_1+1,1}(1-\alpha_{L_1+1,1}))\,\frac{1-\delta^{H-L_1}}{1+\delta} \\ & \quad + \rho_{L_1+1,1}(1-\alpha_{L_1+1,1})\frac{\delta^2(1-\delta^{L_2-L_1-1})}{1+\delta} \\ \geq\ & (1-\rho_{L_1+1,1}) + \rho_{L_1+1,1}\frac{\delta^2(1-\delta^{L_2-L_1-1})}{1+\delta}. \end{aligned} \tag{3.4.10}$$

This inequality holds when

$$\alpha_{L_1+1,1} \geq \frac{\delta^H(1-\rho_{L_1+1,1}) + \delta^{L_2+1}(\delta-\rho_{L_1+1,1}) + (1-\delta)(1+\delta-\rho_{L_1+1,1})\delta^{L_1+1}}{\delta^{L_1}(\rho_{L_1+1,1}-\delta^3) + \delta^{L_2+2} - \delta^H\rho_{L_1+1,1}}. \tag{3.4.11}$$

The right-hand side is at most 1 if and only if

$$\rho_{L_1+1,1} \geq \frac{\delta^{H-L_1}+\delta}{1+\delta^{L_2-L_1+1}+\delta(1-\delta)}. \tag{3.4.12}$$

We discuss in the next section how these beliefs can be guaranteed by mixing either in period 1 or in some cases in period L_1.

In the constructed equilibrium in period L_1 below Player 1 will be offering

$$\left(\frac{1+\delta^{L_2-L_1+2}}{1+\delta}, \frac{\delta(1-\delta^{L_2-L_1+1})}{1+\delta}\right)$$

with positive probability and the low type of Player 2 will be indifferent between accepting and rejecting this. Written in terms of period L_1+1 beliefs, this indifference holds if and only if

$$\begin{aligned}\frac{\delta(1-\delta^{L_2-L_1+1})}{1+\delta} &= \rho_{L_1+1,1}(1-\alpha_{L_1+1,1})\delta^2\frac{\delta(1-\delta^{L_2-L_1-1})}{1+\delta} \\ &+ (1-\rho_{L_1+1,1}(1-\alpha_{L_1+1,1}))\,\delta\frac{1-\delta^{H-L_1}}{1+\delta}.^{6}\end{aligned} \tag{3.4.13}$$

Dividing both sides by δ shows that indifference between accepting and rejecting the equilibrium offer in period L_1 implies that the low type Player 2 is indifferent between offering $\left(\frac{\delta(1+\delta^{L_2-L_1})}{1+\delta}, \frac{1-\delta^{L_2-L_1+1}}{1+\delta}\right)$ or $\left(\frac{\delta(1+\delta^{H-L_1-1})}{1+\delta}, \frac{1-\delta^{H-L_1}}{1+\delta}\right)$ in period L_1+1. Hence, she can mix between the two to make sure that Player 1 has the correct beliefs that allow her to mix in period L_1+2. If we denote the by $\rho^*_{L_1+2,2}$ the solution to the system (3.4.1), (3.4.4), (3.4.5) and (3.4.6) that allows Player 1 to mix in period L_1+2, then the low type Player 2 must be offering the signaling offer $\left(\frac{\delta(1+\delta^{L_2-L_1})}{1+\delta}, \frac{1-\delta^{L_2-L_1+1}}{1+\delta}\right)$ with probability $q_{L_1+1,2}$ that solves Bayesian updating problem:

$$\rho^*_{L_1+2,2} = \frac{\rho_{L_1+1,2}}{\rho_{L_1+1,2}+(1-\rho_{L_1+1,2})(1-q_{L_1+1,2})}.$$

Solving for $q_{L_1+1,2}$ we get that

$$q_{L_1+1,2} = \frac{\rho^*_{L_1+2,2}-\rho_{L_1+1,2}}{\rho^*_{L_1+2,2}(1-\rho_{L_1+1,2})},$$

which is in $[0,1]$ whenever

$$\rho_{L_1+1,2} \in [0, \rho^*_{L_1+2,2}].$$

Hence, the equilibrium exists as long as (3.4.12) is satisfied and the prior in period L_1+1 is at most $\rho^*_{L_1+2,2}$. We concentrate mostly on this case as acceptance behavior in previous periods can be chosen to guarantee that this holds. Whenever, it does not hold one can support an equilibrium where both types of Player 2 pool at offering $(0,1)$ in period L_1+1 and only the low type of Player 1 accepts it. After that the game continues as explained in the previous section.

Summarizing the results so far, we get the following lemma:

Lemma 3.4.2. *If* $\rho_{L_1+1,2} \in \left[\frac{\delta^{H-L_1}+\delta}{1+\delta^{L_2-L_1+1}+\delta(1-\delta)}, \rho^*_{L_1+2,2}\right]$ *where* $\rho^*_{L_1+2,2}$ *is defined as above, then there exists a trembling hand perfect equilibrium for the sub-game starting in period* L_1+1 *where in each even period* L_2-k *both types of Player 2 offer* $\left(\frac{\delta(1+\delta^{H-L_2+k})}{1+\delta}, \frac{1-\delta^{H-L_2+k+1}}{1+\delta}\right)$ *and Player 1 accepts this with a non-zero probability. In odd periods* L_2-k*, Player 1 will offer* $\left(\frac{1+\delta^{k+2}}{1+\delta}, \frac{\delta(1-\delta^{k+1})}{1+\delta}\right)$ *which the high type of Player 2 will always reject and the low type accepts with a positive probability.*

3.4.3 Equilibrium Strategies for Periods $t \leq L_1$

The construction of equilibrium screening strategies for these periods is fairly similar to the symmetric case. Suppose $t \leq L_1$ and that t is even. We will look for an equilibrium where both types of Player 2 make an offer as if Player 1 was known to be the low type:

$$\left(\frac{\delta(1-\delta^{L_1-t+1})}{1+\delta}, \frac{1+\delta^{L_1-t+2}}{1+\delta}\right).$$

The high type Player 1 always rejects this, while the low type accepts it with probability $\alpha_{t,1}$. Deviations to other offers will lead Player 1 to believe that Player 2 is the low type with probability 1. The game will then continue analogously to the one-sided screening equilibrium from Section 3.4.1.

On the equilibrium path, in period $t+1$ both types of Player 1 will offer

$$\left(\frac{1+\delta^{L_2-t+1}}{1+\delta}, \frac{\delta(1-\delta^{L_2-t})}{1+\delta}\right),$$

which corresponds to the Rubinstein offer when Player 2 is known to be the low type and Player 1's type is high. The high type of Player 2 always rejects this while the low type accepts with probability $\alpha_{t,2}$. Similarly to what happens in period t, deviations to other offers lead to Player 2 believing that Player 1 is the low type with probability 1. Given that even Player 2's low type is higher than Player 1's low type, this will either lead to acceptance in period $t+1$ or the game ending in the following period with the offer

$$\left(\frac{\delta(1-\delta^{L_1-t-1})}{1+\delta}, \frac{1+\delta^{L_1-t}}{1+\delta}\right),$$

whichever yields the higher present value for Player 2 in that period.

The low type Player 1 chooses $\alpha_{t,1}$ so that the low type Player 2 is indifferent between accepting and rejecting in period $t-1$. This is true if and only if

$$\begin{aligned} & \frac{\delta(1-\delta^{L_2-t+2})}{1+\delta} \\ = \; & \delta\left(\alpha_{t,1}(1-\rho_{t,1})\frac{1+\delta^{L_1-t+2}}{1+\delta} + (1-\alpha_{t,1}(1-\rho_{t,1}))\frac{\delta^2(1-\delta^{L_2-t})}{1+\delta}\right) \end{aligned}$$

Solving this for $\alpha_{t,1}$ yields

$$\alpha_{t,1} = \frac{(1-\delta^2)\delta^t}{(1-\rho_{t,1})\left((1-\delta^2)\delta^t + \delta^2(\delta^{L_2}+\delta^{L_1})\right)}, \tag{3.4.14}$$

which is less than 1 as long as

$$\rho_{t-1,1} = \rho_{t,1} < \frac{\delta^2(\delta^{L_2}+\delta^{L_1})}{(1-\delta^2)\delta^t + \delta^2(\delta^{L_2}+\delta^{L_1})}. \tag{3.4.15}$$

Notice that the right-hand side converges to 1 as $\delta \to 1$ and hence the condition is satisfied for high levels of patience.

The Bayes' law implies that

$$\rho_{t+1,1} = \frac{\rho_{t,1}}{\rho_{t,1} + (1-\rho_{t,1})(1-\alpha_{t,1})}.$$

Substituting in $\alpha_{t,1}$ from (3.4.14) yields the following recursive equation describing the evolution of Player 2's beliefs:

$$\rho_{t+1,1} = \begin{cases} \rho_{t,1} & t \text{ odd} \\ \left(1 + \frac{\delta^t(1-\delta^2)}{\delta^2(\delta^{L_2}+\delta^{L_1})}\right)\rho_{t,1} & t \text{ even.} \end{cases} \tag{3.4.16}$$

Solving forward for $\rho_{t,1}$ yields

$$\rho_{t,1} = \begin{cases} \rho_{1,1}\prod_{k=1}^{(t-2)/2}\left(1 + \frac{\delta^{2k}(1-\delta^2)}{\delta^2(\delta^{L_2}+\delta^{L_1})}\right) & t \text{ even} \\ \rho_{1,1}\prod_{k=1}^{(t-1)/2}\left(1 + \frac{\delta^{2k}(1-\delta^2)}{\delta^2(\delta^{L_2}+\delta^{L_1})}\right) & t \text{ odd.} \end{cases} \tag{3.4.17}$$

These beliefs can be used with (3.4.14) to fully describe low type Player 1's mixing probability for each period.

Similarly, low type Player 2 chooses $\alpha_{t+1,2}$ to make the low type of Player 1 indifferent between accepting and rejecting the offer made in period t. When $t < L_1 - 1$, the low type Player 1 is indifferent if and only if

$$\begin{aligned} & \frac{\delta(1-\delta^{L_1-t+1})}{1+\delta} \\ = \; & \delta\left(\alpha_{t+1,2}(1-\rho_{t+1,2})\frac{1+\delta^{L_2-t+1}}{1+\delta} + (1-\alpha_{t+1,2}(1-\rho_{t+1,2}))\frac{\delta^2(1-\delta^{L_1-t-1})}{1+\delta}\right). \end{aligned}$$

Solving this for Player 2's mixing probability yields

$$\alpha_{t+1,2} = \frac{(1-\delta^2)\delta^t}{(1-\rho_{t+1,2})\left((1-\delta^2)\delta^t + \delta(\delta^{L_2}+\delta^{L_1})\right)}. \tag{3.4.18}$$

This is less than one as long as

$$\rho_{t,2} = \rho_{t+1,2} < \frac{\delta(\delta^{L_2} + \delta^{L_1})}{(1-\delta^2)\delta^t + \delta(\delta^{L_2} + \delta^{L_1})}. \tag{3.4.19}$$

Combining the mixing probability implied by (3.4.18) for period $t-1$ with Bayes' rule gives us a recursive description of Player 1's equilibrium path beliefs:

$$\rho_{t,2} = \begin{cases} \rho_{t-1,2} & t \text{ odd} \\ \left(1 + \frac{\delta^t(1-\delta^2)}{\delta^3\left(\delta^{L_2}+\delta^{L_1}\right)}\right)\rho_{t-1,2} & t \text{ even.} \end{cases} \tag{3.4.20}$$

Just as above, this can be solved forward to yield:

$$\rho_{t,2} = \begin{cases} \rho_{1,2}\prod_{k=1}^{t/2}\left(1 + \frac{\delta^{2k}(1-\delta^2)}{\delta^3\left(\delta^{L_2}+\delta^{L_1}\right)}\right) & t \text{ even} \\ \rho_{1,2}\prod_{k=1}^{(t-1)/2}\left(1 + \frac{\delta^{2k}(1-\delta^2)}{\delta^3\left(\delta^{L_2}+\delta^{L_1}\right)}\right) & t \text{ odd,} \end{cases} \tag{3.4.21}$$

where we apply the usual convention that the empty product equals 1.

Next consider periods L_1 and $L_1 - 1$. We will look for an equilibrium where in period L_1 the low type of Player 1 submits $(1-\delta, \delta)$ with probability $q_{L_1,1}$ and offers $\left(\frac{1+\delta^{L_2-L_1+2}}{1+\delta}, \frac{\delta(1-\delta^{L_2-L_1+1})}{1+\delta}\right)$ with the remaining probability. The high type chooses the latter offer with probability 1. This $q_{L_1,1}$ is chosen to make sure that in period L_1+1 the low type does not want to deviate to offering $(0,1)$. Let α_{1,L_1+1} be the probability with which the high type Player 1 accepts

$$\left(\frac{\delta(1+\delta^{H-L_1-1})}{1+\delta}, \frac{1-\delta^{H-L_1}}{1+\delta}\right)$$

in period L_1+1. Given the equilibrium for the following subgame, after observing

$$\left(\frac{1+\delta^{L_2-L_1+2}}{1+\delta}, \frac{\delta(1-\delta^{L_2-L_1+1})}{1+\delta}\right)$$

in period L_1, the low type Player 2 is indifferent between accepting and rejecting this offer if and only if

$$\frac{\delta(1-\delta^{L_2-L_1+1})}{1+\delta} = \frac{\rho_{L_1,1}(1-\alpha_{L_1+1,1})}{\rho_{L_1,1}+(1-q_{L_1,1})(1-\rho_{L_1,1})}\delta^2\frac{\delta(1-\delta^{L_2-L_1-1})}{1+\delta} + \left(1-\frac{\rho_{L_1,1}(1-\alpha_{L_1+1,1})}{\rho_{L_1,1}+(1-q_{L_1,1})(1-\rho_{L_1,1})}\right)\delta\frac{1-\delta^{H-L_1}}{1+\delta}. \tag{3.4.22}$$

Whenever this holds, dividing by δ we get that (3.4.10) holds as long as

$$\frac{(1-\delta^{L_2-L_1+1})}{1+\delta} \geq \frac{\rho_{L_1,1}}{\rho_{L_1,1}+(1-q_{L_1,1})(1-\rho_{L_1,1})}\delta\frac{\delta(1-\delta^{L_2-L_1-1})}{1+\delta} + \left(1-\frac{\rho_{L_1,1}}{\rho_{L_1,1}+(1-q_{L_1,1})(1-\rho_{L_1,1})}\right). \tag{3.4.23}$$

Solving for $q_{L_1,1}$ yields

$$q_{L_1,1} \geq 1-\frac{(1-\delta^2)\rho_{L_1,1}}{\delta(1+\delta^{L_2-L_1})(1-\rho_{L_1,1})}. \tag{3.4.24}$$

Notice the right hand is always below one, and at most zero when

$$\rho_{L_1,1} \geq \frac{\delta(1+\delta^{L_2-L_1})}{1+\delta-\delta^2+\delta^{L_2-L_1+1}}. \tag{3.4.25}$$

First, consider the case when (3.4.25) holds. Then one can choose $q_{L_1,1}=0$. Substituting this into (3.4.22) and solving for $\alpha_{L_1+1,1}$ yields

$$\alpha_{L_1+1,1} = \frac{(1-\rho_{L_1,1})(\delta^{L_2+1}-\delta^H)-\rho_{L_1,1}\delta^{L_1}(1-\delta^2)}{(\delta^H-\delta^{L_1}(1-\delta^2)-\delta^{L_2+1})\rho_{L_1,1}} \tag{3.4.26}$$

which is at least zero as long as

$$\rho_{L_1,1} \geq \frac{\delta^{L_2+1}-\delta^H}{\delta^{L_2+1}-\delta^H+\delta^{L_1}(1-\delta^2)}.$$

It is easy to show that this condition is implied by (3.4.25). If the high type Player 1 uses the mixing probability above, neither type of the Player 2 clearly has an incentive to deviate in any previous period from the strategy described earlier in this section.

Mixing is optimal for the high type of Player 1 in period $L_1 + 1$ as long as her beliefs are as described in the previous section. In period L_1 the low type Player 1 is willing to mix between the two offers if the low type of Player 2 accepts with probability $\alpha_{L_1,2}$ after observing

$$\left(\frac{1+\delta^{L_2-L_1+2}}{1+\delta}, \frac{\delta(1-\delta^{L_2-L_1+1})}{1+\delta}\right)$$

such that

$$\alpha_{L_1,2}\frac{1+\delta^{L_2-L_1+2}}{1+\delta} = 1-\delta.$$

In other words, we must have

$$\alpha_{L_1,2} = \frac{1-\delta^2}{1+\delta^{L_2-L_1+2}}. \tag{3.4.27}$$

Notice that this mixing probability also implies that the low type is trivially indifferent between accepting and rejecting $(\delta(1-\delta), 1-\delta(1-\delta))$ being offered in period $L_1 - 1$. Hence, given this mixing probability and the strategy described above for Player 2, the low type will not have an incentive to deviate in period L_1 or any period before it.

Consider next the case when

$$\rho_{L_1,1} < \frac{\delta(1+\delta^{L_2-L_1})}{1+\delta-\delta^2+\delta^{L_2-L_1+1}}. \tag{3.4.28}$$

Then the mixing probability that puts the least restrictions on beliefs in previous periods is given by

$$q_{L_1,1} = 1 - \frac{(1-\delta^2)\rho_{L_1,1}}{\delta(1+\delta^{L_2-L_1})(1-\rho_{L_1,1})}. \tag{3.4.29}$$

Substituting this into (3.4.22) yields

$$\alpha_{L_1+1,1} = \frac{(1-\delta^2)(\delta^H + \delta^{L_1+1})}{\delta(1+\delta^{L_2-L_1})(\delta^{L_1}(1-\delta^2) + \delta^{L_2+1} - \delta^H)},$$

which is always clearly between zero and one.

Since the low type of Player 1 is now offering more to Player 2 in period L_1 than in the previous case, Player 2's continuation value in period $L_1 - 2$ after a rejection will change unless we also adjust Player 1's acceptance probability in period $L_1 - 1$. The low type of Player 2 remains indifferent between accepting and rejecting in period $L_2 - 2$ if

$$\begin{aligned}
&\frac{\delta(1-\delta^{L_2-L_1+4})}{1+\delta} \\
= \;& \delta\Bigg(\alpha_{L_1-1,1}(1-\rho_{L_1-1,1})(1-\delta+\delta^2) + (1-\alpha_{L_1-1,1})(1-\rho_{L_1-1,1})q_{L_1,1}\delta^2 \\
&+ (\rho_{L_1-1,1} + (1-\alpha_{L_1-1,1})(1-\rho_{L_1-1,1})(1-q_{L_1,1}))\,\frac{\delta^2(1-\delta^{L_2-L_1+2})}{1+\delta}\Bigg).
\end{aligned}$$

Combining this with (3.4.29) and solving for $\alpha_{L_1-1,1}$ yields

$$\begin{aligned}
\alpha_{L_1-1,1} \;=\; & \Big((1+\delta^{L_2-L_1})(1-\delta^2-\delta^3) \\
&+ (1+\rho_{L_1-1,1})\delta^{L_2-L_1+4}(1+\delta^{L_2-L_1}) \\
&+ \rho_{L_1-1,1}\delta^2(1+\delta-\delta^2+2\delta^{L_2-L_1+1}-\delta^{L_2-L_1+2})\Big) \\
&\times \left((1-\delta^2)(1+\delta^{L_2-L_1})(1-\rho_{L_1-1,1})\right)^{-1}
\end{aligned} \tag{3.4.30}$$

This is at least zero as long as

$$\rho_{L_1-1,1} \geq \frac{(1+\delta^{L_2-L_1})(\delta^2(1+\delta)-\delta^{L_2-L_1+4}-1)}{\delta^2(1+\delta^{L_2-L_1+1})(1+\delta(1-\delta+\delta^{L_2-L_1}))}.$$

This bound is always below the bound implied by (3.4.28), Bayes' Law and (3.4.30) and hence there exists a range of beliefs for which this is an equilibrium.

From Player 1's perspective the game is just as in the previous case and hence the relevant indifference condition for period L_1 is still given by (3.4.27).

The non-zero $q_{L_1,1}$ lifts the continuation value of the low type Player 2 in period L_1-2 which in turn requires the low type of Player 1 to accept with a lower probability in period L_1-1 to guarantee that Player 2 is willing to accept the offer in L_1-2. If the low type Player 1 is too likely and and Player 2 is sufficiently patient, even setting $\alpha_{L_1-1,1}=0$ is not enough to get the low type of Player 2 to accept in period L_1-2. However, there is an easy fix to this: If Player 2 believes that Player 1 is very likely to be the low type in the beginning of the game, Player 1 can mix in period 1 to guarantee that even condition (3.4.25) holds. Mixing in the first period does not increase the continuation value in future periods (it in fact reduces it) and make rejecting of future offers more attractive. Hence having the right mixing probability in period 1 guarantees an equilibrium of the first type where $q_{L_1,1}=0$.

The following theorem summarizes the on-path properties of the above equilibrium:

Theorem 3.4.3. *As long as the agents are patient enough so that conditions (3.4.15) and (3.4.19) hold for all $t \leq L_1$, then there exists an equilibrium where in each even period $t \leq L_1$ both types of Player 2 offer $\left(\frac{\delta(1-\delta^{L_1-t+1})}{1+\delta}, \frac{1+\delta^{L_1-t+2}}{1+\delta}\right)$, the high type of Player 1 always rejects and the low type accepts with a non-zero probability. Similarly, in every odd period $t \leq L_1$ both types of Player 1 offer $\left(\frac{1+\delta^{L_2-t+2}}{1+\delta}, \frac{\delta(1-\delta^{L_2-t+1})}{1+\delta}\right)$ which the high type of Player 2 always rejects and the low type accepts with a non-zero prob-*

ability. After period L_1 the game continues with the equilibrium outlined in Lemma 3.4.2.

3.5 Discussion of the Properties of the Equilibria

3.5.1 Efficiency Properties of the Equilibria

The most notable feature of both the equilibrium for the symmetric and asymmetric versions of the game are that in a generic round the low types of both players get in expectation what they would get in a game where they were known to be the low type and their opponent was known to be high type. Interestingly, the low type gains no usual information rents from probabilistically revealing her type.

In the symmetric version of the game, when the game ends the high type always gets what she would have gotten in a game with perfect information that started in the resolution period. In this sense the information friction causes screening through waiting which in turn is wasteful even in a strong Pareto sense: If both players could commit to revealing their types ex-ante in the second period, both types of both players would strictly prefer to do so.[7] From the policy point of view this implies that arbitration can imply high efficiency gains if there are mechanisms through which the players can reliably submit their private information to the arbitrator.

In the symmetric version of the model, only the low type of Player 1 can enjoy information rents and this happens only if her prior is that Player 2 is the low type with a too high a probability. In this case, her first period expected utility can be considerably more than what she would expect in a game where she is known to be the low type and her opponent is known to be the high type.

[7]The high type wants to do this to avoid waiting, the low type to benefit from the non-zero probability that also her opponent is the low type.

When one allows for asymmetric deadlines the pictures changes slightly. If the game passes the lower of the two low types' deadlines and nobody has still accepted any offers, the player whose deadline has passed is now known to be the low type. To enable screening she still needs to mix, which in turn requires very specific equilibrium path beliefs after her type is revealed. This implies that the low type of the player whose type is still unknown has to signal her type also with her offers in the period just after first deadline has passed. This in turn may require her to signal her type also in the first period if her opponent's prior favors the low type too heavily. Hence, only the lowest possible type, the type whose deadline is the first enjoys potential information rents and only in the very first period of the game and this happens only if that player has too low a prior on her opponent being the high type.

3.5.2 Trembling Hand Perfection and Conjecture on Uniqueness

Our equilibrium is clearly not the unique perfect Bayesian equilibrium of the game. Most notably, by using harsh off-equilibrium beliefs one can support even pooling equilibria where the game ends in the first period with one of the players walking away with most of the surplus. What is common to these pooling equilibria and many other equilibria that one can construct in the game is that they do not survive if players are relatively patient unless deviations lead to all of the probability mass being placed on the deviator's low type. In the subgame starting from the period prior to a player's deadline, a patient high type has an incentive to reject any offers that yield her less than what is guaranteed by the equilibrium described above as long as any mass is placed on the high type after such rejection. This follows, because in the following period the low type will concede the whole surplus to the opponent and any behavior different from this signals the player's high type.

As long as the players are patient enough, high types being screened in the period prior to the deadline is likely to imply willingness to signal strength by the high types in the previous period. Inductively then the only offer that the high type is going to accept or make are such that only the opponent's low type is willing to accept them. Hence, whenever deviations are required to leave any mass on the deviator's high type we conjecture that, given players that are patient enough, the equilibrium outcome will be identical to the one outlined above. In other words, we expect that for high enough a δ our equilibrium path outcome is the unique trembling hand perfect equilibrium outcome (Selten, 1975).

3.5.3 The Effect of Deadlines on Bargaining Power

Interestingly, the mixing probabilities in the asymmetric version of the game are highly symmetric. Most notably, the low type's mixing probability depends on the deadlines only through the sum $\delta^{L_1} + \delta^{L_2}$. The ordering of the two early deadlines is only reflected in the offers being made to the player and not in the mixing chosen by that player's low type.

The beliefs, however, are reflected in the mixing probabilities. The more weight the opponent puts on the player's high type, the higher must be the acceptance probability of the low type. Notice that this does not imply that if a player is more likely to be the high type, then the player's low type has less bargaining power. As was noted above, the low type's expectation is exactly equal to the equilibrium allocation from a complete information game where she is known to be the low type and her opponent is known to be the high type. Consequently, the low type's expected utility in the equilibrium is fully determined by how close her deadline is.

3.6 Conclusions and Future Research

This paper constructs an intuitive and tractable partial-screening equilibrium in an alternating-offer bargaining game with two-sided asymmetric information. By having the private information be about a deadline, a time after which players will accept all feasible offers, we hope to argue that this class of equilibria is unique with only a mild refinement: trembling-hand perfection.

Our equilibrium exhibits delay through a war-of-attrition structure: players will insist on offers which are made by the highest type, while low types will sometimes relent and accept these offers. We show how the screening through waiting is potentially highly inefficient from the point of view of both types of both players. This result indicates that there is potential for incentive compatible information revelation mechanisms to help us avoid lengthy bargaining impasses. Naturally, designing arbitration mechanisms that allow players to attain the Coasian outcome would yield higher welfare, but without such a mechanism high types of both player may prefer to resort to the costly screening that our equilibrium describes.

Finally, several open questions remain to be investigated. Firstly, the analysis so far has focused on just two types. We believe the basic insights can be extended to an arbitrary number of types, where perhaps with fixed period length some types will prefer to pool together (when the equivalent of the conditions in our theorem 1 fail to be satisfied). This pooling should disappear as we decrease the time between offers to zero and move towards a continuous-time bargaining model. The above intuition comes from the literature on the war of attrition. A second open question is how closely a continuous-time limit of our alternating offer bargaining model will resemble the war of attrition studied by Hendricks et al. (1988) in a way similar to Abreu and Gul (2000). Further research should also address the implications of the trembling-hand perfection refinement in continuous-time bargaining models.

Bibliography

Abbring, J. H., P.-A. Chiappori, and J. Pinquet (2003): "Moral Hazard and Dynamic Insurance Data," *Journal of the European Economic Association*, 1(4), 767–820.

Abbring, J. H., J. J. Heckman, P.-A. Chiappori, and J. Pinquet (2003): "Adverse Selection and Moral Hazard in Insurance: Can Dynamic Data Help to Distinguish?," *Journal of the European Economic Association*, 1(2-3), 512–521.

Abreu, D., and F. Gul (2000): "Bargaining and Reputation," *Econometrica*, 68(1), 85–117.

Adair, J. G. (1984): "The Hawthorne effect: A reconsideration of the methodological artifact," *Journal of Applied Psychology*, 69(2), 334–345.

Adams, W., L. Einav, and J. Levin (2009): "Liquidity Constraints and Imperfect Information in Subprime Lending," *American Economic Review*, 99(1), 49–84.

Bainbridge, W. S. (2007): "The Scientific Research Potential of Virtual Worlds," *Science*, 317(5837), 472–476.

Bertsekas, D. P. (1995): *Dynamic programming and optimal control*. Athena Scientific Belmont, MA.

Briys, E., and H. Schlesinger (1990): "Risk Aversion and the Propensities for Self-Insurance and Self-Protection," *Southern Economic Journal*, 57(2), 458–467.

Cabral, M. (2016): "Claim timing and ex post adverse selection," Unpublished working paper, retrieved from: `http://www.marikacabral.com/Cabral_ExPostAdverseSelection.pdf`.

Castronova, E. (2001): "Virtual Worlds: A First-Hand Account of Market and Society on the Cyberian Frontier," SSRN Scholarly Paper ID 294828, Social Science Research Network, Rochester, NY.

——— (2004): "The Price of Bodies: A Hedonic Pricing Model of Avatar Attributes in a Synthetic World," *Kyklos*, 57(2), 173–196.

Chesney, T., S.-H. Chuah, and R. Hoffmann (2009): "Virtual world experimentation: An exploratory study," *Journal of Economic Behavior & Organization*, 72(1), 618–635.

Chiappori, P., and B. Salanié (2000): "Testing for Asymmetric Information in Insurance Markets," *Journal of Political Economy*, 108(1), 56–78.

Cohen, A. (2005): "Asymmetric Information and Learning: Evidence from the Automobile Insurance Market," *The Review of Economics and Statistics*, 87(2), 197–207.

Cohen, A., and P. Siegelman (2010): "Testing for Adverse Selection in Insurance Markets," *Journal of Risk and Insurance*, 77(1), 39–84.

Cox, D. R. (1972): "Regression Models and Life-Tables," *Journal of the Royal Statistical Society. Series B (Methodological)*, 34(2), 187–220.

——— (1975): "Partial likelihood," *Biometrika*, 62(2), 269–276.

Cramton, P. C. (1992): "Strategic Delay in Bargaining with Two-Sided Uncertainty," *The Review of Economic Studies*, 59(1), 205–225.

Dionne, G., and L. Eeckhoudt (1985): "Self-insurance, self-protection and increased risk aversion," *Economics Letters*, 17(1), 39–42.

Dionne, G., P.-C. Michaud, and M. Dahchour (2013): "Separating Moral Hazard from Adverse Selection and Learning in Automobile Insurance: Longitudinal Evidence from France," *Journal of the European Economic Association*, 11(4), 897–917.

Dobbie, W., and P. M. Skiba (2013): "Information Asymmetries in Consumer Credit Markets: Evidence from Payday Lending," *American Economic Journal: Applied Economics*, 5(4), 256–282.

Edmans, A., and X. Gabaix (2011): "Tractability in Incentive Contracting," *Review of Financial Studies*, 24(9), 2865–2894.

Einav, L., and A. Finkelstein (2010): "Beyond Testing: Empirical Models of Insurance Markets," *Annual review of economics*, 2(1), 311–336.

Einav, L., A. Finkelstein, R. Kluender, and P. Schrimpf (2016): "Beyond Statistics: The Economic Content of Risk Scores," *American Economic Journal: Applied Economics*, 8(2), 195–224.

Einav, L., A. Finkelstein, and P. Schrimpf (2015): "The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D," *The Quarterly Journal of Economics*, 130(2), 841–899.

Einav, L., M. Jenkins, and J. Levin (2012): "Contract Pricing in Consumer Credit Markets," *Econometrica*, 80(4), 1387–1432.

Fang, H., M. Keane, and D. Silverman (2008): "Sources of Advantageous Selection: Evidence from the Medigap Insurance Market," *Journal of Political Economy*, 116(2), 303–350.

Fanning, J. (2013): "Reputational bargaining and deadlines," *Unpublished paper, New York University.*

Fiedler, M., and E. Haruvy (2009): "The Lab Versus the Virtual Lab and Virtual Field – An Experimental Investigation of Trust Games with Communication," *Journal of Economic Behavior & Organization*, 72(2), 716–724.

Fiedler, M., E. Haruvy, and S. X. Li (2011): "Social distance in a virtual world experiment," *Games and Economic Behavior*, 72(2), 400–426.

Finkelstein, A., K. McGarry, and A. Sufi (2005): "Dynamic Inefficiencies in Insurance Markets: Evidence from Long-Term Care Insurance," *American Economic Review*, 95(2), 224–228.

Fuchs, W., and A. Skrzypacz (2013): "Bargaining with Deadlines and Private Information," *American Economic Journal: Microeconomics*, 5(4), 219–243.

Gibbons, R., and K. J. Murphy (1992): "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," *Journal of Political Economy*, 100(3), 468–505.

Golde, S. (2010): "The Economics of Market Thickness," Ph.D. thesis, Stanford University.

Hart, O. (1989): "Bargaining and strikes," *The Quarterly Journal of Economics*, pp. 25–43.

Hendricks, K., A. Weiss, and C. Wilson (1988): "The War of Attrition in Continuous Time with Complete Information," *International Economic Review*, 29(4), 663–680.

Holmström, B. (1979): "Moral Hazard and Observability," *The Bell Journal of Economics*, 10(1), 74–91.

Horton, J. J., D. G. Rand, and R. J. Zeckhauser (2011): "The online laboratory: conducting experiments in a real labor market," *Experimental Economics*, 14(3), 399–425.

Howden, L., and J. Meyer (2011): "Age and Sex Composition: 2010," Discussion paper, US Census Bureau.

Kalbfleisch, J. D., and R. L. Prentice (2011): *The Statistical Analysis of Failure Time Data*. John Wiley & Sons.

Pauly, M. V. (1968): "The Economics of Moral Hazard: Comment," *The American Economic Review*, 58(3), 531–537.

Phelan, C., and R. M. Townsend (1991): "Computing Multi-Period, Information-Constrained Optima," *The Review of Economic Studies*, 58(5), 853–881.

Ridder, G., and I. Tunali (1999): "Stratified partial likelihood estimation," *Journal of Econometrics*, 92(2), 193–232.

Roberts, A. W., and D. E. Varberg (1973): *Convex Functions*. Academic Press Inc, New York.

Rothschild, M., and J. Stiglitz (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *The Quarterly Journal of Economics*, 90(4), 629–649.

Rubinstein, A. (1982): "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50(1), 97–109.

Sandholm, T., N. Vulkan, et al. (1999): "Bargaining with deadlines," in *AAAI/IAAI*, pp. 44–51.

Sannikov, Y. (2008): "A Continuous-Time version of the Principal-Agent Problem," *The Review of Economic Studies*, 75(3), 957–984.

Selten, R. (1975): "Reexamination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory*, 4(1), 25–55.

Simsek, A., and M. Yildiz (2014): "Durability, deadline, and election effects in bargaining," *Unpublished manuscript*.

Ståhl, I. (1972): *Bargaining theory*. Economic Research Institute at the Stockholm School of Economics (EFI), Stockholm.