

NBER WORKING PAPER SERIES

PRIVATE INFORMATION AND INSURANCE REJECTIONS

Nathaniel Hendren

Working Paper 18282

<http://www.nber.org/papers/w18282>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

August 2012

An earlier version of this paper is contained in the first chapter of my MIT graduate thesis. I am very grateful to Daron Acemoglu, Amy Finkelstein, Jon Gruber, and Rob Townsend for their guidance and support in writing this paper. I also thank Victor Chernozhukov, Sarah Miller, Whitney Newey, Ivan Werning, an extensive list of MIT graduate students, and seminar participants at The University of California-Berkeley, Chicago Booth, The University of Chicago, Columbia, Harvard, Microsoft Research New England, Northwestern, The University of Pennsylvania, Princeton, and Stanford for helpful comments and suggestions. I would also like to thank several anonymous insurance underwriters for helpful assistance. Financial support from NSF Graduate Research Fellowship and the NBER Health and Aging Fellowship, under the National Institute of Aging Grant Number T32-AG000186 is gratefully acknowledged. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2012 by Nathaniel Hendren. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Private Information and Insurance Rejections
Nathaniel Hendren
NBER Working Paper No. 18282
August 2012
JEL No. H0,I11,I13

ABSTRACT

Across a wide set of non-group insurance markets, applicants are rejected based on observable, often high-risk, characteristics. This paper argues private information, held by the potential applicant pool, explains rejections. I formulate this argument by developing and testing a model in which agents may have private information about their risk. I first derive a new no-trade result that theoretically explains how private information could cause rejections. I then develop a new empirical methodology to test whether this no-trade condition can explain rejections. The methodology uses subjective probability elicitation as noisy measures of agents' beliefs. I apply this approach to three non-group markets: long-term care, disability, and life insurance. Consistent with the predictions of the theory, in all three settings I find significant amounts of private information held by those who would be rejected; I find generally more private information for those who would be rejected relative to those who can purchase insurance; and I show it is enough private information to explain a complete absence of trade for those who would be rejected. The results suggest private information prevents the existence of large segments of these three major insurance markets.

Nathaniel Hendren
Harvard University
Department of Economics
Littauer Center Room 235
Cambridge, MA 02138
and NBER
nhendren@gmail.com

An online appendix is available at:
<http://www.nber.org/data-appendix/w18282>

1 Introduction

Not everyone can purchase insurance. Across a wide set of non-group insurance markets, companies choose to not sell insurance to potential customers with certain observable, often high-risk, characteristics. In the non-group health insurance market, 1 in 7 applications to the four largest insurance companies in the US were rejected between 2007 and 2009, a figure that excludes those who would be rejected but were deterred from even applying.¹ In US long-term care insurance, 12-23% of 65 year olds have health conditions that would preclude them from being able to purchase insurance (Murtaugh et al. [1995]).²

It is surprising that a company would choose to not offer its products to a certain subpopulation. Although the rejected generally have higher expected expenditures, they still face unrealized risk.³ Regulation does not generally prevent risk-adjusted pricing in these markets.⁴ So why not simply offer them a higher price?

In this paper, I argue that private information, held by the potential applicant pool, explains rejections. In particular, I provide empirical evidence in three insurance market settings that those who have observable conditions that prevent them from being able to purchase insurance also have additional knowledge about their risk beyond what is captured by their observable characteristics. To develop some intuition for this finding, consider the risk of going to a nursing home, one of the three settings that will be studied in this paper. Someone who has had a stroke, which renders them ineligible to purchase long-term care (LTC) insurance, may know not only her personal medical information (which is largely observable to an insurer), but also many specific factors and preferences that are derivatives of her health condition and affect her likelihood of entering a nursing home. These could be whether her kids will take care of her in her condition, her willingness to engage in physical therapy or other treatments that would prevent nursing home entry, or her desire to live independently with the condition as opposed to seek the aid of a nursing home. Such factors and preferences affect the cost of insuring nursing home expenses, but

¹Figures obtained through a formal congressional investigation by the Committee on Energy and Commerce, which requested and received this information from Aetna, Humana, UnitedHealth Group, and WellPoint. Congressional report was released on October 12, 2010. The 1 in 7 figure does not subtract duplicate applications if people applied to more than 1 of these 4 firms.

²Appendix F presents the rejection conditions from Genworth Financial (one of the largest US LTC insurers), gathered from their underwriting guidelines provided to insurance agents for use in screening applicants.

³For example, in long-term care I will show that those who would be rejected have an average five-year nursing home entry rate of less than 25%.

⁴The Civil Rights Act is a singular exception as it prevents purely race-based pricing.

are often difficult an insurance company to obtain and verify. This paper will argue that, because of the private information held by those with rejection conditions, if an insurer were to offer contracts to these individuals, they would be so heavily adversely selected that it wouldn't deliver positive profits, at any price.

To make this argument formally, I begin with a theory of how private information could lead to rejections. The setting is the familiar binary loss environment introduced by Rothschild and Stiglitz [1976], generalized to incorporate an arbitrary distribution of privately informed types. In this environment, I ask under what conditions can anyone obtain insurance against the loss. I derive new a "no-trade" condition characterizing when insurance companies would be unwilling to sell insurance on terms that anyone would accept. This condition has an unraveling intuition similar to the one introduced in Akerlof [1970]. The market unravels when the willingness to pay for a small amount of insurance is less than the pooled cost of providing this insurance to those of equal or higher risk. When this no-trade condition holds, an insurance company cannot offer any contract, or menu of contracts, because they would attract an adversely selected subpopulation that would make them unprofitable. Thus, the theory explains rejections as market segments (segmented by observable characteristics) in which the no-trade condition holds.

I then use the no-trade condition to identify properties of type distributions that are more likely to lead to no trade. This provides a vocabulary for quantifying private information. In particular, I characterize the barrier to trade imposed by a distribution of types in terms of the implicit tax rate, or markup, individuals would have to be willing to pay on insurance premiums in order for the market to exist. The comparative statics of the theory suggests the implicit tax rates should be higher for the rejectees relative to non-rejectees and high enough for the rejectees to explain an absence of trade for plausible values of the willingness to pay for insurance.

I then develop a new empirical methodology to test the predictions of theory. I use information contained in subjective probability elicitation⁵ to infer properties of the distribution of private information. I do not assume individuals can necessarily report their true beliefs. Rather, I use information in the joint distribution of elicitation and the realized events corresponding to these elicitation to deal with potential errors in elicitation.

I proceed with two complementary empirical approaches. First, I estimate the explanatory power of the subjective probabilities on the subsequent realized event, conditional on

⁵A subjective probability elicitation about a given event is a question: "What is the chance (0-100%) that [event] will occur?".

public information. I show that measures of their predictive power provide nonparametric lower bounds on theoretical metrics of the magnitude of private information. In particular, whether the elicitation is predictive at all provides a simple test for the presence of private information. I also provide a test in the spirit of the comparative static of the theory that asks whether those who would be rejected are better able to predict their realized loss.

Second, I estimate the distribution of beliefs by parameterizing the distribution of elicitation given true beliefs (i.e. on the distribution of measurement error). I then quantify the implicit tax individuals would need to be willing to pay in order for an insurance company to be able to profitably sell insurance against the corresponding loss. I then ask whether it is larger for those who would be rejected relative to those who are served by the market and whether it is large (small) enough to explain (the absence of) rejections for plausible values of agents' willingness to pay for insurance.

I apply this approach to three non-group markets: long-term care (LTC), disability, and life insurance. I combine two sources of data. First, I use data from the Health and Retirement Study, which elicits subjective probabilities corresponding to losses insured in each of these three settings and contains a rich set of observable demographic and health information. Second, I construct and merge a classification of those who would be rejected (henceforth "rejectees"⁶) in each market from a detailed review of underwriting guidelines from major insurance companies.

Across all three market settings and a wide set of specifications, I find significant amounts of private information held by the rejectees: the subjective probabilities are predictive of the realized loss conditional on observable characteristics. Moreover, I find that they are more predictive for the rejectees than for the non-rejectees; indeed, once I control for observable characteristics used by insurance companies to price insurance, I cannot reject the null hypothesis of no private information where the market exists in any of the three markets I consider. Quantifying the amount of private information in each market, I estimate rejectees would need to be willing to pay an implicit tax of 80% in LTC, 42% in Life, and 66% in Disability insurance in order for a market to exist. In contrast, I estimate smaller implicit taxes for the non-rejectees that are not statistically different from zero in any of the three market settings.

Finally, not only do the results explain rejections in these three non-group markets, but the pattern of private information about mortality can also explain the *lack of* rejection

⁶Throughout, I focus on those who "would be rejected", which corresponds to those whose choice set excludes insurance, not necessarily the same as those who actually apply and are rejected.

tions in annuity markets. While some individuals are informed about being a relatively high mortality risk, very few are exceptionally informed about having exceptionally low mortality risk. Thus, the population of healthy individuals can obtain annuities without a significant number of even lower mortality risks adversely selecting their contract.

This paper is related to several distinct literatures. On the theoretical dimension, it is, to my knowledge, the first paper to show that private information can eliminate all gains to trade in an insurance market with an endogenous set of contracts. While no trade can occur in the Akerlof [1970] lemons model, this model exogenously restricts the set of tradable contracts, which is unappealing in the context of insurance since insurers generally offer a menu of premiums and deductibles. Consequently, this paper is more closely related to the large screening literature using the binary loss environment initially proposed in Rothschild and Stiglitz [1976]. While the Akerlof lemons model restricts the set of tradable contracts, this literature generally restricts the distribution of types (e.g. “two types” or a bounded support) and generally argues that trade will always occur (Riley [1979]; Chade and Schlee [2011]). But by considering an arbitrary distribution of types, I show this not to be the case. Indeed, not only is no-trade theoretically possible; I argue it is the outcome in significant segments of three major insurance markets.

Empirically, this paper is related to a recent and growing literature on testing for the existence and consequences of private information in insurance markets (Chiappori and Salanié [2000]; Chiappori et al. [2006]; Finkelstein and Poterba [2002, 2004]; see Einav et al. [2010a] and Cohen and Siegelman [2010] for a review). This literature focuses on the revealed preference implications of private information by looking for a correlation between insurance purchase and subsequent claims. While this approach can potentially identify private information amongst those served by the market, my approach can study private information for the entire population, including rejectees. Thus, my results provide a new explanation for why previous studies have not found evidence of significant adverse selection in life insurance (Cawley and Philipson [1999]) or LTC insurance (Finkelstein and McGarry [2006]). The most salient impact of private information may not be the adverse selection of existing contracts but rather the existence of the insurance market.

Finally, this paper is related to the broader literature on the workings of markets under uncertainty and private information. While many theories have pointed to potential problems posed by private information, this paper presents, to the best of my knowledge, the first direct empirical evidence that private information leads to a complete absence of trade.

The rest of this paper proceeds as follows. Section 2 presents the theory and the

no-trade result. Section 3 presents the comparative statics and testable predictions of the model. Section 4 outlines the empirical methodology. Section 5 presents the three market settings and the data. Section 6 presents the empirical specification and results for the nonparametric lower bounds. Section 7 presents the empirical specification and results for the estimation of the implicit tax imposed by private information. Section 8 places the results in the context of existing literature and discusses directions for future work. Section 9 concludes. To keep the main text to a reasonable length, the theoretical proofs and empirical estimation details are deferred to the Online Appendix accompanying this paper.

2 Theory

This section develops a model of private information. The primary result (Theorem 1) is a no-trade condition which provides a theory of how private information can cause insurance companies to not offer any contracts.

2.1 Environment

There exists a unit mass of agents endowed with non-stochastic wealth $w > 0$. All agents face a potential loss of size $l > 0$ that occurs with privately known probability p , which is distributed with c.d.f. $F(p|X)$ in the population, where X is the observable information insurers could use to price insurance (e.g. age, gender, observable health conditions, etc.). For the theoretical section, it will suffice to condition on a particular value for the observable characteristics, $X = x$, and let $F(p) = F(p|X = x)$ denote the distribution of types conditional on this value. I impose no restrictions on $F(p)$; it may be a continuous, discrete, or mixed distribution, and have full or partial support, denoted by $\Psi \subset [0, 1]$.⁷ Throughout the paper, an uppercase P will denote the random variable representing a random draw from the population (with c.d.f. $F(p)$); a lowercase p denote a specific agent's probability (i.e. their realization of P).

Agents have a standard Von-Neumann Morgenstern preferences $u(c)$ with expected

⁷By choosing particular distributions $F(p)$, the environment nests type spaces used in many previous models of insurance. For example, $\Psi = \{p_L, p_H\}$ yields the classic two-type model considered initially by Rothschild and Stiglitz [1976] and subsequently analyzed by many others. Assuming $F(p)$ is continuous with $\Psi = [a, b] \subset (0, 1)$, one obtains an environment similar to Riley [1979]. Chade and Schlee [2011] provide arguably the most general treatment to-date of this environment in the existing literature by considering a monopolists problem with an arbitrary F with bounded support $\Psi \subset [a, b] \subset (0, 1)$.

utility given by

$$pu(c_L) + (1 - p)u(c_{NL})$$

where c_L (c_{NL}) is the consumption in the event of a loss (no loss). I assume $u(c)$ is twice continuously differentiable, with $u'(c) > 0$ and $u''(c) < 0$. An allocation $A = \{c_L(p), c_{NL}(p)\}_{p \in \Psi}$ consists of consumption in the event of a loss, $c_L(p)$, and in the event of no loss, $c_{NL}(p)$ for each type $p \in \Psi$.

2.2 Implementable Allocations

Under what conditions can anyone obtain insurance against the occurrence of the loss? To ask this question in a general manner, I consider the set of implementable allocations.

Definition 1. An allocation $A = \{c_L(p), c_{NL}(p)\}_{p \in \Psi}$ is **implementable** if

1. A is resource feasible:

$$\int [w - pl - pc_L(p) - (1 - p)c_{NL}(p)] dF(p) \geq 0$$

2. A is incentive compatible:

$$pu(c_L(p)) + (1 - p)u(c_{NL}(p)) \geq pu(c_L(\tilde{p})) + (1 - p)u(c_{NL}(\tilde{p})) \quad \forall p, \tilde{p} \in \Psi$$

3. A is individually rational:

$$pu(c_L(p)) + (1 - p)u(c_{NL}(p)) \geq pu(w - l) + (1 - p)u(w) \quad \forall p \in \Psi$$

It is easy to verify that these constraints must be satisfied in most, if not all, institutional environments such as competition or monopoly. Therefore, to ask when agents can obtain any insurance, it suffices to ask when the endowment, $\{(w - l, w)\}_{p \in \Psi}$, is the only implementable allocation.⁸

2.3 The No-Trade condition

The key friction in this environment is that if a type p prefers an insurance contract relative to her endowment, then the pool of risks $P \geq p$ will also prefer this insurance

⁸Focusing on implementable allocations, as opposed to explicitly modeling the market structure, also circumvents problems arising from the potential non-existence of competitive Nash equilibriums, as highlighted in Rothschild and Stiglitz [1976].

contract relative to their endowment. Theorem 1 says that unless some type is willing to pay this pooled cost of worse risks in order to obtain some insurance, there can be no trade. Any insurance contract, or menu of insurance contracts, would be so adversely selected that it would not yield a positive profit.

Theorem 1. *(No Trade). The endowment, $\{(w - l, w)\}$, is the only implementable allocation if and only if*

$$\frac{p}{1-p} \frac{u'(w-l)}{u'(w)} \leq \frac{E[P|P \geq p]}{1-E[P|P \geq p]} \quad \forall p \in \Psi \setminus \{1\} \quad (1)$$

where $\Psi \setminus \{1\}$ denotes the support of P excluding the point $p = 1$.

Conversely, if (1) does not hold, then there exists an implementable allocation which strictly satisfies resource feasibility and individual rationality for a positive mass of types.

Proof. See Appendix A.1⁹ □

The left-hand side of equation (1), $\frac{p}{1-p} \frac{u'(w-l)}{u'(w)}$ is the marginal rate of substitution between consumption in the event of no loss and consumption in the event of a loss, evaluated at the endowment, $(w - l, w)$. It is a type p agent's willingness to pay for an infinitesimal transfer of consumption to the event of a loss from the event of no loss. The actuarially fair cost of this transfer to a type p agent is $\frac{p}{1-p}$. However, if the worse risks $P \geq p$ also select this contract, the cost of this transfer would be $\frac{E[P|P \geq p]}{1-E[P|P \geq p]}$, which is the right hand side of equation (1). The theorem shows that if no agent is willing to pay this pooled cost of worse risks, the endowment is the only implementable allocation.

Conversely, if equation (1) does not hold, there exists an implementable allocation which does not totally exhaust resources and provides strictly higher utility than the endowment for a positive mass of types. So, a monopolist insurer could earn positive profits by selling insurance.¹⁰ In this sense, the no-trade condition (1) characterizes when one would expect trade to occur.¹¹

⁹While Theorem 1 is straightforward, its proof is less trivial because one must show that Condition 1 rules out not only single contracts but also any menu of contracts in which different types may receive different consumption bundles.

¹⁰Also, one can show that a competitive equilibrium, as defined in Miyazaki [1977] and Spence [1978] can be constructed for an arbitrary type distribution $F(p)$ and would yield trade (result available from the author upon request).

¹¹It is easily verified that the no-trade condition can hold for common distributions. For example, if $F(p)$ is uniform on $[0, 1]$, then $E[P|P \geq p] = \frac{1+p}{2}$ so that the no trade condition reduces to $\frac{u'(w-l)}{u'(w)} \leq 2$. Unless individuals are willing to pay a 100% tax for insurance, there can be no trade when $F(p)$ is uniform over $[0, 1]$.

The no-trade condition has an unraveling intuition similar to that of Akerlof [1970]. His model considers a given contract and shows that it will not be traded when its demand curve lies everywhere below its average cost curve, where the cost curve is a function of those who demand it. My model is different in the following sense: while Akerlof [1970] derives conditions under which a *given contract* would unravel and result in no trade, my model provides conditions under which *any contract or menu of contracts* would unravel.¹²

This distinction is important since previous literature has argued that trade must always occur in similar environments with no restrictions on the contract space so that firms can offer varying premium and deductible menus (Riley [1979]; Chade and Schlee [2011]). The key difference in my environment is that I do not assume types are bounded away from $p = 1$.¹³ To see why this matters, recall that the key friction that can generate no trade is the unwillingness of any type to pay the pooled cost of worse risks. This naturally requires the perpetual existence of worse risks. Otherwise the highest risk type, say $\bar{p} = \sup \Psi$, would be able to obtain an actuarially fair full insurance allocation, $c_L(\bar{p}) = c_{NL}(\bar{p}) = w - \bar{p}l$, which would not violate the incentive constraints of any other type. Therefore, the no trade requires some risks be arbitrarily close to $p = 1$.

Corollary 1. *Suppose condition (1) holds. Then $F(p) < 1 \forall p < 1$.*

Corollary 1 highlights why previous theoretical papers have not found outcomes of no trade in the binary loss environment with no restrictions on the contract space; they assume $\sup \Psi < 1$.

The presence of risks near $p = 1$ make the provision of insurance more difficult because it increases the values of $E[P|P \geq p]$ at interior values of p . However, the need for P to have full support near 1 is not a very robust requirement for no trade. In reality, the cost of setting up a contract is nonzero, so that insurance companies cannot offer an infinite set of contracts. Remark 1 shows that if each allocation other than the endowment must attract a non-trivial fraction of types, then risks arbitrarily close to 1 are not required for no trade.

Remark 1. Suppose each consumption bundle (c_L, c_{NL}) other than the endowment must attract a non-trivial fraction $\alpha > 0$ of types. More precisely, suppose allocations $A =$

¹²This is also a difference between my approach and the literature on extreme adverse selection in finance contexts that exogenously restrict the set of tradable assets. Mailath and Noldeke [2008] provide a condition, with similar intuition to the unraveling condition in Akerlof [1970], under which a given asset cannot trade in any nonzero quantity. However, it is easy to verify in their environment that derivatives of the asset could always be traded, even when their no trade condition holds. In contrast, by focusing on the set of implementable allocations, my approach rules out the nonzero trading of any asset derived from the loss.

¹³Both Riley [1979] and Chade and Schlee [2011] assume $\sup \Psi < 1$.

$\{c_L(p), c_{NL}(p)\}_p$ must have the property that for all $q \in \Psi$,

$$\mu(\{p | (c_L(p), c_{NL}(p)) = (c_L(q), c_{NL}(q))\}) \geq \alpha$$

where μ is the measure defined by $F(p)$. Then, the endowment is the only implementable allocation if and only if

$$\frac{p}{1-p} \frac{u'(w-l)}{u'(w)} \leq \frac{E[P|P \geq p]}{1 - E[P|P \geq p]} \quad \forall p \in \hat{\Psi}_{1-\alpha} \quad (2)$$

where $\hat{\Psi}_{1-\alpha} = [0, F^{-1}(1-\alpha)] \cap (\Psi \setminus \{1\})$.¹⁴ Therefore, the no-trade condition need only hold for values $p < F^{-1}(1-\alpha)$.

For any $\alpha > 0$, it is easy to verify that the no trade condition not only does not require types near $p = 1$, but it actually imposes no constraints on the upper range of the support of P .¹⁵ In this sense, the requirement of risks arbitrarily close to $p = 1$ is a theoretical requirement in a world with no other frictions, but not an empirically relevant condition if one believes insurance companies cannot offer contracts that attract an infinitesimal fraction of the population. Going forward, I retain the benchmark assumption of no such frictions or transactions costs, but return to this discussion in the empirical work in Section 7.

In sum, the no-trade condition (1) provides a theory of rejections: individuals with observable characteristics, X , such that the no-trade condition (1) holds are rejected; individuals with observable characteristics, X , such that (1) does not hold are able to purchase insurance. This is the theory of rejections the remainder of this paper will seek to test.

3 Comparative Statics and Testable Predictions

In order to generate testable implications of this theory of rejections, this section derives properties of distributions, $F(p)$, which are more likely to lead to no trade. I provide two such metrics that will be used in the subsequent empirical analysis.

¹⁴If $F^{-1}(1-\alpha)$ is a set, I take $F^{-1}(1-\alpha)$ to be the supremum of this set

¹⁵More precisely, for any $\alpha > 0$ and $\gamma \in (0, 1]$, there exists $u(\cdot)$ and $F(p)$ such that $F(\gamma) = 1$ and the no trade condition in equation(2) holds.

3.1 Two Measures of Private Information

To begin, multiply the no-trade condition (1) by $\frac{1-p}{p}$ yielding,

$$\frac{u'(w-l)}{u'(w)} \leq \frac{E[P|P \geq p]}{1 - E[P|P \geq p]} \frac{1-p}{p} \quad \forall p \in \Psi \setminus \{1\}$$

The left-hand side is the ratio of marginal utilities in the loss versus no loss state, evaluated at the endowment. The right-hand side is independent of the utility function, u , and is the markup that would be imposed on type p if she had to cover the cost of worse risks, $P \geq p$. I define this term the pooled price ratio.

Definition 2. For any $p \in \Psi \setminus \{1\}$, the **pooled price ratio at p** , $T(p)$, is given by

$$T(p) = \frac{E[P|P \geq p]}{1 - E[P|P \geq p]} \frac{1-p}{p} \quad (3)$$

Given $T(p)$, the no-trade condition has a succinct expression.

Corollary 2. (*Quantification of the barrier to trade*) *The no-trade condition holds if and only if*

$$\frac{u'(w-l)}{u'(w)} \leq \inf_{p \in \Psi \setminus \{1\}} T(p) \quad (4)$$

Whether or not there can be trade depends on only two numbers: the agent's underlying valuation of insurance, $\frac{u'(w-l)}{u'(w)}$, and the cheapest cost of providing an infinitesimal amount of insurance, $\inf_{p \in \Psi \setminus \{1\}} T(p)$. I call $\inf_{p \in \Psi \setminus \{1\}} T(p)$ the *minimum pooled price ratio*.

The minimum pooled price ratio has a simple tax rate interpretation. Suppose for a moment that there were no private information but instead a government levies a sales tax of rate t on insurance premiums in a competitive insurance market. The value $\frac{u'(w-l)}{u'(w)} - 1$ is the highest such tax rate an individual would be willing to pay to purchase any insurance. Thus, $\inf_{p \in \Psi \setminus \{1\}} T(p) - 1$ is the implicit tax rate imposed by private information. Given any distribution of risks, $F(p)$, it quantifies the implicit tax individuals would need to be willing to pay so that a market could exist.

Equation (4) leads to a simple comparative static.

Corollary 3. (*Comparative static in the minimum pooled price ratio*) *Consider two market segments, 1 and 2, with pooled price ratios $T_1(p)$ and $T_2(p)$ and common vNM preferences u . Suppose*

$$\inf_{p \in \Psi \setminus \{1\}} T_1(p) \leq \inf_{p \in \Psi \setminus \{1\}} T_2(p)$$

then if the no-trade condition holds in segment 1, it must also hold in segment 2.

Higher values of the minimum pooled price ratio are more likely to lead to no trade. Because the minimum pooled price ratio characterizes the barrier to trade imposed by private information, Corollary 3 is the key comparative static on the distribution of private information provided by the theory.

In addition to the minimum pooled price ratio, it will also be helpful to have another metric to guide portions of the empirical analysis.

Definition 3. For any $p \in \Psi$, define the **magnitude of private information at p** by $m(p)$, given by

$$m(p) = E[P|P \geq p] - p \tag{5}$$

The value $m(p)$ is the difference between p and the average probability of everyone worse than p . Note that $m(p) \in [0, 1]$ and $m(p) + p = E[P|P \geq p]$. The following comparative static follows directly from the no-trade condition (1).

Corollary 4. (*Comparative static in the magnitude of private information*) Consider two market segments, 1 and 2, with magnitudes of private information $m_1(p)$ and $m_2(p)$ and common support Ψ and common vNM preferences u . Suppose

$$m_1(p) \leq m_2(p) \quad \forall p \in \Psi$$

Then if the no-trade condition holds in segment 1, it must also hold in segment 2.

Higher values of the magnitude of private information are more likely to lead to no trade. Notice that the values of $m(p)$ must be ordered for all $p \in \Psi$; in this sense Corollary 4 is a less precise comparative static than Corollary 3.

3.2 Testable Hypotheses

The goal of the rest of the paper is to test whether the no-trade condition (1) can explain rejections by estimating properties of the distribution of private information, $F(p|X)$, for rejectees and non-rejectees. Assuming for the moment that $F(p|X)$ is observable to the econometrician, the ideal tests are as follows. First, do rejectees have private information (i.e. is $F(p|X)$ a non-trivial distribution for the rejectees)? Second, do they have more private information than the non-rejectees, as suggested by the comparative statics in Corollaries 3 and 4? Finally, is the quantity of private information, as measured by the

minimum pooled price ratio, is large (small) enough to explain (the absence of) rejections for plausible values of agents’ willingness to pay, $\frac{u'(w-l)}{u'(w)}$, as suggested by Corollary 2?

Note that these tests do not involve on any observation of adverse selection (i.e. a correlation between insurance purchases and realized losses). Instead, these ideal tests simulate the extent to which private information would afflict a hypothetical insurance market that pays \$1 in the event that the loss occurs and prices policies using the observable characteristics, X .

To implement these tests, one must estimate properties of the distribution of private information, $F(p|X)$, to which I now turn.

4 Empirical Methodology

I develop an empirical methodology to study private information and operationalize the tests in Section 3.2. I rely primarily on four pieces of data. First, let L denote an event (e.g. dying in the next 10 years) that is commonly insured in some insurance market (e.g. life insurance).¹⁶ Second, let Z denote an individual’s subjective probability elicitation about event L (i.e. Z is a response to the question: “What is the chance (0-100%) that L will occur?”). Third, let X continue to denote the set of public information insurance companies would use to price insurance against the event L . Finally, let Θ^{Reject} and $\Theta^{NoReject}$ partition the space of values of X into those for whom an insurance company does and does not offer insurance contracts that provide payment if L occurs (e.g. if L is the event of dying in the next 10 years, Θ^{Reject} would be the values of observables, X , that render someone ineligible to purchase life insurance).

The premise underlying the approach is that the elicitations, Z , are non-verifiable to an insurance company. Therefore, they can be excluded from the set of public information insurance companies would use to price insurance, X , and used to infer properties of the distribution of private information.

I maintain the implicit assumption in Section 2 that individuals behave as if they have true beliefs, P , about the occurrence of the loss, L .¹⁷ But there are many reasons to expect individuals not to report exactly these beliefs on surveys.¹⁸ Therefore, I do not

¹⁶To condense notation, L will denote both a probabilistic event and also the binary random variable equal to 1 if the event occurs and 0 if the event does not occur (i.e. $\Pr\{L\} = \Pr\{L = 1\} = E[L]$).

¹⁷The approach therefore follows the view of personal probability expressed in the seminal work of Savage [1954]. The existence of beliefs P are guaranteed as long as people would behave consistently (in the sense of Savage’s axioms) in response to gambles over L .

¹⁸For example, they may not have the training to know how to answer probabilistic questions; they may intentionally lie to the surveyor; or they may simply be lazy in thinking about their response.

assume $Z = P$. Instead, I use information contained in the joint distribution of Z and L (that are observed) to infer properties about the distribution of P (that is not directly observed).

I conduct two complementary empirical approaches. Under relatively weak assumptions rooted in economic rationality, I provide a test for the presence of private information and a nonparametric lower bound on the average magnitude of private information, $E[m(P)]$. Loosely, this approach asks how predictive the elicitation is of the loss L , conditional on observable information, X . Second, I use slightly stronger structural assumptions to estimate the distribution of beliefs, $F(p|X)$, and the minimum pooled price ratio. I then test whether it is larger for the rejectees and large (small) enough to explain a complete absence of trade for plausible values of $\frac{u'(w-l)}{u'(w)}$, as suggested by Corollary 2.

In this section, I introduce these empirical approaches in the abstract. I defer a discussion of the empirical specification and statistical inference in my particular settings to Sections 6 and 7, after discussing the data and settings in Section 5.

4.1 Nonparametric Lower Bounds

Instead of assuming people necessarily report their true beliefs, I begin with the weaker assumption that people cannot report *more* information than what they know.

Assumption 1. *Z contains no additional information than P about the loss L , so that*

$$\Pr\{L|X, P, Z\} = \Pr\{L|X, P\}$$

This assumption states that if the econometrician were trying to forecast whether or not an agent's loss would occur and knew both the observable characteristics, X and the agent's true beliefs, P , the econometrician could not improve the forecast of L by also knowing the elicitation, Z . All of the predictive power that Z has about L must come from agents' beliefs, P .¹⁹ Proposition 1 follows.

Proposition 1. *Suppose $\Pr\{L|X, Z\} \neq \Pr\{L|X\}$ for a positive mass of realizations of Z . Then, $\Pr\{L|X, P\} \neq \Pr\{L|X\}$ for a positive mass of realizations of P .*

Indeed, existing research suggests the way in which the elicitation is conducted affects the reported belief elicitation (Gigerenzer and Hoffrage [1995], Miller et al. [2008]), which suggests elicitation do not measure true beliefs exactly. Previous literature has also argued that the elicitation in my settings should not be viewed as true beliefs due to excess concentrations at 0, 50%, and 100% (Gan et al. [2005], Hurd [2009]).

¹⁹This assumption would be clearly implied in a model in which agents formed rational expectations from an information set that included X and Z . In this case $\Pr\{L|X, P, Z\} = P$. But, it also allows agents' beliefs to be biased, so that $\Pr\{L|X, P, Z\} = h(P)$ where h is any function not dependent on Z . In particular, $h(P)$ could be an S-shaped function as suggested by Kahneman and Tversky [1979].

Proof. Assumption 1 implies $E[\Pr\{L|X, P\} | X, Z] = \Pr\{L|X, Z\}$. □

Proposition 1 says that if Z has predictive information about L conditional on X , then agents' true beliefs P has predictive information about L conditional on X – i.e. agents have private information. This motivates my test for the presence of private information:

Test 1. (*Presence of Private Information*) *Are the elicitations, Z , predictive of the loss, L , conditional on observable information, X ?*

Although this test establishes the presence of private information, it does not provide a method of asking whether one group has more private information than another. Intuitively, the predictiveness of Z should be informative of how much private information people have. Such a relationship can be established with an additional assumption about how realizations of L relate to beliefs, P .

Assumption 2. *Beliefs P are unbiased:* $\Pr\{L|X, P\} = P$

Assumption 2 states that if the econometrician could hypothetically identify an individual with beliefs P , then the probability that the loss occurs equals P . As an empirical assumption, it is strong, but commonly made in existing literature (e.g. Einav et al. [2010b]); indeed, it provides perhaps the simplest link between the realized loss L and beliefs, P .²⁰

Under Assumptions 1 and 2, the predictiveness of the elicitations form a distributional lower bound on the distribution of P . To see this, define P_Z to be the predicted value of L given the variables X and Z ,

$$P_Z = \Pr\{L|X, Z\}$$

Under Assumptions 1 and 2, it is easy to verify (see Appendix B) that

$$P_Z = E[P|X, Z]$$

so that the true beliefs, P , are a mean-preserving spread of the distribution of predicted values, P_Z . In this sense, the true beliefs are more predictive of the realized loss than are the elicitations.

This motivates my first test of whether rejectees have more private information than non-rejectees. I plot the distribution of predicted values, P_Z , separately for rejectees ($X \in$

²⁰Assumptions 1 and 2 are jointly implied by rational expectations in a model in which agents know both X and Z in formulating their beliefs P . In this case, my approach views Z as a “garbling” of the agent’s true beliefs in the sense of Blackwell ([1951], [1953]).

Θ^{Reject}) and non-rejectees ($X \in \Theta^{NoReject}$). I then assess whether it is more dispersed for the rejectees.

In addition to visual inspection of P_Z , one can also construct a dispersion metric derived from the comparative statics of the theory. Recall from Corollary 4 that higher values of the magnitude of private information, $m(p)$, are more likely to lead to no trade. Consider the average magnitude of private information, $E[m(P)|X]$. This is a non-negative measure of the dispersion of the population distribution of P . If an individual were drawn at random from the population, one would expect the risks higher than him to have an average loss probability that is $E[m(P)|X]$ higher.

Although P is not observed, I construct the analogue using the P_Z distribution. First, I construct $m_Z(p)$ as the difference between p and the average predicted probability, P_Z , of those with predicted probabilities higher than p .

$$m_Z(p) = E_{Z|X}[P_Z|P_Z \geq p, X] - p$$

The $Z|X$ subscript highlights that I am integrating over realizations of Z conditional on X . Then I construct the average magnitude of private information implied by Z in segment X , $E[m_Z(P_Z)|X]$. This is the average difference in segment X between an individual's predicted loss, and the predicted losses of those with higher predicted probabilities. Proposition follows from Assumption 1 and 2.

Proposition 2. (*Lower Bound*) $E[m_Z(P_Z)|X] \leq E[m(P)|X]$

Proof. See Appendix B. □

Proposition 2 states that the average magnitude of private information implied by Z is a lower bound on the true average magnitude of private information. Therefore, using only Assumptions 1 and 2, one can provide a lower bound to the answer to the question: if an individual is drawn at random, on average how much worse are the higher risks?

Given this theoretical measure of dispersion, $E[m_Z(P_Z)|X]$, I conduct a test in the spirit of the comparative statics given by Corollary 4. I test whether rejectees have higher values of $E[m_Z(P_Z)|X]$:

$$\Delta_Z = E[m_Z(P_Z)|X \in \Theta^{Reject}] - E[m_Z(P_Z)|X \in \Theta^{NoReject}] >? 0 \quad (6)$$

Stated loosely, equation (6) asks whether the subjective probabilities of the rejectees better explain the realized losses than the non-rejectees, where “better explain” is measured

using the dispersion metric, $E[m_Z(P_Z)|X]$. I now summarize the tests for more private information for the rejectees relative to the non-rejectees.

Test 2. (*More Private Information for Rejectees*) Are the elicitation, Z , more predictive of L for the rejectees: (a) is P_Z more dispersed for rejectees and (b) Is $\Delta_Z > 0$?

Discussion In sum, I conduct two sets of tests motivated by Assumptions 1 and 2. First, I ask whether the elicitation are predictive of the realized loss conditional on X (Test 1); this provides a test for the presence of private information as long as people cannot unknowingly predict their future loss (Assumption 1). Second, I ask whether the elicitation are more predictive for rejectees relative to non-rejectees (Test 2). To do so, I analyze whether the predicted values, P_Z , are more dispersed for rejectees relative to non-rejectees. In addition to assessing this visually, I collapse these predicted values into the average magnitude of private information implied by Z , $E[m_Z(P_Z)]$ and ask whether it is larger for those who would be rejected relative to those who can purchase insurance (Equation 6).

The approach is nonparametric in the sense that I have made no restrictions on how the elicitation Z relate to the true beliefs P . For example, P_Z and $m_Z(p)$ are invariant to one-to-one transformations in Z : $P_Z = P_{h(Z)}$ and $m_Z(p) = m_{h(Z)}(p)$ for any one-to-one function h . Thus, I do not require that Z be a probability or have any cardinal interpretation. Respondents could all change their elicitation to $1 - Z$ or $100Z$; this would not change the value of P_Z or $E[m_Z(P_Z)|X]$.²¹

But while the lower bound approach relies on only minimal assumptions on how subjective probabilities relate to true beliefs, the resulting empirical test in equation (6) suffers several significant limitations as a test of the theory that private information causes insurance rejections. First, comparisons of lower bounds of $E[m(P)|X]$ across segments do not necessarily imply comparisons of its true magnitude. Second, orderings of $E[m(P)|X]$ does not imply orderings of $m(p)$ for all p , which was the statement of the comparative static in $m(p)$ in Corollary 4. Finally, in addition to having limitations as a test of the comparative static, this approach cannot quantify the minimum pooled price ratio. These shortcomings motivate a complementary empirical approach, which imposes structure on the relationship between Z and P and estimates of the distribution of private information,

²¹In principle, Z need not even be a number. Some individuals could respond to the elicitation question in a crazy manner by saying they like red cars, others that they like Buffy the Vampire Slayer. The empirical approach would proceed to analyze whether a stated liking of red cars versus Buffy the Vampire Slayer is predictive of L conditional on X . Of course, such elicited information may have low power for identifying private information about L .

$F(p|X)$.

4.2 Estimation of the Distribution of Private Information

The second approach estimates the distribution of private information and the minimum pooled price ratio. For expositional ease, fix an observable, $X = x$, and let $f_P(p)$ denote the p.d.f. of the distribution of beliefs, P , given $X = x$, which is assumed to be continuous. For this approach, I expand the joint p.d.f./p.m.f. of the observed variables L and Z , denoted $f_{L,Z}(L, Z)$ by integrating over the unobserved beliefs, P :

$$\begin{aligned} f_{L,Z}(L, Z) &= \int_0^1 f_{L,Z}(L, Z|P = p) f_P(p) dp \\ &= \int_0^1 (\Pr\{L|Z, P = p\})^L (1 - \Pr\{L|Z, P = p\})^{1-L} f_{Z|P}(Z|P = p) f_P(p) dp \\ &= \int_0^1 p^L (1 - p)^{1-L} f_{Z|P}(Z|P = p) f_P(p) dp \end{aligned}$$

where $f_{Z|P}(Z|P = p)$ is the distribution of elicitation given beliefs. The first equality follows by taking the conditional expectation with respect to P . The second equality follows by expanding the joint density of L and Z given P . The third equality follows from Assumptions 1 and 2.

The goal of this approach is to specify a functional form for $f_{Z|P}$, say $f_{Z|P}(Z|P; \theta)$, and a flexible approximation for f_P , say $f_P(p; \nu)$, and estimate θ and ν using maximum likelihood from the observed data on L and Z . To do so, one must impose sufficient restrictions on $f_{Z|P}$ so that θ and ν are identified. Because the discussion of functional form for $f_{Z|P}$ and its identification is more straightforward after discussing the data, I defer a detailed discussion of my choice of specification and the details of identification to Section 7.1. At a high level, identification of the elicitation error parameters, θ , comes from the relationship between L and Z , and identification of the distribution of P is a deconvolution of the distribution of Z , where θ contains the parameters governing the deconvolution. Therefore, a key concern for identification is that the measurement error parameters are well identified from the relationship between Z and L ; I discuss how this is the case in my particular specification in Section 7.1.²²

With an estimate of f_P , the pooled price ratio follows from the identity, $T(p) = \frac{E[P|P \geq p]}{1 - E[P|P \geq p]} \frac{1-p}{p}$. I then construct an estimate of its minimum, $\inf_{p \in [0,1]} T(p)$. Although $T(p)$

²²Indeed, not all distributions $f_{Z|P}$ are identified from data on L and Z since, in general, $f_{Z|P}$ is an arbitrary two-dimensional function whereas L is binary.

can be calculated at each p using estimates of $E[P|P \geq p]$, as p increases, $E[P|P \geq p]$ relies on a smaller and smaller effective sample size. Thus, the minimum of $T(p)$ is not well-identified over a domain including the uppermost points of the support of P . To overcome this extreme quantile estimation problem, I construct the minimum of $T(p)$ over the restricted domain, $\hat{\Psi}_\tau = [0, F_P^{-1}(\tau)] \cap (\Psi \setminus \{1\})$. For a fixed quantile, estimates of the minimum pooled price ratio over $\hat{\Psi}_\tau$ are continuously differentiable functions of the MLE parameter estimates of $f_P(p)$ for $p \leq F_P^{-1}(\tau)$.²³ So, derived MLE estimates of $\inf_{p \in \hat{\Psi}_\tau} T(p)$ are consistent and asymptotically normal, provided $F_P(p)$ is continuous.²⁴ One can assess the robustness to the choice of τ , but the estimates will become unstable as $\tau \rightarrow 1$.

While the motivation for restricting attention to $\hat{\Psi}_\tau$ as opposed to Ψ is primarily because of statistical limitations, Remark 1 in Section 2.3 provides an economic rationale for why $\inf_{p \in \hat{\Psi}_\tau} T(p)$ may not only be a suitable substitute for $\inf_{p \in \Psi \setminus \{1\}} T(p)$ but also may actually be more economically relevant. If contracts must attract a non-trivial fraction $1 - \tau$ of the market in order to be viable, then $\inf_{p \in \hat{\Psi}_\tau} T(p)$ characterizes the barrier to trade imposed by private information.

Given estimates of $\inf_{p \in \hat{\Psi}_\tau} T(p)$ for rejectees and non-rejectees, I test whether it is larger (smaller) for the rejectees (Corollary 3) and whether it is large (small) enough to explain a complete absence of (presence of) trade for plausible values of people's willingness to pay, $\frac{u'(w-l)}{u'(w)}$, as suggested by Corollary 2.

Test 3. (*Quantification of Private Information*) *Is the minimum pooled price ratio larger for the rejectees relative to the non-rejectees; and is it large enough (small enough) to explain an absence of (presence of) trade for plausible values of agents' willingness to pay?*

5 Setting and Data

I ask whether private information can explain rejections in three non-group insurance market settings: long-term care, disability, and life insurance.

²³Non-differentiability could hypothetically occur at points where the infimum is attained at distinct values of p .

²⁴To see this, note if $F_P(p)$ is continuous then $T(p) = \frac{1 - pF_P(p) - \int_0^p F_P(\hat{p})d\hat{p}}{1 - F_P(p)}$, so that $T(p)$ is continuous in the estimated parameters of F_P .

5.1 Short Background on the Three Non-Group Market Settings

Long-term care (LTC) insurance insures against the financial costs of nursing home use and professional home care. Expenditures on LTC represent one of the largest uninsured financial burdens facing the elderly with expenditures in the US totaling over \$135B in 2004. Moreover, expenditures are heavily skewed: less than half of the population will ever move to a nursing home (CBO [2004]). Despite this, the LTC insurance market is small, with roughly 4% of all nursing home expenses paid by private insurance, compared to 31% paid out-of-pocket (CBO [2004]).²⁵

Private disability insurance protects against the lost income resulting from a work-limiting disability. It is primarily sold through group settings, such as one's employer; more than 30% of non-government workers have group-based disability policies. In contrast, the non-group market is quite small. Only 3% of non-government workers own a non-group disability policy, most of whom are self-employed or professionals who do not have access to employer-based group policies (ACLI [2010]).²⁶

Life insurance provides payments to ones' heirs or estate upon death, insuring lost income or other expenses. In contrast to the non-group disability and LTC markets, the private non-group life insurance market is quite big. More than half of the adult US population owns life insurance, 54% of which are sold in the non-group market.²⁷

Previous Evidence of Private Information Previous research has found minimal or no evidence of adverse selection in these three markets. In life insurance, Cawley and Philipson [1999] find no evidence of adverse selection. He [2009] revisits this with a different sample focusing on new purchasers and does find evidence of adverse selection under some empirical specifications. In long-term care, Finkelstein and McGarry [2006] find direct evidence of private information by showing subjective probability elicitation are correlated with subsequent nursing home use. However, they find no evidence that this private information leads to adverse selection: conditional on the observables used to price insurance, those who buy LTC insurance are no more likely to go to a nursing home

²⁵Medicaid pays for nursing home stays provided one's assets are sufficiently low and is a substantial payer of long-term stays.

²⁶In contrast to health insurance where the group market faces significant tax advantages relative to the non-group market, group disability policies are taxed. Either the premiums are paid with after-tax income, or the benefits are taxed upon receipt.

²⁷Life insurance policies either expire after a fixed length of time (term life) or cover one's entire life (whole life). Of the non-group policies in the US, 43% of these are term policies, while the remaining 57% are whole life policies (ACLI [2010]).

than those who do not purchase LTC insurance.²⁸ To my knowledge, there is no previous study of private information in the non-group disability market.

5.2 Data

To implement the empirical approach in Section 4, the ideal dataset contains four pieces of information for each setting:

1. Loss indicator, L , corresponding to a commonly insured loss in a market setting
2. Agents' subjective probability elicitation, Z , about this loss
3. The set of public information, X , which would be observed by insurance companies in the market to set contract terms
4. The classification, Θ^{Reject} and $\Theta^{NoReject}$, of who would be rejected if they applied for insurance in the market setting

The data source for the loss, L , subjective probabilities, Z , and public information X , come from years 1993-2008 of the Health and Retirement Study (HRS). The HRS is an individual-level panel survey of older individuals (mostly over age 55) and their spouses. It contains a rich set of health and demographic information. Moreover, it asks respondents three subjective probability elicitations about future events that correspond to a commonly insured loss in each of the three settings.

Long-Term Care: *"What is the percent chance (0-100) that you will move to a nursing home in the next five years?"*

Disability: *"[What is the percent chance] that your health will limit your work activity during the next 10 years?"*

Life: *"What is the percent chance that you will live to be AGE or more?"* (where $AGE \in \{75, 80, 85, 90, 95, 100\}$ is respondent-specific and chosen to be 10-15 years from the date of the interview)²⁹

Figures 1(a,b,c) display histograms of these responses (divided by 100 to scale to $[0, 1]$).³⁰ These histograms highlight one reason why it would be problematic to view these elicitations as true beliefs. As has been noted in previous literature using these subjective

²⁸They suggest heterogeneous preferences, in which good risks also have a higher valuation of insurance, can explain why private information doesn't lead to adverse selection.

²⁹I construct the corresponding elicitation to be $Z^{die} = 100\% - Z^{live}$ where Z^{live} is the survey elicitation for the probability of living to AGE.

³⁰The histograms use the sample selection described in Subsection (5.2.3)

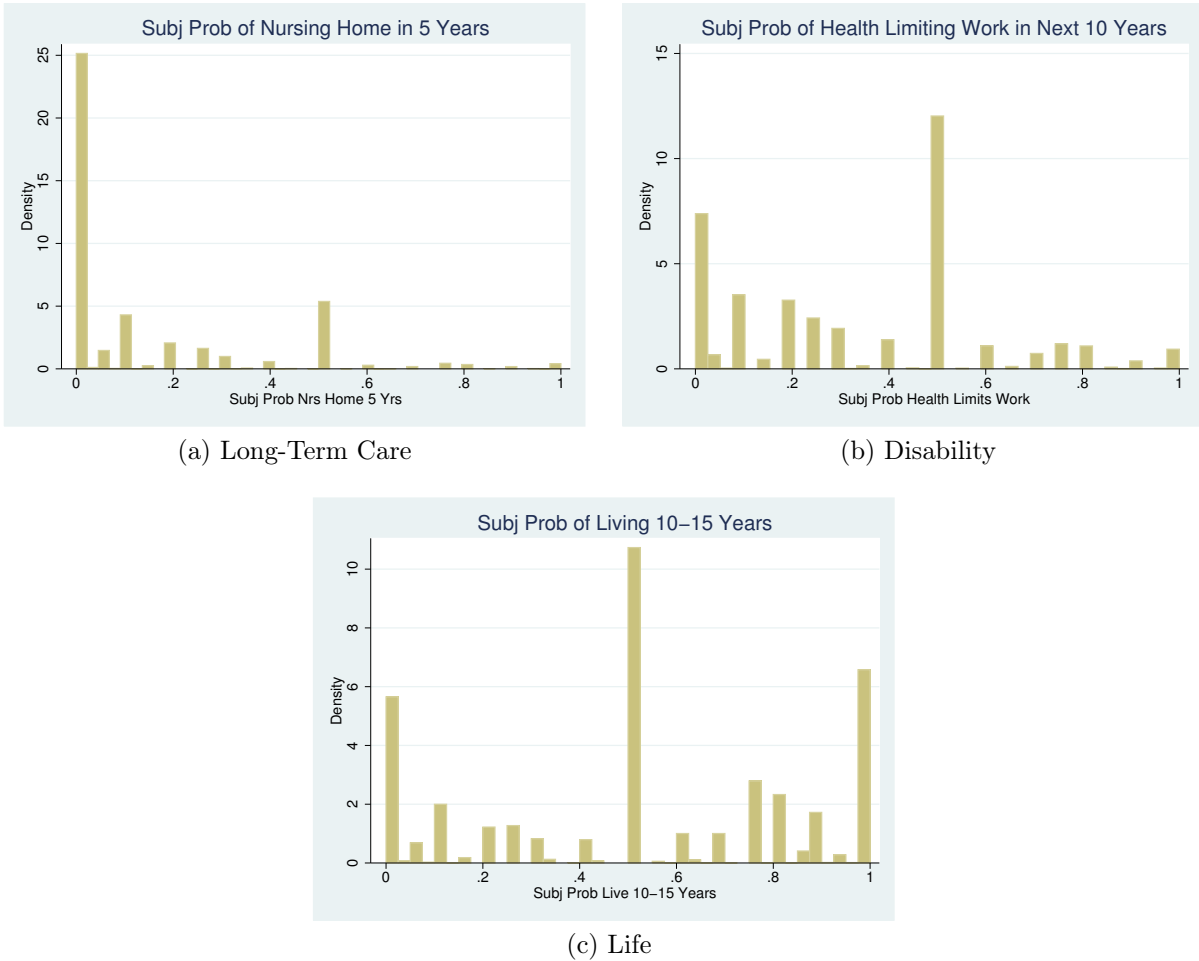


Figure 1: Subjective Probability Histograms

probabilities (Gan et al. [2005]; Finkelstein and McGarry [2006]), many respondents report 0, 50, or 100. Taken literally, responses of 0 or 100 imply an infinite degree of certainty. The lower bound approach remains agnostic on the way in which focal point responses relate to true beliefs. The parametric approach will take explicit account of this focal point response bias in the specification of $f_{Z|P}(Z|P; \theta)$, discussed further in Section 7.1.1.

Corresponding to each subjective probability elicitation, I construct binary indicators of the loss, L . In long-term care, L denotes the event that the respondent enters a nursing home in the subsequent 5 years.³¹ In disability, L denotes the event that the respondent reports that their health limits their work activity in the subsequent 10-11

³¹Although the HRS surveys every two years, I use information from the 3rd subsequent interview (6 years post) which provides date of nursing home entry information to construct the exact 5 year indicator of nursing home entry.

years.³² In life, L denotes the event that the respondent dies before AGE, where AGE $\in \{75,80,85,90,95,100\}$ corresponds to the subjective probability elicitation, which is 10-15 years from the survey date.³³

5.2.1 Public Information

To identify private information, it is essential to control for the public information, X , that would be used by insurance companies to price contracts. For non-rejectees, this is a straightforward requirement which involves analyzing existing contracts. But for rejectees, I must make an assumption about how insurance companies would price these contracts if they were to offer them. My preferred approach is to assume insurance companies price rejectees separately from those to whom they currently offer contracts, but use a similar set of public information. Thus, the primary data requirement is the public information currently used by insurance companies in pricing insurance.

The HRS contains an extensive set of health, demographic, and occupation information that allows me to approximate the set of information that insurance companies use to price insurance. Indeed, previous literature has used the HRS to replicate the observables used by insurance companies to price insurance in LTC and Life (for LTC, see Finkelstein and McGarry [2006] and for Life, see He [2009]), and I primarily follow this literature in constructing this set of covariates. Appendix C.1 provides a detailed listing of the control specifications used in each market setting.

The quality of the approximation to what insurers actually use to price insurance is quite good, but does vary by market. For long-term care, I replicate the information set of the insurance company quite well. For example, perhaps the most obscure piece of information that is acquired by some LTC insurance companies is an interview in which applicants are asked to perform word recall tasks to assess memory capabilities; the HRS conducts precisely this test with survey respondents. In disability and life, I replicate most of the information used by insurance companies in pricing. One caveat is that insurance companies will sometimes perform tests, such as blood and urine tests, which I will not observe in the HRS. Conversations with underwriters in these markets suggest these tests are primarily to confirm application information, which I can approximate quite well with

³²The loss is defined as occurring when the individual reports yes to the question: “Does your health limit your work activity?” over the subsequent five surveys, which is 10 years for all waves except AHEAD wave 2, which corresponds to a time interval of 11 years because of a slightly different survey spacing. Although the HRS has other measures of disability (e.g. SSDI claims), I use this measure because the wording corresponds exactly to the subjective probability elicitation, which will be important for the structural assumptions made to estimate the minimum pooled price ratio.

³³The HRS collects date of death information that allows me to establish the exact age of death.

the HRS. But, I cannot rule out the potential that there is additional information which can be gathered by insurance companies in the disability and life settings.³⁴

While the preferred specification attempts to replicate the variables used by insurance companies in pricing, I also assess the robustness of the estimates to larger and smaller sets of controls.³⁵ As a baseline, I consider a specification with only age and gender. As an extension, I also consider an extended controls specification that adds a rich set of interactions between health conditions and demographic variables that could be, but are not currently, used in pricing insurance. I conduct the lower bound approach for all three sets of controls. For brevity, I focus exclusively on the preferred specification of pricing controls for the parametric approach.

5.2.2 Rejection Classification

Not everyone can purchase insurance in these three non-group markets. To identify conditions that lead to rejection, I obtain underwriting guidelines used by underwriters and provided to insurance agents for use in screening applicants. An insurance company’s underwriting guidelines list the conditions for which underwriters are instructed to not offer insurance at any price and for which insurance agents are expected to discourage applications. These guidelines are generally viewed as a public relations liability and are not publicly available.³⁶ Thus, the extent of my access varies by market: In long-term care, I obtained a set of guidelines used by an insurance broker from 18 of the 27 largest long-term care insurance companies comprising a majority of the US market.³⁷ In disability and life, I obtained several underwriting guidelines and supplement this information with interviews with underwriters at several major US insurance companies. Appendix F provides several pages from the LTC underwriting guideline from Genworth Financial,

³⁴In LTC, insurance companies are legally able to conduct tests, but it is not common industry practice.

³⁵While it might seem intuitive that including more controls would reduce the amount of private information, this need not be the case. To see why, consider the following example of a regression of quantity on price. Absent controls, there may not exist any significant relationship. But, controlling for supply (demand) factors, price may have predictive power for quantity as it traces out the demand (supply) curve. Thus, adding controls can increase the predictive power of another variable (price, in this case). Of course, conditioning on additional variables X' which are uncorrelated with L or Z has no effect on the population value of $E[m(P) | X \in \Theta]$.

³⁶An example of these guidelines is presented in Appendix F and a collection of these guidelines is available on my website. Also, many underwriting guidelines are available via internet searches of “underwriting guideline not-for-public-use pdf”. These are generally left on the websites of insurance brokers who leave them electronically available to their sales agents and, potentially unknowingly, available to the general public.

³⁷I thank Amy Finkelstein for making this broker-collected data available.

one of the largest LTC insurers in the US.³⁸

I then use the detailed health and demographic information available in the HRS to identify individuals with these rejection conditions. While the HRS contains a relatively comprehensive picture of respondents' health, sometimes the rejection conditions are too precise to be matched to the HRS. For example, individuals with advanced stages of lung disease would be unable to purchase life insurance, but some companies will sell policies to individuals with a milder case of lung disease; however, the HRS only provides information for the presence of a lung disease.

Instead of attempting to match all cases, I construct a third classification in each setting, "Uncertain", to which I classify those who may be rejected, but for whom data limitations prevent a solid assessment. This allows me to be relatively confident in the classification of rejectees and non-rejectees. For completeness, I present the lower bound analysis for all three classifications.

Table 1 presents the list of conditions for the rejection and uncertain classification, along with the frequency of each condition in the sample (using the sample selection outlined below in Section 5.2.3). LTC insurers generally reject applicants with conditions that would make them more likely to use a nursing home in the relatively near future. Activity of daily living (ADL) restrictions (e.g. needs assistance walking, dressing, using toilet, etc.), a previous stroke, any previous home nursing care, and anyone over the age of 80 would be rejected regardless of health status. Disability insurers reject applicants with back conditions, obesity ($BMI > 40$), and doctor-diagnosed psychological conditions such as depression or bi-polar disorder. Finally, life insurers reject applicants who have had a past stroke or currently have cancer.

Table 1 also lists the conditions which may lead to rejection depending on the specifics of the disease. People with these conditions are allocated into the Uncertain classification.³⁹ In addition to health conditions, disability insurers also have stringent income and job characteristic underwriting. Individuals earning less than \$30,000 (or wages below \$15/hr) and individuals working in blue-collar occupations are often rejected regardless of health condition due to their employment characteristics. I therefore allocate all such individuals to the uncertain category in the disability insurance setting.

Given these classifications, I construct the Reject, No Reject, and Uncertain samples

³⁸A collection of underwriting guidelines from these three markets are available from the author upon request and are posted on my website.

³⁹I also attempt to capture the presence of rarer conditions not asked in the HRS (e.g. Lupus would lead to rejection in LTC, but is not explicitly reported in the HRS). To do so, I allocate to the uncertain classification individuals who report having an additional major health problems which was not explicitly asked about in the survey.

by first taking anyone who has a known rejection condition in Table 1 and classify them into the Reject sample in each setting. I then classify anyone with an uncertain rejection condition into the Uncertain classification, so that the remaining category is the set of people who can purchase insurance (the No Reject classification).

5.2.3 Sample Selection

For each sample, I begin with years 1993-2008 of the HRS. The selection process varies across each of the three market settings due to varying data constraints. Appendix C.2 discusses the specific data construction details for each setting. The primary sample restrictions arise from requiring the subjective elicitation be asked (e.g. only individuals over age 65 are asked about future nursing home use) and needing to observe individuals in the panel long enough to construct the loss indicator, L in each setting.⁴⁰ For LTC, the sample consists of individuals aged 65 and older; for disability the sample consists of individuals aged 60 and under⁴¹; and for life, the sample consists of individuals over age 65. Table 2 presents the summary statistics for each sample. I include multiple observations for a given individual (which are spaced roughly two years apart) to increase power.⁴²

There are several broad patterns across the three samples. First, there is a sizable sample of rejectees in each setting. Because the HRS primarily surveys older individuals, the sample is older, and therefore sicker, than the average insurance purchaser in each market. Obtaining this large sample size of rejectees is a primary benefit of the HRS; but it is important to keep in mind that the fraction of rejectees in the HRS is not a measure of the fraction of the applicants in each market that are rejected.

Second, many rejectees own insurance. These individuals could (and perhaps should) have purchased insurance prior to being stricken with their rejection condition. Also, they may have been able to purchase insurance in group markets through their employer, union, or other group which has less stringent underwriting requirements than the non-group market.

⁴⁰Note that death during this subsequent time horizon does not exclude an individual from the sample; I classify the event of dying before the end of the time horizon as $L = 0$ for the LTC and Disability settings as long as an individual did not report the loss (i.e. nursing home entry or health limiting work) prior to death.

⁴¹The disability question is asked of individuals up to age 65, but I exclude individuals aged 61-65 because of the near presence of retirement. Ideally, I would focus on a sample of even younger individuals, but unfortunately the HRS contains relatively few respondents below age 55.

⁴²All standard errors will be clustered at the household level. Because the multiple observations within a person will always have different X values (e.g. different ages), including multiple observations per person does not induce bias in the construction of $F(p|X)$.

However, the fact that some own insurance raises the concern that moral hazard could generate heterogeneity in loss probabilities from differential insurance ownership. Therefore, I also perform robustness checks in LTC and Life on samples that exclude those who currently own insurance.⁴³ Since Medicaid also pays for nursing home use, I also exclude Medicaid enrollees from this restricted LTC sample. Unfortunately, the HRS does not ask about disability insurance ownership, so I cannot conduct this robustness check for the disability setting.

Finally, although the rejectees have, on average, a higher chance of experiencing the loss than the non-rejectees, it is not certain that they would experience the loss. For example, only 22.5% of rejectees in LTC actually end up going to a nursing home in the subsequent 5 years. This suggests there is substantial unrealized risk amongst the rejectees.

5.2.4 Relation to Ideal Data

Before turning to the results, it is important to be clear about the extent to which the data resembles the ideal dataset in each market setting. In general, I approximate the ideal dataset quite well, aside from the necessity to classify a relatively large fraction of the sample to the Uncertain rejection classification. In Disability and in Life, I classify a smaller fraction of the sample as rejected or not rejected as compared with LTC. Also, for Disability and Life I rely on a smaller set of underwriting guidelines (along with underwriter interviews) to obtain rejection conditions, as opposed to LTC where I obtain a fairly large fraction of the underwriting guidelines used in the market. In Disability and Life I also do not observe medical tests that may be used by insurance companies to price insurance (although conversations with underwriters suggest this is primarily to verify application information, which I approximate quite well using the HRS). In contrast, in LTC I classify a relatively large fraction of the sample, I closely approximate the set of public information, and I can assess the robustness of the results to the exclusion of those who own insurance to remove the potential impact of a moral hazard channel driving any findings of private information. While re-iterating that all three of the samples approximate the ideal dataset quite well, the LTC sample is arguably the best of the three samples.

⁴³Since rejection conditions are generally absorbing states, this rules out the path through which insurance contract choice could generate heterogeneity for the rejectees. For the non-rejectees, this removes the heterogeneity induced by current contract choice; but it does not remove heterogeneity introduced from expected future purchase of insurance contracts. But, for my purposes this remaining moral hazard impact only biases against finding more private information amongst the rejectees.

6 Lower Bound Estimation

I now turn to the estimation of the distribution of P_Z and the lower bounds of the average magnitude of private information, $E[m_Z(P_Z)|X]$, outlined in Section 4.1.

6.1 Specification

All of the empirical estimation is conducted separately for each of the settings and rejection classifications within each setting. Here I provide an overview of the preferred specification, which controls for the variables used by insurance companies to price insurance. I defer a detailed discussion of all three control specifications to Appendix D.1.

I estimate the distribution of $P_Z = \Pr\{L|X, Z\}$ using a probit specification

$$\Pr\{L|X, Z\} = \Phi(\beta X + \Gamma(\text{age}, Z))$$

where X are the control variables (i.e. the pricing controls listed in Table A1) and $\Gamma(\text{age}, Z)$ captures the relationship between L and Z , allowing it to depend on age. With this specification, the null hypothesis of no private information, $\Pr\{L|X, Z\} = \Pr\{L|X\}$, is tested by restricting $\Gamma = 0$.⁴⁴ I choose a flexible functional form for $\Gamma(\text{age}, Z)$ that uses full interactions of basis functions in age and Z . For the basis in Z , I use second-order Chebyshev polynomials plus separate indicators for focal point responses at $Z = 0, 50$, and 100 . For the basis in age , I use a linear specification.

With infinite data, one could estimate $E[m_Z(P_Z)|X]$ at each value of X . However, the high-dimensionality of X requires being able to aggregate across values of X . To do this, I assume that conditional on ones' age and rejection classification, the distribution of $P_Z - \Pr\{L|X\}$ does not vary with X . This allows the rich set of observables to flexibly affect the mean loss probability, but allows for aggregation of the dispersion of the distribution across values of X .⁴⁵

I then estimate the conditional expectation, $m_Z(p) = E[P_Z|P_Z \geq p, X] - p$ using the estimated distribution of $P_Z - \Pr\{L|X\}$ within each age grouping and rejection classification. After estimating $m_Z(p)$, I use the estimated distribution of P_Z to construct its average, $E[m_Z(P_Z)|X \in \Theta]$, where Θ is a given sample (e.g. LTC rejectees). I construct

⁴⁴At various points in the estimation I require an estimate of $\Pr\{L|X\}$, which I obtain with the same specification as above, but restricting $\Gamma = 0$.

⁴⁵Note also that I only impose this assumption *within* a setting/rejection classification – I do not require the dispersion of the rejectees to equal that of the non-rejectees. Also, note that this assumption is only required to arrive at a point estimate for $E[m_Z(P_Z)|X \in \Theta]$, and is not required to test for the presence of private information (i.e. whether $\Gamma = 0$).

the difference between the reject and no reject estimates,

$$\Delta_Z = E [m_Z (P_Z) | X \in \Theta^{Reject}] - E [m_Z (P_Z) | X \in \Theta^{NoReject}]$$

and test whether I can reject a null hypothesis that $\Delta_Z \leq 0$.

6.2 Statistical Inference

Statistical inference for $E [m_Z (P_Z) | X \in \Theta]$ for a given sample Θ and for Δ_Z is straightforward, but requires a bit of care to cover the possibility of no private information. In any finite sample, estimates of $E [m_Z (P_Z) | X \in \Theta]$ will be positive (Z will always have some predictive power in finite samples). Provided the true value of $E [m_Z (P_Z) | X \in \Theta]$ is positive, the bootstrap provides consistent, asymptotically normal, standard errors for $E [m_Z (P_Z) | X \in \Theta]$ (Newey [1997]). But, if the true value of $E [m_Z (P_Z) | X \in \Theta]$ is zero (as would occur if there were no private information amongst those with $X \in \Theta$), then the bootstrap distribution is not asymptotically normal and does not provide adequate finite-sample inference.⁴⁶ Therefore, I supplement the bootstrap with a Wald test that restricts $\Gamma (age, Z) = 0$.⁴⁷ The Wald test is the key statistical test for the presence of private information, as it tests whether Z is predictive of L conditional on X . I report results from both the Wald test and the bootstrap.

I conduct inference on Δ_Z in a similar manner. To test the null hypothesis that $\Delta_Z \leq 0$, I construct conservative p-values by taking the maximum p-value from two tests: 1) a Wald test of no private information held by the rejectees, $E [m_Z (P_Z) | X \in \Theta^{Reject}] = 0$, and 2) the p-value from the bootstrapped event of less private information held by the rejectees, $\Delta \leq 0$.⁴⁸

6.3 Results

I begin with graphical evidence of the predictive power of the subjective probability elicitation in each sample. Figures 2(a,b,c) plot the estimated distribution of $P_Z - E [P_Z | X]$ aggregated by rejection classification for the rejectees and non-rejectees, using

⁴⁶In this case, $\hat{\Gamma} \rightarrow 0$ in probability, so that estimates of the distribution of $P_Z - E [P_Z | X]$ converge to zero in probability (so that the bootstrap distribution converges to a point mass at zero).

⁴⁷The event $\Gamma (age, Z) = 0$ in sample Θ is equivalent to both the event $\Pr \{L | X, Z\} = \Pr \{L | X\}$ for all $X \in \Theta$ and the event $E [m_Z (P_Z) | X \in \Theta] = 0$.

⁴⁸More precise p-values would be a weighted average of these two p-values, where the weight on the Wald test is given by the unknown quantity $\Pr \{E [m_Z (P_Z) | X \in \Theta^{Reject}] = 0 | \Delta \leq 0\}$. Since this weight is unknown, I use these conservative p-values that are robust to any weight in $[0, 1]$.

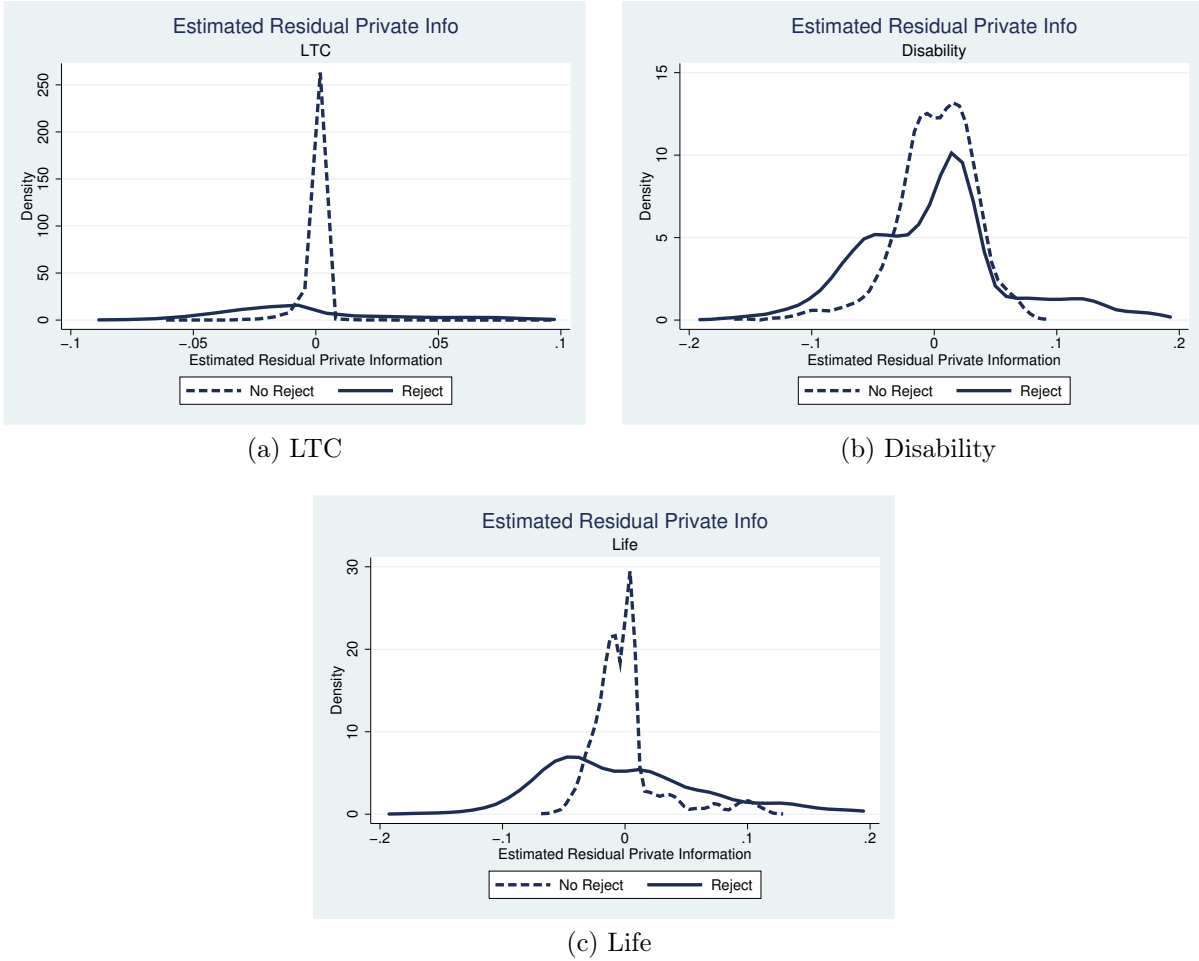


Figure 2: Distribution of $P_Z - \Pr\{L|X\}$

the preferred pricing control specification.⁴⁹

Across all three market settings, the distribution of $P_Z - \Pr\{L|X\}$ appears more dispersed for the rejectees relative to non-rejectees. In this sense, the subjective probability elicitation contains more information about L for the rejectees than for the non-rejectees.

Table 3 presents the measurements of this dispersion using the average magnitude of private information implied by Z . The first set of rows, labelled “Reject”, presents the estimates for the rejectees in each setting and control specification. Across all settings and control specifications, I find significant evidence of private information amongst the rejectees ($p < 0.001$); the subjective probabilities are predictive of the realized loss, conditional on the set of insurance companies used to price insurance and also are predictive

⁴⁹Subtracting $E[P_Z|X]$ or equivalently, $\Pr\{L|X\}$, allows for simple aggregation across X within each sample.

conditional on the baseline controls (age and gender) and the extended controls.

In addition, the estimates provide an economically significant lower bound on the average magnitude of private information. For example, the estimate of 0.0358 for the LTC price controls specification indicates that if a rejectee was drawn at random, one would expect the average probability of higher risks (with the same observables, X) to be at least 3.58pp higher, which is 16% higher than the mean loss probability of 22.5% for LTC rejectees.

The third set of rows in Table 3 provides the estimates of Δ_Z . Again, across all specifications and market settings, I estimate larger lower bounds on the average magnitude of private information for the rejectees relative to those served by the market. These differences statistically significant at the 1% level in LTC and life, and positive (but not significant at standard levels) in disability.⁵⁰

Not only do I find smaller amounts of private information for the non-rejectees, but I cannot actually reject the null hypothesis of no private information amongst this group once one includes the set of variables insurers use to price insurance, as indicated by the second set of rows in Table 3.⁵¹ This provides a new explanation for why previous research has not found significant amounts of adverse selection of insurance contracts in LTC (Finkelstein and McGarry [2006]) and Life insurance (Cawley and Philipson [1999]). The practice of rejections by insurance companies limits the extent to which private information manifests itself in adverse selection of contracts.

⁵⁰The estimated magnitudes for the uncertain classification generally fall between the estimates for the rejection and no rejection groups, as indicated by the bottom set of rows in Table 3. In general, the theory does not have a prediction for the uncertain group. However, if $E[m_Z(P_Z)|X]$ takes on similar values for all rejectees (e.g. $E[m_Z(P_Z)|X] \approx m^R$) and non-rejectees (e.g. $E[m_Z(P_Z)|X] \approx m^{NR}$), then linearity of the expectation implies

$$E[m_Z(P_Z)|X \in \Theta^{\text{Uncertain}}] = \lambda m^R + (1 - \lambda) m^{NR} \quad (7)$$

where λ is the fraction in the uncertain group who would be rejected. Thus, it is perhaps not unreasonable to have expected $E[m_Z(P_Z)|X \in \Theta^{\text{Uncertain}}]$ to lie in between the estimates for the rejectees and non-rejectees, as I find. Nevertheless, there is no theoretical reason to suppose the average magnitude of private information is constant within rejection classification; thus this should be viewed only as a potential rationalization of the results, not as a robust prediction of the theory.

⁵¹Of course, the difference between the age and gender specification and the price controls specification is not statistically significant. Also, the inability to reject a null of no private information is potentially driven by the small sample size in the Disability setting; but the LTC sample of non-rejectees is quite large (>9K) and the sample of non-rejectees in Life is larger than the sample of rejectees.




Figure 3: Magnitude of Private Information by Age & Rejection Classification

6.4 Age 80 in LTC insurance

LTC insurers reject applicants above age 80 regardless of health status. This provides an opportunity for a finer test of the theory by exploring whether those without rejection health conditions start to obtain private information at age 80. To do so, I construct a series of estimates of $E[m_Z(P_Z)]$ by age for the set of people who do not have a rejection health condition and thus would only be rejected if their age exceeded 80.⁵²

Figure 3 plots the results for those without health conditions (hollow circles), along with a comparison set of results for those with rejection health conditions (filled circles).⁵³ The figure suggests that the subjective probability elicitations of those without rejection health conditions become predictive of L right around age 80 – exactly the age at which insurers choose to start rejecting applicants based on age, regardless of health status. Indeed, from the perspective of $E[m_Z(P_Z)]$, a healthy 81 year old looks a lot like a 70 year old who had a stroke. This is again consistent with the theory that private information limits the existence of insurance markets.

⁵²To ensure no information from those with rejection health conditions is used in the construction of $E[m_Z(P_Z)]$ for those without health conditions above age 80, I split the Reject sample into two groups: those who do not have a rejection health condition (and thus would only be rejected because their age is above 80) and those who do have a rejection condition. I estimate P_Z separately on these two samples using the pricing specification outlined in Section 6.1.

⁵³The graph presents bootstrapped 95% confidence intervals adjusted for bias using the non-accelerated procedure suggested in Efron and Gong [1983]. These are appropriate confidence intervals as long as the true magnitude of private information is positive; In the aggregate sample of rejectees, I reject the null hypothesis of no private information (see Table 3). However, for any particular age, I am unable to reject a null hypothesis of no private information using the Wald test.

6.5 Robustness

Moral Hazard As discussed in Section 5.2.3, one alternative hypothesis is that the private information I estimate is the result of moral hazard from insurance contract choice, not an underlying heterogeneity in loss probabilities. To assess whether this is driving any of the results, I re-estimate the average magnitude of private information implied by Z on samples in LTC and Life that exclude those who currently own insurance. For LTC, I exclude those who own private LTC insurance along with those who are currently enrolled in Medicaid, since it pays for nursing home stays. As shown in Table 2, this excludes 20.6% of the sample of rejectees and 19.5% of non-rejectees. For Life, I exclude those with any life insurance policy. Unfortunately, this excludes 63% of the rejectees and 65% of the non-rejectees; thus the remaining sample is quite small.

Table 4 presents the results. For LTC, I continue to find significant amounts of private information for the rejectees ($p < 0.001$), that is significantly more than for the non-rejectees ($\Delta_Z = 0.0313$, $p < 0.001$), and cannot reject the null hypothesis of no private information for the non-rejectees ($p = 0.8325$). For Life, I estimate marginally significant amounts of private information for the rejectees ($p = 0.0523$) of a magnitude similar to what is estimated on the full sample (0.0491 versus 0.0587). I estimate more private information for the rejectees relative to the non-rejectees, however the difference is no longer statistically significant ($\Delta_Z = 0.011$, $p = 0.301$), which is arguably a result of the reduced sample size. I also continue to be unable to reject the null hypothesis of no private information for the non-rejectees ($p = 0.2334$). In short, the results suggest moral hazard is not driving my findings of private information for the rejectees and more private information for the rejectees relative to the non-rejectees.

Additional Robustness Checks Appendix D.2 contains a couple of additional robustness checks. I present the age based plots, similar to Figure 3, for the Disability and Life settings and show that I generally find larger amounts of private information across all age groups for the rejectees in each setting. I also present an additional specification in life insurance that includes additional cancer controls, discussed in Appendix C.1, that are available for a smaller sample of the HRS data; I show that the estimates are similar when introducing these additional controls.

6.6 Summary

In all three market settings, I estimate a significant amount of private information held by the rejectees that is robust to a wide set of controls for public information. I find more private information held by the rejectees relative to the non-rejectees; and I cannot reject a null hypothesis of no private information held by those actually served by the market. Moreover, a de-aggregated analysis of the practice of LTC insurers rejecting all applicants above age 80 (regardless of health) reveals that healthy individuals begin to have private information right around age 80 – precisely the age chosen by insurers to stop selling insurance. In sum, the results are consistent with the theory that private information leads to insurance rejections.

7 Estimation of Distribution of Private Information

While the lower bound results, and in particular the stark pattern of the presence of private information, provides support for the theory that private information would afflict a hypothetical insurance market for the rejectees, it does not establish whether the amount of private information is sufficient to explain why insurers don't sell policies to the rejectees. This requires an estimate of the minimum pooled price ratio, and hence an estimate of the distribution of private information, $F(p|X)$. To do so, I follow the second approach, outlined in Section 4.2: I impose additional structure on the relationship between elicitation, Z , and true beliefs, P , that allows for a flexible estimation of $F(p|X)$.

7.1 Empirical Specification

7.1.1 Elicitation Error Model

Elicitations Z may differ from true beliefs P in many ways. They may be systematically biased, with values either higher or lower than true beliefs. They may be noisy, so that two individuals with the same beliefs may have different elicitation. Moreover, as shown in Figures 1(a,b,c) and recognized in previous literature (e.g. Gan et al. [2005]), people may have a tendency to report focal point values at 0, 50, and 100%. My model of elicitation will capture all three of these forms of elicitation error.

To illustrate the model, first define the random variable \tilde{Z} by

$$\tilde{Z} = P + \epsilon$$

where $\epsilon \sim N(\alpha, \sigma^2)$. The variable \tilde{Z} is a noisy measure of beliefs with bias α and noise variance σ^2 where the error follows a normal distribution. I assume there are two types of responses: focal point responses and non-focal point responses. With probability $1 - \lambda$, an agent gives a non-focal point response, Z^{nf} ,

$$Z^{nf} = \begin{cases} \tilde{Z} & \text{if } \tilde{Z} \in [0, 1] \\ 0 & \text{if } \tilde{Z} < 0 \\ 1 & \text{if } \tilde{Z} > 1 \end{cases}$$

which is \tilde{Z} censored to the interval $[0, 1]$. These responses are continuously distributed over $[0, 1]$ with some mass at 0 and 1.

The second type of responses are focal point responses. With probability λ an agent reports Z^f given by:

$$Z^f = \begin{cases} 0 & \text{if } \tilde{Z} \leq \kappa \\ 0.5 & \text{if } \tilde{Z} \in (\kappa, 1 - \kappa) \\ 1 & \text{if } \tilde{Z} \geq 1 - \kappa \end{cases}$$

where $\kappa \in [0, .5)$ captures the focal point window. With this structure, focal point responses have the same underlying structure as non-focal point responses, but are reported on a scale of low, medium, and high as opposed to a continuous scale on $[0, 1]$.⁵⁴ As a result, non-focal point responses will contain more information about P than will focal point responses. Therefore, most of the identification for the distribution of P will come from those reporting non-focal point values.

Given this model, I have four elicitation parameters to be estimated: $\{\alpha, \sigma, \kappa, \lambda\}$, which will be estimated separately in each market setting and classification. This allows for the potential that rejectees have a different elicitation error process than non-rejectees.

7.1.2 Flexible Approximation for the Distribution of Private Information

With infinite data, one could flexibly estimate $f(p|X)$ separately for every possible value of X and p . Faced with finite data and a high dimensional X , this is not possible. Since the minimum pooled price ratio is essentially a function of the shape of the distribution of $f(p|X)$ across values of p , I choose a specification that allows for considerable flexibility

⁵⁴Note that I do assume the act of providing a focal point response is not informative of P (λ is not allowed to be a function of P). Ideally, one would allow focal point respondents to have differing beliefs from non-focal point respondents; yet the focal point bias inherently limits the extent of information that can be extracted from their responses.

across p . In particular, I assume $f(p|X)$ is well-approximated by a mixture of beta distributions,

$$f(p|X) = \sum_i w_i \text{Beta}(p|a_i + \Pr\{L|X\}, \psi_i) \quad (8)$$

where $\text{Beta}(p|\mu, \psi)$ is the p.d.f. of the beta distribution with mean μ and shape parameter ψ .⁵⁵ With this specification, $\{w_i\}$ governs the weights on each beta distribution, $\{a_i\}$ governs the non-centrality of each beta distribution, and ψ_i governs the dispersion of each beta distribution. The flexibility of the beta distributions ensures that I impose no restrictions on the size of the minimum pooled price ratio.⁵⁶ For the main specification, I include 3 beta distributions.⁵⁷ Additional details of the specification are provided in Appendix E.1.

7.1.3 Pooled Price Ratio (and its Minimum)

With an estimate of $f(p|X)$ the pooled price ratio is easily constructed as $T(p) = \frac{E[P|X] \frac{1-p}{p}}{1-E[P|X]}$ for each p , where $E[P|P \geq p, X]$ is computed using the estimated $f(p|X)$. Throughout, I focus on estimates evaluated for a mean loss characteristic, $\Pr\{L|X\}$. In principle, one could analyze the pooled price ratio across all values of X ; but given the specification, focusing on differing values of X or $\Pr\{L|X\}$ does not yield an independent test of the theory. In Appendix E.2, I show the results are generally robust to focusing on values of $\Pr\{L|X\}$ at the 20, 50, and 80th percentiles of its distribution.

As described in Section 4.2, I estimate the analogue to the minimum pooled price

⁵⁵The p.d.f. of a beta distribution with parameters α and β is given by

$$\text{beta}(x; \alpha, \beta) = \frac{1}{B(\alpha, \beta)} x^{\alpha-1} x^{\beta-1}$$

where $B(\alpha, \beta)$ is the beta function. The mean of a beta distribution with parameters α and β is given by $\mu = \frac{\alpha}{\alpha+\beta}$ and the shape parameter is given by $\psi = \alpha + \beta$.

⁵⁶In principle, the event of no private information is captured with $\psi_1 \rightarrow \infty$, $a_1 = 0$, and $w_1 = 1$. For computational reasons, I need to impose a cap on ψ_i in the estimation. In the initial estimation, this cap binds for the central most beta distribution in both the LTC No Reject and Disability No Reject samples. Intuitively, the model wants to estimate a large fraction of very homogenous individuals around the mean. Therefore, for these two samples, I also include a point-mass distribution with weight w_0 in addition to the three beta distributions. This allows me to capture a large concentration of mass in a way that does not require integrating over a distribution $f(p|X)$ with very high curvature. Appendix E.1 provides further details.

⁵⁷While equation 8 allows for a very flexible shape of $f(p|X)$ across p ; it is fairly restrictive in how this shape varies across values of X . Indeed, I do not allow the distribution parameters to vary with X . This is a practical necessity due to the size of my samples and the desire to allow for a very flexible shape for $f(p|X)$. Moreover, it is important to stress that I will still separately estimate $f(p|X)$ for the rejectees and the non-rejectees using the separate samples.

ratio, $\inf_{p \in \hat{\Psi}_\tau} T(p)$, for the restricted domain $\hat{\Psi}_\tau = [0, F^{-1}(\tau)]$. My preferred choice for τ is 0.8, as this ensures at least 20% of the sample (conditional on q) is used to estimate $E[P|P \geq p]$ and produces estimates that are quite robust to changes in the number of approximating beta distributions. For robustness, I also present results for $\tau = 0.7$ and $\tau = 0.9$ along with plots of the pooled price ratio for all p below the estimated 90th quantile, $F^{-1}(0.9)$.

7.1.4 Identification

Before turning to the results, it is important to understand the sources of identification for the model. As discussed above, much of the model is identified from the non-focal point responses. If the elicitation error parameters were known, then identification of the distribution of P is a deconvolution of the distribution of Z^{nf} ; thus, the empirical distribution of non-focal elicitation provides a strong source of identification for the distribution of P conditional on having identified the elicitation error parameters.⁵⁸

To identify the elicitation error parameters, the model relies on the relationship between Z^{nf} and L . To see this, note that Assumptions 1 and 2 imply

$$E[Z^{nf} - P] = E[Z^{nf}] - E[L]$$

so that the mean elicitation bias is the difference between the mean elicitation and the mean loss probability. This provides a strong source of identification for α .⁵⁹ In practice, the model calculates α jointly with the distribution of P to adjust for the fact that the non-focal elicitation are not censored over $[0, 1]$.

To identify σ , note that Assumptions 1 and 2 imply

$$\text{var}(Z^{nf}) - \text{cov}(Z^{nf}, L) = \text{var}(Z^{nf} - P) + \text{cov}(Z^{nf} - P, P) \quad (9)$$

where $\text{var}(Z^{nf} - P)$ is the variance of the non-focal elicitation error and $\text{cov}(Z^{nf} - P, P)$ is correction term that accounts for the fact that I allow non-focal elicitation are censored

⁵⁸If Z^{nf} were not censored on $[0, 1]$, then P would be non-parametrically identified from the observation of the distribution of $Z^{nf} = \tilde{Z}$ (this follows from the completeness of the exponential family of distributions). However, since I have modeled the elicitation as being censored at 0 and 1, some distributions of P , especially those leading to a lot of censored values, may not be non-parametrically identified solely from the distribution of Z^{nf} and may also rely on moments of the joint distribution of Z^{nf} and L for identification.

⁵⁹Indeed, if Z^{nf} were not censored on $[0, 1]$ this quantity would equal α .

on $[0, 1]$.⁶⁰ The quantity $\text{var}(Z^{nf}) - \text{cov}(Z^{nf}, L)$ is the variation in Z that is not explained by L . Since the primary impact of changing σ is to change the elicitation error variance of $Z^{nf} - P$, the value of $\text{var}(Z^{nf}) - \text{cov}(Z^{nf}, L)$ provides a strong source of identification for σ .⁶¹ Finally, the fraction of focal point respondents, λ , and the focal point window, κ , are identified from the distribution of focal points and the loss probability at each focal point.

7.1.5 Statistical Inference

Bootstrap delivers appropriate confidence intervals for the estimates of $\inf_{p \in [0, F^{-1}(\tau)]} T(p)$ and the values of $f_P(p|X)$ and $F_P(p|X)$ as long as the estimated parameters are in the interior of their potential support. This assumption is violated in the potentially relevant case in which there is no private information. In this case, $\psi_1 \rightarrow \infty$, $w_1 = 1$, and $a_1 = 0$. As with the lower bound approach, the problem is that in finite samples one may estimate a nontrivial distribution of P even if the true P is only a point mass. Because the parameters are at a boundary, one cannot use bootstrapped estimates to rule out the hypothesis of no private information.

To account for the potential that individuals have no private information, I again use the Wald test from the lower bound approach (see Table 3) that tests whether $\Pr\{L|X, Z\} = \Pr\{L|X\}$ for all X in the sample (by restricting $\Gamma = 0$).⁶² I construct 5/95% confidence intervals for $\inf_{p \in \hat{\Psi}_\tau} T(p)$ by combining bootstrapped confidence inter-

⁶⁰To see this, note that

$$\text{var}(Z^{nf}) = \text{var}(Z^{nf} - P) + \text{var}(P) + 2\text{cov}(Z^{nf} - P, P)$$

and

$$\text{cov}(Z^{nf}, L) = \text{cov}(Z^{nf} - P, P) + \text{cov}(P, L) = \text{cov}(Z^{nf} - P, P) + \text{var}(P)$$

where the latter equality follows from $\Pr\{L|P\} = P$. Subtracting these equations yields equation 9.

⁶¹More generally, Assumptions 1 and 2 impose an infinite set of moment conditions that can be used to identify the elicitation parameters:

$$E[P^N | L = 1] \Pr\{L\} = E[P^{N+1}]$$

It is easy to verify that $N = 0$ provides the source of identification for α mentioned above and $N = 1$ provides the source of identification for σ . This expression suggests one could in principle allow for a richer specification of the elicitation error; I leave the interesting but difficult question of the nonparametric identification conditions on the elicitation error for future work.

⁶²This test also has the advantage that mis-specification of $f_{Z|P}$ will not affect the test for private information. But in principle, one could use the structural assumptions made on $f_{Z|P}$ to generate a more powerful test for the presence of private information. Such a test faces technical hurdles since it involves testing whether $F(p|q)$ lies along a boundary of the set of possible distributions and must account for sample clustering (which makes a likelihood ratio test inappropriate). Andrews [2001] provides a potential method for constructing an appropriate test; but this is left for future work.

vals and extending the 5% boundary to 1 in the event that I cannot reject a null hypothesis of no private information at the 5% level. Given the results in Table 3, this amounts to extending the 5/95% CI to include 1 for the non-rejectees in each of the three settings.

I will also present graphs of the estimated p.d.f., $f_P(p|X)$, c.d.f., $F_P(p|X)$, and pooled price ratio, $T(p)$, evaluated at the mean characteristic, $\Pr\{L|X\} = \Pr\{L\}$, in each sample. For these, I present the 95% confidence intervals and do not attempt to incorporate information from the Wald test. The reader should keep in mind that one cannot reject $F(p|X) = 1 \{p \leq \Pr\{L|X\}\}$ at the 5% level for the non-rejectees in any of the three settings.⁶³ Also, for the estimated confidence intervals of $F_P(p|X)$, I impose monotonicity in a conservative fashion by defining $F_P^5(p|X) = \min_{\hat{p} \leq p} \hat{F}_P^5(p|X)$ and $F_P^{95}(p|X) = \max_{\hat{p} \geq p} \hat{F}_P^{95}(p|X)$ where $\hat{F}_P^5(p|X)$ and $\hat{F}_P^{95}(p|X)$ are the estimated point-wise 5/95% confidence thresholds from the bootstrap.


7.2 Estimation Results

Qualitatively, no trade is more likely for distributions with a thick upper tail of high risks, the presence of which inhibit the provision of insurance to lower risks by raising the value of $E[P|P \geq p]$. In each market setting, I find evidence consistent with this prediction. Figure 4 presents the estimated p.d.f. $f_P(p|X)$ and c.d.f. $F_P(p|X)$ for each market setting, plotted for a mean characteristic within each sample using the price controls, X .⁶⁴ The solid line presents estimates for the rejectees; the dotted line for non-rejectees. Across all three settings, there is qualitative evidence of a thick upper tail of risks as $p \rightarrow 1$ for the rejectees. In contrast, for the non-rejectees, there is less evidence of such an upper tail.

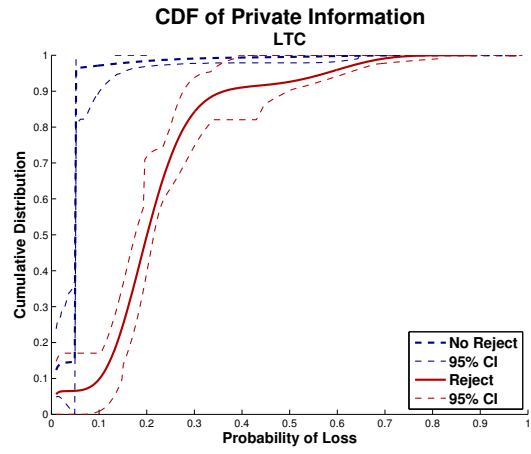
Figure 4 translates these estimates into their implied pooled price ratio, $T(p)$, for $p \leq F^{-1}(0.8)$, and Table 5 presents the estimated minimums over this same region, $\inf_{p \in [0, F^{-1}(0.8)]} T(p)$. Across all three market settings, I estimate a sizable minimum pooled price ratio for the rejectees: 1.82 in LTC (5/95% CI [1.657, 2.047]), 1.66 in Disability

⁶³Estimates of the p.d.f., c.d.f., and minimum pooled price ratio exhibited considerable bias in the bootstrap estimation, especially among the life and disability settings since they have smaller samples. To be conservative, I present confidence intervals that are the union of bias-corrected confidence intervals (Efron and Gong [1983]) and the more traditional studentized-t confidence intervals. In practice, the studentized-t confidence intervals tended to be wider than the bias-corrected confidence intervals for the disability and life estimates. However, the use of either of these methods does not affect the statistical conclusions.

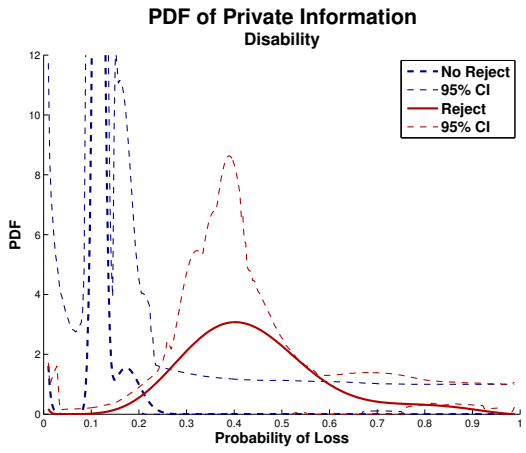
⁶⁴This involves setting $\Pr\{L|X\} = \Pr\{L\}$ in equation (8) within each sample (e.g. $\Pr\{L\} = 0.052$ for the LTC No Reject sample - the other means are reported in Table 2). Appendix E.2 shows the general conclusions are robust to focusing on other values of $\Pr\{L|X\}$ in each sample; I focus on the mean since it is the most in-sample estimate.



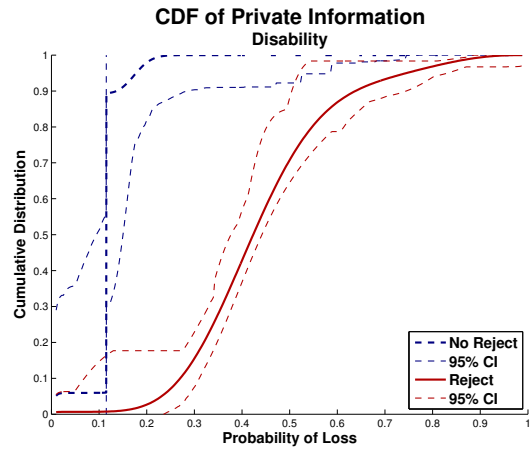
(a)




(b)




(c)



(d)



(e)



(f)

Figure 4: Distribution of Private Information

(5/95% CI [1.524, 1.824]), and 1.42 in Life (5/95% CI [1.076,1.780]). In contrast, in all three market settings I estimate smaller minimum pooled price ratios for the non-rejectees. Moreover, consistent with the prediction of Corollary 3, the estimated differences between rejectees and non-rejectees are large and significant in both LTC and Disability (roughly 59%); for Life the difference is positive (8%) but not statistically different from zero.

The estimates suggests that an insurance market cannot exist for the rejectees unless they are willing to pay a 82% implicit tax in LTC, a 66% implicit tax in Disability and a 42% implicit tax in Life. These implicit taxes are large enough relative to the magnitudes of willingness to pay found in existing literature and those implied by simple models of insurance. For LTC, there is no exact estimate corresponding to the willingness to pay for a marginal amount of LTC insurance, but Brown and Finkelstein [2008] suggests most 65 year olds are not willing to pay more than a 60% markup for existing LTC insurance policies.⁶⁵ For disability, Bound et al. [2004] calibrates the marginal willingness to pay for an additional unit of disability insurance to be roughly 46-109%. This estimate is arguably an over-estimate of the willingness to pay for insurance because the model calibrates the insurance value using income variation, not consumption variation, which is known to be less variable than income. Nonetheless, the magnitudes are of a similar level to the implicit tax of 66% for the disability rejectees.⁶⁶ Finally, if a loss incurs a 10% drop in consumption and individuals have CRRA preferences with coefficient of 3, then $\frac{u'(w-l)}{u'(w)} = 1.372$, so that individuals would be willing to pay a 37.2% markup for insurance, a magnitude that roughly rationalizes the pattern of trade in all three market settings.⁶⁷ In short, the size of the estimated implicit taxes suggest the barrier to trade imposed by private information is large enough to explain a complete absence of trade for the rejectees.

⁶⁵More specifically, the results of Brown and Finkelstein [2008] imply that an individual at the 60-70th percentile of the wealth distribution is willing to pay roughly a 27-62% markup for existing LTC insurance policies This is not reported directly, but can be inferred from Figure 1 and Table 2. Figure 2 suggests the break-even point for insurance purchase is at the 60-70th percentile of the wealth distribution. Table 2 shows this corresponds to individuals being willing to pay a tax of 27-62%. Their model would suggest that those above the 80th percentile of the wealth distribution are willing to pay a substantially higher implicit tax; however Lockwood [2012] shows that incorporating bequest motives significantly reduces the demand for LTC insurance in the upper income distribution.

⁶⁶See column 6 of Table 2 in Bound et al. [2004]. The range results from differing samples. The lowest estimate is 46% for workers with no high school diploma and 109% for workers with a college degree. The sample age range of 45-61 is roughly similar to the age range used in my analysis.

⁶⁷To the best of my knowledge, there does not exist a well-estimated measure of the marginal willingness to pay for an additional unit of life insurance.

Robustness to choice of τ The results in Table 5 focus on the results for $\tau = 80\%$. Table 6 assesses the robustness of the findings to the choice of τ by also presenting results for $\tau = 0.7$ and $\tau = 0.9$. In general, the results are quite similar. For LTC and Disability, both the minimums for the rejectees and non-rejectees are obtained at an interior point of the distribution, so that the estimated minimum is unaffected by the choice of τ in the region $[0.7, 0.9]$. For Life, the minimums are obtained at the endpoints, so that changes in τ do affect the estimated minimum. At $\tau = 0.7$, the minimum pooled price ratio rises to 1.488 for the rejectees and 1.423 for the non-rejectees; at $\tau = 0.9$ the minimum pooled price ratio drops to 1.369 for the rejectees and 1.280 for the non-rejectees. In general, the results are similar across values of τ .

Additional Robustness Checks The results in Tables 5 and 6 evaluate the minimum pooled price ratio for a characteristic, X , corresponding to a mean loss probability within each sample, $\Pr\{L|X\} = \Pr\{L\}$. In Appendix E.2, I show that the estimates are quite similar if, instead of evaluating at the mean, one chooses X such that $\Pr\{L|X\}$ lies at the 20th, 50th or 80th quantile of its within sample distribution.⁶⁸ The minimum pooled price ratio for rejectees ranges from 1.77 to 2.09 in LTC, 1.659 to 1.741 in Disability, and 1.416 to 1.609 in Life. For the non-rejectees I estimate significantly smaller magnitudes in LTC and Disability and the estimated differences between rejectees and non-rejectees for Life remain statistically indistinct from zero.


8 Discussion

The results shed new light on many existing patterns found in existing literature and pose new questions for future work.


8.1 Annuities

There are no rejections in annuity markets. Indeed, annuity companies generally post the same prices to all applicants based solely on their age and gender. At first glance, it may seem odd that I find evidence of private information about mortality that, I argue,


⁶⁸Because of the choice of functional form for $f_P(p|X)$, these should not be considered separate statistical tests of the theory. The functional form is restrictive in the extent to which the shape of the distribution can vary across values of X within a rejection classification. But, nonetheless it is important to ensure that the results do not change simply by focusing on different levels of the index, $\Pr\{L|X\}$.



(a)



(b)



(c)

Figure 5: Pooled Price Ratio

leads to rejections in life insurance. But annuities, which provide a fixed income stream regardless of one's length of life, insure the same (yet opposing) risk of living too long.

However, the pattern of private information found in this paper can explain not only why applicants for annuities are not rejected, but also why previous literature has found adverse selection in annuity markets (Finkelstein and Poterba [2002, 2004]) but not life insurance markets (Cawley and Philipson [1999]). My results suggests that although some people, namely those with health conditions, know that they have a relatively higher than average mortality risk, *few people know that they have an exceptionally lower than average mortality risk*. There's only one way to be healthy but many (unobservable) ways to be sick. Thus, annuity companies can sell to an average person without any major health conditions without the risk of it being adversely selected by an even healthier subset of the population. Annuities may be adversely selected, as the sick choose not to buy them, but by reversing the direction of the incentive constraints, rejections no longer occur.⁶⁹

8.2 Welfare

My results suggest that the practice of rejections by insurers is *constrained* efficient. Insurance cannot be provided without relaxing one of the three implementability constraints. Either insurers must lose money or be subsidized (relax the resource constraint), individuals must be convinced to be irrational (relax the incentive constraint), or agents' outside option must be adjusted via mandates or taxation (relax the participation constraint). However, policymakers must ask whether they like the constraints. Indeed, the first-best utilitarian allocation is full insurance for all, $c = W - E[p]L$, which could be obtained through subsidies or mandates that use government conscription to relax the participation constraints.

However, literal welfare conclusions based on the stylized model in this paper should be highly qualified. The model abstracts from many realistic features such as preference heterogeneity, moral hazard, and the dynamic aspect of insurance purchase. Indeed, the latter may be quite important for understanding welfare. Although my analysis asks why the insurance market shuts down, I do not address why those who face rejection did not purchase a policy before they obtained the rejection condition. Perhaps they don't value insurance (in which case mandates may lower welfare) or perhaps they face credit constraints (in which case mandates may be beneficial). Unpacking the decision of when to purchase insurance in the presence of potential future rejection is an interesting

⁶⁹Moreover, the presence of private information amongst those with health conditions may explain why annuity companies are generally reluctant to offer discounts to those with health conditions.

direction for future work.

8.3 Group Insurance Markets

Although this paper focuses on non-group insurance markets, much insurance is sold in group markets, often through one's firm. For example, more than 30% of non-government US workers have group-based disability insurance; whereas just 3% of workers a non-group disability policy ([ACLI 2010]). Similarly, in health insurance 49% of the US population has an employer-based policy, whereas only 5% have a non-group policy.⁷⁰

While it is commonplace to assume that the tax advantage status for employer-sponsored health insurance causes more insurance to be sold in group versus non-group health insurance markets, tax advantages cannot explain the same pattern in disability insurance. Disability benefits are always taxed regardless of whether the policy is sold in the group or non-group market.⁷¹ This suggests group markets may be more prevalent because of their ability to deal with informational asymmetries. Indeed, group markets can potentially relax participation constraints by subsidizing insurance purchase for its members. Identifying and quantifying this mechanism is an important direction for future work, especially for understanding the impact of government policies that attempt to promote either the individual or the group-based insurance market.

8.4 Private Information versus Adverse Selection

There is a recent and growing literature seeking to identify the impact of private information on the workings of insurance markets. Generally, this literature has searched for adverse selection, asking whether those with more insurance have higher claims. Yet my theoretical and empirical results suggest this approach is unable to identify private information precisely in cases where its impact is most severe: where the insurance market completely shuts down. This provides a new explanation for why previous literature has found mixed evidence of adverse selection and, in cases where adverse selection is found, estimated small welfare impacts (Cohen and Siegelman [2010], Einav et al. [2010a]).

Existing explanations for the oft-absence of adverse selection focus on preference heterogeneity (see Finkelstein and McGarry [2006] in LTC, Fang et al. [2008] in Medigap, and Cutler et al. [2008] for a broader focus across five markets). At a high level, these papers suggest that in some contexts the higher risk (e.g. the sick) may have a lower

⁷⁰Figures according to Kaiser Health Facts, www.statehealthfacts.org.

⁷¹If premiums are paid with after-tax income, then benefits are not taxed. If premiums are paid with pre-tax income (as is often the case with an employer plan), then benefits are taxed.

preference for insurance. Unfortunately, this paper cannot directly shed more light on whether those with different beliefs have different utility functions, u . Indeed, I do not estimate demand and instead assume u is constant throughout the population. But future work could merge my empirical approach to identify beliefs with traditional revealed preference approaches to identify demand, thereby identifying the distribution of preferences for insurance conditional on beliefs and further exploring the role of preference heterogeneity in insurance markets.

But it is important to note that my results raise concerns about the empirical conclusion that the sick have lower demand for insurance; such studies generally have not considered the potential that the supply of insurance to the sick, especially those with *observable* health conditions, is limited through rejections.⁷² It may not be that the sick don't want insurance, but rather that the insurers don't want the sick.

9 Conclusion

This paper argues private information leads insurance companies to reject applicants with certain observable, often high-risk, characteristics. In short, my findings suggest that if insurance companies were to offer any contract or set of contracts to those currently rejected, they would be too adversely selected to yield a positive profit. More generally, the results suggest that the most salient impact of private information may not be the adverse selection of existing contracts, but rather the existence of the market itself.

References

- G Akerlof. The market for lemons: Qualitative uncertainty and the market mechanism. *Quarterly journal of economics*, 84(3):488–500, 1970.
- D Andrews. Testing when a parameter is on the boundary of the maintained hypothesis. *Econometrica*, 69:683–734, 2001.
- D Blackwell. Comparison of experiments. *Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102, Jan 1951.

⁷²Finkelstein and McGarry [2006] note that rejections could pose an issue and provide a specification that excludes individuals with health conditions leading to rejection (Table 4), but they do not exclude individuals over age 80 who would be rejected solely based on age. Fang et al. [2008] does not discuss the potential that rejections limits the extent of adverse selection in Medigap. Although Medigap insurers are not allowed to reject applicants during a 6-month open enrollment period at the age of 65, beyond this grace period rejections are allowed and are common industry practice in most states.

- D Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, Jan 1953.
- J Bound, J B Cullen, A Nichols, and L Schmidt. The welfare implications of increasing disability insurance benefit generosity. *Journal of Public Economics*, 88(12):2487–2514, 2004.
- J Brown and A Finkelstein. The interaction of public and private insurance: Medicaid and the long-term care insurance market. *The American Economic Review*, 98:1083–1102, 2008.
- J Cawley and T Philipson. An empirical examination of information barriers to trade in insurance. *The American economic review*, 89(4):827–846, 1999.
- H Chade and E Schlee. Optimal insurance with adverse selection. *Theoretical Economics*, Forthcoming, May 2011.
- P Chiappori and B Salanié. Testing for asymmetric information in insurance markets. *Journal of Political Economy*, pages 56–78, 2000.
- P Chiappori, B Jullien, B Salanié, and F Salanié. Asymmetric information in insurance: General testable implications. *RAND Journal of Economics*, pages 783–798, 2006.
- A Cohen and P Siegelman. Testing for adverse selection in insurance markets. *The Journal of Risk and Insurance*, 77:39–84, 2010.
- D Cutler, A Finkelstein, and K McGarry. Preference heterogeneity and insurance markets: Explaining a puzzle of insurance. *American Economic Review*, 98(2):157–62, 2008.
- B Efron and G Gong. A leisurely look at the bootstrap, the jackknife, and cross-validation. *The American Statistician*, 37(1):36–48, 1983.
- L Einav, A Finkelstein, and J Levin. Beyond testing: Empirical models of insurance markets. *Annual Review of Economics*, 2(1):311–336, Sep 2010a.
- L Einav, A Finkelstein, and P Schrimpf. Optimal mandates and the welfare cost of asymmetric information: evidence from the uk annuity market. *Econometrica*, 78(3): 1031–1092, 2010b.
- H Fang, M Keane, and D Silverman. Sources of advantageous selection: Evidence from the medigap insurance market. *Journal of Political Economy*, 116(2):303–350, 2008.

- A Finkelstein and K McGarry. Multiple dimensions of private information: Evidence from the long-term care insurance market. *American Economic Review*, 96(4):938–958, 2006.
- A Finkelstein and J Poterba. Selection effects in the united kingdom individual annuities market. *The Economic Journal*, Jan 2002.
- A Finkelstein and J Poterba. Adverse selection in insurance markets: Policyholder evidence from the uk annuity market. *Journal of Political Economy*, Jan 2004.
- L Gan, M Hurd, and D McFadden. Individual subjective survival curves. *Analyses in the Economics of Aging*, ed. D Wise, Aug 2005.
- G Gigerenzer and U Hoffrage. How to improve bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102:684–704, 1995.
- D He. The life insurance market: Asymmetric information revisited. *Journal of Public Economics*, Jan 2009.
- M Hurd. Subjective probabilities in household surveys. *Annual Review of Economics*, pages 543–562, 2009.
- D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.
- L Lockwood. Incidental bequests: Bequest motives and the choice to self-insure late-life risks. *Working Paper*, 2012.
- G. J. Mailath and G. Noldeke. Does competitive pricing cause market breakdown under extreme adverse selection? *Journal of Economic Theory*, 140:97–125, 2008.
- S. Miller, A. Kirlik, A. Kosorukoff, and J. Tsai. Supporting joint human-computer judgement under uncertainty. *Proceedings of the Human Factors and Ergonomics Society 52nd Annual Meeting*, pages 408–412, 2008.
- H Miyazaki. The rat race and internal labor markets. *The Bell Journal of Economics*, 8 (2):394–418, 1977.
- C Murtaugh, P Kemper, and B Spillman. Risky business: Long-term care insurance underwriting. *Inquiry*, 35(3):204–218, Jan 1995.

- W K Newey. Convergence rates and asymptotic normality for series estimators. *Journal of Econometrics*, 79(1):147–168, 1997.
- American Council of Life Insurers. *Life Insurers Fact Book 2010*. American Council of Life Insurers, Nov 2010.
- Congressional Budget Office. *Financing Long-Term Care for the Elderly*. Congressional Budget Office, Apr 2004.
- JG Riley. Informational equilibrium. *Econometrica*, 47(2):331–359, 1979.
- M Rothschild and J Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, pages 629–649, 1976.
- L. J. Savage. *The Foundations of Statistics*. John Wiley & Sons Inc., New York, 1954.
- A.M Spence. Product differentiation and performance in insurance markets. *Journal of Public Economics*, 10(3):427–447, 1978.

Table 1: Rejection Classification

Classification	Long-Term Care		Disability		Life	
	Condition	% Sample	Condition	% Sample	Condition	% Sample
Rejection	Any ADL/IADL Restriction	7.5%	Back Condition	22.7%	Cancer ⁴ (Current)	13.1%
	Past Stroke	8.3%	Obesity (BMI > 40)	1.7%	Stroke (Ever)	7.3%
	Past Nursing/Home Care	13.6%	Psychological Condition	6.3%		
	Over age 80	20.0%				
Uncertain	Lung Disease	10.7%	Arthritis	36.9%	Diabetes	13.8%
	Heart Condition	29.6%	Diabetes	7.7%	High Blood Pressure	50.7%
	Cancer (Current)	15.4%	Lung Disease	5.1%	Lung Disease	10.9%
	Hip Fracture	1.3%	High Blood Pressure	31.3%	Cancer (Ever, not current)	12.1%
	Memory Condition ¹	0.9%	Heart Condition	6.9%	Heart Condition	26.5%
	Other Major Health Problems ²	26.8%	Cancer (Ever Have)	4.6%	Other Major Health Problems ²	23.5%
			Blue-collar/high-risk Job ³	23.3%		
			Wage < \$15 or income < \$30K	65.5%		
			Other Major Health Problems ²	16.2%		

¹Memory conditions generally lead to rejection, but were not explicitly asked in waves 2-3; I classify memory conditions as uncertain for consistency, since they would presumably be considered an "other" condition in waves 2-3.

²Wording of the question varies slightly over time, but generally asks: "Do you have any other major/serious health problems which you haven't told me about?"

³I define blue collar/high-risk jobs as non-self employed jobs in the cleaning, foodservice, protection, farming, mechanics, construction, and equipment operators

⁴Basal cell (skin) cancers are excluded from the cancer classification

Note: percentages will not add to the total fraction of the population classified as rejection and uncertain because of people with multiple conditions

Table 2: Sample Summary Statistics

	Long-Term Care			Disability			Life		
	No Reject	Reject	Uncertain	No Reject	Reject	Uncertain	No Reject	Reject	Uncertain
Subj. Prob (mean) ¹ (std dev)	0.112 (0.195)	0.171 (0.252)	0.132 (0.207)	0.276 (0.245)	0.385 (0.264)	0.335 (0.263)	0.366 (0.313)	0.556 (0.341)	0.491 (0.337)
Loss	0.052 (0.222)	0.225 (0.417)	0.073 (0.26)	0.115 (0.32)	0.441 (0.497)	0.286 (0.452)	0.273 (0.446)	0.572 (0.495)	0.433 (0.496)
Demographics									
Age	71.7 (4.368)	79.7 (6.961)	72.3 (4.319)	54.6 (4.066)	55.0 (4.016)	55.3 (3.795)	70.4 (7.627)	75.3 (7.785)	72.9 (7.548)
Female	0.618 (0.486)	0.619 (0.486)	0.557 (0.497)	0.453 (0.498)	0.602 (0.49)	0.590 (0.492)	0.595 (0.491)	0.564 (0.496)	0.588 (0.492)
Health Status Indicators									
Arthritis	0.479 (0.5)	0.613 (0.487)	0.551 (0.497)	0.000 (0)	0.553 (0.497)	0.346 (0.476)	0.351 (0.477)	0.435 (0.496)	0.443 (0.497)
Diabetes	0.141 (0.348)	0.181 (0.385)	0.150 (0.357)	0.000 (0)	0.090 (0.287)	0.082 (0.274)	0.000 (0)	0.163 (0.369)	0.185 (0.388)
Heart Condition	0.000 (0)	0.401 (0.49)	0.432 (0.495)	0.000 (0)	0.083 (0.275)	0.061 (0.24)	0.000 (0)	0.375 (0.484)	0.332 (0.471)
Sample Size									
Observations (Ind x wave)	9,027	11,259	10,976	763	2,216	5,534	2,689	2,362	6,800
Unique Individuals	4,379	3,587	5,291	391	1,280	3,018	1,720	1,371	4,270
Unique Households	3,206	2,887	3,870	290	975	2,362	1,419	1,145	3,545
Fraction Insured ²	14.0%	10.5%	14.6%				65.1%	63.3%	64.2%
Fraction Insured (Incl Medicaid)	19.5%	20.6%	19.7%						

¹I transform the life insurance variable to 1-Pr{living to AGE} to correspond to the loss definition

²Calculated based on full sample prior to excluding individuals who purchased insurance

Table 3: Lower Bound Results

Classification	LTC			Disability			Life		
	Age & Gender	Price Controls	Extended Controls	Age & Gender	Price Controls	Extended Controls	Age & Gender	Price Controls	Extended Controls
Reject	0.0336*** (0.0038)	0.0358*** (0.0037)	0.0313*** (0.0036)	0.0727*** (0.0092)	0.0512*** (0.0086)	0.0504*** (0.0083)	0.0759*** (0.0088)	0.0587*** (0.0083)	0.0604*** (0.0078)
s.e. ¹	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
p-value ²									
No Reject	0.0048 (0.0018)	0.0049 (0.0018)	0.0041 (0.0018)	0.036 (0.0116)	0.024 (0.009)	0.023 (0.0072)	0.031** (0.0076)	0.025 (0.007)	0.021 (0.0066)
s.e. ¹	0.2557	0.3356	0.3805	0.6843	0.8525	0.9324	0.0102	0.1187	0.2395
p-value ²									
Difference: Δ_z	0.0288***	0.0309***	0.0272***	0.0365*	0.027	0.0274*	0.0449***	0.0338***	0.0397***
s.e. ¹	(0.0041)	(0.0041)	(0.0039)	(0.0146)	(0.0127)	(0.0109)	(0.0112)	(0.0107)	(0.0103)
p-value ³	0.000	0.000	0.000	0.091	0.121	0.092	0.000	0.000	0.001
Uncertain	0.009***	0.0086***	0.0079***	0.0506***	0.0409***	0.0363***	0.0463***	0.0294***	0.028***
s.e. ¹	(0.0024)	(0.0025)	(0.0024)	(0.0058)	(0.0047)	(0.0051)	(0.0058)	(0.0054)	(0.0051)
p-value ²	0.0001	0.0014	0.0001	0.0000	0.0000	0.0000	0.0000	0.0001	0.0001

¹Bootstrapped standard errors computed using block re-sampling at the household level (results shown for N=1000 repetitions)

²p-value for the Wald test which restricts coefficients on subjective probabilities equal to zero

³p-value is the maximum of the p-value for the rejection group having no private information (Wald test) and the p-value for the hypothesis that the difference is less than or equal to zero, where the latter is computed using bootstrap

*** p<0.01, ** p<0.05, * p<0.10

Table 4: Robustness to Moral Hazard: No Insurance Sample

	LTC, Price Controls		Life, Price Controls	
	Primary Sample	Excluding Insured	Primary Sample	Excluding Insured
Reject	0.0358***	0.0351***	0.0587***	0.0491*
s.e. ¹	(0.0037)	(0.0041)	(0.0083)	(0.0115)
p-value ²	0.0000	0.0000	0.0000	0.0523
No Reject	0.0049	0.0038	0.0249	0.0377
s.e. ¹	(0.0018)	(0.0019)	(0.007)	(0.0107)
p-value ²	0.3356	0.8325	0.1187	0.2334
Difference: Δ_z	0.0309***	0.0313***	0.0338***	0.011
s.e. ¹	(0.0041)	(0.0046)	(0.0107)	(0.0157)
p-value ³	0.000	0.000	0.000	0.301
Uncertain	0.0086***	0.0064	0.0294***	0.0269
s.e. ¹	(0.0025)	(0.0024)	(0.0054)	(0.0078)
p-value ²	0.0014	0.1130	0.0001	0.1560

¹Bootstrapped standard errors computed using block re-sampling at the household level (results shown for N=1000 repetitions)

²p-value for the Wald test which restricts coefficients on subjective probabilities equal to zero

³p-value is the maximum of the p-value for the rejection group having no private information (Wald test) and the p-value for the hypothesis that the difference is less than or equal to zero, where the latter is computed using bootstrap

*** p<0.01, ** p<0.05, * p<0.10

Table 5: Minimum Pooled Price Ratio

	LTC	Disability	Life
Reject	1.827	1.661	1.428
5% ¹	1.657	1.524	1.076
95%	2.047	1.824	1.780
No Reject	1.163	1.069	1.350
5% ¹	1.000	1.000	1.000
95%	1.361	1.840	1.702
Difference	0.664	0.592	0.077
5% ²	0.428	0.177	-0.329
95%	0.901	1.008	0.535

Note: Minimum Pooled Price Ratio evaluated for X s.t. $\Pr\{L|X\} = \Pr\{L\}$ in each sample

¹5/95% CI computed using bootstrap block re-sampling at the household level (N=1000 Reps); 5% level extended to include 1.00 if p-value of F-test for presence of private information is less than .05; Bootstrap CI is the union of the percentile-t bootstrap and bias corrected (non-accelerated) percentile intervals from Efron and Gong (1983).

²5/95% CI computed using bootstrap block re-sampling at the household level (N=1000 Reps); 5% level extended to include 1.00 if p-value of F-test for presence of private information for the rejectees is less than .05; Bootstrap CI is the union of the percentile-t bootstrap and bias corrected (non-accelerated) percentile intervals from Efron and Gong (1983).

Table 6: Minimum Pooled Price Ratio: Robustness to τ

Quantile Region: Ψ_τ	LTC			Disability			Life		
	0-70%	0-80%	0-90%	0-70%	0-80%	0-90%	0-70%	0-80%	0-90%
Reject	1.827	1.827	1.827	1.661	1.661	1.661	1.488	1.428	1.369
5% ¹	1.661	1.657	1.624	1.518	1.524	1.528	1.124	1.076	1.000
95%	2.250	2.047	2.030	1.824	1.824	1.795	1.815	1.780	1.754
No Reject	1.163	1.163	1.163	1.069	1.069	1.069	1.423	1.350	1.280
5% ¹	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
95%	1.361	1.361	1.366	1.918	1.840	1.728	1.750	1.702	1.665
Difference	0.664	0.664	0.664	0.592	0.592	0.592	0.065	0.077	0.089
5% ²	0.430	0.428	0.407	0.158	0.177	0.215	-0.344	-0.329	-0.340
95%	1.026	0.901	0.922	1.026	1.008	0.970	0.505	0.535	0.558

¹5/95% CI computed using bootstrap block re-sampling at the household level (N=1000 Reps); 5% level extended to include 1.00 if p-value of F-test for presence of private information is less than .05; Bootstrap CI is the union of the percentile-t bootstrap and bias corrected (non-accelerated) percentile intervals from Efron and Gong (1983).

²5/95% CI computed using bootstrap block re-sampling at the household level (N=1000 Reps); 5% level extended to include 1.00 if p-value of F-test for presence of private information for the rejectees is less than .05; Bootstrap CI is the union of the percentile-t bootstrap and bias corrected (non-accelerated) percentile intervals from Efron and Gong (1983).