

Misperceptions and Contract Distortions in Insurance Markets

Adam Solomon*

April 11, 2021

Abstract

It is well known that private information can impair the functioning of insurance markets, and widely documented that individuals misperceive their private information, but these two facts are rarely analyzed jointly. Theoretically, I find that whether low or high types make upward or downward misperceptions, as well as the assumed market structure, can determine whether the misperceptions increase or decrease welfare relative to the perfect private information benchmark. I then use panel data from a survey of the elderly to compare risk type and risk perception. I find a strong covariance between risk type and risk misperception: Individuals tend to shade their perceptions toward the mean. These analyses suggests that some insurance markets are impaired by individuals not correctly perceiving their riskiness and their value of insurance, whilst others are improved through less than perfect private information.

*Department of Economics, Massachusetts Institute of Technology. Email: adamsol@mit.edu. I thank Jim Poterba, Amy Finkelstein, Stephen Morris, Dmitry Taubinsky, Nathan Hendren, Hunt Allcott, Rafael Veiel and Rob Townsend for very helpful discussions and comments.

Contents

1	Introduction	3
1.1	Literature	6
2	Theory	7
2.1	Equilibrium with Misperceptions	8
2.1.1	Welfare Consequences of Misperceptions	11
2.2	Wilson-Miyazaki equilibrium	13
2.3	The Planner’s Problem	16
3	Empirical Analysis	18
3.1	Ideal Data and Quantities of Interest	18
3.2	Data	19
3.3	Prediction of True Risk	20
3.4	Accounting for Rounding Error	23
3.5	Results	26
3.5.1	Evidence from Wave 1	27
3.5.2	Evidence from the Panel Data	28
4	Discussion and Conclusion	29
	Bibliography	30
A	Theoretical Extensions	34
A.1	Monopoly	34
A.2	Heterogeneity in misperceptions within type classes	36
A.3	Partial pooling in the Rothschild and Stiglitz (1976) framework	38
B	Proofs	39
B.1	Proof of Proposition 1	39
B.2	Proof of Proposition 2	41
B.3	Proof of Proposition 3	42
B.4	Proof of Proposition 4	44
B.5	Proof of Lemma 1	47
B.6	Proof of Proposition 5	48
B.7	Proof of Proposition 7	48
B.8	Proof of Proposition 8	54
B.9	Proof of Proposition 9	54
C	Dataset Construction	55
C.1	Full set of mortality prediction covariates	55

1 Introduction

Accurately understanding and insuring risk is vital to economic well-being. A perpetual problem in the smooth functioning of insurance markets is the presence of private information held by individuals. Yet there is substantial evidence that individuals widely and systematically misperceive their risk in a variety of contexts.⁽¹⁾ In this paper I study how this imperfect knowledge of an individual’s private information affects market equilibrium and welfare.

It is not obvious how inaccurate private information will affect the market. If everyone thinks they are average, then the adverse selection problem might dissipate. But if the high risk think they are a little higher risk than in truth, will this worsen the adverse selection problem, or alleviate it?

To this problem I bring to bear two related insights:

First, the contracts offered in the equilibrium of a private insurance market depend on the informational primitives and potential misperceptions. Assuming, as in Akerlof (1970) and Einav et al. (2010), that the set of tradeable contracts is fixed and the market can only clear on price shuts down this channel.

Second, not all misperceptions are equal in their effect. Misperceptions by high risk types have different effects than misperceptions by lower risk types. Optimistic misperceptions have different effects than pessimistic misperceptions. This is because certain incentive constraints are more important than others in equilibrium, and so a misperception in one direction by one type may loosen a constraint and improve welfare, the same misperception by a different risk type may reduce it.

To develop intuition, consider the annuity market. There are no rejections, minimal pricing on observables and relatively transparent products in the annuity market. This is probably because “there is one way to be healthy, but many (unobservable) ways to be sick” (Hendren (2013)) . Beyond knowing they have no immediate threat to their mortality, individuals likely do not have very accurate information about whether they might live to 75, 85 or 95. The market doesn’t completely collapse through rejections or death spirals, but a more subtle yet equally pernicious distortion is present.

Consider the highest risk in annuity markets - the very long lived. If they do not internalize just how long lived they have the propensity to be, they will find annuity products priced as if they do to be bad value. These individuals will overstate the probability of dying the initial years of an annuity, and understate the probability of survival to very old ages. If they find traditional annuities at actuarially fair prices to be bad value, then at best they can be sold an annuity with a guarantee period, that provides much less longevity insurance and much more bond per dollar annuitized.

Then consider the individuals of moderate longevity risk - neither presently sick nor with nonagenarian parents. They would not like to buy the annuity sold to the longest lived at prices the longest live would. So to separate the high risk from the moderate risk, when the high risk refuse to buy actuarially fair full insurance, even less insurance is sold to the moderate types, who perhaps

⁽¹⁾See, for example, Handel et al. (2020), Handel and Kolstad (2015), Mueller et al. (2018) and Handel et al. (2018).

end up with annuities with a 20 year guarantee period that provides 90 cents of bond and 10 cents of insurance for each dollar. And so on it goes down the risk distribution: When some individuals at the top of the market refuse to buy actuarially fair insurance, the insurance for all is distorted downwards.

This offers a novel explanation for the puzzle of very low annuitization rates: Many marginal types might not find the set of products attractive at the offered prices since some of the riskier are not correctly perceiving their risk, thereby distorting down the amount of insurance per dollar every lower risk type can receive. Faced then with pure annuities that are priced for very long-lived, or guaranteed annuities that are essentially just bonds priced for the quite long-lived, the moderately long-lived rationally choose to forgo participating in this distorted market.

This type of reasoning is not unique to the annuity market in particular or mortality markets in general. In any selection market where the cost of serving a consumer and their willingness-to-pay both depend on private information, which may not be perceived correctly, these dynamics are relevant. In this paper I make this intuition precise both theoretically and empirically.

I first develop the theory as to how misperceptions affect welfare when the contract menu is endogenous. The primary setting is the binary loss framework of Rothschild and Stiglitz (1976) generalized to allow for misperceptions of risk types between and within risk classes. I study how the equilibrium changes with the introduction of risk misperceptions and what the welfare consequences are. As a preview, I find that misperceptions by high risk types have a material effect on the equilibrium, while those by low risk types do not. I find that pessimistic misperceptions by high types may be welfare increasing, whilst optimistic misperceptions by high types (thinking they are closer to the mean) are welfare decreasing. To be sure these conclusions are not sensitive to the particular assumptions of Rothschild and Stiglitz (1976), I study how misperceptions affect the optimal menu of the monopolist and the planner. In both cases I find that misperceptions differ in their impact based on the direction of the misperception and the identity of the misperceiver.

These theoretical results highlight shortcomings of prior work that fixes the contract menu when performing comparative statics. By missing this important channel there is an undue focus on the marginal consumer and interventions aimed at them. When a market can clear only on price with the distribution of true risks fixed, any policy that reduces the willingness-to-pay of the infra-marginal types, or increases the willingness-to-pay of the marginal types, will marginally alleviate the under-supply of insurance due to adverse selection. But when a market clears on both contract type and price, as here, reducing the willingness-to-pay of the infra-marginal types may increase the informational cost of separating the higher and lower types into different contracts. Moreover, increasing the willingness-to-pay of the lower under-insured types need not be an effective or efficient way to increase their insurance, if their under-insurance is due to the supply side needing to separate types and not due to a lack of demand.

Further, by fixing the contract menu the implications for informational interventions are distorted and incomplete. With a fixed contract, an intervention that reduced spread in subjective beliefs between objectively different risk types reduces adverse selection and increases welfare. With

endogenous contracts there is a non-monotonicity: If the different risk types come to believe they are identical to each other, they can then be pooled and adverse selection overcome. But if their subjective belief convergence is anything but total, then the cost of separating these different risk types will add to adverse selection and decrease welfare further.

With these theoretical conclusions in hand I proceed to empirically analyse the pattern of misperceptions that are present within mortality risk. Using data from the Health and Retirement Study I compare an elicited subjective probability of mortality with an objective prediction of that same risk. I find that whilst there is mild overall pessimism in individuals mortality (consistent with the literature) this is not uniform across risk types. High mortality risk individuals are optimistic about their risk, whilst low mortality risk individuals are pessimistic. This indicates that the distortion hypothesized in the annuity market is likely present.

An empirical challenge is that the subjective elicitation from which I infer individuals perception of their risk are highly bunched at focal responses such as 0%, 50% and 100%. More generally, I want to do inference on any gap between individuals true risk and their demand-relevant subjective perception of that risk. The latter quantity may not be accurately reported in a zero-stakes survey such as the HRS, due to noise introduced by the survey that would not be present in a higher stakes demand-relevant setting, rounding to focal responses is the leading example. To show robustness to accounting for this elicitation error due to the survey, I proceed in three ways.

First, using individuals' responses I allocate them conservatively into rounding categories that imply by how much their rounded elicitation could differ from their true subjective belief. I show that allowing for some of the gap between predicted objective risk and elicited subjective risk to be due to rounding, even the conservative residual difference between 'de-rounded' elicitation and the predicted true risk generates the same qualitative pattern: pessimism at the bottom, optimism at the top.

Second, in the first wave of the HRS answers could only be given on an integer scale out of 10, the unit of rounding is clearer. I can then de-round elicitation in the same way and establish, subject to sample size and precision deteriorations, that the qualitative facts go through.

Third, using the panel structure of the data, any individually fixed elicitation errors can be removed and the within-person variation in objective and subjective risk can be isolated. That variation compares whether changes over time in objective risk are correctly reacted to in the subjective elicitation. Even generously allowing for rounding, and restricting the sample to those who actively adjust their elicitation over time, the evidence for misperceptions within person is strong.

This body of evidence offers a qualitative, if not quantitative insight: Individuals misperceive their risk away from the tails, which is a fact widely observed in behavioural economics. Whilst others have studied the correlation of misperceptions with non-risk covariates, I focus on the joint distribution of risk and risk misperception. The most novel implication concerns the resulting distortion of the equilibrium contract menu due to this pattern of risk and misperception.

The intention of this paper is to encourage discussion of this aspect of insurance market function.

Policy recommendation in the future should consider both direct price and quantity interventions (e.g. subsidies and mandates) but also regulation of the policy space (such as occurs in the market-places established by the Affordable Care Act) to counteract distortion in the private market due to these informational frictions.

1.1 Literature

This paper contributes to multiple literatures.

The first contribution is to behavioural mechanism design, specifically to the study of departures from perfect rationality in an insurance context. The problem under perfect rationality was first studied in the canonical papers by Akerlof (1970), Rothschild and Stiglitz (1976), Stiglitz (1977), and to overcome the possible non-existence of equilibria alternative equilibrium notions were proposed by Riley (1979), Wilson (1977), Spence (1978), Miyazaki (1977) and most recently by Azevedo and Gottlieb (2017).

Working within the Rothschild and Stiglitz (1976) framework, there has been work that relaxes the informational assumptions therein. Young and Browne (2000) shows that the qualitative features remain in a non-EU framework, Chassagnon and Villeneuve (2005) characterizes the constrained-efficient frontier when risk misperceptions similar to this paper are made, and Sandroni and Squintani (2007) study a version of this model where optimistic high types are indistinguishable from low types and so cannot be screened apart.

In contrast, assuming a monopolistic rather than competitive market structure as in Akerlof (1970), Jeleva and Villeneuve (2004) studies misperceptions of risk as in this paper, and obtain a monopolistic analogue to the pooling result I obtain in section A.3. Sandroni and Squintani (2013) extend their earlier work to the monopolistic setting. A key departure of this paper is that I do not ex-ante rule out over-insurance resulting from misperceptions and as such obtain the possibility of welfare enhancing misperceptions.

A focus on the inefficiency of the private equilibrium and the scope for welfare-improving government intervention is long-standing. See for example: Crocker and Snow (1992), Veiga (2020), Bisin and Gottardi (2006), Hoy (1982) and Finkelstein and McGarry (2006) . In different contexts the efficiency and welfare implications of misperceptions are studied by Spinnewijn (2013), Spinnewijn (2015), Arya and Mittendorf (2004), Eliaz and Spiegler (2006), Köszegi (2006), Bénabou and Tirole (2002), Spiegler (2006), Gabaix and Laibson (2006), Landier and Thesmar (2009), de la Rosa (2011) , Santos-Pinto (2007) and surveyed by Koszegi (2014). The common but not universal finding is that misperception leads to inefficiency and/or exploitation by better informed parties.

More recently, based on the influential framework by Einav et al. (2010), a literature has documented empirically and theoretically the impact of choice frictions and misperceptions within a setting with fixed contracts. For example, in the setting of health insurance, studies include Abaluck and Gruber (2016a), Abaluck and Gruber (2016b), ?, Handel et al. (2015), Handel et al. (2018), Ericson et al. (2020), Einav et al. (2015), Bhargava et al. (2017), Ketcham et al. (2012), Fang et al. (2008), in the context of unemployment insurance Landais et al. (2017) and Mueller et al. (2018),

and with regards to tax design Allcott et al. (2019) and Stantcheva (2020).

Of particular note are Mueller et al. (2018) and Handel et al. (2020) who study how choice frictions and risk misperceptions vary across the population. The former documents how optimism is concentrated in those with a high risk of long spells of unemployment and derives implications for the optimal public unemployment insurance system. The latter demonstrates that choice quality within a healthcare framework differs markedly across demographic, economic and occupational strata. They focus on the distributional implications that these frictions introduce. In these settings the contract menu is fixed, and so this channel by which misperceptions can affect the equilibrium is closed down by assumption. In this paper I focus squarely on that channel and ask how the distribution of risk misperceptions with respect to risk affect the contract menu, choices and welfare that prevail in equilibrium.

Finally, this work speaks to the accuracy and predictive content in subjective elicitation, particularly within the Health and Retirement Study. This is a literature initiated by Hurd and McGarry (1995) and Hurd and McGarry (2002) that shows that short-run survival outcomes are predicted by responses to the longer-term mortality questions asked in the HRS. A drawback of the early papers was that long-run survival outcomes such as those asked about in the HRS (e.g. survival to 75, or 85) were not observable. A contribution of this paper is to examine how accurate those predictions were ex-post. Also, there is a literature using various econometric techniques to de-round elicitation, for example Gan et al. (2003) and Kleinjans and Soest (2014). The typical finding is that even after a structural model to de-round elicitation is estimated, the quantitative changes are minimal and, in the context of the models explored in those papers, the qualitative changes are essentially null.

The rest of this paper is organized as follows. Section 2 describes the theoretical framework that accounts for the impact of risk misperceptions on the equilibrium contracts and welfare. Different market structures are explored in sections 2.3, 2.2 and A.1. Section 3 presents the empirical analysis, within which 3.1 describes the ideal data, 3.2 details and compares the actual data, 3.3 explains how the prediction algorithm for risk works and 3.5 presents the empirical results. Finally, section 4 discusses the implications of the results for insurance markets and concludes.

2 Theory

There is a unit mass of agents, all of whom are endowed with wealth $w > 0$ and face a stochastic loss of size $l > 0$. There are two types of agents. The fraction $\alpha_H \in (0, 1)$ of the population are high risk types, for whom the probability of the loss occurring is $p_H \in (0, 1)$. The complementary $\alpha_L = 1 - \alpha_H$ of the agents are low risk for whom the loss occurs with probability $p_L < p_H$.⁽²⁾

The risk of loss occurring for each agent is private information unobservable to the insurer. Unlike in the canonical model of Rothschild and Stiglitz (1976) where their private risk of loss is perfectly known by each agent, here the agents may misperceive this probability. High type agents

⁽²⁾All of the analysis that follows extends straightforwardly to any finite number of types. For clarity of exposition I discuss only the two type case throughout.

believe their risk of loss to be q_H where we allow $q_H \neq p_H$, and similarly for low type agents. For a given consumption bundle in the *loss* (L) and *no loss* (NL) states of the world, denoted $\mathbf{c} = (c^{NL}, c^L)$, each agent with *perceived* risk $q_i \in \{q_L, q_H\}$ receives value

$$V(q, \mathbf{c}) = (1 - q_i)u(c_{NL}) + q_i u(c_L),$$

where $u(\cdot)$ is the VnM utility function assumed to be twice continuously differentiable, with $u'(c) > 0$ and $u''(c) < 0$.

The supply side of this market consists of a continuum of identical insurers. Each contract offered by an insurer specifies consumption in each state of the world. A typical contract is labelled $\mathbf{c} = (c^{NL}, c^L)$. Selling such a contract to a consumer with *true* risk $p_i \in \{p_L, p_H\}$ results in a profit of

$$\begin{aligned} \pi(p_i, \mathbf{c}) &= (1 - p_i)(W - c^{NL}) - p_i(c_L - (W - l)) \\ &= W - (1 - p_i)c^{NL} - p_i(c^L + l). \end{aligned}$$

Firms know the true and perceived risks, p_i and q_i as well as the relative sizes α_i , for each $i \in \{H, L\}$ of each group in the population, but they cannot identify which group a particular individual falls into.

As shorthand, write the marginal rate of substitution as (the absolute value of) the slope of the indifference curve of an individual with (real or perceived) risk q indifferent to contract \mathbf{c} as

$$MRS(q, \mathbf{c}) = \frac{1 - q}{q} \frac{U'(c^{NL})}{U'(c^L)},$$

and note the slope of the zero profit line when individuals of average type p are insured is $\frac{1-p}{p}$.

2.1 Equilibrium with Misperceptions

The equilibrium concept we use, following Rothschild and Stiglitz (1976), Sandroni and Squintani (2007), Sandroni and Squintani (2013), is of a *locally-competitive* equilibrium. This captures the idea that firms experiment with new policies that are similar to the current market offering.

Definition 1. *We say a set of contracts \mathcal{C} is a **locally-competitive equilibrium** when: 1. No contract $\mathbf{c} \in \mathcal{C}$ makes strictly negative profits; 2. There exists $\epsilon > 0$ such that every possible contract \mathbf{c}' with $\|\mathbf{c} - \mathbf{c}'\| < \epsilon$ for each $\mathbf{c} \in \mathcal{C}$ does not make strictly positive profits if offered in addition to \mathcal{C} .*

For the rest of the paper, when I write equilibrium I will mean locally-competitive equilibrium, unless otherwise specified. In this context with single-dimensional types and no further heterogeneity between agents, the (typically unique) local equilibrium coincides with the alternative equilibrium relaxations used by Riley (1979) and Azevedo and Gottlieb (2017). Moreover, the local equilibrium

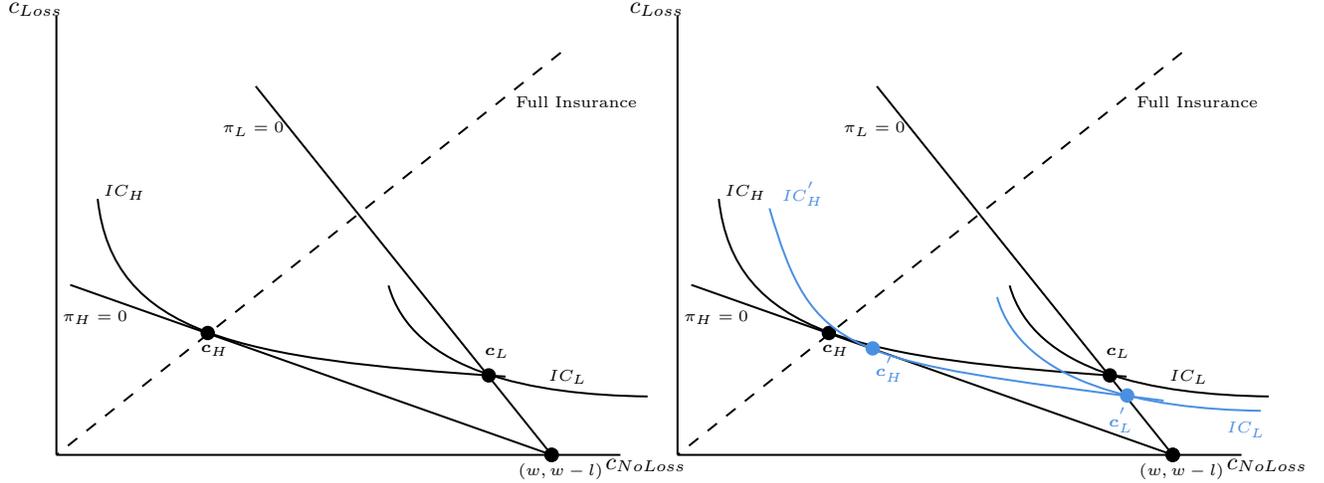


Figure 1: Rothschild Stiglitz Equilibrium

Figure 2: Equilibrium with misperceptions

is the unique Nash equilibrium, when the latter exists. A typical sufficient condition for existence is that the mass of low types is not too large, although I will not explore that issue further.

To understand the forces at work once misperceptions are introduced, the natural point of comparison is the equilibrium in the model of Rothschild and Stiglitz (1976) which assumes no misperceptions, that is, in this terminology, $q_H = p_H, q_L = p_L$. The unique equilibrium is illustrated in figure ??.

The high types purchase the contract marked c_H , which offers full insurance whilst earning the insurers zero profits. The high types indifference curve through c_H is marked IC_H . The low types receive contract c_L defined by the intersection of IC_H with the zero profit line for the low type.

The key qualitative features of the RS equilibrium are: They high risk types receive their first best, full insurance contract (subject to the contract breaking even). The low risk types receive only partial insurance, with this distortion due to the information rents that accrue to the high types due to private information. In order for the high types to self-select into the contract designed for them, the low-type contract must be made sufficiently unattractive. For a high risk type who values insurance highly, a downward distortion in the low risk contract will successfully deter the high types from swapping contracts.

These forces are not specific to the RS model. As shown in subsection A, the incentive compatibility constraint always distorts down the contract received by the low types. In essentially any competitive model this incentive compatibility constraint must be present. Only if the government compelled everyone to reveal their types, or compelled insurers to offer just one defined contract, could an incentive compatibility constraint be avoided.

Now with the primary forces in hand I move to an analysis of misperceptions. The equilibrium characterization is stated in the following proposition. All proofs are in the appendix.

Proposition 1. *Suppose $\|\mathbf{p} - \mathbf{q}\|_\infty < \xi$, with $q_H > q_L$. For small enough ξ ,⁽³⁾ in the unique*

⁽³⁾When I assume ξ is small (the exact bound is given in the proof), qualitatively I am assuming that the contract

locally-competitive equilibrium, high risk individuals buy the contract \mathbf{c}_H that solves:

$$MRS(q_H, \mathbf{c}'_H) = \frac{1 - p_H}{p_H} \text{ and } \pi(p_H, \mathbf{c}'_H) = 0,$$

and low risk individuals receive the contract \mathbf{c}_L that solves

$$V(q_H, \mathbf{c}'_H) = V(q_H, \mathbf{c}'_L) \text{ and } \pi(p_L, \mathbf{c}'_L) = 0.$$

The equilibrium with misperceptions is illustrated in comparison to the equilibrium without misperceptions in figure 2.1. Figure 2.1 illustrates the resulting equilibrium contracts when the high type makes a downward, or optimistic error, by believing that their risk is lower than in truth: $q_H < p_H$. The equilibrium contracts with such a misperception are labeled $\mathbf{c}'_H, \mathbf{c}'_L$ in figure 2.1, without misperceptions are labelled as $\mathbf{c}_H, \mathbf{c}_L$. Per proposition 1, \mathbf{c}'_H is defined as the point of tangency between the zero profit line and the high type's indifference curve IC'_H according to their *perceived* risk q_H . Then \mathbf{c}'_L is defined by the intersection of IC'_H and the zero profit line for the low types. This ensures that the high type is indifferent between the two contracts offered in equilibrium, as proposition 1 specifies.

The effects of misperceptions in this model are felt through two channels. First, an individual who misperceives their risk will, directly, have a higher or lower perceived utility from a given contract. In this case, the optimism of the high types means, loosely, they have a lesser demand for insurance. In equilibrium this lesser demand leads to an equilibrium contract featuring less insurance than without a misperception.

Second, an individual who misperceives their risk will find contracts designed for the other type to be more or less attractive. Depending on the incentive constraints, the other types contract may then change in response to this misperception. In this case, the optimism of the high types leads them to find the original low type contract, \mathbf{c}_L , more attractive than before. To maintain a separating equilibrium, this leads to the low types new contract, \mathbf{c}'_L being distorted away from full insurance even further. This is because the binding incentive constraint is the high type's. I refer to this effect as the **externality** of a misperception, and in the illustrated case it is a negative externality.

So the misperception by the high type has both a direct and indirect effect on the menu of contracts offered in equilibrium. On the other hand, the low type might themselves have made a misperception, yet per proposition 1 this will have no effect on the equilibrium contract menu. That is, only the misperceptions of the high types have a direct or indirect effect on the equilibrium contract menu, with the externality of that indirect effect flowing downwards and affecting the low types.⁽⁴⁾

defined by the intersection of IC_H and $\pi_L = 0$ is the low type's preferred contract out of all separating contracts that break even and the high type weakly doesn't prefer. A simplified version of the case where IC_H doesn't bind and the low type can receive a contract they prefer to \mathbf{c}'_L without tempting the high types to pretend to be low types is treated in Sandroni and Squintani (2007) and Sandroni and Squintani (2013). I will not pursue it further here.

⁽⁴⁾Proposition 1 easily generalizes to multiple types, maintaining the assumption that their perceived risks are

Proposition 1 establishes that the effect of a misperception on the equilibrium contract menu depends crucially on: Who makes the misperception (high or low type) and what type of misperception is made (optimism or pessimism). An important consequence of proposition 1 is that knowing only the level of aggregate misperception in the population, or knowing the misperception made by the marginal type (in a context of insurance purchase) is not sufficient for welfare calculations. In the following section I will make these consequences precise.

2.1.1 Welfare Consequences of Misperceptions

The introduction of misperceptions into this model of competitive insurance purchase has four key implications for welfare that this section will explore. First, that misperceptions of risk may increase or may decrease welfare, depending on the who makes the misperceptions and in what directions those misperceptions are. Second, that knowing the level of average level of misperception in a population is not welfare-sufficient. Third, knowing the misperception made by (almost) any particular type is not welfare-sufficient. Fourth, to increase insurance coverage by a type of people, sometimes interventions (risk-rating, education, subsidies etc) should be directed to infra-marginal types.

To make these statements precise, I define a utilitarian welfare function:

$$W(\mathbf{c}_L, \mathbf{c}_H) = \sum_{i \in \{L, H\}} \alpha_i V(p_i, \mathbf{c}_i). \quad (2.1)$$

The core assumptions here are that each individuals utility is equally weighted, and that the planner evaluates welfare according to *realized* expected utilities (i.e. according to p , not q). This is justified if a planner cares about utilities that are experienced when the risk is realized (or not). It would not be correct if the planner cared about utility received during a forward-looking evaluation of that risk.

The primary result of this section that drives all others are that some misperceptions impair welfare while others improve it. These two situations are illustrated in the left and right panels, figures 4 and 3.

In figure 4, the high types make an optimistic misperception. They think the probability of loss is lower than it actually is. As illustrated, the new equilibrium, relative to the pre-misperception, features the high types buying \mathbf{c}'_H that offers less insurance than their no-misperception contract \mathbf{c}_H . But their original contract \mathbf{c}_H was first-best for them, and their new contract will deliver a loss of realized utility due to the optimistic error. Due to the binding incentive compatibility constraint of the high type, the low type's new contract \mathbf{c}'_L also offers less insurance relative to \mathbf{c}_L , whilst both contract still just break-even. This is even further from full insurance for the low type, and so the high type's error also causes a decrease in the utility that the low type will realize, despite their not making any error.

ordered in the same way as their true risks. With multiple types, and 'small' misperceptions, only the highest type can have a direct effect on the entire equilibrium, whereas all types but the lowest exert an indirect effect on the contract offered to all lower risk types than themselves.

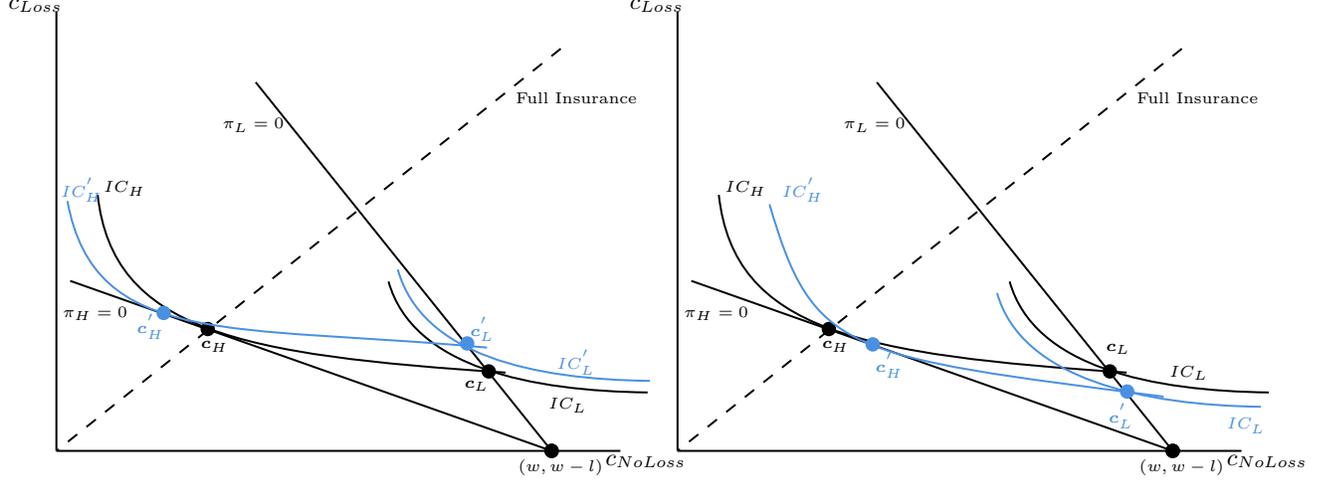


Figure 3: Pessimistic Error by the High Types Figure 4: Optimistic Error by the High Types

In figure 3, the high types make a small pessimistic misperception. They think their probability of loss is slightly higher than it truly is. Inversely to figure ??, they are now over-insured relative to their no-misperception first-best contract. This causes a loss of realized utility, but only a second-order loss by the envelope theorem, since along the zero profit line c_H was a global maximum. In contrast, the incentive compatibility constraint distorting down the low types coverage has been loosened, such that the new contract c'_L delivers strictly higher utility than c_L . Since c_L is away from the first-best, this improvement is first order. Overall, a first order improvement in the low type's utility outweighs a second order decline in the high type's utility, leading to an overall welfare gain.

Hence, there exist misperceptions that strictly increase, strictly decrease, or have no effect at all on, the overall welfare in the market. The above discussion is summarized in the following proposition.

Proposition 2. *Suppose the high types perceived risk is $q_H = p_H + \xi$. Then $\frac{\partial W}{\partial \xi} > 0$ when evaluated at (c_L, c_H) , the undistorted contracts. In particular, welfare increases for small positive ξ , and decreases for negative ξ .*

In contrast to the effects, positive or negative, of errors by the high type, any misperceptions by the low types have no effect at all on the equilibrium menu of contracts, nor on welfare. This is because the low types contract is pinned down by a zero-profit condition according to true probabilities, and the incentive compatibility constraint of the *high* type. Neither of these are affected by a misperception by the low type. This observation, in addition to proposition 2, makes the following immediate. Define the average misperception in the population as

$$\bar{\xi} = \alpha_H \xi_H + \alpha_L \xi_L$$

where $\xi_H = q_H - p_H, \xi_L = q_L - p_L$.

Corollary 1. *For any given (small) $\bar{\xi}$ there exist (ξ_L, ξ_H) and (ξ'_L, ξ'_H) , both of which generate an average error of $\bar{\xi}$, but with the former causing a welfare loss, and the latter a welfare gain, relative to $\xi = \mathbf{0}$.*

There are important implications of this result. Some studies⁽⁵⁾ document misperceptions that occur on average in the population. Statements about welfare are then typically made with reference to these average misperceptions. The result above shows that such welfare implications are possibly wrong, and definitely unfounded. The welfare implications depend on, at least, which risk class makes which type of error, and analysis that does not take this into account might be missing an important channel.

Were this analysis to be embedded in a larger framework of demand for a particular type of insurance, one could define which risk-type was marginal. If the marginal type is found to be optimistic (hence under-demanding insurance), one might suggest an intervention be targeted at the marginal type. For example, an educational or informational campaign that aims to better inform the marginal type as to their true risk. The analysis above shows that this targeting of an intervention might be misguided. In particular, should the marginal type be a low(er) type, no amount of information provision to that type will affect the insurance they purchase in equilibrium. This is because, for all but the highest type, the insurance purchased in equilibrium depends on the misperceived risk *only of the higher types*. So if high and low types are both making a misperception, and a planner corrects the latter but not the former, then the contract purchased by the latter in equilibrium will not change.

The general implication, that non-local or non-marginal interventions are sometimes necessary to for the contract and utility of the marginal types to change, is summarized in the following.

Corollary 2. *Suppose a planner can correct the misperception of some types. The equilibrium contract received by the low types will change if and only if the misperception of the high type is changed.*

2.2 Wilson-Miyazaki equilibrium

Two concerns that have been repeatedly levelled at the Rothschild and Stiglitz (1976) equilibrium concept are that sometimes no equilibrium exists and even when the equilibrium does exist it need not be pareto optimal. To address these concerns a loosening of the equilibrium notion was developed by Wilson (1977) and subsequently further studied by Miyazaki (1977), Spence (1978) and Netzer and Scheuer (2014) amongst others. When deciding whether a given menu is an equilibrium, instead of allowing any profitable deviations, the Wilson-Miyazaki (WM) concept only considers deviations that continue to make a non-negative profit even after all contracts rendered unprofitable by the initial deviation are withdrawn. Thus, firms anticipate best responses to their deviations in this limited manner. A formal definition is given in Wilson (1977).

⁽⁵⁾For example, O’Dea and Sturrock (2020)

From now on I focus on the maximization problem that defines the MW equilibrium. Although there is no planner performing this maximization, the market behaves as if the welfare of the low risk type is being maximized subject to constraints. Consider the program:

$$\max_{\mathbf{c}_H, \mathbf{c}_L} V(q_L, \mathbf{c}_L) \tag{2.2}$$

$$\text{such that} \tag{2.3}$$

$$V(q_H, \mathbf{c}_H) \geq V(q_H, \mathbf{c}_L) \tag{2.4}$$

$$V(q_L, \mathbf{c}_L) \geq V(q_L, \mathbf{c}_H) \tag{2.5}$$

$$V(q_H, \mathbf{c}_H) \geq V(q_H, \mathbf{c}_H^{RS}) \tag{2.6}$$

$$\alpha_H \pi(p_H, \mathbf{c}_H) + \alpha_L \pi(p_L, \mathbf{c}_L) = 0 \tag{2.7}$$

$$(\mathbf{c}_H, \mathbf{c}_L) \in A = \left\{ \mathbf{c}_H, \mathbf{c}_L : MRS_H \leq \frac{1 - p_H}{p_H} \wedge c_L^{NL} \geq c_L^L \right\}. \tag{2.8}$$

This generalizes the program in Miyazaki (1977) and Netzer and Scheuer (2014) by allowing for q_H and q_L to differ from p_H and p_L . In addition I make one more restriction on the equilibrium. The usual restriction on the contract space is that $c_{NL} \geq c_L$, for both types and is called an indemnity constraint (as in Netzer and Scheuer (2014)) such that the loss cannot make an individual better off. I wish to keep this in spirit - not allowing individuals to harm themselves through perceived over-insurance - but I do not want to shut down the wedge between subjective and objectively defined optimal insurance. So I define the set of acceptable contracts, A , such that \mathbf{c} can only allow for small over-insurance for the high type, and since I will assume that errors are small, such that the only over-insurance that can be bought is that which improves perceived utility for the high type. This assumption rules out non-infinitesimal over-insurance for the low types, so the standard indemnity constraint can remain for them.

This rules out equilibria where, according to the misperceived risk, a type could be made better off with a profit-neutral reduction of insurance. If no errors are made, $\mathbf{q} = \mathbf{p}$, this reduces to the standard indemnity constraint. If one omits this restriction, uniqueness of the equilibria that follow might be lost, and the new equilibria will be qualitatively similar but possibly inverted.

Proposition 3. *For small errors $\xi = \|\mathbf{q} - \mathbf{p}\|$, the unique solution to the program (2.2) has (2.7), (2.4) and (2.8) binding and coincides with the the unique MW equilibrium also constrained by (2.8) and (??).*

There is a continuum of contract pairs that satisfy the three constraints in the above program. They differ only in their degree of cross-subsidization from low-type to high type (never the other way). It will be useful to think in terms of this cross-subsidization, so I define:

$$\chi = p_H c_H^L + (1 - p_H) c_H^{NL} - (w - lp_H) \geq 0.$$

And so program (2.2) can be written as a univariate maximization over χ where the constraints

above and definition of χ implicitly define the contracts as functions of χ , e.g. $c_H^{NL} = c_H^{NL}(\chi)$ and so on. The solution can then be thought of as the optimal amount of cross-subsidization from low types to high types to maximise the perceived utility of the former.

When $\chi = 0$ this reduces to the RS equilibrium. When $\chi > 0$ the equilibrium features a profit-making contract offered to the low types and a loss-making contract offered to the high types that makes both better off relative to $\chi = 0$. RS is the particular case where the low types are not willing to cross-subsidize the high types so that the high types will receive more in both states of the world and the low types can move closer to full insurance owing to the relaxed IC constraint. This happens, most intuitively, when there are too many high types, so that the cross-subsidization is too expensive to justify. On the other hand when there are many low types, the RS equilibrium, even if it existed, would be inefficient as the high types could all pay a small subsidy to get a relatively large relaxation of the high types IC constraint.

The prior sections study the comparative statics when $\chi = 0$. Here I focus on the case where, prior to errors being made, there is a non-zero amount of cross-subsidization: $\chi > 0$.

It will turn out that in that situation the movement of χ with q_H or q_L is welfare sufficient. In particular, welfare increases with cross-subsidization. And so if some error increases the cross-subsidization in the equilibrium then welfare will increase, and vice versa. And so if I sign $\frac{\partial}{\partial q_H} \chi^*(q_H)$ or similarly for q_L the welfare implications will follow.

Proposition 4. *Suppose without errors $\chi^* > 0$. Then if a small error is made: welfare increases in χ , $\frac{\partial}{\partial q_H} \chi^*(q_H) < 0$ and $\frac{\partial}{\partial q_L} \chi^*(q_L) > 0$.*

To understand the intuition for this result, one should keep in mind that the χ^* in the MW equilibrium is pushed higher and higher so long as the low types are still willing to give a dollar to the high types to relax the incentive constraint. The question is how an error by either type changes the incentive for the low types to give the final dollar of cross-subsidy to the high types.

First suppose the high types make a small upward error, $q_H > p_H$. Then, holding the no-error χ^* fixed, the high types contract will move up and left on the same iso-profit line, as in figure 4 (except with the iso-profit line not necessarily earning zero). The low types will then also receive closer to full insurance and be better off even before the cross-subsidy is adjusted. But because the low type is now closer to full insurance, the marginal utility from relaxing the cross-subsidy is now lessened, and so the final dollar of cross-subsidy that was optimal with no errors is now longer optimal with this error. There is a countervailing force: that the marginal relaxation of the incentive constraint with a dollar of cross-subsidy is larger with the higher q_H (as the indifference curves are flatter) but this is dominated by the first effect. As such the new $\chi^* |_{q_H > p_H} < \chi^* |_{q_H = p_H}$. But welfare is evaluated according the objective probabilities, and $\chi^* |_{q_H = p_H}$ was optimal according to objective probabilities, and so the final foregone dollar of cross-subsidy is welfare decreasing.

Conversely, but similarly, if the low types make a downward error, $q_L < p_L$, they are no longer willing to pay for the final dollar of cross-subsidy, even though it is welfare optimal. This explains the sign of $\frac{\partial}{\partial q_L} \chi^*(q_L) > 0$.

These results reverse the logic from the RS equilibrium: errors that make a type think they are closer to the other type are welfare enhancing, errors that spread the difference between perceived types are welfare reducing. The intuition here is that **type-spreading errors are substitutes for cross-subsidies**. That is, in the RS equilibrium where cross-subsidies are impossible, the only mechanism by which the incentive constraint can be loosened is through errors that spread the types apart. This is what was found in proposition 2. In the MS equilibrium, cross-subsidies are present without errors. Those cross-subsidies loosen the incentive constraint by as much as the low type is willing to pay. When we add an error which also loosens the incentive constraint, the low types cross-subsidize a little less. In this sense when cross-subsidies are present they are crowded out, suboptimally, when errors are introduced.

This highlights a crucial distinction: the impact of errors depends on whether or not contracts are currently cross-subsidizing each other.

2.3 The Planner's Problem

The contrasting implications in the prior two sections are problematic. In this section I study the planner's problem and show that the patterns obtained identically track the Rothschild and Stiglitz (1976) style formulation from section 2.1. This, in combination with similar conclusions in the monopolist case in section A.1, restore confidence that it is the MW equilibrium concept that generates outlier predictions. This is due to MW allowing for a particular type of cross-subsidization that is not welfare maximizing.

I continue to assume that the planner has preferences as in (2.1) and can implement any incentive compatible and budget-balanced menu of contracts. Their problem is then:

$$\max_{\mathbf{c}_H, \mathbf{c}_L} \sum_{i \in \{L, H\}} \alpha_i V(p_i, \mathbf{c}_i) \quad (2.9)$$

$$\text{such that} \quad (2.10)$$

$$V(q_H, \mathbf{c}_H) \geq V(q_H, \mathbf{c}_L) \quad (2.11)$$

$$V(q_L, \mathbf{c}_L) \geq V(q_L, \mathbf{c}_H) \quad (2.12)$$

$$\alpha_H \pi(p_H, \mathbf{c}_H) + \alpha_L \pi(p_L, \mathbf{c}_L) = 0 \quad (2.13)$$

$$(\mathbf{c}_H, \mathbf{c}_L) \in A. \quad (2.14)$$

Equations (2.11) and (2.12) are incentive compatibility constraints for the high and low risk types respectively, and equation (2.13) is an aggregate resource constraint which allows for cross-subsidization between types, unlike in the competitive equilibrium in which each contract individually broke even. I also re-emphasize that whilst incentive constraints are evaluated according to subjectively perceived probabilities q_i , welfare and profit are evaluated according to objective probabilities p_i . This is in contrast to the optimization problem in the MW framework, in which the low-type's utility is maximized according to their subjective probability.

Consider again the cross-subsidization from low types to high types χ . Denote the cross-subsidization in the planner's optimum by $\chi_{Planner}^* > 0$ and in the MW problem by $\chi_{MW}^* > 0$.

Lemma 1. *Suppose there are no errors, $\mathbf{q} = \mathbf{p}$. If $\chi_{MW}^* > 0$ then $\chi_{planner}^* > \chi_{MW}^*$.*

This says that the planner's incentives to cross-subsidize types are weakly higher the hypothetical optimizer of the MW problem, and strictly so when the MW equilibrium features strictly positive cross-subsidies. The intuition for this is straightforward: The MW equilibrium keeps increasing the cross-subsidy so long as the low types are willing to pay an extra dollar to the high types to relax the latter's IC constraint. This process ends when the benefit to the low type from the marginal dollar is zero. But at that point the benefit to the high type of an extra dollar is strictly positive, so that incentives to the planner are for at least a little more cross-subsidization.

The MW equilibrium cannot feature $\chi < 0$ because it would be an allowable deviation for a different firm to offer the RS $\chi = 0$ contracts to the high types. Contrastingly, the planner's solution might feature cross-subsidization from the high type to the low type. Even though this is a cross-subsidization that would benefit the objective function of the MW problem, the nature of competition precludes it. It is natural that competition between firms, as opposed to the fiat of the planner, leads to some inefficiency.

Now I introduce small errors in the perceived probabilities. The planner's optimum has the following features.

Proposition 5. *Suppose there are small errors, $|\mathbf{q} - \mathbf{p}| < \xi$. The optimal contracts vary with the high type error q_H , and welfare increases in q_H . Neither the optimal contracts nor welfare change with low type errors q_L .*

Once one has internalized the logic of the RS equilibrium this conclusion is unsurprising. In the planner's no-error optimum the high type's incentive constraint binds and the high type receives their optimal contract (full insurance) conditional on the level of cross-subsidization. Once a small error is introduced, over-insuring the high type is now possible so that the contract can be of the form depicted in figure 3. That is, a first order gain can be given to the low types with only a second order loss to the high types.

That the low types error doesn't matter for the optimum or welfare is also intuitive. The low type's perceived risk q_L appears only in the incentive constraint for the low type, which is slack optimally, and so small changes to it will not be material. In the MW case, it appeared directly in the objective, and so an increase in q_L meant that there was a little more incentive for cross-subsidization, and because the no-error MW cross subsidization was inefficiently low this helped welfare. This required a third best situation: some cross-subsidization was effectively allowed but not the efficient amount, and the (imaginary) optimizer in the MW equilibrium cared about perceived not experienced utility. These two inefficiencies were partially cancelled out by a low type error, which would not otherwise have had an effect if either of the two inefficiencies were not present.

This analysis partially resolves the ambiguity around the conflicting conclusions of the RS and MW models. That the planner’s problem displays identical comparative statics with respect to q_H and q_L with respect to the former, and the monopolists (see section A.1 agreement in the q_H result, lends credibility to the qualitative insights from the core RS model. I proceed to the empirical analysis with those qualitative insights taken as correct.

3 Empirical Analysis

3.1 Ideal Data and Quantities of Interest

The novel insight in the preceding analysis is that it is not sufficient to measure the misperception averaged over the whole population. We need to know, for each true risk type p_i , what their perceived risk q_i is. So we are interested in the distribution of $q_i | p_i$. The following analysis studies key parameters of this conditional distribution.

In this paper, I specialize the analysis and discussion to consider mortality risk, as relevant to the life insurance and annuity markets. Specifically, I will consider the probability of death by the age of 75. But the techniques developed are much more general, and can be used to analyze any other risk market such as health, disability or property.

The first key datum is a measure of an individuals perception of their mortality risk, q_i . Ideally, I would observe this perception as relevant to an actual life insurance or annuity purchase. Unfortunately, the best I can do is to observe this quantity in setting of a non-incentivized survey, which is detailed below. The key issue with this datum is that there might be ‘elicitation error’ due only to the survey that wouldn’t be present in an actual insurance purchase. I develop methods to circumvent this.

Secondly, I need a measure of the individuals true propensity for the risk to occur, p_i . That is, need a measure of the true probability of death by age 75. This is a binary risk only realized once per person, and so is in some sense unknowable at an individual level. But at a group or population level it is coherent to assign propensities that are on average correct. This propensity needs to be predicted in a non-degenerate way so that the conditional distribution of $q_i | p_i$ can be analyzed.

To make that prediction, I first need as rich a set of mortality-relevant covariates as possible. Natural candidates include medical history data, longevity information about the individuals parents and other close family, employment history, income and wealth information. Then I need to know how this risk (death by 75) was realized for at least some of the population of interest. More details are given in section ??, but by splitting the sample into a test set and training set, valid out-of-sample probabilities of death by 75 can be retroactively predicted. These probabilities will be (my noisy estimate of) p_i . They can then be compared to subjective assessments of that same probability q_i to do inference on moments of the distribution $q_i | p_i$.

3.2 Data

The data come from the Health and Retirement Study (HRS), a biennial panel survey of Americans who are (mostly) older than 50 years. It contains the three key ingredients described above.

1. The subjective probability of mortality by age 75 is elicited from all respondents aged 65 and younger. The survey asks: "*What is the percent chance that you will live to be 75 or more*". This will form the raw estimate of q_i .
2. The survey has run since 1992, with relatively low attrition.⁽⁶⁾ As such, the outcome of the risk of death by age 75 is known for approximately half of the respondents. Those who attrited before age 75 or who are still younger than age 75 are excluded. Those that remain are assigned a death indicator $D_i = 1$ if they died before age 75 and $D_i = 0$ if not.
3. The survey collects a rich variety of health data, employment data, family history data and others relevant to mortality prediction. The entire set of data used in the prediction algorithm is specified in appendix C.1. The various prediction algorithms are described in the next section.

I pause here to note some advantages and disadvantages of the HRS. The primary features are the availability of the three key classes of data above. Few collect data on subjective elicitation of medium run mortality risk. Fewer still have gone on for long enough that the realization of the mortality risk is known. Whilst some progress could be made using, for example, hazard modelling if only short-run mortality was known, a distinction needs to be made between subjective elicitation with high short run accuracy (e.g. if an individual has a terminal disease their prediction is likely to be very accurate and resolved soon) and higher long run accuracy. This is one, to my knowledge, one of the first studies to assess the accuracy of long-run mortality predictions over the entire interval to which the prediction pertains. For no other reason than the lack of time, this has not been possible until quite recently.

A key disadvantage is the lack of incentivization for giving accurate or thoughtful answers. There are some checks in place to make sure the respondents are not answering randomly, but the survey is still entirely cheap talk. It is this uncertainty as to how well the elicitation from the HRS correspond to true private estimates of risk that presents a central empirical problem in the analysis that follows.

Some of the key data are summarized in table 1. Each observation is an individual-wave, and so each individual typically appears multiple times. I separate out those observations for which the respondent ended up dying by age 75, and those for which the respondent survived. As expected, in both observable characteristics (e.g. major medical problems) and in subjective elicitation, those who die differ from those who survive.

⁽⁶⁾See, for example, Banks et al. (2011), who finds low attrition and no correlation between attrition and major medical conditions.

Variable	Survive to 75	Death before 75
Health		
Cancer	0.06 (0.23)	0.11 (0.32)
Heart Disease	0.10 (0.29)	0.25 (0.43)
Lung Disease	0.04 (0.18)	0.05 (0.35)
Diabetes	0.09 (0.28)	0.25 (0.43)
Doctor Visits (past year)	6.5 (11.1)	11.6 (25.6)
Self-health Rating (1=Excellent, 5=Poor)	2.51 (1.09)	3.30 (1.20)
Private Health Insurance Plans	0.89 (0.56)	0.72 (0.69)
Demographics		
Age	60.13 (3.60)	59.32 (3.76)
Female	0.57 (0.5)	0.45 (0.5)
Sample Size		
Observations (Ind \times wave)	25,027	10,940

Table 1: Selected Summary Statistics

Amongst those who remain in the sample in multiple waves, the general trend is for the probability of death by 75 to decrease, simply because each wave means there are two fewer years to survive until reaching 75. The evidence I present below is broadly consistent with individuals not sufficiently updating their survival probabilities as they survive for two more years. I could instead have analysed the probability of surviving 10 more years from the current age, as the HRS (almost) asks this question. I leave that for separate research.

3.3 Prediction of True Risk

Since the true probability of death p_i by age 75 is unobservable to the researcher (indeed, to anyone) the best I can do is make a prediction \hat{p}_i that has favourable properties for inference. I begin by taking the set of individuals for whom death or survival to 75 is observed, and randomly splitting them into a training set and a test set with equal probability. Note this is done on the individual level, not the individual times wave level. I train various models to predict death before 75 as a function of covariates. Inference is then performed on all the observations that were left out of the training set, i.e. the test set. Sample-splitting of this form is needed to avoid overfitting.

There are two prediction methods that I will refer to at various points. The first is a classical

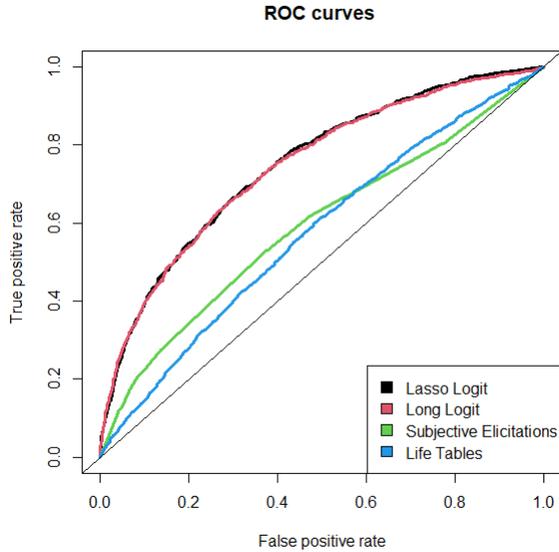


Figure 5: Operator Receiver Curves for different predictors

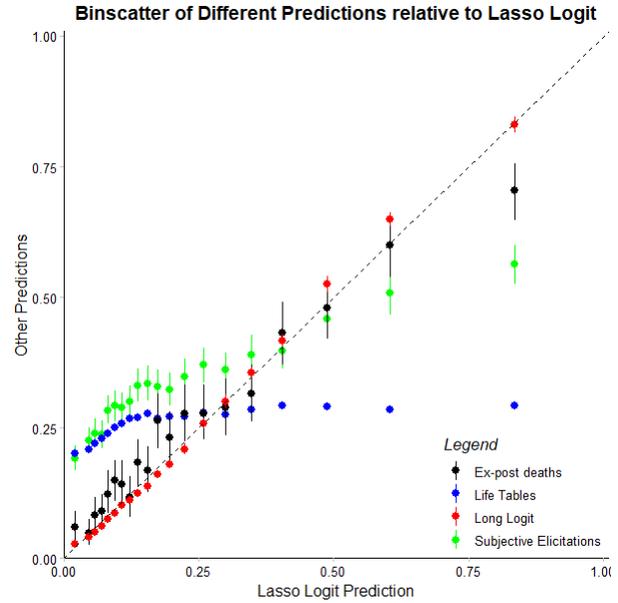


Figure 6: Binscatter of different predictions relative to Lasso Logit

logit model with over 400 variables and interactions that seemed, ex-ante, relevant to mortality prediction. I will refer to this as the 'logit' or 'long logit' model. The second is a LASSO logit model. The LASSO adds to the logit's log-likelihood function a 'penalty' or 'regularization' term of the form $\lambda |\beta|$ where λ is a tuning parameter to be chosen by cross-validation and β is the vector of coefficients to be estimated. The LASSO 'zeroes out' coefficients of variables with minimal predictive power which leads to a less noisy and better prediction.

Figures 5 and 6 are two different illustrations of the quality of the prediction from the logit LASSO model and the long logit model. Natural comparisons for the prediction accuracy are to the using life tables that only use gender and age, and using just the subjective elicitation. All these four different predictions are contrasted in the figures.

Figure 5 shows the receiver operating characteristic curves (AUC), or the tradeoff between false positives and false negatives. The 45 degree line corresponds to a random classifier. The blue life tables line is the receiver operating characteristic curve corresponding to assigning each individual's probability of death by 75 to be whatever the population life tables for their sex and gender would imply. Naturally, this is not a very good predictor, with an area under the curve of 0.57. The green line is the curve obtained when individuals raw subjective elicitations are used as predictors of death. These are only slightly better than the life table probabilities, with an area under the receiver operating characteristic curve (AUROC) of 0.59. On the other hand, the logistic classifiers I trained does much better at predicting death by 75 than the subjective probabilities or the life tables. The AUROCs are 0.746 and 0.749 for the long logit and lasso logit curves respectively.

Figure 6 is a bin-scatter with the best prediction algorithm - the lasso logit - on the horizontal axis, and all the other prediction algorithms, and ex-post deaths on the vertical axis. It is

constructed as follows: all observations are sorted by the prediction outputted by the lasso logit model, and then put into 20 bins. In each bin the average ex-post deaths, prediction from the life tables, prediction from the long logit, and subjective elicitation, are all computed and plotted. So for example, looking at the bins that are above 0.50 on the horizontal axis: of those in that bin with an average prediction from the lasso logit of 0.5, the average ex-post deaths is exactly on the 45 degree line and hence agrees exactly with the prediction of lasso logit, the long logit is a little above (0.52) and the subjective elicitations are a little below (0.48).

There are a few interesting patterns revealed by figure 6. First, the two lasso models predict very well on average as the red bins (long logit) and the 45 degree line (lasso logit) are both close to the black dots (actual deaths). The life tables do very poorly, as expected. The subjective elicitations are overly high for those with low predicted risk, and overly low for those with high predicted risk. Whilst these approximately wash out, so that on average over the entire sample subjective elicitations are quite close to predicted objective risk (see the following histograms), this hides the large and disparate gaps between subjective elicitation and predicted objective risk at the top and bottom of the risk distribution. This is a key distinction as the theory indicates: the errors made by the high risk types matter much more than the errors made by the low risk types, and the average error in the population is not very informative.

Recall here we do not know objective risk $p_i = 1 / (1 + \exp - x_i \beta) \equiv S(x_i \beta)$ perfectly, we just have a prediction \hat{p}_i of it from the logistic routines. In the case of ordinary logit, the estimator minimizes the log-likelihood function with score given by:

$$\nabla_{\beta} l(\beta, y, X) = \sum_i [y_i - S(x_i \beta)] x_i.$$

As such, mechanically it is the case that the residuals satisfy:

$$E [[y_i - S(x_i \beta)] x_i] = 0.$$

And so even though we do not have a perfect prediction of risk, we have

$$p_i = \hat{p}_i + \eta_i, \text{ where } E(\eta_i \hat{p}_i) = 0.$$

So in the case of ordinary logit, if I assume I exactly know each individuals subjective belief q_i , then a simple OLS regression of \hat{p}_i on q_i will deliver a correct estimate of the conditional mean function $E(q_i | p_i) = 0$. This will then allow for inference about the average error made by the high risk or the low risk, as the theory demands.

The lasso logit is more desirable from the point of view of prediction, as the variable selection is data-driven and the out-of-sample properties better (although only marginally so, as figures 5 and 6 demonstrate). However, for the second stage inference about $E(q_i | p_i)$, there are two issues. The first is potential bias due to the regularization in the lasso prediction first stage, as in Chernozhukov et al. (2017), Singh and Sun (2020) and Chernozhukov et al. (2018). The second is the lack of

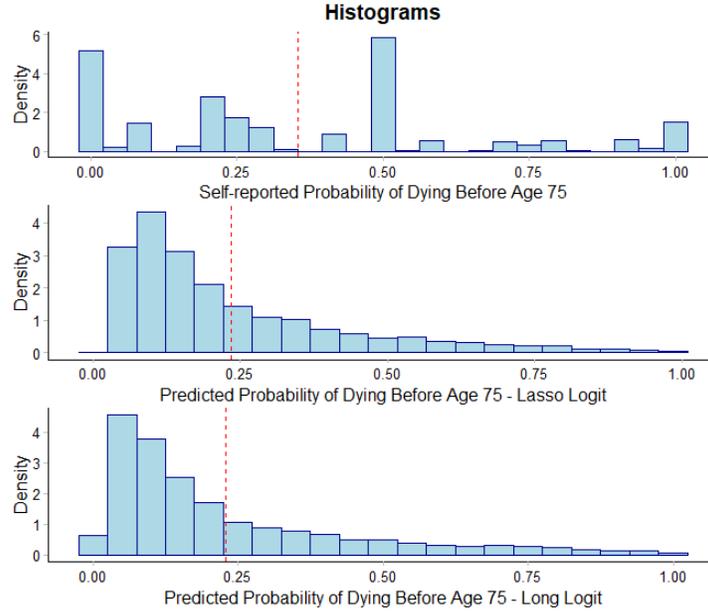


Figure 7: Histograms of self-reported and two predicted probabilities of dying before 75

orthogonality between the predicted \hat{p}_i and the residual values from prediction, as the log-likelihood and score functions will have an extra term that comes from the 1-norm regularization lasso. As such the second-stage is less reliable. So I will present all estimates using both the lasso logit that has superior prediction and the long lasso that has more desirable properties for the second stage inference. Histograms of the predictions from the lasso logit and the long logit are in the upper two panels of figure 7 respectively.

The next issue to address is q_i , the subjective perception of the risk. A histogram of the raw responses to the HRS question is in the third panel of figure 7. The clearest issue is the rounding. That is, even if individuals would have more precise beliefs about their probability when they enter an insurance market, in this low-stakes setting perhaps they simply report their belief rounded to the nearest 10,25 or even 50%. It is possible that individuals rounded beliefs are biased relative to their true risk, but their unrounded demand-relevant beliefs are accurate. I account for this in various ways in the next sections.

3.4 Accounting for Rounding Error

When individuals are faced with an HRS interviewer and asked to consider their mortality prospects, there are broadly two types of errors they might make. One is a genuine misperception, of the kind I am interested in throughout this paper. That is, they might truly think their mortality prospects are better or worse than they actually are, and this misperception would persist even in a perfect elicitation of their belief. These I will call **true errors** or **true misperceptions**. On the other hand, the response given in the HRS might be infected with what I will call **elicitation noise** or **elicitation error**. By this I mean a discrepancy between the individuals 'true' perception of their

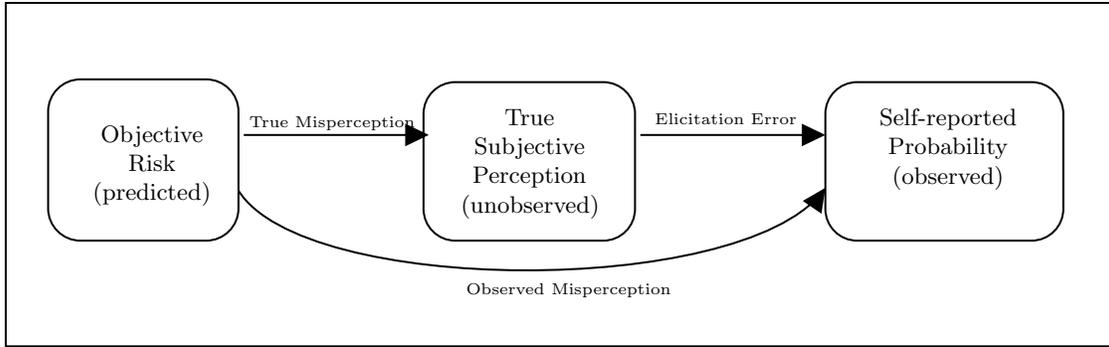


Figure 8: Schematic view of true misperception versus elicitation error

risk and their reported perception of their risk that is due only to the artefact of the survey and its nature. In particular, there might be errors made by the individual only because the survey is unincentivized, or the individual is tired and so rounds their answer, or because no actual insurance purchase is being made here. If such errors could be corrected, hypothetically, if we observed the individual's perception in a demand-relevant incentivized high stakes and otherwise perfect scenario, then I deem such errors elicitation noise of elicitation error.

The challenge is that I only want to do inference on the true misperceptions after filtering out the elicitation noise (primarily rounding) but neither of these quantities can be disentangled. Only the sum of both is observed by looking at the difference between an individual's (predicted) true risk and their elicited perception of that risk from the HRS. A summary of this is in figure 8. I observe the 'observed misperception' but wish to filter out the 'elicitation error' to be able to do inference only on the 'true misperception'

I will attempt to overcome this in multiple ways. First, from the individuals response I will classify them into a rounding category, and then attribute as much of the observed misperception to elicitation error as is consistent with the rounding category. I will do this in a conservative way and then argue that the residual (conservatiely small) true misperception still generates the same qualitative patterns as the observed misperception. Second, I will leverage a change in scale on which the probabilities were elicited in the first wave of the HRS. In that wave, the unit to which everyone is rounding is unambiguous and so there will be much less uncertainty that the rounding category is correct. Third, I will study within person variation across time, so that any individually fixed elicitation errors are removed, as well as applying the rounding procedure to be described in detail, so that plausibly only time varying true-misperception remains. Using all these methods the same qualitative conclusions will obtain.

Before proceeding I should note that even though actual insurance purchases are available in the data for some insurance markets (e.g. life insurance), I will not use them to infer private beliefs. This is because insurance markets are already distorted in multiple ways such that private perceptions of risk cannot be easily extracted from demand information. For example, in the life insurance market there is a litany of health conditions that would lead to rejection if one tried to

buy insurance. As such, life insurance holding is *negatively* correlated with elicited risk perception, contrary to the theory.⁽⁷⁾ Put another way, the insurance markets are in a second-best state due to possibly many factors. To isolate the impact of the one factor under study - misperceptions in risk - we cannot look at the current equilibrium and isolate which distortion is due to this factor and which is due to other factors on which we do not have data.

I now describe the de-rounding procedure. As can be seen in the top panel of figure 7, many of the self-reports are multiples of 50%, 20% or 10%. I take the view that a report of 40% likely means that the individual has rounded to the nearest 20%. So even if the 40% represents rounding due to impatience or inability to narrow it down more precisely, I can be quite sure the true subjective belief, if perfectly reported, lies between 30% and 50%. So to adjust the 40% report, I consider what the (predicted) true risk for that individual is. If the individuals (predicted) true risk is $\hat{p}_i \in [30\%, 50\%]$, then I set their "de-rounded elicitation" \tilde{q}_i to be $\tilde{q}_i = \hat{p}_i$ such that, after accounting for rounding in this way, their true misperception is zero. Alternatively, if their (predicted) true risk is $\hat{p} < 30\%$ then I set $\tilde{q}_i = 30\%$, and if their (predicted) true risk is $\hat{p} > 50\%$ then I set $\tilde{q}_i = 50\%$.

In words: If the gap between an individual's raw elicitation and their (predicted) true risk is fully explainable by rounding error, then I assume their true misperception is zero. If the gap between an individual's raw elicitation and their (predicted) true risk is not fully explainable by rounding error, then I attribute as much as consistent with their rounding category to be due to rounding, and call the remainder (a conservative bound on) true misperception.

Finally, I categorize individuals into rounding categories as follows. If an individual is designated an $x\%$ rounder, then this means I allow for up to $x/2\%$ error in each direction, as in the example detailed above. I categorize:

- If $q_i \in \{0, 50\%, 100\%\}$, then assume 50% rounder.
- If $q_i \in \{25\%, 75\%\}$, then assume 25% rounder.
- If $q_i \in \{10\%, 20\%, 30\%, 40\%, 60\%, 70\%, 80\%, 90\%\}$, then assume 20% rounder.
- If $q_i \in \{5\%, 15\%, \dots, \%95\}$, then assume 10% rounder.
- Everyone else, assume no rounding.

According to this procedure, the de-rounded perceptions \tilde{q}_i are computed. In the next section inference will be performed as if either subjective beliefs are truly q_i or as if they are truly \tilde{q}_i . Likely the former will understate the elicitation error and the latter will overstate it, with the true private beliefs lying somewhere in between. But what we will find is that the qualitative conclusions hold in both cases, and thus likely would be true if we could observe the true subjective perceptions devoid of any elicitation noise.

⁽⁷⁾Note though that conditional on not having a rejectable condition - life insurance holding is correlated with perception of risk, indicating that there is demand-relevant private information in the HRS elicitation. See, for example, Hendren (2013).

Quantity	Lasso Logit		Long Logit	
	Raw q_i	Derounded \tilde{q}_i	Raw q_i	Derounded \tilde{q}_i
$E(q_i p_i = 0)$	0.23*** (0.01)	0.13*** (0.005)	0.25*** (0.01)	0.16*** (0.005)
β	0.46*** (0.003)	0.70*** (0.02)	0.38*** (0.03)	0.58*** (0.02)
$E(q_i p_i = 1)$	0.69*** (0.02)	0.83*** (0.01)	0.63*** (0.02)	0.74*** (0.01)
Observations (Ind \times wave)	5,305	5,305	4,736	4,736
Individuals	2,227	2,227	2,090	2,090

Table 2: Results. *** means significant at the 1% level against 0 (first row) or 1 (second and third rows).

3.5 Results

It has been quite widely established⁽⁸⁾ that individuals are on average pessimistic about their mortality beliefs. They underestimate the chances of surviving for a long time. What I seek to interrogate here is how this misperception is distributed over the risk distribution. That is, is everyone uniformly pessimistic, is there some variation uncorrelated with risk, or is there a clear relationship between risk and risk misperception.

I estimate the following four OLS specifications:

$$\begin{aligned}
 q_i &= \beta \hat{p}_i^{LassoLogit} + \epsilon_i, \\
 q_i &= \beta \hat{p}_i^{LongLogit} + \epsilon_i, \\
 \tilde{q}_i &= \beta \hat{p}_i^{LassoLogit} + \epsilon_i, \\
 \tilde{q}_i &= \beta \hat{p}_i^{LongLogit} + \epsilon_i.
 \end{aligned}$$

So the dependent variable is either the raw elicitation of subjective risk q_i or the de-rounded version of that elicitation \tilde{q}_i . The independent variable is the predicted true risk, predicted by either the lasso logit or the long logit specification. The quantities of interest are $E(q_i | p_i = 0)$, $E(q_i | p_i = 1)$ and $\beta \approx \frac{\partial q_i}{\partial p_i}$, the average error at the top of the risk distribution, the average error at the bottom of the risk distribution, and the slope of the conditional mean, interpretable approximately as the partial effect on the subjective elicitation of a change in the predicted true risk. The estimates for these three quantities for each of the four specifications are in table ??? below. Standard errors are clustered by respondent throughout and only the test set is included, never the training set.

Figure 7 shows that on average I observe pessimism, with the mean of the self-reports above the

⁽⁸⁾See for example, O’Dea and Sturrock (2020)

mean of the predicted true probabilities. But table 2 indicates that the pessimism is not constant over the risk distribution. The lowest mortality risk individuals are pessimistic, overstating their mortality prospects by 13-26% depending on the specification. But the highest mortality risk individuals are optimistic, understating their mortality prospects by 17-37%. As we move from the bottom to the top of the risk distribution, an extra 1% of predicted true risk is associated with only 0.38-0.7% greater subjective elicitation. With pessimism at the bottom, optimism at the top and attenuated response in the middle, at most risk levels there is error on average, but not uniform error as is often assumed.

Naturally when the derounded subjective elicitations are used the estimates more closer to unbiasedness. This is by construction: In all cases \tilde{q}_i is closer to \hat{p}_i than is q_i . But even generously allowing for much of the observed error to be simply rounding error (over half the sample is classified as 50% rounders), the core facts hold up: pessimism at the bottom, optimism at the top, and attenuated slope in the middle. Although this correction was primarily motivated by the observed pattern of rounding, it captures any sort of elicitation noise that might have been introduced by the survey instrument and process.

3.5.1 Evidence from Wave 1

A concern with the evidence just presented might be that the rounding categories into which I have assigned individuals are incorrect. I have attempted to be conservative and err on the side that would be adversarial to my results (by assigning to coarser categories whenever in doubt), but nevertheless this concern might persist. To circumvent this concern, I offer evidence below that comes only from the responses in the very first HRS wave.

In the first wave of the HRS, the same question was asked of survival to age 75, but the responses were constrained to be an integer from 0-10 inclusive. These responses are then aligned with later HRS waves by treating the integer response of x as signifying a response of x in 10, or $10x$ percent⁽⁹⁾. So a response of 4 (out of 10) in wave one is equivalent to a response of 40% in any subsequent wave. Interpreted this way, the unit of rounding is clearer. I assume that an individual considered all possibilities and chose the appropriate response. As such, I assume the unit of rounding is to the nearest 10% for everyone.

The same prediction and inference exercise is performed as above, except now restricted to the individuals for whom a prediction and subjective elicitation is obtained in wave 1. Naturally this restricts the sample size considerably, such that there is more noise in the prediction and inference. The results, again with the four different specifications (raw and derounded elicitations, long logit and logit lasso predictions) are presented in table 3.

The patterns are primarily the same as table 2. In all specifications we observe pessimism at the bottom of the risk distribution, optimism at the top and attenuation as we move from the bottom to the top. I reject the hypothesis of constant misperception (in particular, of zero misperception).

⁽⁹⁾Study (2018)

Quantity	Lasso Logit		Long Logit	
	Raw q_i	Derounded \tilde{q}_i	Raw q_i	Derounded \tilde{q}_i
$E(q_i p_i = 0)$	0.23*** (0.01)	0.19*** (0.01)	0.26*** (0.01)	0.23*** (0.01)
β	0.45*** (0.04)	0.53*** (0.04)	0.31*** (0.04)	0.38*** (0.03)
$E(q_i p_i = 1)$	0.67*** (0.03)	0.71*** (0.03)	0.57*** (0.03)	0.61*** (0.03)
Observations (Individuals)	1,449	1,449	1,420	1,420

Table 3: Results from Wave One. *** means significant at the 1% level against 0 (first row) or 1 (second and third rows).

3.5.2 Evidence from the Panel Data

The HRS is a panel data set that has been collected for almost 30 years. Up to this point I have pooled all observations without meaningfully using the panel structure, except when clustering standard errors. In this section I offer within-person evidence using the panel data to reinforce the conclusions so far. The questions such evidence can answer include: Is the within-person time variation in subjective perceptions concordant with the within-person time variation in predicted objective risk? Is this particularly true of those who are actively changing their subjective elicitation?

The latter question speaks to the hypothesis that within the sample there are a group of people giving thoughtful answers that update over time as risk changes, and a group of people who do not wish to engage with the question as so repeat the same number every wave. Perhaps amongst the group of active changers (the 'non-inert') there is closer to unbiasedness.

Now I will introduce a time subscript t to make clear the panel structure. I model the relationship between objective risk and subjective perception as:

$$q_{it} = \alpha + \beta p_{it} + u_i + \epsilon_{it}. \quad (3.1)$$

I estimate (3.1) by time-demeaning within person, such that u_i drops out and β is identified using only within person variation. In particular, any time invariant optimism or unobservable factor drops out. It is very unlikely there is a true time invariant unobservable. Suppose an individual had some disease that was unobserved to the HRS. If at age 60 say this disease decreased the probability of survival by 10%, it is unlikely at age 62 the decreased probability of survival is still 10%, because the individual is now 2 years closer to 75. So the primary purpose of the fixed effect is to remove a hypothetical baseline optimism or pessimism, such that an individual always inflates or deflates their perception by a constant amount.

The first specification in table 4 is the fixed effect regression of (3.1) with raw elicitation q_i used. The second specification is again (3.1) with de-rounded \tilde{q}_i used. Finally, the third specification

Quantity	Raw q_i	Derounded \tilde{q}_i	Derounded \tilde{q}_i and $\Delta_t q_{it} \neq 0$
β	0.08** (0.03)	0.42*** (0.02)	
β'			0.47*** (0.07)
Observations (Ind \times waves)	4,736	4,736	1,140
Individuals	2,090	2,090	751

Table 4: Fixed Effect Results. *** means significant at the 1% level against 0 (first row) or 1 (second and third rows).

explicitly selects the sample such that $\Delta_t q_{it} \neq 0$. To emphasize that this does not estimate the same β , I denote that slope parameter by β' . Throughout I use only the long logit prediction of p_{it} to avoid any machine-learner plug-in issues.

As before similar patterns obtain. I reject the hypothesis that changes in objective risk are reacted to in an unbiased way over time. Moreover, this is not just due to inertia: even amongst those who actively change their elicitation and accounting for rounding error in periods t and $t - 1$, still a 1% change in true risk leads to less than half a percent change in elicitation.

4 Discussion and Conclusion

The robust findings presented above show that individual private perceptions bend toward the mean. The high risk are optimistic and the low risk are pessimistic. On average, although this depends on which part of the risk distribution is surveyed, there is pessimism. In this section I discuss the broad implications of considering the extra margins, risk misperceptions and contract distortion, relative to the existing literature. I will then specialize to a discussion of the what the pattern in mortality risk implies about private annuity markets in the context of retirement saving.

First, some common findings must be reconsidered in light of this evidence of misperception. A strand of the literature attributes market unravelling to very accurate and predictive private information. For example, Hendren (2013) and Hendren (2017) argue that life insurance, disability insurance, long term care insurance and unemployment insurance cannot be profitably to be provided to select groups (in the latter case the whole population) owing to the private information about risk that these groups possess. Critical to those conclusions are the assumption that individuals private perception of their risk is equal to the econometricians prediction of that risk. That is, once elicitation error is removed, private beliefs are unbiased.

I have presented evidence against this assumption. As such, it is possible that the unravelling of the, for example, the life insurance market for those with a rejectable condition is not entirely explained by too much private information. Indeed, section 3.5 shows that those with high mortality risk are unduly optimistic and the distribution of private information likely has a thinner right tail than estimated. So the no-trade condition in Hendren (2013), which says no one is willing to pay

the average cost of all those riskier than themselves, is driven less by a fat right tail of private information about risk and more by the wedge between objective risk (according to which profit is earned) and subjective perceptions (according to which insurance choices are made). This does not rule out that the subjective distribution of risk is sufficiently fat-right-tailed as to necessitate unravelling, but simply that this difference in levels between objective and subjective risk must also be considered.

Second, most explanations of insurance choice fix the contract space. A consequence of the former is that the manner in which the contract offering is distorted by informational primitives is neglected. I have shown that when those near the top of the risk distribution are not willing to buy objectively priced insurance, not only will they be sold a contract with less insurance, but all the contracts offered to those further down the risk distribution will also be distorted down.

With regards to the annuity market in particular, the pattern suggests that the high longevity risk are not willing to buy fair insurance. They are then sold annuities with less longevity insurance (e.g a guarantee period or a nominal premium return period). And so on down the risk distribution, to the point where the annuities offered to the medium risk both feature essentially no insurance. Faced with this distorted menu, it is no wonder that many choose not to participate in the annuity market. This sheds new light on the 'annuity puzzle': the well documented lack of participation is at least in part due to distorted contract offering. Incentives to the marginal participants are less important than interventions at the top of risk distribution, so that the externality on lower contracts can be removed. This channel is entirely missed when the contract space is fixed.

Finally, almost all studies of risk misperception have calculated only average misperception in the population or in some demographic subgroup. As the theory shows, the key covariate is with risk itself. This covariance determines any distortion in the contract menu, and even with a fixed contract menu determines the misperception of the marginal type. The typical pattern of mortality pessimism on average misses the mortality optimism at the top of the risk distribution. The implications for other forms of insurance and retirement financing in general are multiple but left for future work.

Similar analyses could be performed for all selection and risk markets. Of particular interest are markets where the risk is realized repeatedly, such as annual medical expenditure, in which for obvious opportunities for learning and improving the accuracy of one's subjective risk perception are present. The implications of differential errors in that setting, or indeed whether they dissipate over time, is a an interesting avenue for future research.

Bibliography

- Abaluck, J. and J. Gruber (2016a, August). Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance. *The American economic review* 106(8), 2145–2184.
- Abaluck, J. and J. Gruber (2016b, December). Improving the Quality of Choices in Health Insurance Markets. Technical report.

- Akerlof, G. A. (1970). The Market for Lemons: Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics* 84(3), 488–500.
- Allcott, H., B. B. Lockwood, and D. Taubinsky (2019, August). Regressive Sin Taxes, with an Application to the Optimal Soda Tax*. *The Quarterly Journal of Economics* 134(3), 1557–1626.
- Arya, A. and B. Mittendorf (2004, December). Benefits of a slanted view: a discussion of ‘disclosure bias’. *Journal of Accounting and Economics* 38, 251–262.
- Azevedo, E. M. and D. Gottlieb (2017). Perfect Competition in Markets With Adverse Selection. *Econometrica* 85(1), 67–105.
- Banks, J., A. Muriel, and J. Smith (2011). Attrition and health in ageing studies: evidence from elsa and hrs. *Longitudinal and Life Course Studies* 2(2), 101–126.
- Bhargava, S., G. Loewenstein, and J. Sydnor (2017, August). Choose to Lose: Health Plan Choices from a Menu with Dominated Option*. *The Quarterly Journal of Economics* 132(3), 1319–1372.
- Bisin, A. and P. Gottardi (2006, June). Efficient Competitive Equilibria with Adverse Selection. *Journal of Political Economy* 114, 485–516.
- Bénabou, R. and J. Tirole (2002, August). Self-Confidence and Personal Motivation*. *The Quarterly Journal of Economics* 117(3), 871–915.
- Chassagnon, A. and B. Villeneuve (2005). Optimal Risk-Sharing under Adverse Selection and Imperfect Risk Perception. *The Canadian Journal of Economics / Revue canadienne d’Economie* 38(3), 955–978.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, and W. Newey (2017, May). Double/Debiased/Neyman Machine Learning of Treatment Effects. *American Economic Review* 107(5), 261–265.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018, February). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* 21(1), C1–C68.
- Crocker, K. J. and A. Snow (1992). The Efficiency Effects of Categorical Discrimination in the Insurance Industry. Huebner International Series on Risk, Insurance and Economic Security, pp. 444–467. Dordrecht: Springer Netherlands.
- de la Rosa, L. E. (2011, November). Overconfidence and moral hazard. *Games and Economic Behavior* 73(2), 429–451.
- Dubey, P. and J. Geanakoplos (2002). Competitive pooling: Rothschild-stiglitz reconsidered. *The Quarterly Journal of Economics* 117(4), 1529–1570.
- Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating Welfare in Insurance Markets Using Variation in Prices. *The Quarterly Journal of Economics* 125(3), 877–921.
- Einav, L., A. Finkelstein, and P. Schrimpf (2015, April). The Response of Drug Expenditure to Non-Linear Contract Design: Evidence from Medicare Part D. *The Quarterly Journal of Economics* 130.

- Eliasz, K. and R. Spiegel (2006, July). Contracting with Diversely Naive Agents. *The Review of Economic Studies* 73(3), 689–714.
- Ericson, K. M., P. Kircher, J. Spinnewijn, and A. Starc (2020, May). Inferring Risk Perceptions and Preferences Using Choice from Insurance Menus: Theory and Evidence. *The Economic Journal* (ueaa069).
- Fang, H., M. P. Keane, and D. Silverman (2008). Sources of Advantageous Selection: Evidence from the Medigap Insurance Market. *Journal of Political Economy* 116(2), 303–350.
- Finkelstein, A. and K. McGarry (2006, September). Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market. *American Economic Review* 96(4), 938–958.
- Gabaix, X. and D. Laibson (2006, May). Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets*. *The Quarterly Journal of Economics* 121(2), 505–540.
- Gan, L., M. Hurd, and D. Mcfadden (2003, March). Individual Subjective Survival Curves. *Analyses in the Economics of Aging*.
- Handel, B., I. Hendel, and M. D. Whinston (2015). Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk. *Econometrica* 83(4), 1261–1313.
- Handel, B. R. and J. T. Kolstad (2015, August). Health Insurance for Humans: Information Frictions, Plan Choice, and Consumer Welfare. *American Economic Review* 105(8), 2449–2500.
- Handel, B. R., J. T. Kolstad, T. Minten, and J. Spinnewijn (2020, September). The Social Determinants of Choice Quality: Evidence from Health Insurance in the Netherlands. Technical report.
- Handel, B. R., J. T. Kolstad, and J. Spinnewijn (2018, November). Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets. *The Review of Economics and Statistics* 101(2), 326–340.
- Hendren, N. (2013). Private Information and Insurance Rejections. *Econometrica* 81(5), 1713–1762.
- Hendren, N. (2017, July). Knowledge of Future Job Loss and Implications for Unemployment Insurance. *American Economic Review* 107(7), 1778–1823.
- Hoy, M. (1982, May). Categorizing Risks in the Insurance Industry*. *The Quarterly Journal of Economics* 97(2), 321–336.
- Hurd, M. and K. McGarry (2002, February). The Predictive Validity of Subjective Probabilities of Survival. *Economic Journal* 112, 966–985.
- Hurd, M. D. and K. McGarry (1995). Evaluation of the subjective probabilities of survival in the health and retirement study. *The Journal of Human Resources* 30, S268–S292.
- Jeleva, M. and B. Villeneuve (2004, May). Insurance contracts with imprecise probabilities and adverse selection. *Economic Theory* 23(4), 777–794.
- Ketcham, J. D., C. Lucarelli, E. J. Miravete, and M. C. Roebuck (2012, May). Sinking, Swimming, or Learning to Swim in Medicare Part D. *American Economic Review* 102(6), 2639–2673.

- Kleinjans, K. J. and A. V. Soest (2014). Rounding, Focal Point Answers and Nonresponse to Subjective Probability Questions. *Journal of Applied Econometrics* 29(4), 567–585.
- Koszegi, B. (2014, December). Behavioral Contract Theory. *Journal of Economic Literature* 52(4), 1075–1118.
- Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association* 4(4), 673–707.
- Landais, C., A. Nekoei, P. Nilsson, D. Seim, and J. Spinnewijn (2017, October). Risk-Based Selection in Unemployment Insurance: Evidence and Implications. Technical report, Centre for Economic Performance, LSE.
- Landier, A. and D. Thesmar (2009). Financial contracting with optimistic entrepreneurs. *The Review of Financial Studies* 22(1), 117–150.
- Miyazaki, H. (1977). The rat race and internal labor markets. *The Bell Journal of Economics* 8(2), 394–418.
- Mueller, A. I., J. Spinnewijn, and G. Topa (2018, November). Job seekers’ perceptions and employment prospects: Heterogeneity, duration dependence and bias. Working Paper 25294, National Bureau of Economic Research.
- Netzer, N. and F. Scheuer (2014). A game theoretic foundation of competitive equilibria with adverse selection. *International Economic Review* 55(2), 399 – 422.
- O’Dea, C. and D. Sturrock (2020, August). Survival pessimism and the demand for annuities. Working Paper 27677, National Bureau of Economic Research.
- Riley, J. (1979). Informational Equilibrium. *Econometrica* 47(2), 331–59.
- Rothschild, M. and J. Stiglitz (1976, November). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *The Quarterly Journal of Economics* 90(4), 629–649.
- Sandroni, A. and F. Squintani (2007, December). Overconfidence, Insurance, and Paternalism. *American Economic Review* 97(5), 1994–2004.
- Sandroni, A. and F. Squintani (2013, September). Overconfidence and asymmetric information: The case of insurance. *Journal of Economic Behavior & Organization* 93, 149–165.
- Santos-Pinto, L. (2007, February). Positive self-image and incentives in organizations. Technical report, University Library of Munich, Germany.
- Shi, S. Y. (1988, January). Rothschild-Stiglitz competitive insurance market under quasilinear preferences. *Economics Letters* 27(1), 27–30.
- Singh, R. and L. Sun (2020, June). De-biased Machine Learning in Instrumental Variable Models for Treatment Effects. Technical report, arXiv.org.
- Spence, M. (1978, December). Product differentiation and performance in insurance markets. *Journal of Public Economics* 10(3), 427–447.

- Spiegler, R. (2006). Competition over agents with boundedly rational expectations. *Theoretical Economics* 1(2), 207–231.
- Spinnewijn, J. (2013). Insurance and Perceptions: How to Screen Optimists and Pessimists. *The Economic Journal* 123(569), 606–633.
- Spinnewijn, J. (2015). Unemployed But Optimistic: Optimal Insurance Design with Biased Beliefs. *Journal of the European Economic Association* 13(1), 130–167.
- Stantcheva, S. (2020, August). Understanding Tax Policy: How Do People Reason? Technical report.
- Stiglitz, J. E. (1977). Monopoly, Non-Linear Pricing and Imperfect Information: The Insurance Market. *The Review of Economic Studies* 44(3), 407–430.
- Stiglitz, J. E., J. Yun, and A. Kosenko (2018, June). Characterization, existence, and pareto optimality in insurance markets with asymmetric information with endogenous and asymmetric disclosures: Revisiting rothschild-stiglitz. Working Paper 24711, National Bureau of Economic Research.
- Study, H. . R. (2018). public use dataset. produced and distributed by the university of michigan with funding from the national institute on aging (grant number nia u01ag009740).
- Veiga, A. (2020, February). Community Rating in Markets for Lemons. SSRN Scholarly Paper ID 3545479, Rochester, NY.
- Wilson, C. (1977, December). A model of insurance markets with incomplete information. *Journal of Economic Theory* 16(2), 167–207.
- Young, V. R. and M. J. Browne (2000, December). Equilibrium in Competitive Insurance Markets Under Adverse Selection and Yaari’s Dual Theory of Risk. *The Geneva Papers on Risk and Insurance Theory* 25(2), 141–157.

A Theoretical Extensions

A.1 Monopoly

As another robustness check that the particular equilibrium concept is not driving the qualitative results, here I consider the provision of insurance by a monopolist instead of a competitive market or planner. As standard, the monopolists problem is to, writing the endowment as $\mathbf{0} = (w, w - l)$,

$$\max_{\mathbf{c}_H, \mathbf{c}_L} \sum_{i \in \{L, H\}} \alpha_i \pi(p_i, \mathbf{c}_i) \quad (\text{A.1})$$

$$\text{such that} \quad (\text{A.2})$$

$$V(q_H, \mathbf{c}_H) \geq V(q_H, \mathbf{c}_H) \quad (\text{A.3})$$

$$V(q_L, \mathbf{c}_L) \geq V(q_L, \mathbf{c}_H) \quad (\text{A.4})$$

$$V(q_H, \mathbf{c}_H) \geq V(q_H, \mathbf{0}) \quad (\text{A.5})$$

$$V(q_L, \mathbf{c}_L) \geq V(q_L, \mathbf{0}). \quad (\text{A.6})$$

The equilibrium configuration is a direct adaptation of the various parts of Jeleva and Villeneuve (2004) and is summarized below. Using their terminology: when I say one type receives optimal insurance means that the type is receiving their first best contract **fixing the level of profit** from that types contract. Or geometrically, it means that the type's indifference curve is tangent to the iso-profit line at their equilibrium contract.

Proposition 6 (Jeleva and Villeneuve (2004)). *Suppose $q_H > q_L$ and errors are small. In the unique optimum:*

1. (A.6) and (A.3) bind.
2. Types are always seperated.
3. *H* types obtain optimal insurance. If $q_H > p_H$ this is over-insurance, if $q_H < p_H$ then under insurance.
4. *L* types are sub-optimal partial insurance, perhaps none.
5. There is a threshold $\bar{\gamma}_H$ such that low types receive no insurance if $\gamma_H > \bar{\gamma}_H$ and are served if not.

The case where the low type is not served is not particularly interesting. In such a case low type errors clearly have no impact. The high types participation constraint binds according to their perceived probability, so a pessimistic error is bad and an optimistic error good according to their objective welfare as calculated by the planner.

The more interesting dynamics occur when both types are served. The monopolist faces a trade-off: They can extract more profit from the low types by offering them more and more insurance but this requires less (perhaps negative) profits to be earned from the high types due to the binding incentive constraint. Conversely, they can extract more profits from the high types by raising the price of optimal insurance and pushing the high types closer to their outside option, but for the same incentive reason the low types contract must then feature even less insurance and profit (perhaps zero of both) so as to not tempt the high type. These forces will help with interpreting the welfare implications of errors.

First, consider a small upward misperception by the high type with no error by the low type. The low type will continue to receive utility equal to their outside option and so their welfare will not change. The high types incentive constraint has now loosened. By how much it has loosened depends on where their risk was to begin with. When they were already high risk, the loosening is greater (at the limit, consider flat indifference curves) and so per dollar of extra profit to be extracted from the low type the incentive cost to the high type is lower. This motivates the (sufficient) condition $2p_H > 1 + p_L$. The higher risk the high type is, the higher the maginal loosening in incentive cost is when they make a small error, and so the temptation to earn more profits from the low types whilst making the high types slightly better off is stronger.

On the other hand, suppose the low type makes a small error. The effect on their welfare is unambiguous. If the low types think they are riskier than they are, the monopolist will sell them insurance they don't need, still restrict them to zero utility - but according to their perceived pessimistic probability - which the planner will then determine as detrimental to their welfare as defined by the objective probability. On the other hand, if the low types are optimistic, their zero perceived utility earned will actually be an objective improvement in the planners eyes.

But the effect of a low type error on a high type is more subtle. To charge more for insurance to the low types (their IR curve has shifted left), the high types must be made better off. If the high

types are very different from the low types in risk, then the insurer can charge the low types more for the same loss payout whilst not affecting their profit from the high types much at all (at the limit, again, consider high types having horizontal indifference curves). If the high types are closer in risk to the low types, then extracting profit from the low type will require a greater marginal loss from the high types. Hence, the (sufficient) condition $\frac{p_H(1-p_L)}{p_L(1-p_H)} < \frac{u'(w)}{u'(w-l)}$ says that when p_H is small and p_L large, i.e. their difference is small, the low types naivete can be exploited at lower change to high types profit.

Proposition 7. *For a small error, the effects on welfare for the high and low types are:*

For an error by the high type $\frac{\partial V_H(p_H, \mathbf{c}^)}{\partial q_H} |_{\mathbf{q}=\mathbf{p}} > 0$ if $2p_H > 1 + p_L$, $\frac{\partial V_L(p_L, \mathbf{c}^*)}{\partial q_H} |_{\mathbf{q}=\mathbf{p}} = 0$. For an error by the low type $\frac{\partial V_H(p_H, \mathbf{c}^*)}{\partial q_L} |_{\mathbf{q}=\mathbf{p}} > 0$ if $\frac{p_H(1-p_L)}{p_L(1-p_H)} < \frac{u'(w)}{u'(w-l)}$, $\frac{\partial V_L(p_L, \mathbf{c}^*)}{\partial q_L} |_{\mathbf{q}=\mathbf{p}} < 0$.*

As compared to the competitive market dynamics and the planners optimal choice problem, the key difference here is that errors by both types matter for welfare. But the main qualitative insight holds: errors that make types more disparate (which, loosely speaking, might 'increase' the amount of private information) can actually be good for welfare by weakening the incentive constraints. When there is no endogenous contract adjustment, this channel is closed down and the welfare conclusions might be changed.

A.2 Heterogeneity in misperceptions within type classes

Amongst the simplifications in the basic model in section 2 is that everyone in a given risk class makes the same misperception. Clearly this is unrealistic: For any given level of true risk we would expect some people to be optimistic, some to be pessimistic, and others accurate. In this section we demonstrate how the equilibrium notion generalizes in when there is heterogeneity of misperception with a risk class. The key forces that determine the equilibrium, and hence that determine the welfare consequences of the misperceptions, remain qualitatively unchanged. Throughout we focus on the case where there are two risk classes, and within each 2 different classes of misperception are made. The analysis naturally extends to any finite number of risk classes and within those any finite number of perceived risk classes, as will be clear.

Amongst those with true risk p_H , we define $q_{H,O}$ and $q_{H,P}$ as two different perceived risk classes: The former being relative optimists, the latter relative pessimists, such that $q_{H,O} < q_{H,P}$. Similarly amongst those with true risk p_L we analogously define $q_{L,O}$ and $q_{L,P}$. I write $\alpha_{H,O}, \alpha_{H,P}, \alpha_{L,O}, \alpha_{L,P}$ for the proportion of the population that falls into each of these classes. In this section I study the case in which

$$q_{H,P} > q_{H,O} > q_{L,P} > q_{L,O}. \quad (\text{A.7})$$

This leads to qualitatively similar outcomes as in the core model. In the following subsection A.3, we study the case in which $q_{H,O} < q_{L,P}$ and find this to be a novel mechanism for pooling in equilibrium - a rare theoretical foundation for a ubiquitous observed phenomenon.

The equilibrium that obtains if we assume (A.7) is illustrated in figure 9. It is formalized in the following proposition.

Proposition 8. *Suppose $\|\mathbf{p} - \mathbf{q}\|_\infty < \xi$, and A.7 holds. For small enough ξ , in the unique locally-competitive equilibrium, high risk pessimistic individuals buy the contract $\mathbf{c}_{H,P}$ that solves:*

$$MRS(q_{H,P}, \mathbf{c}_{H,P}) = \frac{1 - p_H}{p_H} \text{ and } \pi(p_H, \mathbf{c}_{H,P}) = 0,$$

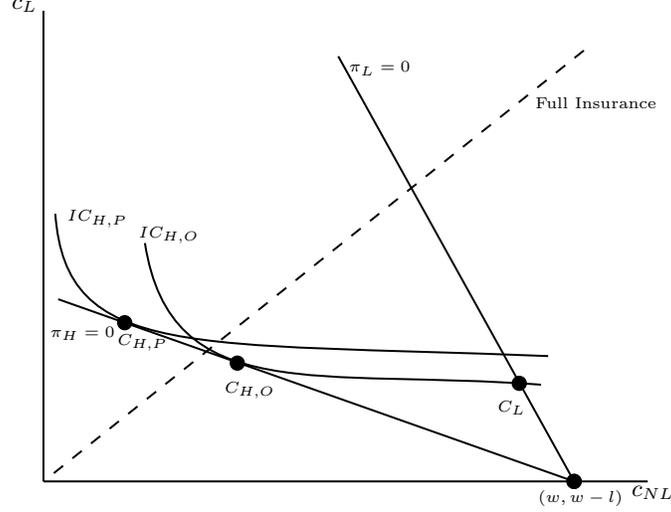


Figure 9: Heterogenous perceptions within risk class, seperating equilibrium

high risk optimistic individuals buy the contract $c_{H,O}$ that solves:

$$MRS(q_{H,O}, c_{H,O}) = \frac{1 - p_H}{p_H} \text{ and } \pi(p_H, c_{H,O}) = 0,$$

and low risk individuals receive the contract c_L that solves

$$V(q_{H,O}, c_{H,O}) = V(q_{H,O}, c_L) \text{ and } \pi(p_L, c_L) = 0.$$

The high risk types, both optimistic and pessimistic, obtain their perceived first best contract subject to zero profit being earned. Naturally, $c_{H,P}$ features more insurance than $c_{H,O}$ as the pessimists think the risk is more likely then the optimists and so have a greater demand for insurance.

The low risk types all receive the same insurance, and as in 1 this equilibrium contract features partial insurance distorted away from the first best owing to the information rents that accrue to the high types. As I noted above, misperceptions by the low-types have no effect on the equilibrium menu, and hence heterogeneity in misperceptions by the low types also have no effect.

The equilibrium menu is separating in the sense that all true high types are separated from all true low types, as opposed to the following subsection. The incentive constraint that binds to determine the low types contract is that of the most optimistic high type. But this is optimistic only relative to the pessimistic high types, not necessarily relative to the truth. That is, if $p_H > q_{H,O}$ (as is illustrated in figure 9) in which case the misperception will certainly create a welfare loss in the market in the same way as proposition 2. Or, if $p_H < q_{H,O} < q_{H,P}$ then if $q_{H,O} - p_H$ is small enough then the misperceptions can lead to a welfare improvement as also discussed in proposition 2.

Broadly, when the heterogeneity within risk classes is ordered as in (A.7), all the dynamics apply as in the core model and proofs therein, with the slight caveat that the relevant incentive compatibility constraint for determining a lower types contract is that of the most optimistic amongst the higher types.

A.3 Partial pooling in the Rothschild and Stiglitz (1976) framework

The assumption (A.7) states that despite heterogeneity within risk classes, the most pessimistic low type still thinks they are lower risk than the most optimistic high type. As the empirical analysis speaks to, this is likely unrealistic, with each risk class having a large amount of heterogeneity in perceived risk, such that the most pessimistic low types perceive their risk to be higher than the most optimistic high types. In this section I study the effect of such an ordering. Specifically, I assume:

$$q_{H,P} > q_{L,P} > q_{H,O} > q_{L,O}. \quad (\text{A.8})$$

I will find that under a mild assumption, formalized below, that in equilibrium those with perceived risk $q_{L,P}$ are pooled together with those with perceived risk $q_{H,O}$. Define the average risk if those with perceived risk $q_{L,P}$ and $q_{H,O}$ are pooled together:

$$p_{HL} = \frac{\alpha_{L,P}p_L + \alpha_{H,O}p_H}{\alpha_{L,P} + \alpha_{H,O}}.$$

The partial-pooling equilibrium that obtains is:

Proposition 9. *Suppose A.8 holds, $\alpha_{H,P} \geq \alpha_{H,O}$ and $MRS(q_{L,P}, \mathbf{c}_{HL}) < \frac{1-p_{HL}}{p_{HL}}$.⁽¹⁰⁾ In the unique locally-competitive equilibrium, high risk pessimistic individuals buy the contract $\mathbf{c}_{H,P}$ that solves:*

$$MRS(q_{H,P}, \mathbf{c}_{H,P}) = \frac{1-p_H}{p_H} \text{ and } \pi(p_H, \mathbf{c}_{H,P}) = 0,$$

*high risk optimistic **and** low risk pessimistic individuals buy the contract \mathbf{c}_{HL} that solves:*

$$V(q_{H,P}, \mathbf{c}_{H,P}) = V(q_{H,P}, \mathbf{c}_{HL}) \text{ and } \pi(p_{HL}, \mathbf{c}_{HL}) = 0,$$

and low risk optimistic individuals receive the contract $\mathbf{c}_{L,O}$ that solves

$$V(q_{H,O}, \mathbf{c}_{HL}) = V(q_{H,O}, \mathbf{c}_{L,O}) \text{ and } \pi(p_L, \mathbf{c}_{L,O}) = 0.$$

This equilibrium is illustrated in Figure 10. The pessimistic high types still receive their perceived first-best contract subject to zero profits being made. Then, most notably, the pessimistic low types and the optimistic high types are pooled together. The contract they are pooled into is defined so that the pessimistic high types are indifferent between in and their contract, as standard.

It is worth considering why this specific heterogeneity (equation (A.8) in misperception allows for the fundamental no-pooling result in Rothschild and Stiglitz (1976) to not hold here. Typically, any pooling contract is not an equilibrium because a deviating contract offering slightly less insurance, attracting only the lower risk type would make a profit. Here due to the fact that $q_{H,O} < q_{L,P}$, a contract offering less insurance is more attractive to the *higher* risk type, due to their large misperception. So you cannot 'cream-skim' in the normal way by offering less insurance than the pooling contract. Here, inversely, to separate the two pooled types and only attract the lower (true) type, a contract with more insurance than the pooled contract needs to be offered. This is because the lower of the two true types being pooled has a higher perceived type, per (A.8).

⁽¹⁰⁾This assumption rules out a deviation to a contract that offers less insurance, is preferred by both pooled types (but no one else) and makes a positive profit. Alternatively we could use the Riley notion since such a deviation is not 'safe'. This is a minor assumption unrelated to the substance of these two different true risk classes being pooled together, as the discussion makes clear.

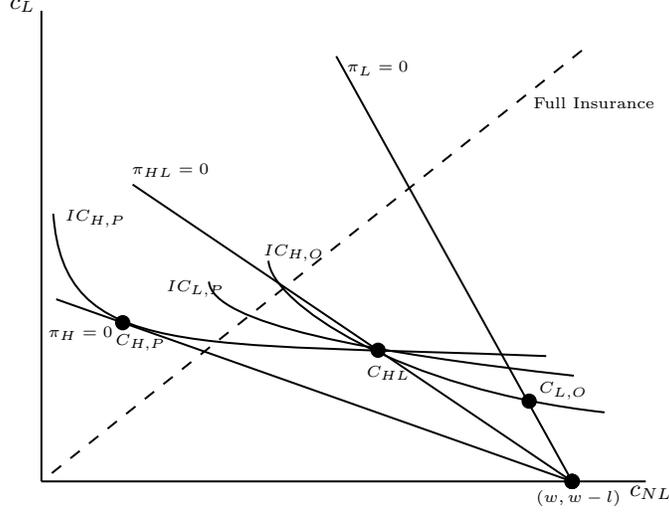


Figure 10: Heterogenous perceptions within risk class, pooling equilibrium

So skimming the lower type out of the pool can be done with a contract offering more insurance. , But such a deviation will also attract the pessimistic high types, who are even higher perceived risk than the optimistic low type being skimmed out. The assumption that $\alpha_{H,P} > \alpha_{H,O}$ means that this cream skimming contract replaces the optimistic high types with a larger mass of pessimistic high types, thereby making a loss.

This result resolves the empirically unsatisfactory no-pooling prediction of the Rothschild-Stiglitz model. It shows that misperceptions of risk that cause 'overlap' of perceived risk types as in (A.8) can generate equilibrium pooling. Alternative mechanisms that generate pooling in this framework include: departing from the expected utility framework as in Shi (1988); quantity caps as in Dubey and Geanakoplos (2002); or endogenous sharing of information by firms that allows for partial non-exclusivity of contracts (e.g. Stiglitz et al. (2018)).

Even in this differently structured equilibrium, the qualitative insights of the core model hold true. Misperceptions made by the higher types still cascade down and exert an externality on the equilibrium contract offered to lower types. The welfare consequences are still ambiguous, with the pooled contract being worse for the pessimistic low type and better for the optimistic high type than the no-misperception equilibrium.

B Proofs

B.1 Proof of Proposition 1

Proof. TO DO: check notation lines up. Remove all references to medium types. Add in explicit condition that defines bound on ξ .

Step 1. No pooling

First, note an intuitive but repeatedly used fact: for types i, j if $q_i > q_j$ then the equilibrium contract that i buys (and perhaps others) must offer weakly more insurance than the equilibrium contract j buys: $\mathbf{c}_i \succeq \mathbf{c}_j$. If i and j pool this is trivial. If not, then either the contracts are not comparable with respect to \succeq , in which case one dominates the other and an IC constraint must break, or $\mathbf{c}_j \succ \mathbf{c}_i$. In that case, we must have $V(q_j, \mathbf{c}_j) \geq V(q_j, \mathbf{c}_i)$, but also that $V(q_i, \mathbf{c}_j) > V(q_i, \mathbf{c}_i)$ since $MRS(q_i, \mathbf{c}_j) < MRS(q_j, \mathbf{c}_j)$ and $\mathbf{c}_j \succ \mathbf{c}_i$. This contradicts the i types IC constraint. Hence

higher types must receive weakly more insurance than lower types in equilibrium. Denote this observation by (*)

The impossibility of pooling between H and M types, and between M and L types, and between all three types, in a RS equilibrium, and hence in a locally-competitive equilibrium, follows by almost identical arguments to R-S. We need only check that the deviating contract that attracts the relatively lower type in each situation does not attract the currently un-pooled type, if relevant. For the latter two cases this is obvious. For the first, where we conjecture one contract that pools and one that attracts only L , $\mathbf{c}_{HM}, \mathbf{c}_L$, note simply that the low types IC constraint cannot bind. If it did, then the medium types IC constraint wouldn't bind (as the indifference curves are distinct), and so a deviation to offering $\mathbf{c}_L + (-\epsilon, \epsilon r)$ with $r \in (MRS(q_L, \mathbf{c}_L), \frac{1-p_L}{p_L})$ will by construction improve utility, earn positive profits, yet not attract any M or H types for small enough $\epsilon > 0$. So the low types IC constraint does not bind, so a contract marginally more attractive to the medium type will not attract the low type, and hence the pooling between high and medium types cannot be sustained. ⁽¹¹⁾

Finally, H and S cannot pool in any locally competitive equilibrium. If they did, then by the observation (*) above the contracts must satisfy $\mathbf{c}_{HS} \succ \mathbf{c}_M$ or the opposite. In the former case, we contradict (*) with S types being more insured than M types. In the latter, we contradict the same (*) with M types being more insured than H types. This establishes that pooling cannot be sustained in any configuration.

Step 2. Any allocation except the solution to the maximization problem in the Proposition cannot be a local equilibrium

Step 1 shows that an equilibrium menu must feature three separating contracts with $\mathbf{c}_H \succ \mathbf{c}_M \succ \mathbf{c}_L$.

We first verify that the incentive constraints cannot bind upwards. Assume that \mathbf{c} underinsures.⁽¹²⁾ For a contradiction suppose $V(q_M, \mathbf{c}_M) = V(q_M, \mathbf{c}_H)$. This implies $V(q_H, \mathbf{c}_M) < V(q_H, \mathbf{c}_H)$. Hence, for small enough $\epsilon > 0$, there is a new contract $\mathbf{c}'_M = \mathbf{c}_M + (-\epsilon, \epsilon r)$ with

$r \in \left(MRS(q_M, \mathbf{c}_M), \min \left\{ \frac{1-q_M}{q_M}, MRS(q_L, \mathbf{c}_L) \right\} \right)$ that strictly improves on the utility of the medium types, yet makes positive profits as it doesn't attract any other types. Hence the IC constraint preventing the medium type from reporting themselves as high doesn't bind in any equilibrium.

Similarly, we cannot have $V(q_L, \mathbf{c}_L) = V(q_L, \mathbf{c}_M)$. Assume that \mathbf{c}_L features underinsurance.⁽¹³⁾ If we did have $V(q_L, \mathbf{c}_L) = V(q_L, \mathbf{c}_M)$, then we must have $V(q_M, \mathbf{c}_L) < V(q_M, \mathbf{c}_M)$. This also implies that $V(q_H, \mathbf{c}_L) < V(q_H, \mathbf{c}_M)$ as the MRS curves are always shallower for H types. Hence with $r \in \left(MRS(q_L, \mathbf{c}_L), \frac{1-q_L}{p_L} \right)$, then $\mathbf{c}_L + (-\epsilon, \epsilon r)$ will make low types strictly better off, earning strictly positive profits by just attracting low types so long as $\epsilon > 0$ is small enough.

Then, suppose the H types received any contract except the specific contract whilst maintaining zero profits. Then we must have $MRS(q_H, \mathbf{c}_H) \neq \frac{1-p_H}{p_H}$. In either case there is a contract slightly north-west or south-east of the current contract (depending on whether the inequality is $>$ or $<$,

⁽¹¹⁾Note an $r \in (MRS(q_L, \mathbf{c}_L), \frac{1-p_L}{p_L})$ exists when q_L receives less than full insurance. If L has full insurance, their IC constraint cannot bind as they are at their first-best. If they are strictly over-insured then they can move toward full insurance with a positive profit made.

⁽¹²⁾If not, then since $MRS(q_M, \mathbf{c}_M) > \frac{1-p_M}{p_M}$ for overinsured contracts, the M types indifference curve is always above $(1-p_M)(W - c_M^{NL}) + p_M(c_M^L - (W-l)) = 0$ for any contract that features more insurance than \mathbf{c}_M , and hence cannot intersect with $(1-p_H)(W - c_H^{NL}) + p_H(c_H^L - (W-l)) = 0$ at such contract, and hence the IC constraint for the medium type cannot bind.

⁽¹³⁾If not, then the logic of the footnote above applies

that offers positive profits and strictly higher utility to the H types without attracting any lower types when $\epsilon > 0$ is small since their IC constraints do not bind upwards.

Then, suppose that $V(q_H, \mathbf{c}_H) > V(q_H, \mathbf{c}_M)$. Then the deviation from we considered M 's IC constraint (pretending to be H) is a profitable, utility improving deviation here as well (assuming underinsurance at \mathbf{c}_M)⁽¹⁴⁾. Similarly, suppose that $V(q_M, \mathbf{c}_M) > V(q_M, \mathbf{c}_L)$. Then the deviation from when we considered L s IC constraint suffices here as well. Note that that deviation relies on L being underinsured, which must be the case since they can always be offered less insurance without affecting the IC constraints.

This shows that the given menu is the only possible equilibrium configuration.

Step 3. Existence

It remains to verify that the proposed menu is indeed a locally competitive equilibrium.

Consider first \mathbf{c}_H . At \mathbf{c}_H we have $MRS(q_H, \mathbf{c}_H) = \frac{1-p_H}{p_H}$. Further, since $U'' < 0$ these curves are tangent. This directly means that any contract attracting only H types cannot both make a non-negative profit and also strictly improve upon H 's utility. As above, the IC constraints of the lower types do not bind upward, and so no local deviation will attract anyone but the H types. It follows there are no local deviations near \mathbf{c}_H .

Consider next \mathbf{c}_M . Recall that $MRS(q_M, \mathbf{c}_M) < \frac{1-p_M}{p_M}$ for small enough ξ . Consider a deviation in the direction $(-1, r)$ from \mathbf{c}_M with $r \in \left[MRS(q_M, \mathbf{c}_M), \frac{1-p_M}{p_M} \right]$. Since $MRS(q_M, \mathbf{c}_M) > MRS(q_M, \mathbf{c}_M)$ this contract attracts high and medium types, but for small ϵ this will be loss making since $p_{HM} > p_M$ and so the pooled contract will make profits arbitrarily close to $(1 - p_{HM})(W - c_M^{NL}) + p_{HM}(c_M^L - (W - l)) < (1 - p_M)(W - c_M^{NL}) + p_M(c_M^L - (W - l)) = 0$. Similarly, any contract deviations with $r \in [MRS(q_M, \mathbf{c}_M), MRS(q_M, \mathbf{c}_M)]$ will attract only high types and also earn negative profits. With $r > MRS(q_M, \mathbf{c}_M)$ no one will demand the contract.

Then if we consider deviations of the form $(1, r)$ from \mathbf{c}_M , if $r > -\frac{1-p_M}{p_M}$ then the contract will make negative profit if it attracts medium types, and hence also if it attracts medium and high or just high types. If $r < -\frac{1-p_M}{p_M}$ in particular $r < -MRS(q_M, \mathbf{c}_M) < MRS(q_H, \mathbf{c}_H)$ for small ξ and so no medium or high types will swap to it, and for small enough deviations neither will low types. This shows that there are no local deviations from \mathbf{c}_M . Similar working rules out local deviations from \mathbf{c}_L .⁽¹⁵⁾

From this we can conclude that the proposed menu is a locally competitive equilibrium, and from step 2 that there are no others. This concludes the proof. ■

B.2 Proof of Proposition 2

Proof. First we suppose that the error made is a small upward error by the high types. Proposition 2 shows that the high types contract can be found as the solution to the constrained maximization problem:

$$\hat{V}_H(\mathbf{q}) \equiv \max_{\mathbf{c}} V(q_H, \mathbf{c}) \quad \text{s.t. } \pi_H(p_h, \mathbf{c}) = 0.$$

By the envelope theorem, for the multiplier $\lambda > 0$ we have

⁽¹⁴⁾If not, and \mathbf{c}_M is overinsured, then and since L s IC constraint does not bind, M can be moved toward full insurance slightly at a profit and utility gain since $MRS(q_M, \mathbf{c}_M) > \frac{1-p_M}{p_M}$ for overinsured \mathbf{c}_M

⁽¹⁵⁾This all assumes that \mathbf{c}_M features underinsurance, which is true for small enough ξ .

$$\frac{\partial \hat{V}}{\partial q_H} = \frac{\partial V}{\partial q_H} + \lambda \frac{\partial g}{\partial q_H} = \frac{\partial V}{\partial q_H} = u(c_H^L) - u(c_H^{NL}).$$

In particular, at $q_H = p_H$ we have $c_H^L = c_H^{NL}$ and hence :

$$\frac{\partial \hat{V}}{\partial \xi} \Big|_{q_H=p_H} = \frac{\partial \hat{V}}{\partial \xi} \Big|_{q_H=p_H} = 0.$$

For low types, we can write their equilibrium contract, for a small error, as the solution to the (degenerate) maximization program:

$$\begin{aligned} \hat{V}_L(\mathbf{q}) \equiv \max_{\mathbf{c}} V(q_L, \mathbf{c}) \quad \text{s.t } \pi_L(p_L, \mathbf{c}) = 0 \\ \text{and } h_L \equiv V(q_H, \mathbf{c}_H) - V(q_H, \mathbf{c}) = 0. \end{aligned}$$

This is degenerate in the sense that the constraints exactly pin down the solution, and hence the maximization is superfluous. Nevertheless, with multiplier $\lambda > 0$ on the IC constraint we immediately get that:

$$\frac{\partial \hat{V}_L}{\partial q_H} = u(c_H^L) - u(c_H^{NL}) - (u(c_L^L) - u(c_L^{NL})).$$

At $q_H = p_H$ we have $c_H^L = c_H^{NL}$ and $c_L^L < c_L^{NL}$ and so

$$\frac{\partial \hat{V}_L}{\partial q_H} \Big|_{q_H=p_H} = \frac{\partial \hat{V}_L}{\partial \xi} \Big|_{q_H=p_H} > 0.$$

It follows that for a small change ξ , we have

$$\frac{\partial W}{\partial \xi} > 0 \text{ when } \xi > 0, \text{ and } \frac{\partial W}{\partial \xi} < 0 \text{ when } \xi < 0,$$

which completes the proof. ■

B.3 Proof of Proposition 3

Proof. First I show that each of the constraints must hold else the allocation cannot be an MW equilibrium.

First the budget constraint (2.7). If negative profit is earned the contract will be withdrawn. If positive profit is made and the equilibrium is pooling, then consumption in each state of the world can be increased, and everyone will swap. If the equilibrium is separating, then increase consumption for both contracts to make everyone better off whilst keeping the IC constraint binding. The old contracts will be withdrawn, the new contracts will be demanded by the appropriate types. Hence (2.7) must hold in the MW equilibrium.

Next, suppose that (??) does not hold. Write the level of profit earned from each type as $\pi_H = \pi(p_H, \mathbf{c}_H), \pi_L = \pi(p_L, \mathbf{c}_L)$. First, suppose IC_H binds and the equilibrium is separating. This means IC_L does not bind. By assumption (2.8), this means that \mathbf{c}_H is down and left of the tangency point of iso-profit line for π_H with the appropriate H indifference curve, Then suppose a firm deviated to offer the contract $\mathbf{c}' = \mathbf{c} - (\epsilon, \frac{\epsilon(1-p_H)}{p_H})$. The high types are better off, as the slope of their utility function is monotone along an iso-profit line with a maximum at the point

of tangency, whilst the low types still buy their old contract, and total profit has not changed by construction. So no contracts are withdrawn, and this is a valid deviation. So (??) must hold. Suppose the equilibrium is pooling, then the deviation to each types optimal contract subject to the pooling profit iso-profit line still breaks even and makes the high type strictly better off, and the low type weakly better off. All types change, so no contracts are withdrawn, so this is a valid deviation. Hence (??) must hold here by contradiction.

Suppose IC_L binds and the equilibrium is separating. As c_H is down and right of the optimal insurance point on the π_H isoprofit line, then $MRS(q_H, c_H) < \frac{1-p_H}{p_H}$ and also $MRS(q_H, c_H) < MRS(q_L, c_H)$ and so there exists a slope $\psi \in ([MRS(q_H, c_H), \frac{1-p_H}{p_H}) \cap (MRS(q_H, c_H) < MRS(q_L, c_H))$ such that the contract $c + \epsilon(-1, \xi)$ is preferred by H types to c_H , not preferred by low types to c_L and costs less. This deviation shows the initial menu was not an MW equilibrium and hence (??) must hold.

Next, suppose that (2.6) does not hold, i.e. $MRS(q_H, c_H) < V(q_H, c_H^{RS})$. If a firm then offers a contract arbitrarily close (but earning positive profit) to c_H^{RS} all the high types will swap to it, and even if the low types prefer (perhaps after other contracts are withdrawn) it then profit will be even higher. So this is a valid deviation and hence (2.6) must hold.

Finally, suppose that (2.4) does not hold. Then as in RS the low types can be made better off with positive profit without inducing the high types to swap. So it must be that (2.4).

Now I show that the maximization has a unique solution. After which I will show that the MW equilibrium is exactly that unique solution. The three binding constraints mean that this is reducible to a one-dimensional problem. In particular, we can index the problem by the cross-subsidy χ . The constraint (2.6) implies $\chi \geq 0$. Moreover, an upper bound on feasible χ is the maximum that which induces pooling: $\chi_1 = (p_H - P_L)l$ since for any higher χ the incentive constraint for the high type cannot possibly bind, or that which hits the indemnity constraint for the low type (??), χ_2 . Write $\bar{\chi} = \max\{\chi_1, \chi_2\}$. Moreover, a choice of $\chi \in [0, \bar{\chi}]$ uniquely defines c_H and c_L . To see this: once a χ is chosen, the iso-profit for each type is pinned down. The high types iso-profit, together with (2.8), uniquely pins down c_H as the MRS is monotonic along an iso-profit line. Then (2.4) and the low types isoprofit line intersect at two places. One of them is up and left of c_H and violates (??). Hence only the intersection that offers the low type less insurance remains. So c_L is pinned down.

Given all this, the problem becomes:

$$\max_{\chi \in [0, \bar{\chi}]} V(q_L, c_L(\chi)) \tag{B.1}$$

so that all the constraints hold with c_H, c_L implicit functions of χ . The existence of a solution follows by Weierstrass' theorem. Uniqueness follows directly from the proof of uniqueness in Netzer and Scheuer (2014), the fact that errors are small and so the solutions to this problem are arbitrarily close to the problem without errors, and continuity of the second derivative. In particular, they showed that the second derivative is globally concave, which then holds here.

This shows that any MW equilibrium must satisfy the given constraints, and that a unique solution to the optimization problem exists. It remains to show that only the amongst the allocations that satisfy the constraints, only that which maximizes the objective is an MW equilibrium. Suppose not. Then suppose a firm does offer the menu that satisfies the constraints and maximizes the objective. By definition the low types will prefer to switch to the new contract. That contract makes weakly positive profit, by constraint (2.6). The high types will either remain in their old contract or swap to the newly offered contract for them. In the former case the deviating firm will make zero profits, in the latter they will make positive profits. No withdrawal of contracts will

make the low types switch as their subjective utility is being maximized. This shows there is a MW deviation, and so the supposition was false. Hence the objective must be maximized. This concludes the proof. ■

B.4 Proof of Proposition 4

Proof. The argument above establishes that to sign $\frac{\partial}{\partial q_H} \chi^*(q_H)$ when without errors it is the case that $\chi^* > 0$ the comparative statics with respect to q_H can be computed in the following way. (The result for q_L is entirely similar and omitted for brevity.)

The following binding constraints (and a definition) implicitly define $c_H^{NL} = c_H^{NL}(\chi)$ and similarly for c_H^L, c_L^L, c_L^{NL} . :

$$\frac{u'(\chi)}{u'(c_H^L(\chi))} = \frac{(1 - p_H)q_H}{p_H(1 - q_H)} \quad (\text{B.2})$$

$$q_H u(c_H^L(\chi)) + (1 - q_H)u(c_H^{NL}(\chi)) = q_H u(c_L^L(\chi)) + (1 - q_H)u(c_L^{NL}(\chi)) \quad (\text{B.3})$$

$$w - l(p_H + p_L) = (1 - \alpha_H)p_L c_L^L(\chi) + (1 - \alpha_H)(1 - p_L)c_L^{NL}(\chi) + \alpha_H(l(w - p_H) + \chi) \quad (\text{B.4})$$

$$\chi = p_H c_H^L(\chi) + (1 - p_H)c_H^{NL}(\chi) - (w - lp_H). \quad (\text{B.5})$$

Differentiating each with respect to χ and solving yields

$$(c_H^{NL})'(\chi) = \frac{u''(c_H^L(\chi))u'(c_H^{NL}(\chi))}{p_H u'(c_H^L(\chi))u''(c_H^{NL}(\chi)) - (p_H - 1)u''(c_H^L(\chi))u'(c_H^{NL}(\chi))} \quad (\text{B.6})$$

$$(c_H^L)'(\chi) = \frac{u'(c_H^L(\chi))u''(c_H^{NL}(\chi))}{p_H u'(c_H^L(\chi))u''(c_H^{NL}(\chi)) - (p_H - 1)u''(c_H^L(\chi))u'(c_H^{NL}(\chi))} \quad (\text{B.7})$$

$$(c_L^I)'(\chi) = \left[u'(c_H^L(\chi))u''(c_H^{NL}(\chi)) (q_H(\alpha_H(-p_L) + \alpha_H + p_L - 1)u'(c_H^L(\chi)) + \alpha_H p_H(q_H - 1)u'(c_L^{NL}(\chi))) \right. \\ \left. + \alpha_H(p_H(-q_H) + p_H + q_H - 1)u''(c_H^L(\chi))u'(c_H^{NL}(\chi))u'(c_L^{NL}(\chi)) \right. \\ \left. + (\alpha_H - 1)(p_L - 1)(q_H - 1)u''(c_H^L(\chi))u'(c_H^{NL}(\chi))^2 \right] \quad (\text{B.8})$$

$$+ \alpha_H(p_H(-q_H) + p_H + q_H - 1)u''(c_H^L(\chi))u'(c_H^{NL}(\chi))u'(c_L^{NL}(\chi)) \quad (\text{B.9})$$

$$+ (\alpha_H - 1)(p_L - 1)(q_H - 1)u''(c_H^L(\chi))u'(c_H^{NL}(\chi))^2 \quad (\text{B.10})$$

$$/ \left[(\alpha_H - 1) ((p_H - 1)u''(c_H^L(\chi))u'(c_H^{NL}(\chi)) - p_H u'(c_H^L(\chi))u''(c_H^{NL}(\chi))) \right. \\ \left. \times ((p_L - 1)q_H u'(c_L^I(\chi)) - p_L(q_H - 1)u'(c_L^{NL}(\chi))) \right] \quad (\text{B.11})$$

$$\times ((p_L - 1)q_H u'(c_L^I(\chi)) - p_L(q_H - 1)u'(c_L^{NL}(\chi))) \quad (\text{B.12})$$

$$(c_L^{NL})'(\chi) = \left[u''(c_H^L(\chi))u'(c_H^{NL}(\chi)) ((\alpha_H - 1)p_L(q_H - 1)u'(c_H^{NL}(\chi)) - \alpha_H(p_H - 1)q_H u'(c_L^I(\chi))) \right. \\ \left. + q_H u'(c_H^L(\chi))u''(c_H^{NL}(\chi)) ((p_L - \alpha_H p_L)u'(c_H^L(\chi)) + \alpha_H p_H u'(c_L^I(\chi))) \right] \quad (\text{B.13})$$

$$+ q_H u'(c_H^L(\chi))u''(c_H^{NL}(\chi)) ((p_L - \alpha_H p_L)u'(c_H^L(\chi)) + \alpha_H p_H u'(c_L^I(\chi))) \quad (\text{B.14})$$

$$= / \left[(\alpha_H - 1) ((p_H - 1)u''(c_H^L(\chi))u'(c_H^{NL}(\chi)) - p_H u'(c_H^L(\chi))u''(c_H^{NL}(\chi))) \right. \\ \left. \times ((p_L - 1)q_H u'(c_L^I(\chi)) - p_L(q_H - 1)u'(c_L^{NL}(\chi))) \right]. \quad (\text{B.15})$$

$$\times ((p_L - 1)q_H u'(c_L^I(\chi)) - p_L(q_H - 1)u'(c_L^{NL}(\chi))). \quad (\text{B.16})$$

Being overly explicit, the objective can be written:

$$V_L(q_L, \chi) = q_L u(c_L^L(\chi(q_H), q_H)) + (1 - q_L)u(c_L^{NL}(\chi(q_H), q_H))$$

By the implicit function theorem, the sign of $\chi'(q_H)$ is the same as the sign of $\partial^2 V_L(q_L, \chi) / \partial \chi \partial q_L$. So computing:

$$\frac{\partial^2 V_L(q_L, \chi)}{\partial \chi \partial q_H} = q_L \frac{\partial c_L^L \chi}{\partial q_H} u'(c_L^L(\chi(q_H), q_H)) + (1 - q_L) \frac{\partial c_L^{NL}(\chi(q_H), q_H)}{\partial q_H} u'(c_L^{NL}(\chi(q_H), q_H)) \quad (\text{B.17})$$

$$+ q_L (c_L^L)'(q_H) (c_L^L)'(\chi) u''(c_L^L(\chi(q_H), q_H)) + (1 - q_L) (c_L^{NL})'(q_H) (c_L^{NL})'(\chi) u''(c_L^{NL}(\chi(q_H), q_H)) \quad (\text{B.18})$$

where $(c_L^{NL})'(q_H)$ is shorthand for the partial derivative with respect to q_H when χ is fixed (i.e. only the explicit dependence is taken into account, not the implicit dependence through $\chi(q_H)$). Such quantities can be found by considering each consumption quantity to be a function of q_H , not χ in (B.2) - (B.5), differentiating them all and solving. The solution to this is, suppressing

dependence on q_H for brevity:

$$\begin{aligned}
(c_H^{NL})'(q_H) &= \frac{(p_H - 1)u'(c_H^L(q_H))^2}{(q_H - 1)^2 ((p_H - 1)u''(c_H^L(q_H))u'(c_H^{NL}) - p_H u'(c_H^L)u''(c_H^{NL}))} \\
(c_H^L)'(q_H) &= \frac{(p_H - 1)^2 u'(c_H^L)^2}{p_H(q_H - 1)^2 ((p_H - 1)u''(c_H^L)u'(c_H^{NL}) - p_H u'(c_H^L)u''(c_H^{NL}))} \\
(c_L^L)'(q_H) &= - \left[(p_L - 1)(p_H^2(q_H - 1)^2 u'(c_H^L)u''(c_H^{NL})(u(c_H^L) - u(c_H^{NL}) - u(c_L^L) + u(c_L^{NL})) \right. \\
&\quad - (p_H - 1)p_H(q_H - 1)^2 u''(c_H^L)u'(c_H^{NL})(u(c_H^L) - u(c_H^{NL}) - u(c_L^L) + u(c_L^{NL})) \\
&\quad \left. + (p_H - 1)p_H(q_H - 1)u'(c_H^L)^2 u'(c_H^{NL}) - (p_H - 1)^2 q_H u'(c_H^L)^3 \right] \\
&\quad / \left[p_H(q_H - 1)^2 ((p_L - 1)q_H u'(c_L^L) - p_L(q_H - 1)u'(c_L^{NL})) ((p_H - 1)u''(c_H^L)u'(c_H^{NL}) - p_H u'(c_H^L)u''(c_H^{NL})) \right] \\
(c_L^{NL})'(q_H) &= \left[p_L(-p_H^2(q_H - 1)^2 u'(c_H^L)u''(c_H^{NL})(u(c_H^L) - u(c_H^{NL}) - u(c_L^L) + u(c_L^{NL})) \right. \\
&\quad + (p_H - 1)p_H(q_H - 1)^2 u''(c_H^L)u'(c_H^{NL})(u(c_H^L) - u(c_H^{NL}) - u(c_L^L) + u(c_L^{NL})) \\
&\quad \left. + p_H(p_H(-q_H) + p_H + q_H - 1)u'(c_H^L)^2 u'(c_H^{NL}) + (p_H - 1)^2 q_H u'(c_H^L)^3 \right] \\
&\quad / \left[p_H(q_H - 1)^2 ((p_L - 1)q_H u'(c_L^L) - p_L(q_H - 1)u'(c_L^{NL})) ((p_H - 1)u''(c_H^L)u'(c_H^{NL}) - p_H u'(c_H^L)u''(c_H^{NL})) \right]
\end{aligned}$$

The appropriate second order partials with respect to χ can be found by differentiating (B.6)-(B.13) with respect to χ once more. Substituting these all in, evaluating at $\mathbf{q} = \mathbf{p}$ and hence at $c_H^{NL} = c_H^L = c$ yields

$$\begin{aligned}
\frac{\partial}{\partial q_H} \chi'(q_H) |_{\mathbf{q}=\mathbf{p}} &= (p_L - 1)p_L((u(c_L^L) - u(c_L^{NL}))) \\
&\times \left[\frac{(-(\alpha_H - 1)u'(c)((p_L - 1)u''(c_L^L) - p_L u''(c_L^{NL})) + \alpha_H(p_H - 1)u''(c_L^L)u'(c_L^{NL}) - \alpha_H p_H u'(c_L^L)u''(c_L^{NL}))}{(\alpha_H - 1)(p_H(p_L - 1)u'(c_L^L) - (p_H - 1)p_L u'(c_L^{NL}))^2} \right. \\
&\left. + \frac{-(u'(c_L^L) - u'(c_L^{NL}))((\alpha_H - 1)u'(c)((p_L - 1)u'(c_L^L) - p_L u'(c_L^{NL})) + \alpha_H u'(c_L^L)u'(c_L^{NL}))}{(\alpha_H - 1)(p_H(p_L - 1)u'(c_L^L) - (p_H - 1)p_L u'(c_L^{NL}))^2} \right] \\
&< 0.
\end{aligned}$$

The denominator is negative and the numerator is positive as $c_L^L < c_L^{NL}$, hence $u(c_L^L) < u(c_L^{NL})$ and $u'(c_L^L) < u'(c_L^{NL})$. This shows $\frac{\partial}{\partial q_H} \chi'(q_H) |_{\mathbf{q}=\mathbf{p}} < 0$. Similar working shows that for a small error by the low type:

$$\frac{\partial}{\partial q_L} \chi'(q_L) |_{\mathbf{q}=\mathbf{p}} = \frac{p_L ((\alpha_H - 1)u'(c) ((p_L - 1)u'(c_L^L) - p_L u'(c_L^{NL})) + \alpha_H u'(c_L^L)u'(c_L^{NL}))}{(\alpha_H - 1) (p_H(p_L - 1)u'(c_L^L) - (p_H - 1)p_L u'(c_L^{NL}))} > 0.$$

once one notices that $p_H > p_L$, $1 - p_L > 1 - p_H$ and $u'(c_L^L) > u'(c_L^{NL})$ and hence the denominator is positive.

Lastly it remains to show that welfare increases in χ . Recall welfare is given by

$$\alpha_H(p_H u(c_H^L(\chi)) + (1 - p_H)u(c_H^{NL}(\chi))) + (1 - \alpha_H)(p_L u(c_L^L(\chi)) + (1 - p_L)u(c_L^{NL}(\chi))).$$

Differentiating with respect to χ and substituting in $W = (B.8)$ and $(B.13)$, and then evaluating at $\mathbf{p} = \mathbf{q}$, yields

$$\frac{\partial W}{\partial \chi} = \frac{u'(c) \left((p_L - 1)(\alpha_H p_H - \alpha_H p_L + p_L)u'(c_L^L) + p_L(-\alpha_H p_H + (\alpha_H - 1)p_L + 1)u'(c_L^{NL}) \right)}{p_H(p_L - 1)u'(c_L^L) - (p_H - 1)p_L u'(c_L^{NL})} \quad (B.19)$$

$$+ \frac{\alpha_H(p_H - p_L)u'(c_L^L)u'(c_L^{NL})}{p_H(p_L - 1)u'(c_L^L) - (p_H - 1)p_L u'(c_L^{NL})}. \quad (B.20)$$

Note that the first order condition here reads:

$$p_L \frac{\partial c_L^L}{\partial \chi} + (1 - p_L) \frac{\partial c_L^{NL}}{\partial \chi} = 0,$$

and substituting this in and simplifying yields such that the change in welfare is just given by the effect on the high types utility:

$$\frac{\partial W}{\partial \chi} = \alpha_H u'(c) > 0. \quad \blacksquare$$

B.5 Proof of Lemma 1

Proof. First notice that with no errors, and fixing a level of cross-subsidization χ , the planner's optimal contracts coincide with the MW equilibrium contracts. Those contracts are: full insurance for the high type, with the low type's contract determined by the binding IC constraint of the high type. The proof of this follows from the proof of proposition 3, but also follows from Netzer and Scheuer (2014).

As such, the planner's problem is then to pick a level of cross-subsidization knowing that the MW contracts will be optimal conditional on χ .

Recall from the proof of Proposition 4 that the rate of change of welfare with respect to χ at $\mathbf{p} = \mathbf{q}$ is given by

$$\frac{\partial W}{\partial \chi} = \frac{u'(c) \left((p_L - 1)(\alpha_H p_H - \alpha_H p_L + p_L)u'(c_L^L) + p_L(-\alpha_H p_H + (\alpha_H - 1)p_L + 1)u'(c_L^{NL}) \right)}{p_H(p_L - 1)u'(c_L^L) - (p_H - 1)p_L u'(c_L^{NL})} \quad (B.21)$$

$$+ \frac{\alpha_H(p_H - p_L)u'(c_L^L)u'(c_L^{NL})}{p_H(p_L - 1)u'(c_L^L) - (p_H - 1)p_L u'(c_L^{NL})}. \quad (B.22)$$

Defining the first order condition of the MW optimization problem as

$$FOC_{MW}(\chi) = p_L \frac{\partial c_L^L}{\partial \chi} + (1 - p_L) \frac{\partial c_L^{NL}}{\partial \chi}$$

we can rewrite the change in welfare as

$$\frac{\partial W}{\partial \chi} = FOC_{MW}(\chi) + \alpha_H u'(c). \quad (\text{B.23})$$

From this we immediately conclude that when the MW equilibrium features $\chi_{MW}^* > 0$ such that $FOC_{MW}(\chi) = 0$ then the planner will choose $\chi_{planner}^* > \chi_{MW}^* > 0$. ■

B.6 Proof of Proposition 5

Proof. First, suppose IC_H doesn't bind. If the optimum is pooling then this is an immediate contradiction. So the optimum must separate the types. Then IC_L must bind else one type can be made better off holding cross-subsidization fixed (since both cannot be at their tangency point, by the assumption of small errors, as then both individuals tangency points are arbitrarily close to full insurance and one will be strictly preferred to the other by both types).

If $c_L \succ c_H$ we have an immediate contradiction as the indifference curve of the high type is strictly shallower than the low type, so if $V_L(c_L, q_L) = V_L(c_H, q_L)$ then it must be that $V_H(c_L, q_H) > V_L(c_H, q_H)$, contracting IC_H . On the other hand, if $c_H \succ c_L$ then, since $(c_H, c_L) \in A$ and the contracts are separate, the low type must receive less than full insurance, $c_L^{NL} > c_L^L$. But, fixing a level of cross-subsidization, their optimal contract from the planner's perspective is full insurance. They can then be made better off by moving along the profit line $\pi_L = -\chi$ toward full insurance. That IC_H will continue to be slack for small enough movements, giving a contradiction to welfare optimality. Hence IC_H must bind.

Then suppose there is a small error by the high type only. Whereas the no-error optimum for the high type featured $c_H^{NL} = c_H^L$, now slight over-insurance $c_H^{NL} + \epsilon = c_H^L$ for some small $\epsilon > 0$ is feasible, according to the constraints A . Holding constant the level of cross-subsidization, the proof of Proposition 2 shows that this slight over-insurance for the high types and a correspondingly better contract for the low types owing to the relaxed IC_H strictly improves welfare. This proves the first part of the theorem.

Now if the low type makes a small error, that IC_H binds and IC_L binds at the optimum immediately implies that the optimal contracts or welfare will not change, since such alternative contracts would have been feasible before the error, since IC_L can't have bound and q_L only matters there.

Finally, if the no-error optimum is pooling at the $\pi_{HL} = 0$ full insurance line, then the same deviation as in the separating case shows that when $q_H > p_H$ welfare can be strictly increased. ■

B.7 Proof of Proposition 7

First I calculate the change in welfare with respect to a small high type error. The process for a low type is similar and not explicated at length.

Proof. The three binding constraints are respectively the tangency constraint for the high type, the incentive constraint for the high type, and the participation constraint for the low type:

$$\frac{(1 - q_H)u'(c_H^{NL})}{q_H u'(c_H^L(c_H^{NL}))} = \frac{1 - p_H}{p_H} \quad (\text{B.24})$$

$$\text{IC} = q_H u(c_H^L(c_H^{NL})) + (1 - q_H)u(c_H^{NL}) = q_H u(c_L^L(c_H^{NL})) + (1 - q_H)u(c_L^{NL}(c_H^{NL})) \quad (\text{B.25})$$

$$\text{IR} = q_L u(c_L^L(c_H^{NL})) + (1 - q_L)u(c_L^{NL}(c_H^{NL})) = q_L u(w - l) + (1 - q_L)u(w). \quad (\text{B.26})$$

Given these, I can implicitly express each of c_H^L, c_L^L, c_L^{NL} in terms of c_H^{NL} in which case the monopolists maximization problem is solely a function of c_H^{NL} . As such, from the implicit function theorem I have

$$\frac{\partial c_H^{NL}}{\partial q_H} = - \frac{\partial^2 \Pi / \partial q_H \partial c_H^{NL}}{\partial^2 \Pi / \partial^2 c_H^L}.$$

Recalling that

$$\Pi(c_H^{NL}) = \alpha_H [(1 - p_H)(W - c_H^{NL}) - p_H(c_H^L(c_H^{NL}) - (W - l))] + \alpha_H [(1 - p_L)(W - c_L^{NL}(c_H^{NL})) - p_L(c_L^L(c_H^{NL}) - (W - l))]$$

differentiating we have

$$\begin{aligned} \frac{\partial^2 \Pi(c_H^{NL})}{\partial q_H \partial c_H^{NL}} &= \alpha_H \left[p_H \frac{\partial^2}{\partial q_H \partial c_H^{NL}} (c_H^L(c_H^{NL})) \right] \\ &+ \alpha_H \left[(1 - p_L) \left(\frac{\partial^2}{\partial q_H \partial c_H^{NL}} c_L^{NL}(c_H^{NL}) \right) - p_L \left(\frac{\partial^2}{\partial q_H \partial c_H^{NL}} c_L^L(c_H^{NL}) \right) \right] \end{aligned}$$

and

$$\frac{\partial^2 \Pi(c_H^{NL})}{\partial (c_H^{NL})^2} = \alpha_H \left[p_H \frac{\partial^2}{\partial (c_H^{NL})^2} (c_H^L(c_H^{NL})) \right] + \alpha_H \left[(1 - p_L) \left(\frac{\partial^2}{\partial (c_H^{NL})^2} c_L^{NL}(c_H^{NL}) \right) - p_L \left(\frac{\partial^2}{\partial (c_H^{NL})^2} c_L^L(c_H^{NL}) \right) \right]$$

To calculate $\frac{\partial c_H^L(c_H^{NL})}{\partial c_H^{NL}}, \frac{\partial c_L^L(c_H^{NL})}{\partial c_H^{NL}}, \frac{\partial c_L^{NL}(c_H^{NL})}{\partial c_H^{NL}}$ differentiate each of the binding constraints with respect to c_H^{NL} to get respectively:

$$0 = \frac{(q_H - 1)(u''(c_H^{NL})u'(c_H^L(c_H^{NL})) - (c_H^L)'(c_H^{NL})u'(c_H^{NL})u''(c_H^L(c_H^{NL})))}{q_H u'(c_H^L(c_H^{NL}))} \quad (\text{B.27})$$

$$q_H (c_H^L)'(c_H^{NL})u'(c_H^L(c_H^{NL})) + (1 - q_H)u'(c_H^{NL}) = q_H (c_L^L)'(c_H^{NL})u'(c_L^L(c_H^{NL})) + (1 - q_H)(c_L^{NL})'(c_H^{NL})u'(c_L^{NL}(c_H^{NL})) \quad (\text{B.28})$$

$$0 = q_L (c_L^L)'(c_H^{NL})u'(c_L^L(c_H^{NL})) + (1 - q_L)(c_L^{NL})'(c_H^{NL})u'(c_L^{NL}(c_H^{NL})) \quad (\text{B.29})$$

Solving this system yields

$$\frac{\partial c_L^L(c_H^{NL})}{\partial c_H^{NL}} = -\frac{(q_L - 1) \left((q_H - 1)u'(c_H^{NL}) - \frac{q_H u''(c_H^{NL})u'(c_H^L(c_H^{NL}))^2}{u'(c_H^{NL})u''(c_H^L(c_H^{NL}))} \right)}{(q_L - q_H)u'(c_L^L(c_H^{NL}))} \quad (\text{B.30})$$

$$\frac{\partial c_L^{NL}(c_H^{NL})}{\partial c_H^{NL}} = -\frac{q_L \left((q_H - 1)u'(c_H^{NL}) - \frac{q_H u''(c_H^{NL})u'(c_H^L(c_H^{NL}))^2}{u'(c_H^{NL})u''(c_H^L(c_H^{NL}))} \right)}{(q_L - q_H)u'(c_L^{NL}(c_H^{NL}))} \quad (\text{B.31})$$

$$\frac{\partial c_H^L(c_H^{NL})}{\partial c_H^{NL}} = \frac{u''(c_H^{NL})u'(c_H^L(c_H^{NL}))}{u'(c_H^{NL})u''(c_H^L(c_H^{NL}))} \quad (\text{B.32})$$

To get the cross partials, differentiate with respect to q_H :

$$\frac{\partial^2}{\partial q_H \partial c_H^{NL}} c_L^L(c_H^{NL}) = \frac{(q_L - 1) (q_L u''(c_H^{NL})u'(c_H^L(c_H^{NL}))^2 - (q_L - 1)u'(c_H^{NL})^2 u''(c_H^L(c_H^{NL})))}{(q_H - q_L)^2 u'(c_H^{NL})u''(c_H^L(c_H^{NL}))u'(c_L^L(c_H^{NL}))} \quad (\text{B.33})$$

$$\frac{\partial^2}{\partial q_H \partial c_H^{NL}} c_L^{NL}(c_H^{NL}) = \frac{q_L (q_L u''(c_H^{NL})u'(c_H^L(c_H^{NL}))^2 - (q_L - 1)u'(c_H^{NL})^2 u''(c_H^L(c_H^{NL})))}{(q_H - q_L)^2 u'(c_H^{NL})u''(c_H^L(c_H^{NL}))u'(c_L^{NL}(c_H^{NL}))} \quad (\text{B.34})$$

$$\frac{\partial^2}{\partial q_H \partial c_H^{NL}} c_H^L(c_H^{NL}) = 0. \quad (\text{B.35})$$

Hence we have

$$\begin{aligned} \frac{\partial^2 \Pi(c_H^{NL})}{\partial q_H \partial c_H^{NL}} &= \frac{\alpha_L (q_L u''(c_H^{NL})u'(c_H^L(c_H^{NL}))^2 - (q_L - 1)u'(c_H^{NL})^2 u''(c_H^L(c_H^{NL})))}{(q_H - q_L)^2 u'(c_H^{NL})u''(c_H^L(c_H^{NL}))u'(c_L^L(c_H^{NL}))u'(c_L^{NL}(c_H^{NL}))} \\ &\quad \times ((p_L - 1)q_L u'(c_L^L(c_H^{NL})) - p_L (q_L - 1)u'(c_L^{NL}(c_H^{NL}))) \end{aligned}$$

On the other hand, differentiating (B.30) with respect to c_H^{NL} again, and using the values from (B.30) yields

$$\frac{\partial^2}{\partial(c_H^{NL})^2} c_L^L(c_H^{NL}) = \left[(q_L - 1)((1 - q_L)u''(c_H^L(c_H^{NL}))u''(c_L^L(c_H^{NL}))(q_H u''(c_H^{NL})u'(c_H^L(c_H^{NL})))^2 \right. \quad (\text{B.36})$$

$$\left. - (q_H - 1)u'(c_H^{NL})^2 u''(c_H^L(c_H^{NL})) \right)^2 \quad (\text{B.37})$$

$$- (q_H - q_L)u'(c_L^L(c_H^{NL}))^2 (-q_H u^{(3)}(c_H^L(c_H^{NL}))u''(c_H^{NL})^2 u'(c_H^L(c_H^{NL})))^3 \quad (\text{B.38})$$

$$+ q_H u'(c_H^L(c_H^{NL}))^2 u''(c_H^L(c_H^{NL}))^2 (u^{(3)}(c_H^{NL})u'(c_H^{NL}) + u''(c_H^{NL})^2) \quad (\text{B.39})$$

$$\left. - (q_H - 1)u'(c_H^{NL})^2 u''(c_H^{NL})u''(c_H^L(c_H^{NL}))^3 \right] \quad (\text{B.40})$$

$$/ [(q_H - q_L)^2 u'(c_H^{NL})^2 u''(c_H^L(c_H^{NL}))^3 u'(c_L^L(c_H^{NL}))^3] \quad (\text{B.41})$$

$$\frac{\partial^2}{\partial(c_H^{NL})^2} c_L^{NL}(c_H^{NL}) = \left[-q_L((q_H - q_L)u'(c_L^{NL}(c_H^{NL}))^2 (-q_H u^{(3)}(c_H^L(c_H^{NL}))u''(c_H^{NL})^2 u'(c_H^L(c_H^{NL})))^3 \right. \quad (\text{B.42})$$

$$\left. + q_H u'(c_H^L(c_H^{NL}))^2 u''(c_H^L(c_H^{NL}))^2 (u^{(3)}(c_H^{NL})u'(c_H^{NL}) + u''(c_H^{NL})^2) \right) \quad (\text{B.43})$$

$$- (q_H - 1)u'(c_H^{NL})^2 u''(c_H^{NL})u''(c_H^L(c_H^{NL}))^3 \quad (\text{B.44})$$

$$\left. + q_L u''(c_H^L(c_H^{NL}))u''(c_L^{NL}(c_H^{NL}))(q_H u''(c_H^{NL})u'(c_H^L(c_H^{NL}))^2 - (q_H - 1)u'(c_H^{NL})^2 u''(c_H^L(c_H^{NL})))^2 \right] \quad (\text{B.45})$$

$$/ [(q_H - q_L)^2 u'(c_H^{NL})^2 u''(c_H^L(c_H^{NL}))^3 u'(c_L^{NL}(c_H^{NL}))^3] \quad (\text{B.46})$$

$$\frac{\partial^2}{\partial(c_H^{NL})^2} c_H^L(c_H^{NL}) = \frac{u'(c_H^L(c_H^{NL}))(u^{(3)}(c_H^{NL})u'(c_H^{NL})u''(c_H^L(c_H^{NL}))^2 - u^{(3)}(c_H^L(c_H^{NL}))u''(c_H^{NL})^2 u'(c_H^L(c_H^{NL})))}{u'(c_H^{NL})^2 u''(c_H^L(c_H^{NL}))^3} \quad (\text{B.47})$$

And so the second order total derivative of profit is

$$\frac{\partial^2 \Pi(c_H^{NL})}{\partial (c_H^{NL})^2} = \frac{\gamma H(p_H - 1)u'(c_H^L(c_H^{NL}))(u^{(3)}(c_H^{NL})u'(c_H^{NL})u''(c_H^L(c_H^{NL}))^2 - u^{(3)}(c_H^L(c_H^{NL}))u''(c_H^{NL})^2u'(c_H^L(c_H^{NL})))}{u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3} \quad (\text{B.48})$$

$$- \frac{(1 - \alpha_H)p_L(q_L - 1)(1 - q_L)u''(c_H^L(c_H^{NL}))u''(c_L^L(c_H^{NL}))(q_H u''(c_H^{NL})u'(c_H^L(c_H^{NL}))^2)}{(q_H - q_L)^2u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3u'(c_L^L(c_H^{NL}))^3} \quad (\text{B.49})$$

$$- \frac{-(q_H - 1)u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^2}{(q_H - q_L)^2u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3u'(c_L^L(c_H^{NL}))^3} \quad (\text{B.50})$$

$$- \frac{(q_H - q_L)u'(c_L^L(c_H^{NL}))^2(-q_H u^{(3)}(c_H^L(c_H^{NL}))u''(c_H^{NL})^2u'(c_H^L(c_H^{NL}))^3)}{u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3} \quad (\text{B.51})$$

$$- \frac{q_H u'(c_H^L(c_H^{NL}))^2u''(c_H^L(c_H^{NL}))^2(u^{(3)}(c_H^{NL})u'(c_H^{NL}))}{u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3} \quad (\text{B.52})$$

$$+ \frac{u''(c_H^{NL})^2 - (q_H - 1)u'(c_H^{NL})^2u''(c_H^{NL})u''(c_H^L(c_H^{NL}))^3}{u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3} \quad (\text{B.53})$$

$$+ \frac{(1 - \alpha_H)(1 - p_L)q_L q_L u''(c_H^L(c_H^{NL}))u''(c_L^{NL}(c_H^{NL}))(q_H u''(c_H^{NL})u'(c_H^L(c_H^{NL}))^2)}{(q_H - q_L)^2u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3u'(c_L^{NL}(c_H^{NL}))^3} \quad (\text{B.54})$$

$$+ \frac{(q_H - 1)u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^2}{(q_H - q_L)^2u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3u'(c_L^{NL}(c_H^{NL}))^3} \quad (\text{B.55})$$

$$+ \frac{(q_H - q_L)u'(c_L^{NL}(c_H^{NL}))^2 - q_H u^{(3)}(c_H^L(c_H^{NL}))u''(c_H^{NL})^2u'(c_H^L(c_H^{NL}))^3}{u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3} \quad (\text{B.56})$$

$$+ \frac{q_H u'(c_H^L(c_H^{NL}))^2u''(c_H^L(c_H^{NL}))^2(u^{(3)}(c_H^{NL})u'(c_H^{NL}))}{u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3} \quad (\text{B.57})$$

$$+ \frac{u''(c_H^{NL})^2 - (q_H - 1)u'(c_H^{NL})^2u''(c_H^{NL})u''(c_H^L(c_H^{NL}))^3}{u'(c_H^{NL})^2u''(c_H^L(c_H^{NL}))^3}. \quad (\text{B.58})$$

Now, knowing how a change in q_H affects the optimized value of c_H^{NL} , I calculate how this will affect the optimized values of c_H^L, c_L^L, c_L^{NL} . The optimized values solve the same constraints of course, now written as dependant on q_H . Differentiating them with respect to q_H and substituting the known expression for $\frac{\partial^2 \Pi(c_H^{NL})}{\partial (c_H^{NL})^2}$ yields a system that implicitly defines $\frac{\partial (c_H^L)^*}{\partial q_H}, \frac{\partial (c_L^{NL})^*}{\partial q_H}, \frac{\partial (c_L^L)^*}{\partial q_H}$.

Solving these gives, evaluated now at $\mathbf{p} = qvec$:

$$\frac{\partial(c_L^L)^*}{\partial q_H} = \frac{(p_L - 1)}{(p_H - p_L)u'(c_L^L)} \quad (\text{B.59})$$

$$\times \left(\frac{u'(c)^2 u'(c_L^L)^2 u'(c_L^{NL})^2 (u'(c_L^{NL}) - u'(c_L^L))}{u'(c_L^{NL})^2 ((p_H - p_L)u''(c)u'(c_L^L)^2 (u'(c_L^L) - u'(c_L^{NL})) - (p_L - 1)u'(c)^2 u''(c_L^L)u'(c_L^{NL})) + p_L u'(c)^2 u'(c_L^L)^3 u''(c_L^{NL})} \right) \quad (\text{B.60})$$

$$- \frac{u'(c)^2}{(p_H - 1)u''(c)} + u(c_L^L) - u(c_L^{NL}) \quad (\text{B.61})$$

$$\frac{\partial(c_L^{NL})^*}{\partial q_H} = \frac{p_L}{(p_H - p_L)u'(c_L^{NL})} \left(u'(c)^2 u'(c_L^L)^2 u'(c_L^{NL})^2 (u'(c_L^{NL}) - u'(c_L^L)) \right) \quad (\text{B.62})$$

$$/ \left(u'(c_L^{NL})^2 ((p_H - p_L)u''(c)u'(c_L^L)^2 (u'(c_L^L) - u'(c_L^{NL})) - (p_L - 1)u'(c)^2 u''(c_L^L)u'(c_L^{NL})) \right) \quad (\text{B.63})$$

$$+ p_L u'(c)^2 u'(c_L^L)^3 u''(c_L^{NL}) - \frac{u'(c)^2}{(p_H - 1)u''(c)} + u(c_L^L) - u(c_L^{NL}) \quad (\text{B.64})$$

$$\frac{\partial(c_H^L)^*}{\partial q_H} = \left((p_H^2 - p_L)u'(c)u''(c)u'(c_L^L)^2 u'(c_L^{NL})^2 (u'(c_L^L) - u'(c_L^{NL})) + u'(c)^3 \left[p_L u'(c_L^L)^3 u''(c_L^{NL}) \right. \right) \quad (\text{B.65})$$

$$\left. - (p_L - 1)u''(c_L^L)u'(c_L^{NL})^3 \right] \quad (\text{B.66})$$

$$/ \left((p_H - 1)p_H u''(c)(u'(c_L^{NL})^2 ((p_H - p_L)u''(c)u'(c_L^L)^2 (u'(c_L^L) - u'(c_L^{NL})) \right) \quad (\text{B.67})$$

$$- (p_L - 1)u'(c)^2 u''(c_L^L)u'(c_L^{NL})) + p_L u'(c)^2 u'(c_L^L)^3 u''(c_L^{NL}) \quad (\text{B.68})$$

Now we can evaluate the change in welfare at the new optimum versus the old optimum: Directly from the constraints we have

$$\frac{\partial W_L^*}{\partial q_H} \Big|_{\mathbf{q}=\mathbf{p}} = 0.$$

For the high type:

$$\frac{\partial W_H^*}{\partial q_H} \Big|_{q=pvec} = \alpha_H p_H (c_H^L)'(q_H) u'(c_H^L) + \alpha_H (1 - p_H) (c_H^{NL})'(q_H) u'(c_H^{NL}) \quad (\text{B.69})$$

$$= \left(\alpha_H (2p_H - p_L - 1) u'(c)^2 u''(c) u'(c_L^L)^2 u'(c_L^{NL})^2 (u'(c_L^L) - u'(c_L^{NL})) \right) \quad (\text{B.70})$$

$$+ \alpha_H u'(c)^4 (p_L u'(c_L^L)^3 u''(c_L^{NL}) - (p_L - 1) u''(c_L^L) u'(c_L^{NL})^3) \quad (\text{B.71})$$

$$/ \left((p_H - 1) u''(c) (u'(c_L^{NL})^2 \left[(p_H - p_L) u''(c) u'(c_L^L)^2 (u'(c_L^L) - u'(c_L^{NL})) \right. \right. \quad (\text{B.72})$$

$$\left. \left. - (p_L - 1) u'(c)^2 u''(c_L^L) u'(c_L^{NL}) + p_L u'(c)^2 u'(c_L^L)^3 u''(c_L^{NL}) \right] \right) \quad (\text{B.73})$$

The denominator can be seen to be negative, since $u'(c_L^L) - u'(c_L^{NL}) > 0$ and utility is concave. As such, the welfare change will be positive when, the numerator is negative. A sufficient condition for this is $2p_H > 1 + p_L > 0$, which proves the claim.

The calculations for a change in the low types error are completely analogous and omitted for brevity. ■

B.8 Proof of Proposition 8

Proof. Essentially identical to the proof of proposition 1. ■

B.9 Proof of Proposition 9

Proof. Essentially identical to the proof of proposition 1 except that it must be established that the optimistic high types and pessimistic low types must be pooled together.

Suppose not. First, suppose that they are not pooled together. The pessimistic high types must then have their own contract that earns zero profits.⁽¹⁶⁾ As such both true low types must individually earn zero profits. But then we cannot have both that $V(q_{L,P}, \mathbf{c}_{LP}) \geq V(q_{H,O}, \mathbf{c}_{HO})$ and the reverse since we have $\mathbf{c}_{HO} \succsim \mathbf{c}_{LP}$ whilst $q_{H,O} < q_{L,P}$ contradicting the fact established at the beginning of proposition 1.

So they must be pooled together and the contract must earn zero profits. If not another firm will offer a slightly more attractive contract. That the constraints must bind in the way described is clear and arguments as in proposition 1 apply.

It remains to show that the contract defined by the zero profit line from the pool of high optimists and low pessimists and the high pessimists binding IC constraint is in fact an equilibrium. The usual deviation away from pooling is to just attract the lower risk type by offering slightly less insurance. This doesn't work here as offering slightly less insurance attracts the high optimists but not the low pessimists as the latter think they are higher risk than the former, despite being objectively lower risk. So to attract the objectively low risk pessimists, slightly more insurance needs to be offered. But given that $V(q_{H,O}, \mathbf{c}_{H,O}) = V(q_{H,O}, \mathbf{c}_L)$, offering slightly more insurance that attracts the low

⁽¹⁶⁾This is because they cannot be pooled with the high pessimistic types since a deviation along $\pi_H = 0$ but tangent for the high optimistic types will attract just them as the low pessimistic types and high pessimistic types prefer the high pessimistic types contract. They cannot be pooled with the lowest types for the standard reasons.

pessimists will also attract the high optimists, since (A.8) holds. Then, because of the additional assumption that $\alpha_{H,P} > \alpha_{H,O}$ such a contract will make less than the zero profit from c_{HL} since the group of high optimists is being replaced by the larger pool of high pessimists. Hence there is no local deviation and the equilibrium is as posited. ■

C Dataset Construction

C.1 Full set of mortality prediction covariates

The set of variables that appear in the long logit specification, and that were the starting point for the variable selection in the lasso logit specification, are as follows:

1. The respondents subjective elicitation.
2. Sex and age(nearest year)
3. Dummies for diabetes, cancer, lung disease, heart disease, stroke, arthritis, and the first time difference for all of these.
4. Dummies for whether the respondent has ever had any of these diseases.
5. The full set of interactions of 2. and 3.
6. BMI, dummies for being married, seperated, divorced, never marraid. The length of the current marriage.
7. A dummy for whether the respodnents mother and father are alive, and their current age (or age of death).
8. The subjective self health assessment and the first time difference.
9. The number of overnight hospital visits, nursing home visits, doctor visits and episodes of home care since the previous interview. Out of pocket medical expenditures and the first time difference.
10. Indices for activities of daily living (and a time difference), mobility, large muscle, gross motor skills, fine motor skills, and instrumental activities of daily living.
11. Indicators for whether depression was experienced, whether the respondent was happy, the number of days each week in which alchohol was drunk and the first time difference, whether the respondent now smokes or ever smoked, an indicator for high blood pressure.
12. Indicators for past or current memory problems, whether the repondent has public or private health insurance (if the latter how many plans), has life insurance.
13. The number of children, an indicator for whether the respondent is retired, current income and the first time difference in income.