Heterogeneity, Demand for Insurance and Adverse Selection

Johannes Spinnewijn*
London School of Economics

December 15, 2011
COMMENTS VERY WELCOME.

Abstract

Recent empirical work finds that surprisingly little variation in the demand for insurance is explained by heterogeneity in risks. The welfare and policy conclusions are substantially different when the residual demand variation is due to heterogeneity in risk perceptions and other noisy determinants rather than to heterogeneous preferences, as previously assumed. This heterogeneity induces a systematic difference between the revealed and actual value of insurance as a function of the insurance price, which is used to extend the sufficient statistics approach to the welfare analysis of adverse selection. The source of heterogeneity is essential for the effectiveness of insurance subsidies and mandates, information policies and risk-adjusted pricing.

1 Introduction

Adverse selection due to heterogeneity in risks has been considered a prime reason for governments to intervene in insurance markets. The classic argument is that the presence of higher risk types increases insurance premia and drives lower risk types out of the market (Akerlof 1970). However, empirical work has found surprisingly little evidence supporting the existence of adverse selection in insurance markets.[1] Most of the variation in the demand for insurance can not be attributed to heterogeneity in risks. The attribution to heterogeneity in preferences in recent empirical work (Cohen and Einav 2007, Einav, Finkelstein and Cullen 2010a, Einav, Finkelstein and Schrimpf 2010b) implies that the estimated welfare cost of inefficient pricing due to adverse selection is very small; the value of insurance for the uninsured is estimated to be small, as is the pool of inefficiently uninsured individuals. Heterogeneity in preferences thus undermines the role for policy interventions in insurance markets.

*Department of Economics, STICERD R515, LSE, Houghton Street, London WC2A 2AE, United Kingdom (email: j.spinnewijn@lse.ac.uk, web: http://personal.lse.ac.uk/spinnewi/).

[1] See Cohen and Siegelman (2010) for an overview.

An alternative explanation why risks do not explain the demand for insurance is the discrepancy between perceived and actual risks. A large literature documents the biases and heuristics in forming risk perceptions and the inherently subjective nature of people's formed risk perceptions.[2] Neighbors in a coastal area may have very different perceptions about the risk of a natural disaster damaging their property, even though they face the same actual risk. In general, risk perceptions are likely to be only a noisy measure of one's actual risk, which drives a wedge between the *actual* value of insurance and the value of insurance as *revealed* by an individual's demand. Recent empirical evidence on inertia (Handel 2011) and inconsistencies (Abaluck and Gruber, 2011) in insurance choices and on the substantial role played by cognitive ability (Fang, Keane and Silverman 2008) confirm the importance of this wedge. To the extent that we care about the actual value rather than the revealed value of insurance, the *non-welfarist* heterogeneity underlying the demand for insurance changes earlier welfare and policy conclusions.

This paper presents a framework with different dimensions of heterogeneity underlying the variation in the demand for insurance and the related adverse selection. A simple selection argument implies that non-welfarist heterogeneity drives a systematic difference between the actual and revealed value of insurance along the demand curve. The paper analyzes and calibrates the consequence for welfare and policy analysis in the presence of adverse selection, extending the sufficient statistics approach by Einav et al. (2010a). The welfare cost of adverse selection is substantially higher than previously estimated if some of the unexplained heterogeneity in insurance choices is due to non-welfarist heterogeneity. Moreover, the effectiveness of information policies and standard policy interventions in insurance markets crucially depends on the source of heterogeneity.

I consider a simple model where individuals are heterogeneous in risks, preferences and perceptions (or any other non-welfarist noise), and decide whether or not to buy insurance. Although perceptions are accurate on average, a simple selection argument implies that the insured tend to overestimate, while the uninsured tend to underestimate the value of insurance and this at any price. That is, as overly pessimistic beliefs encourage individuals to buy insurance, individuals buying insurance are more likely to be too pessimistic and vice versa.[3] This result applies when the noise determining people's risk perceptions is independent of the actual risk, but generally extends when the correlation between the perceived and actual risks is imperfect and the dispersion in perceived risk exceeds the dispersion in actual risk.

The welfare implication is that the demand curve overstates the surplus for the insured individuals and understates the potential surplus for the uninsured individuals.

[2] See Tversky and Kahneman (1974) and Slovic (2000) for the seminal contributions to this literature.

[3] The selection effect is structurally similar to the mechanisms underlying for example the winner's curse, regression towards to the mean, and choice-driven optimism (Van Den Steen 2004), conditioning an expected value on a particular choice or outcome.

When taking the demand curve at face value, the evaluation of policy interventions, targeting the insured and uninsured respectively, will be unambiguously biased in opposite directions. For instance, the welfare gain of a universal mandate is unambiguously higher than the demand for insurance would suggest. In order to extend the results from infra-marginal individuals to marginal individuals and thus to evaluate more targeted policy interventions, information on the thickness of the market is required. I derive conditions under which the *value curve*, depicting the actual value of insurance for the marginally insured at a given price, is a counter-clockwise rotation of the demand curve.[4] With normal heterogeneity, the slope of the value curve will become smaller relative to the slope of the demand curve, the lower the correlation between the perceived and actual risk and the larger the variance in perceived risks relative to the variance in actual risks.

I use this systematic relation between the value and demand curve to extend the sufficient statistics approach by Einav, Finkelstein and Cullen (2010a) for non-welfarist heterogeneity. When the demand reveals the actual value, the demand and cost curves are sufficient statistics for welfare analysis. In the presence of non-welfarist heterogeneity, the one additional statistic that is required captures the extent to which heterogeneous choices - left unexplained by heterogeneity in risks - are explained by heterogeneous risk perceptions (or other noise) rather than by heterogeneous preferences. Using the empirical estimates of the demand and cost curves from Einav et al. (2010a) for employer-provided health insurance, I find that the actual cost of adverse selection would be thirty percent higher when ten percent of the unexplained variation is driven by variation in perceptions and four times as high when this share increases to fifty percent.

I also the framework to analyze and evaluate the impact of non-welfarist heterogeneity on standard government interventions in insurance markets. First, the heterogeneity introduces a disconnect between price and quantity policies, since price policies aiming to induce more insurance coverage are still constrained by individuals' perceived valuations, while the welfare effect of an increase in insurance coverage solely depends on their actual valuations. Second, policies that reduce the wedge between the perceived and actual values have an ambiguous effect on welfare. They increase expected welfare to the extent that they induce the purchase of insurance by individuals for whom the net-value of insurance rather the provision cost of insurance is higher. The latter would increase the equilibrium price and cause the market to be more adversely selected.[5] The framework with multi-dimensional heterogeneity allows to disentangle the two effects.

[4] Johnson and Myatt (2006) analyze how marketing and advertizing rotate a demand curve. Notice that the distribution underlying the rotated demand curve is derived, but independent of the distribution underlying the original demand curve. Here, the value curve however depends on a distribution that is explicitly correlated with the distribution underlying the original demand curve.

[5] Condon, Kling and Mullanaithan (2011) also discuss the potential welfare loss when people are better informed about their risks. Handel (2010) provides an empirical welfare analysis of a similar trade-off for a nudging policy when people's decisions are subject to switching costs or inertia.

An immediate policy implication is that improving individuals' information about the expected risk decreases welfare, while improving individuals' information about the variance of the risk increases welfare. Finally, adjusting the pricing of insurance contracts for the particular risk type of the buyer is often argued to reduce adverse selection and thus increase inefficiency. However, I show that the adjustment for risks may in fact decrease efficiency when individuals do not perceive these risks accurately.

1.1 Related Literature

Starting with the work by Chiappori and Salanié (1996, 2000), several papers have tested for the presence of adverse selection in different insurance markets, using the testable implication that the correlation between insurance coverage and risk is positive. The mixed evidence, with some insurance markets being advantageously rather than adversely selected (Cohen and Siegelman 2010), inspired a new series of studies to estimate the heterogeneity in risk preferences jointly with the heterogeneity in risk types, using data on insurance choices and claim rates (Cohen and Einav, 2007; Einav et al. 2010a, 2010b, 2010c). The estimated welfare cost of inefficient pricing due to adverse selection is small, as the estimated heterogeneity in preferences dominates the heterogeneity in risks.

A parallel literature argues that insurance behavior cannot be adequately explained with standard preferences and risk perceptions. Cutler and Zeckhauser (2004) argue that distorted risk perceptions are one of the main reasons why some insurance markets do not work efficiently. Recent empirical evidence confirms the tenuous relation between the value of insurance and insurance choices. For instance, Fang et al. (2008) find that heterogeneity in cognitive ability is important (relative to risk aversion) in explaining the choice of eldery to buy Medigap. Abaluck and Gruber (2011) identify explicit inconsistencies in their choices, also documenting dominated choices with an alternative plan offering better risk protection at a lower cost available. A number of recent papers have a more explicit focus on the importance risk perceptions for insurance choices, but estimates a uniform model to explain insurance behavior rather than attributing heterogeneous choices to different sources of heterogeneity. For example, Barseghyan, Molinari, O'Donoghue and Teitelbaum (2010) find that a model with nonlinear probability weighting rather than with standard risk aversion performs better in explaining deductible choices in auto and house insurance.[6] Some studies also analyze the stability or generality of an individual's risk preference across choices (Barseghyan, Prince and Teitelbaum 2010, Einav, Finkelstein, Pascu and Cullen 2011). Despite the presence of some commonality, choices in different insurance domains seem not only explained by risks and risk preferences.

This paper fills a gap between these two strands of literature by analyzing welfare and policy interventions without assuming that the heterogeneity in insurance choices

[6] Other recent examples are Sydnor 2010, Snowberg and Wolfers 2010, Bruhin et al. 2010.

- unexplained by heterogeneity in risks - is only driven by heterogeneity in preferences. The welfare analysis relates to two recent trends in public economics by considering non-standard decision makers in the welfare analysis (Bernheim and Rangel 2009, Congdon et al. 2011) and expressing the optimal policy in terms of sufficient statistics (Chetty 2010). In a similar spirit, Chetty, Kroft and Looney (2009) extend the sufficient statistics approach to tax policy for tax salience and Spinnewijn (2009) extends the sufficient statistics approach to unemployment policy for biased perceptions of employment prospects.

The remainder of the paper is as follows. Section 2 introduces a simple model of the demand for insurance and characterizes the difference between actual and perceived insurance values along the demand curve. Section 3 introduces heterogeneity in risk types and preferences to analyze and calibrate the cost of adverse selection. Section 4 analyzes the effectiveness of different government interventions depending on the importance of non-welfarist heterogeneity. Section 5 concludes.

2 Demand and Welfare

This section introduces a simple model of insurance demand and analyzes the systematic difference between the value of insurance, as revealed by an individual's demand for insurance, and the actual value of insurance as a function of the insurance price. I thus deviate from the revealed preference paradigm and assume that the variation in insurance decisions may be driven by heterogeneity in *non-welfarist* variables, unrelated to the actual value of insurance. While this may imply a particular stance on what the welfare criterion should be for non-standard decision makers, the analysis does not only apply to heterogeneity in *behavioral* variables like misperceptions, inattention, cognitive inability or inertia, but also to heterogeneity in *economic* variables, like liquidity constraints or adjustment costs, which are also restricting people's ability to buy insurance regardless of the value of insurance for those individuals. The analysis is general, but I will mostly interpret the source of the non-welfarist heterogeneity as coming from differences between perceived and actual risks.

2.1 Simple Model

Individuals decide whether or not to buy insurance against a risk. I assume that only one contract is provided and all individuals can buy this contract at a variable price p. Individuals may differ in several dimensions and these different characteristics are captured by a vector ζ. Examples of characteristics are individuals' risk preferences, risk types, perceptions of their risk types, cognitive ability, wealth and liquidity constraints,... I distinguish between the *true* value of insurance $v(\zeta)$ and the *perceived* value of insurance $\hat{v}(\zeta)$ for an individual with characteristics ζ. The true value refers to the actual value of the insurance contract for a given individual and is relevant for

evaluating welfare and policy intervention. The perceived value, however, refers to the value as perceived by this individual and determines his or her demand for insurance.

I assume that the perceived value equals the sum of the true value and a noise term ε,

$$\hat{v}(\zeta) = v(\zeta) + \varepsilon(\zeta) \text{ with } E_{\zeta}(\varepsilon) = 0.$$

The noise reduces the correlation between the perceived and true value in the population. When capturing misperceptions of the risk, the noise term would be positive when an individual overestimates the risk she is facing and negative when the individual underestimates that risk. I assume that on average the noise cancels out across the entire population so that the average true and perceived value coincide. However, since the demand for insurance only depends on the perceived value, the true and perceived value may differ substantially conditional on the insurance decision.

An individual with characteristics ζ will buy an insurance contract if her perceived value exceeds the price, $\hat{v}(\zeta) \geq p$. The demand for insurance at price p equals $D(p) = 1 - F_{\hat{v}}(p)$. As well known, the demand curve reflects the marginal willingness to pay of marginal buyers at different prices. That is, the price reveals the perceived value for the marginal buyers at that price, $p = E_{\zeta}(\hat{v}|\hat{v} = p)$. However, to evaluate welfare, we would like to know the (average) true value for the marginal buyers, which I denote by $MV(p) \equiv E_{\zeta}(v|\hat{v} = p)$.[7] The central question is thus to what extent the true value covaries with the perceived value. A central statistic capturing this co-movement is the ratio of the covariance between the true and perceived value to the variance of the perceived value,

$$\frac{cov(v, \hat{v})}{var(\hat{v})} = \frac{var(v) + cov(v, \varepsilon)}{var(v) + var(\varepsilon) + 2cov(v, \varepsilon)}.$$

Graphically, one can construct the *value curve*, depicting the expected true value for the marginal buyers for any level of insurance coverage $q = D(p)$, and compare this to the demand curve, depicting the perceived value at any level of insurance coverage q. The mistake made by a *naive* policy maker who uses the demand curve to evaluate policies depends on the wedge between the two curves. For policies that target marginal individuals, the difference in levels between the two curves is relevant. For policies that target infra-marginal individuals, either the insured or the uninsured, the difference in the areas below the two curves is relevant. I analyze the systematic nature of these differences along the demand curve.

2.2 Infra-marginally Insured and Uninsured

I start by analyzing the average insurance value for the infra-marginal individuals, given by $E(v|\hat{v}(\zeta) \geq p)$ for the insured and $E(v|\hat{v}(\zeta) < p)$ for the uninsured. For the insured,

[7] Individuals with the same perceived value may have very different actual values. I take the unweighted average of the insurance value to evaluate welfare. This utilitarian approach implies that in the absence of noise, total welfare is captured by the consumer surplus.

the average value of insurance determines the actual consumer surplus generated in the insurance market and thus the value of any policy affecting all insured individuals, like banning an insurance product. For the uninsured, this determines the value of any policy affecting all uninsured individuals, like a universal mandate to buy insurance.

Random Noise I first consider the case where the noise determining the perceived value is independent of the true value. The implied co-movement of the actual and perceived value only depends on the relative variances of the true value and the noise term,

$$\frac{cov\left(v,\hat{v}\right)}{var\left(\hat{v}\right)}=\frac{var\left(v\right)}{var\left(v\right)+var\left(\varepsilon\right)}.$$

Not surprisingly, an increase in the perceived value is less indicative of an increase in the actual value if noise is more important. Moreover, since the noise term determines the perceived value of insurance, the expected noise realization will be different among those who buy and do not buy insurance.

Proposition 1 *If the true value v and the noise term ε are independent, the demand curve overestimates the value of the insured and underestimates the value of the uninsured,*

$$E\left(\varepsilon|\hat{v}\left(\zeta\right)\geq p\right)\geq 0\geq E\left(\varepsilon|\hat{v}\left(\zeta\right)<p\right).$$

The Proposition relies on a simple selection effect; characteristics that affect the decision to buy insurance will be differently represented among the insured and the uninsured. Even if these characteristics average out over the entire population, they do not conditional on the decision to buy insurance. For example, since optimistic beliefs discourage individuals from buying insurance and pessimistic beliefs encourage individuals to buy insurance, those who buy insurance are more likely to be too pessimistic, while those who do not buy insurance are more likely to be too optimistic, even when beliefs are accurate on average. This simple argument has important policy consequences. The selection effect unambiguously signs the mistake *naive* policy makers make by using the demand curve to evaluate welfare consequences of policy interventions targeting either all the insured or uninsured. They overestimate the surplus generated in the insurance market and underestimate the potential value of insurance for the uninsured. As a consequence, universal insurance mandates, often central in the insurance policy debate, are always underappreciated.

Normal Heterogeneity Random noise decreases the correlation between the perceived and true value of insurance and increases the dispersion in the perceived value relative to the dispersion in the actual value. Both a reduction in the correlation and an increase in the relative dispersion decrease the extent to which the true value covaries with the perceived value. For tractability, I only illustrate this here for normal

distributions, allowing for correlation between the true value and the noise term, but I extend this insight for more general distributions in Appendix. Denote the mean value of insurance by $\mu_{\hat{v}} = \mu_v$ and the covariance matrix for the joint distribution by

$$\Sigma = \begin{pmatrix} \sigma_v^2 & \rho_{v,\hat{v}}\sigma_v\sigma_{\hat{v}} \\ \rho_{v,\hat{v}}\sigma_v\sigma_{\hat{v}} & \sigma_{\hat{v}}^2 \end{pmatrix}.$$

Proposition 2 *If the true and perceived value are normally distributed, the demand curve overestimates the value of the insured and underestimates the potential value for the uninsured if and only if* $\rho_{v,\hat{v}} \times \dfrac{\sigma_v}{\sigma_{\hat{v}}} \leq 1$.

The proposition shows that naive policy makers underestimate the value of insurance for the uninsured when either the correlation between the perceived and actual value of insurance or the dispersion in the perceived values is relatively high. When risk perceptions drive the wedge between the actual and perceived value, the correlation is reduced when individuals confound their risk types and thus misperceive who is facing a higher risk. When individuals also exaggerate the differences in their risk types, this increases the dispersion in perceived values relative to the dispersion in true values and thus increases the wedge even further. However, when individuals underappreciate the differences in their actual risks, the wedge is reduced and could even be opposite. Expressing this in terms of noise, the condition states that the result derived under independence is robust as long as the correlation between the noise term and the true value is not too negative. That is, the correlation $\rho_{v,\varepsilon}$ cannot be more negative than the ratio of the standard deviations, $\rho_{v,\varepsilon} \leq -\frac{\sigma_\varepsilon}{\sigma_v}$.

2.3 Marginally Insured

I continue by analyzing the average value of insurance for the marginal buyers, who are indifferent about buying insurance at a price p. From the selection argument before, we expect that, *on average*, people with high perceived value are more likely to overestimate the value of insurance than people with low perceived value. However, to have that higher perceived values *always* signal stronger overestimation of the true values, we require more structure corresponding to the monotone likelihood ratio property (Milgrom 1981).

Proposition 3 *If* $f(\hat{v}|\varepsilon)$ *satisfies the monotone likelihood ratio property,* $\frac{f(\hat{v}_H|\tilde{\varepsilon})}{f(\hat{v}_H|\varepsilon)} \geq \frac{f(\hat{v}_L|\tilde{\varepsilon})}{f(\hat{v}_L|\varepsilon)}$ *for any* $\hat{v}_H \geq \hat{v}_L$, $\tilde{\varepsilon} \geq \varepsilon$, *the difference between the true and perceived value of insurance is increasing in the price,*

$$\frac{\partial}{\partial p} E\left(\varepsilon|\hat{v}\left(\zeta\right) = p\right) \geq 0.$$

The immediate policy implication is that a naive policy maker underestimates the value of an increase in insurance coverage more, the thicker the market is. Moreover,

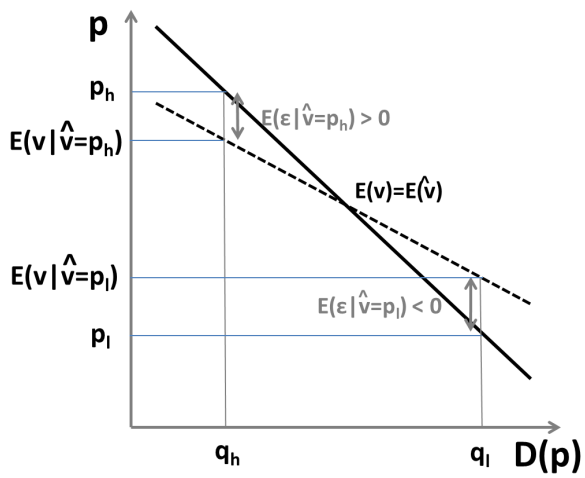

Figure 1: The Demand Curve and the Value Curve.

since the population averages of the perceived and actual value are assumed to be equal, the demand curve and thus the naive policy maker overestimate the true value of additional insurance if and only if the market is sufficiently thin.

Graphically, the Proposition implies that the value curve is a counter-clockwise rotation of the demand curve, as shown in Figure 1. The value curve lies below the demand curve when prices are high and above the demand curve when prices are low, and the difference between the two curves is monotone in the price. If both the perceived and true values are symmetrically distributed, the intersection of the demand and value curve will be exactly where the price equals the median value, which coincides with the average value. The counter-clockwise rotation naturally implies that the area to the left of any q is larger below the demand curve than below the value curve, while to the right of any q it is smaller, which implies Proposition 1. When the true and perceived value are normally distributed, the condition for the value curve to be a counter-clockwise rotation of the demand curve is $\rho_{v,\hat{v}} \times \frac{\sigma_v}{\sigma_{\hat{v}}} \leq 1$, exactly the same as in Proposition 2.

3 Adverse Selection

I now introduce risk heterogeneity into the analysis and consider the supply of insurance contracts. Particular to insurance markets is that an individual's risk type influences not only her demand for insurance, but also the cost to the insurer of providing insurance to that individual. I decompose a type's valuation of insurance in a risk component and a preference component with only the former determining the cost of insuring that type. Following the approach by Einav et al. (2010a), I derive a sufficient statistics formula to evaluate the welfare cost of inefficient pricing due to adverse selection. This formula shows the mistake made by a *naive* policy maker when determining the efficient price and estimating the cost of adverse selection, by ignoring

the non-welfarist heterogeneity underlying the heterogeneous choices.

3.1 Heterogeneity in the Simple Model

The true value of insurance $v(\zeta)$ for an individual with characteristics ζ depends on a risk term, denoted by $\pi(\zeta)$, and a preference term, denoted by $r(\zeta)$,

$$v(\zeta) \equiv \pi(\zeta) + r(\zeta).$$

The risk term also determines the expected cost for the insurance company of providing the insurance contract $c(\zeta) \equiv \pi(\zeta)$. I assume that the risk term is the only determinant of the insurer's cost, ignoring for example potential heterogeneity in loads or enforcement costs.

The model thus captures heterogeneity in three dimensions: heterogeneity in risk types, risk preferences and risk perceptions. Notice that the additivity is not restrictive without restrictions on the distribution of the heterogeneity in the different dimension. Moreover, it would naturally arise in a model when individuals have CARA preferences and face a normally distributed risk x. In this particular case, the actual value of full insurance equals the sum of the expected risk, $\pi(\zeta) = E(x|\zeta)$, and the risk premium, $r(\zeta) = \frac{\eta(\zeta)Var(x|\zeta)}{2}$, where $\eta(\zeta)$ is the individual's parameter of absolute risk aversion. Hence, in the decomposition above, the preference term should be interpreted as the net value of insurance, i.e., the valuation that is not related to the cost of insurance.

3.2 Cost of Adverse Selection

The expected cost of an insurance contract depends on the types who decide to buy the contract. The market is adversely selected when lower risk types are less willing to buy insurance and thus are increasingly willing to buy insurance as the price decreases. The average and marginal cost of providing a contract at price p equal respectively

$$AC(p) = E(\pi|\hat{v}(\zeta) \geq p) \text{ and } MC(p) = E(\pi|\hat{v}(\zeta) = p).$$

Adverse selection results when the marginal cost is an increasing function of the price, in which case the average cost function is increasing and above the marginal cost function. The less an individual's risk affects her insurance choice, the less the marginal cost will depend on the price. If individuals with higher risk are less likely to buy insurance, the market will be *advantageously* selected and the average cost function will be below rather than above the increasing marginal cost function.

In a competitive equilibrium, as Einav et al. (2010a) argue, the equilibrium price p^c equals the average cost of providing insurance given that competitive price,

$$AC(p^c) = p^c.$$

However, this price is (constrained) efficient only if the marginal cost of insurance equals the marginal actual value of insurance.[8] That is, the efficient price p^* solves

$$MC(p^*) = MV(p^*) \;\; (= E(r + \pi | \hat{v}(\zeta) = p)).$$

When the market is adversely selected and the marginal cost is thus below the average cost ($MC(p) < AC(p)$), the equilibrium price is inefficiently high if the price reflects the marginal value of insurance. When the market is sufficiently thick and the demand curve thus underestimates the value of insurance ($p < MV(p)$), the inefficiency is further increased.

The total cost of adverse selection equals the area between the welfare curve and the marginal cost curve - from the equilibrium to the eficient level of insurance,

$$\Gamma = \int_{p^*}^{p^c} [MV(p) - MC(p)] \, dD(p).$$

As shown in the Figure below, the equilibrium level is given by the intersection of the demand and average cost curve, while the efficient coverage level is given by the intersection of the welfare curve and the marginal cost curve. As illustrated by Einav et al. (2010), the demand and cost functions can be estimated using data on insurance choices and claims. However, when the perceived and actual values do not coincide, the demand and cost curves are no longer sufficient to determine the cost of adverse selection. A *naive* policy maker mistakenly beliefs that the efficient price p^n is given by

$$MC(p^n) = p^n$$

and evaluates the inefficiency comparing the wedge between the price and the associated marginal cost. The policy maker thus misestimates this welfare cost as he (1) misidentifies the pool of individuals who should be insured and (2) misestimates the welfare loss for the adversely uninsured,

$$\Gamma = \Gamma^n + \underbrace{\int_{p^*}^{p^n} [MV(p) - MC(p)] \, dD(p)}_{(1)} + \underbrace{\int_{p^n}^{p^c} [MV(p) - p] \, dD(p)}_{(2)},$$

where $\Gamma^n = \int_{p^n}^{p^c} [p - MC(p)] \, dD(p)$ denotes the naively estimated welfare cost of adverse selection. The difference between Γ and Γ^n depends on the thickness of the market ($MV(p) \gtrless p$) and whether the market is adversely selected or not ($p^c \gtrless p^n$). The Figure below illustrates the difference between the actual and naively estimated inefficiency cost for a thick and adversely selected market. The inefficiency is higher than a naive policy maker thinks, both because the extent of underinsurance is worse

[8] In the unconstrained efficient allocation, an individual buys insurance if and only if $r \geq 0$.

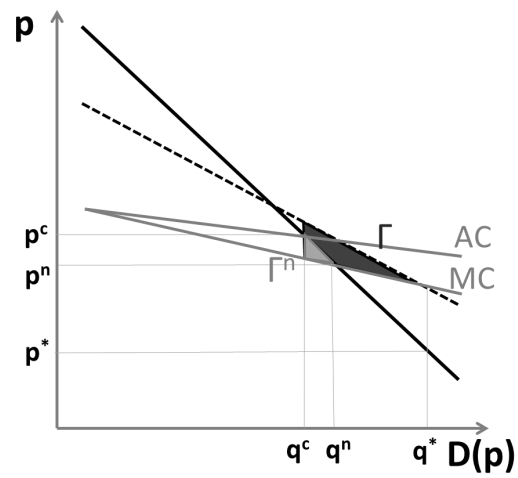

Figure 2: Adverse Selection: the naively estimated cost $\hat{\Gamma}$ vs. the actual cost Γ.

and the welfare loss of underinsurance at a given price is larger than expected.

3.3 Sufficient Statistics Formula

I assume that all three variables determining the perceived value, $\hat{v} = \pi + r + \varepsilon$, are normally distributed. I put no restrictions on the covariance matrix and use notation as before. The expected value of any variable $z \in \{\pi, r, \varepsilon\}$, conditional on the perceived value, equals

$$E\left(z|\hat{v}\left(\zeta\right) = p\right) = \frac{cov\left(z, \hat{v}\right)}{var\left(\hat{v}\right)}\left[p - \mu_{\hat{v}}\right] + \mu_z.$$

The covariance between any the perceived value and any term determining the perceived value indicates how much this term moves with the price. For individuals with the same perceived value, this value can easily be decomposed in the expected value of the risk term, the preference term and the noise term, depending on the respective covariances. If all terms are independent, the covariance of each term with the perceived value is equal to the variance of that term.

The first reason for misestimating the cost of adverse selection is the under- or overestimation of the welfare cost of being uninsured. The size of this misestimation depends on the wedge between the true and perceived marginal surplus of insurance. With normal heterogeneity, this depends on the relative magnitude of the covariances between the perceived value and the risk preference and the perceived value and the noise term,

$$\frac{E\left(\varepsilon|\hat{v} = p\right)}{E\left(r|\hat{v} = p\right) - \mu_r} = \frac{cov\left(\varepsilon, \hat{v}\right)}{cov\left(r, \hat{v}\right)}.$$

The sign of the misestimation depends not only on whether the expected noise is increasing or decreasing in the perceived value ($cov\left(\varepsilon, \hat{v}\right) \gtrless 0$), but also on whether the market is thick or thin ($p \gtrless \mu_{\hat{v}}$) . The second reason for misestimating the cost of

adverse selection is the misidentification of the pool of adversely selected. This mistake depends on the difference between the price that is perceived to be efficient and the price that is actually efficient. With normal heterogeneity, this difference equals

$$p^n - p^* = \frac{cov\left(\varepsilon, \hat{v}\right)}{cov\left(r + \varepsilon, \hat{v}\right)} \mu_r,$$

which again depends on the relative magnitude of the covariances.[9] By linearizing the demand curve through (p^n, q^n) and (p^c, q^c) and linearizing the associated cost and value curve, we obtain the following result.

Proposition 4 *With normal heterogeneity, the bias in welfare cost estimation equals*

$$\frac{\Gamma}{\Gamma^n} \cong \frac{[1 + \frac{cov\left(\varepsilon, \hat{v}\right)}{cov\left(r, \hat{v}\right)} \mathcal{P}]^2}{1 + \frac{cov\left(\varepsilon, \hat{v}\right)}{cov\left(r, \hat{v}\right)}} \text{ where } \mathcal{P} \equiv \frac{\mu_{\hat{v}} - p^n}{p^c - p^n}.$$

The misestimation by a naive policy maker thus crucially depends on the covariance ratio $cov\left(\varepsilon, \hat{v}\right) / cov\left(r, \hat{v}\right)$. This ratio captures the extent to which the variation in demand is explained by noise rather than by preferences. If all terms are independent, the ratio equals the variance in noise relative to the variance in preferences. The impact of a positive covariance ratio depends on the price ratio $\mathcal{P} = \frac{\mu_{\hat{v}} - p^n}{p^c - p^n}$, which is known from the intersections of the demand curve with the average and marginal cost curve, and the value curve. If this price ratio is larger than one, the policy maker unambiguously underestimates the efficiency cost of selection. This is the case if the market is thick and adversely selected, $\mu_{\hat{v}} \geq p^c \geq p^n$, as all adversely uninsured are underestimating the value of insurance on average.[10] For an equilibrium price equal to the average valuation in the market (i.e., $p^c = \mu_{\hat{v}}$) so that half of the market is covered, the misestimation increases approximately linearly in the covariance ratio,

$$\frac{\Gamma}{\Gamma^n} \cong 1 + \frac{cov\left(\varepsilon, \hat{v}\right)}{cov\left(r, \hat{v}\right)}.$$

That is, if the variation in preferences explains 90% of the variation in perceived values, unexplained by variation in risks, and noise explains the remaining 10%, the cost of adverse selection is approximately 11% higher than estimates based on the demand curve suggest. For a thicker market, this bias is larger and increases at faster rate. If the equilibrium price drops below the average valuation to $p^c = \frac{\mu_v + p^n}{2}$, the cost of adverse selection is about 34% higher than naively estimated. The welfare gain of increasing the level of market coverage may thus be much larger than naively estimated if the role of noise is substantial.

[9] Notice that the sign of the price differential is reversed for a thin market, implied by $\mu_r < 0$.

[10] Notice that the price ratio $\mathcal{P}$ is also larger than one if the market is thin and advantageously selected, $\hat{\mu}_v \leq p^n \leq p^n$.

Notice if the market is sufficiently thin ($p^c > p^n \geq \mu_v$), the price ratio $\mathcal{P}$ is smaller than one and the policy maker underestimates the inefficiency cost of selection. The thinner the market, the more of the adversely uninsured are overestimating rather than underestimating the value of insurance, but this reverses if the role of noise is sufficiently important such that p^* decreases below μ_v. Not surprisingly, if the market is thin, but adversely selected, the efficient price may be above the equilibrium price such that it becomes welfare improving to decrease rather than increase the level of market coverage.

3.4 Calibration

In order to assess the potential importance of the bias, I build on the empirical analysis of employer-provided health insurance by Einav, Finkelstein and Cullen (2010), henceforth EFC. Using the health insurance options, choices, and medical insurance claims of the employees of Alcoa, a multi-national producer of aluminium, EFC estimate the demand for additional insurance coverage and the associated cost of providing additional insurance.[11,12] They find that the marginal cost is increasing in the price, but the increase is very small. The increase indicates the existence of adverse selection, but the small magnitude of the increase suggests that very little heterogeneity in insurance choices is explained by heterogeneity in risks. They assume that the residual heterogeneity in insurance choices is due to heterogeneity in (welfarist) preferences and estimate a very small welfare cost of adverse selection, equal to \$9.55 per employee per year, with a 95% confidence interval ranging from \$1 to \$40 per employee. Relative to the average price of \$463.5 - the maximum amount of money at stake - this suggest a welfare cost of only 2.2%. Relative to the estimated surplus at efficient pricing, this suggests a welfare cost of only only 3%.

I relax the assumption that the demand curve reveals the actual value of insurance and use the estimates in EFC to illustrate how welfare conclusions are affected by potentially non-welfarist noise explaining people's insurance choices. Table 1 shows how the bias in the estimate of the welfare cost increase as a function of $cov\left(\varepsilon, \hat{v}\right) / cov\left(r + \varepsilon, \hat{v}\right)$. I apply the formula derived in Proposition 4, which was derived for a linear approximation of the demand and cost curves with normal heterogeneity. Since EFC estimate a linear demand and cost system, the formula is exact when the welfare curve is a rotation of the demand curve like in the case with normal heterogeneity.[13] Using the earlier interpretation, if 1 percent of the residual variation is explained by noise, the actual cost of adverse selection is 3% higher than estimated when using the demand function.

[11] EFC argue that the price variation is exogenous, as business unit managers set the prices for a menu of different health insurance options, offered to all employees within their business unit.

[12] In particular, they consider a sample of 3,779 salaried employees, who chose one of the two modal health insurance choices, where one option is more expensive, but provides more coverage.

[13] I thus assume that the welfare curve has slope $\frac{cov(\pi+\varepsilon,\hat{v})}{var(\hat{v})} p'\left(q\right)$ and crosses the demand curve at $q = 0.5$.

If this share increases to a reasonable 10 percent, the actual cost of adverse selection is already 30 percent higher. If half of the residual variation is explained by noise, the actual cost of adverse selection is more than 4 times higher than estimated based on the demand function. This would imply that rather than \$9.6 per employee per year, the cost of adverse selection would be \$38.5 per employee per year, corresponding to 12 percent of the surplus generated in this market at the efficient price.[14]

4 Policy Interventions

Adverse selection is considered to be the prime reason for governments to intervene in insurance markets. If the cost of inefficient pricing due to adverse selection is small, the net gain from intervening in the insurance market will be small as well. The presence of non-welfarist heterogeneity may not only impact the gain, but also the cost from different policy interventions. In this section, I show that the nature of the heterogeneity driving the demand for insurance and the adverse selection is crucial for the evaluation and comparison of different policy interventions. I discuss insurance subsidies and mandates, information policies and risk-adjustments of insurance premia. To focus the discussion, I assume that the market is adversely selected and the value curve is a counter-clockwise rotation of the demand curve. I continue to assume normal heterogeneity in all dimensions.

4.1 Price vs. Quantity

The two most common interventions to tackle adverse selection in practice are price subsidies and insurance mandates. A price subsidy provides financial incentives encouraging individuals to buy insurance. Individuals still decide, based on their perceived value of insurance, whether or not to buy the contract. While the actual value determines the welfare of a price policy that induces an individual to buy insurance, the perceived value determines how big the price incentives need to be. Hence, inducing an individual to buy insurance through a price policy will be more costly the more the individual underestimates the value of insurance. In contrast, a universal mandate forces an individual to buy insurance, regardless of her perceived value of the contract. In principle, the cost of implementation does not depend on the perceived values.[15] The wedge between the perceived and actual values thus affects the relative effectiveness of price and quantity policies, providing another reason to expect that price and quantity policies are not equivalent (Weitzman 1974).

An increase in the variance of perceived values decreases the value of insurance as perceived by the uninsured in a thick market. The required lower price inducing

[14] Notice that the actual efficient allocation is bounded by complete market coverage. The calculations take this into account.

[15] Notice that people's resistance to a mandate or the benefit of implementing a non-universal mandate will depend on the perceived values as well.

the purchase of insurance for a targeted group of uninsured increases the budgetary cost of that price policy. Consider the clockwise rotation of the demand curve around the equilibrium (q^c, p^c) for invariant cost and value curves. The equilibrium remains unchanged, while the demand rotation decreases the efficient price such that a larger subsidy is required to achieve efficient coverage. The welfare benefit of the price policy is, however, unaffected. For a universal mandate, both the welfare benefit and cost are unaffected. The demand rotation above is obtained when the variance of the perceived values increases, but the correlation with the actual values and the insurance cost remains unchanged. In case of independence, this is simply implied by an increase in the variance of the noise term.

Proposition 5 *In an adversely selected market with equilibrium price $p^c = \mu_v$, an increase in the variance of the perceived values $\sigma^2_{\hat{v}}$, keeping $\rho_{\hat{v},r}$ and $\rho_{\hat{v},\pi}$ unchanged, reduces the net welfare gain of the efficient price policy, while leaving the net welfare gain from a universal mandate unaffected.*

Notice that if the market is thick ($p^c = \mu_v$), the demand rotation increases the equilibrium price as the decreasing average cost curve remains unchanged. This implies that both the welfare benefit and the welfare cost of the efficient price policy increase. The net welfare gain of the price subsidy may thus increase, but will increase by less compared to a universal mandate for which the same higher welfare benefit is realized, but the cost remains unchanged.

Calibration EFC evaluate the welfare gains and losses from different policy interventions, using their estimated demand and cost curves. Table 2 shows how the implied estimates for an efficient price subsidy and a universal mandate would change when the importance of noise increases.[16] To realize the foregone surplus due to adverse selection, an efficient price subsidy needs to reduce the equilibrium price p^c the employees pay to the efficient price p^*. The social cost of such a subsidy is given by $0.3 \times q^* \times (p^c - p^*)$, where 0.3 is a standard estimate of the cost of public funds. EFC estimate the social cost of the efficient price subsidy to be \$45 per employee per year, almost five times as large as the social gain. Interestingly, as the importance of noise increases, the net gain from the efficient price subsidy initially decreases, despite the increasing social gain. The reason is that the social cost of required subsidy increases by even more, as the willingness to pay for insurance of the employee for whom insurance is marginally efficient decreases. In contrast, the net gain from a universal mandate unambiguously increases as the role of noise increases in importance. Ignoring the role of noise leads to two mistakes when evaluating the net welfare gain from a universal mandate; first by understating the welfare gain Γ from forcing those for whom insurance is efficient

[16] Notice that we take the estimated demand curve as given, but consider a counter-clockwise rotation of the associated value curve. In the Proposition, we keep the value curve unchanged, but consider a clockwise rotation of underlying demand curve.

to buy insurance and second by overstating the welfare cost Φ from forcing employees for whom insurance is inefficient to buy insurance. In fact, if more than 17 percent of the unexplained heterogeneity is due to noise, a universal mandate would be welfare increasing, while the estimates of EFC imply that a universal mandate is welfare decreasing.

4.2 Information Policies

When individuals take misguided choices, a natural government intervention is to provide them with information to reduce or overcome their mistakes. However, such information policy does not unambiguously increase welfare as the provision of information may affect the pool of insured and thus the equilibrium price. While an individual is always better off when her perceived valuation is closer to her actual valuation for a given price, the higher price when more costly individuals buy insurance reduces coverage in equilibrium. This implies an important trade-off for information policies. Better information induces people with high actual rather than high perceived value to buy insurance, but for welfare purposes the key concern is whether people with high net-value buy insurance. If better informed types with more costly risks tend to buy insurance due to the information policy, the market becomes more adversely selected and welfare decreases. If types with higher risk preference tend to buy insurance due to the information policy, the market selects individuals with higher net-value and welfare increases.

To disentangle these two effects, I consider two information policies; a first policy that increases the correlation between the actual risk π and the *perceived* risk $\hat{\pi} \equiv \pi + \varepsilon$, a second policy that increases the correlation between the actual preference r and the *perceived* preference $\hat{r} \equiv r + \varepsilon$. In both cases, everything else remains unchanged.[17] As the variance in the perceived values is unaffected by the information policies, the same number of individuals buy insurance at a given price, but the selection of individuals buying insurance will depend on the policy.

The first policy will induce individuals with high risk π rather than individuals with high perceived risk $\hat{\pi}$ to buy insurance. The average expected cost of the individuals buying insurance at a given price level increases, which increases the equilibrium prices as the demand function is unaffected. However, the expected net-value of the individuals buying insurance at a given price is still the same. The same welfare surplus is generated for those buying insurance, but less individuals buy insurance. Hence, welfare is unambiguously lower.

[17] An alternative interpretation is that the information policy reduces the variance in the noise term, where the noise term is independent of the one term x, but negatively related to the other term y (i.e., $\rho_{\varepsilon,x} = 0$ and $\rho_{\varepsilon,y} = -\frac{1}{2}\frac{\sigma_\varepsilon}{\sigma_y}$ for $x = r, \pi$, $y = \pi, r$). In this interpretation, ε could be interpreted as a misperception of y where the dispersion of the perceived and true y is the same.

Proposition 6 *An information policy that increases the correlation between the actual and perceived risks, ceteris paribus, unambiguously reduces welfare.*

The second policy has an opposite effect. While the same number of individuals buy insurance, a higher welfare surplus is generated for those buying insurance. The information policy induces people with a high net-value r to buy insurance, but the equilibrium price remains unchanged as the expected cost of the individuals buying insurance is not affected. Hence, welfare unambiguously increases.

Proposition 7 *An information policy that increases the correlation between the actual and perceived preferences, ceteris paribus, unambiguously increases welfare.*

Better information induces people to make better decisions, but increases the scope for adverse selection. This potential trade-off can be relaxed by providing the right type of information. Information regarding the cost-related value of information will be detrimental, as it only affects the market price, while information regarding the net-value of insurance will be beneficial, as it only affects the selection of the insured pool. Notice that next to an individual's risk preference, the variance of the risk they are facing also affects their value of insurance, without affecting the expected cost for a risk-neutral insurance company. For instance, in the earlier case of CARA-preferences and normal risk, the net-value of insurance equals the risk premium, which depends on both the risk-aversion and the variance of the normal risk. Hence, while providing information about one's expected risk may be a bad idea, providing information about the variance of the risk one is facing may be a very good idea. Similar recommendations apply for policies that tackle other factors driving a wedge between the perceived and actual value. If liquidity constraints prevent individuals from buying insurance, a policy that helps these individuals overcome their liquidity constraints will always be welfare-improving, unless the liquidity-constrained have particularly high risk-types.

Calibration [TBC] I again use the empirical analysis in EFC to evaluate the potential impact of information policies. While in the previous calibrations the estimated demand and cost functions were taken as given, information policies will change the selection of employees buying insurance contracts. Like with normal heterogeneity, I assume that the marginal cost curve has slope $\frac{cov(\pi,\hat{v})}{var(\hat{v})}p'(q)$, but is linear as estimated by EFC. We start from a benchmark case where $\frac{cov(\varepsilon,\hat{v})}{cov(r+\varepsilon,\hat{v})} = 0.10$.

4.3 Risk-Adjusted Pricing

The uniform price that individuals pay for insurance drives a wedge between the average and marginal cost borne by the insurer, causing the equilibrium price to be inefficient. Adjusting the insurance premium to reflect an individidual's risk could reduce adverse selection and thus increase efficiency. The adjustment, however, introduces inequality

between higher and lower risks, since the former will face higher prices for the same insurance contract. In practice, risk-adjusted pricing is often regulated, like the recent ban on gender discrimination in insurance pricing by the European Court of Justice. The policy maker faces a difficult trade-off between efficiency and equity. However, the efficiency gain from adjusting premia to individuals' risks crucially depends on the individuals appreciating those differences. If not, the risk-adjustment will lower the net surplus rather than increase the net surplus generated in equilibrium.

Consider the adjustment $\alpha(\pi)$ to the insurance premium for an individual with risk π. Since the average risk is trivially reflected in equilibrium price, the risk-adjustment will depend on an individual's risk relative to the average risk, with $\alpha(\pi)$ weakly increasing in π and equal to 0 if $\pi = \mu_\pi$. Perfect risk-adjusted pricing is achieved when $\alpha(\pi) = \pi - \mu_\pi$, but in general, the risk-adjustment can be only based on observable dimensions of the risk (e.g., $\pi = \pi_{unobs.} + \pi_{obs.}$ and $\alpha(\pi) = \pi_{obs.} - \mu_{\pi_{obs.}}$). An individual with characteristics ζ buys insurance if and only if

$$\hat{v}(\zeta) \geq p + \alpha(\pi(\zeta)) \Longleftrightarrow \tilde{v}(\zeta) \geq p$$

where $\tilde{v}(\zeta)$ denotes the perceived value of insurance adjusted for the risk adjustment. The cost for the insurer, accounting for the risk adjustment of the premium now equals

$$\begin{aligned} \widetilde{AC}(p) &= E(\pi - \alpha(\pi) | \tilde{v}(\zeta) \geq p). \\ \widetilde{MC}(p) &= E(\pi - \alpha(\pi) | \tilde{v}(\zeta) = p). \end{aligned}$$

Given these adjusted expressions, we can apply the equilibrium and welfare analysis from before.

The risk-adjustment of the insurance price affects both the equilibrium price and the selection of individuals buying insurance at the new equilibrium price. Pricing the risk or part of the risk mechanically reduces the difference between the average cost of providing insurance and the marginal cost, conditional on the demand for insurance. That is, the difference between the unpriced risk among the insured and the unpriced risk for the marginal individual is reduced. With perfect risk-adjusted pricing, the average cost becomes independent of the price p and coincides with the marginal cost curve. The resulting equilibrium price equals the average risk, which is (constrained) efficient. Pricing the risk also makes high risk types less likely to buy insurance and low risk types more likely to buy insurance. This selection change lowers the average cost curve and thus the equilibrium price even further. However, the change also affects the surplus generated at a given price p,

$$E_\zeta(r | \tilde{v} \geq p).$$

Intuitively, the insurance surplus will be higher the less any heterogeneity other than

the heterogeneity in preferences drives the demand for insurance. That is, the surplus is higher the more the preference term r covaries with the risk-adjusted perceived value $\tilde{v}$, $\frac{cov(r,\tilde{v})}{var(\tilde{v})}$. The risk type π does not affect the net value of insurance, as it also determines the cost of providing insurance. Hence, when the preference term is independently distributed, reducing the role that one's risk plays in the decision to buy insurance increases the equilibrium surplus. The problem is that when perceived and actual risks differ, risk-adjusted pricing does not decrease this role as expected and, in fact, may even increase this role. It only increases the surplus if

$$var\left(\pi+\varepsilon\right) \geq var\left(\pi+\varepsilon-\alpha\left(\pi\right)\right).$$

With perfect risk-adjusted pricing, this condition simplifies to $\rho_{\varepsilon,\pi} > -\frac{1}{2}\frac{\sigma_\pi}{\sigma_\varepsilon}$. Hence, if the correlation between the risk and noise term is sufficiently negative, the introduction of risk-adjusted pricing reduces the consumer surplus at any given price. The following Proposition considers two extreme cases to evaluate the effect on the consumer surplus in equilibrium.

Proposition 8 *When the perceived and actual risk types coincide $\hat{\pi} = \pi$, perfect risk-adjusting pricing unambiguously increases the equilibrium surplus. Without heterogeneity in perceived risks $\hat{\pi} = E\left(\pi\right)$, perfect risk-adjusted pricing unambiguously decreases the equilibrium surplus.*

Consider first the case with accurate risk perceptions. When facing a uniform price, an individual buys insurance if $\pi + r \geq p$. When the price instead reflects the individual's risk $p + \pi - \mu_\pi$, her decision only depends on her risk preference. In equilibrium, the price equals the average cost, $p = \mu_\pi$, and an individual only buys insurance if $r \geq 0$, which is exactly efficient. When an individual does not appreciate her particular risk type, she will buy insurance if $r \geq p - \mu_\pi$. Again the equilibrium price equals $p = \mu_\pi$ and an individual only buys insurance if $r \geq 0$. However, when the price reflects her risk, she will only buy insurance if $r - \pi \geq p - \mu_\pi$. The risk-adjustment thus introduces the inefficiency that it is supposed to eliminate. The two considered cases are extreme, but make the policy implication very clear. The impact of risk-adjusted policies very much depend on how these risks are perceived.

Calibration [TBC] I simulate the new equibrium when prices are adjusted for observable risks (e.g., male vs. female) and contrast the welfare conclusions depending on the difference between the actual and perceived risks.

5 Discussion

The analysis shows that the welfare cost of adverse selection and the effectiveness of potential government interventions depend crucially on the sources of heterogeneity

driving differences in the demand for insurance. In this section, I discuss the robustness of the welfare results and some empirical studies in light of this welfare analysis.

Robustness The assumption that heterogeneity in risk perceptions perfectly reflects heterogeneity in risk types seems strong, even for rational individuals. Rationality may restrict individuals to be Bayesian, but it puts no restrictions on priors themselves, which are primitives of the model (Van Den Steen 2004). While one's prior expectation $\hat{\pi}$ is likely to be different from one's actual expected risk π, learning about one's risk will increase the correlation. However, learning may well be incomplete and any correlation lower than one induces the average actual value of insurance to be lower than the average perceived value of insurance above any price. This suggests that the discussed bias in evaluating infra-marginal policies is robust.

The wedge is also affected by the relative dispersion of the perceived and actual risk types. The wedge is further increased when the perceived expected risks are more dispersed than the actual expected risks, while it is reduced and potentially reversed if the perceived risks are less dispersed than the actual risks. In Appendix, I extend the insights that both a decreased correlation and increase in the relative dispersion induce a wedge in the same direction for general, discrete distributions of risk types. A potential model of the relation between perceived and actual risk types has $\hat{\pi} = \mu_\pi + \alpha(\pi - \mu_\pi)$.[18] This models allows the dispersion in the perceived risk types to be too large ($\alpha > 1$) or too small ($\alpha < 1$), such that the sign of the wedge would only depend on the sign of $\alpha - 1$, since $\frac{cov(\pi,\hat{\pi})}{var(\hat{\pi})} = \frac{1}{\alpha}$. Notice that this model assumes perfect correlation between the perceived and actual risk types. With imperfect correlation, the dispersion in perceived risk should be sufficiently smaller ($\alpha << 1$) to reverse the wedge. The extent to which actual risks do correlate with perceived risks can be analyzed empirically. In general, estimating one's actual risk type is challenging. However, when observing for each individual a realization of that individual's risk x_i, which is a draw from F_{π_i}, and her perceived risk $\hat{\pi}_i$, the statistic of interest $\frac{cov(\pi,\hat{\pi})}{var(\hat{\pi})}$ is estimated by regressing x on $\hat{\pi}$. Some empirical studies, discussed in Hurd (2009) find a significant positive relation between risk realizations and individual's surveyed risk perceptions. However, the estimated relation is small.[19] For example, Finkelstein and McGarry (2006) find estimates smaller than 0.10 when estimating a linear probability model of whether an individual went into a nursing home in the five years between 1995 and 2000 on his 1995 self-reported beliefs of this probability. The self-reported probabilities have some predictive value, but an increase in the risk perceptions is associated with a much smaller average increase in the actual risk realizations. Unfortunately, the

[18] This specification is considered in considered in Einav et al. (2011). One could interpret this model as coming from the updating of a uniform prior, equal to the population distribution, using individual signals. The updating is non-Bayesian if indeed $\alpha \neq 1$.

[19] For an empirical analysis of the relation between subjective life expectations and survival probabilities, see for instance (Hamermesh 1985, Hurd and McGarry 1995 and 2002).

self-reported probability does not measure the perceived probability $\hat{\pi}$ that affects the insurance demand without error, which would also attenuate the regression estimate below 1.[20,21]

Rotations vs. Shifts The analysis assumes that the population averages of the actual value of insurance and the value as revealed by the demand are the same. Regarding risk perceptions, various studies suggest that people are too optimistic or too pessimistic on average, depending on the context, the size of the probability, the own control, etc (see Tversky and Kahneman 1974, Slovic 2000, Weinstein 1980, 1982 and 1984). The presence of an average wedge between actual and perceived values does not affect the nature of the insights regarding the impact of heterogeneity. While heterogeneity affects how the wedge changes a long the demand curve, the sign of the wedge is relevant for evaluating policy interventions. For instance, heterogeneous risk perceptions induce the uninsured to be more optimistic than the average individual. However, if the average individual is too optimistic, the underappreciation of insurance for the uninsured will be smaller. Graphically, heterogeneity in perceptions induces a rotation of the value curve relative to the demand curve, while an average optimistic or pessimistic bias introduces a shift and thus changes the intersection of the demand and the value curve.[22] Similarly, if an individual does not buy or switch insurance contracts because of liquidity constraints or inertia, the demand curve will underestimate the actual value of insurance. Heterogeneity in liquidity constraints or inertia causes the bias to be particularly large for those not buying insurance relative to those buying insurance.

Empirical Implementation The analysis shows that disentangling the sources of heterogeneity driving the insurance demand is essential, but this may be challenging. The empirical evidence that risk plays a minor role in explaining insurance coverage has led to some recent studies analyzing the role of heterogeneity in risk preferences. Some papers use the heterogeneity in insurance choices, left unexplained by heterogeneity in risks, to estimate heterogeneity in risk preferences (for example Cohen and Einav 2007), but few papers use explicit measures of risk preferences to explain insurance choices, like Cutler et al. (2008). An overview of the empirical literature by Cohen and Spiegelman (2009) suggests that also these preference measures explain only a minor part of the variation in insurance demand. While measuring the heterogeneity in preferences across individuals is of course difficult, Barseghyan et al. (2011) and Einav

[20] Notice that Finkelstein and McGarry (2006) find a positive relationship between the self-reported probability and insurance coverage, but no significant relationship between the actual risk and insurance coverage.

[21] Kircher and Spinnewijn (2011) suggest an alternative approach to use price variation to disentangle perceived risks from risk preferences.

[22] Also changes in the symmetry of the distribution of the different components of the perceived value would make that demand and value curve do not exactly interact at the median perceived value.

et al (2011b) analyze individuals' choices across different insurance domains. These analyses suggest again that the role of an individual-specific, but domain-general component is modest at most.[23] While these findings seem suggestive, they are definitely not sufficient to conclude that choices do not reveal an individual's preferences and that $cov\left(r,\hat{v}\right)/var\left(\hat{v}\right)$ would be small. Empirical work is needed to provide direct evidence of the role of non-welfarist heterogeneity and the potential importance of $cov\left(\varepsilon,\hat{v}\right)/var\left(\hat{v}\right)$. One natural approach is to identify a variable which is considered not to affect the actual value of insurance, but can be shown to affect the insurance demand. For example, Fang et al. (2008) find that cognitive ability is a strong predictor of Medigap insurance coverage, while cognitive ability is unlikely to be related to the actual value of Medigap insurance. In general, wealth, income and education are often found to be predictors of insurance purchases. While it is harder to argue that these variables are unrelated to the true value of insurance, empirical evidence (Choi et al. 2011) suggests that they are strongly related to the quality of the decisions. In a similar spirit, one could try to identify a selection of individuals for which the importance of non-welfare noise is expected to be smaller to uncover the value function associated with the observed demand. An alternative approach relates to the estimated relation between actual and perceived risks $cov\left(\pi,\hat{\pi}\right)/var\left(\hat{\pi}\right)$.[24] Notice that in combination with the price variation used in EFC to estimate the relation between the insurance demand and the risk realizations, $cov\left(\pi,\hat{v}\right)/var\left(\hat{v}\right)$, one could recover the noise in perceptions underlying the demand for insurance,

$$\frac{cov\left(\pi,\hat{v}\right)}{var\left(\hat{v}\right)}\Big/\frac{cov\left(\pi,\hat{\pi}\right)}{var\left(\hat{\pi}\right)}=\left[\frac{cov\left(\hat{\pi},\hat{v}\right)}{var\left(\hat{v}\right)}-\frac{cov\left(\hat{\pi},r\right)}{var\left(\hat{v}\right)}\right]\times\left[\frac{cov\left(\pi,\hat{\pi}\right)}{cov\left(\pi,\hat{\pi}\right)}+\frac{cov\left(\pi,r\right)}{cov\left(\pi,\hat{\pi}\right)}\right].$$

If the covariance between risk preferences and risk perceptions is small, we thus find

$$\frac{cov\left(\pi,\hat{v}\right)}{var\left(\hat{v}\right)}\Big/\frac{cov\left(\pi,\hat{\pi}\right)}{var\left(\hat{\pi}\right)}-\frac{cov\left(\pi,\hat{v}\right)}{var\left(\hat{v}\right)}\cong\frac{cov\left(\varepsilon,\hat{v}\right)}{var\left(\hat{v}\right)}.$$

6 Conclusion

What drives the heterogeneity in the demand for insurance? This difficult question has been central in a recent, but already prominent empirical literature. The question remains open, but has led to a seeming contradiction. While a number of recent empirical studies suggest that what drives the selection into insurance contracts is often unrelated to the actual value of these contracts (Fang et al. 2008, Abaluck and Gruber 2011, Handel 2011), the papers analyzing the importance of adverse selection in insurance markets, have mostly evaluatesd potential government interventions under

[23] Surprisingly, the individuals' (predicted) risks in the different domains again seem to play only a minor role in explaining the insurance choices.

[24] An alternative to measuring perceived risks is to provide information about risks in a controlled experiment and analyze the effect on the demand for insurance and the associated costs.

the assumption that individual's choices reveal the actual value of insurance. This paper provides a simple framework to analyze the consequences of heterogeneity in the differences of the actual and revealed value of insurance. The analysis presents a simple selection argument that shows that even without an average bias in the valuation, the welfare conclusions will be systematically biased. Not only the welfare cost of adverse selection, but also the relative welfare gains from standard policy intervention in insurance markets depends on the source of the heterogeneity undering the demand for insurance. A calibration exercise illustrates that for small (and thus plausible) differences between the actual and perceived value of insurance, the policy conclusions will be substantially different.

7 References

Abaluck, J., and J. Gruber, 2011. Choice Inconsistencies among the Elderly: Evidence from Plan Choice in the Medicare Part D Program. American Economic Review. forthcoming

Akerlof, G., 1970. The Market for "Lemons": Quality Uncertainty and the Market Mechanism. The Quarterly Journal of Economics, 84(3), 488-500.

Arrow, K. J., 1963. Uncertainty and the Welfare Economics of Medical Care. The American Economic Review, 53(5), 941–973.

Barseghyan, L., Molinari, F., O'Donoghue, T., and J. Teitelbaum, 2011. The Nature of Risk Preferences: Evidence from Insurance Choices. mimeo

Barseghyan, L., Prince, J., and J. Teitelbaum, 2011. Are Risk Preferences Stable Across Contexts? Evidence from Insurance Data. American Economic Review, 101(2), 591-631.

Bernheim, D., and A. Rangel, 2009. Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics. Quarterly Journal of Economics, 124(1), 51-104.

Bruhin, A., Fehr-Duda, H. and T. Epper, 2010. Risk and Rationality: Uncovering Heterogeneity in Probability Distortion. Econometrica, 78(4), 1375-1412.

Chetty, R., 2009. Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods. Annual Review of Economics, 1, 451-488.

Chetty, R., Kroft, L., and A. Looney 2009. Salience and Taxation: Theory and Evidence. American Economic Review, 99(4), 1145-1177.

Chiappori, P., and B. Salanié, 1997. Empirical Contract Theory: The Case of Insurance Data. European Economic Review, 41(3), 943-950.

Chiappori, P., and B. Salanié, 2000. Testing for asymmetric information in insurance markets. Journal of Political Economy, 108(1), 56-78.

Chiappori, P.-A., Jullien, B., Salanié, B., Salanié, F., 2006. Asymmetric Information in Insurance: General Testable Implications. Rand Journal of Economics 37,

783-798.

Choi, S., Kariv, S., Muller, W. and D. Silverman, 2011. Who is (More) Rational? mimeo

Cohen, A., Einav, L., 2007. Estimating Risk Preferences from Deductible Choice. American Economic Review 97(3), 745-788.

Cohen, A., Siegelman, P., 2010. Testing for Adverse Selection in Insurance Markets. Journal of Risk and Insurance 77(1), 39-84.

Congdon, W., Kling, J. and S. Mullanaithan 2011. Policy and Choice: Public Finance Through the Lens of Behavioral Economics. Brookings Institute.

Cutler, D., Finkelstein, A., and K. McGarry, 2008. Preference Heterogeneity and Insurance Markets: Explaining a Puzzle of Insurance. American Economic Review Papers and Proceedings.

Cutler, D., and R., Zeckhauser, 2004. Extending the Theory to Meet the Practice of Insurance. In R. Litan and R. Herring, eds., Brookings-Wharton Papers on Financial Service. Washington, D.C.: Brookings Institution.

Einav, L., Finkelstein, A., Cullen, M., 2010a. Estimating Welfare in Insurance Markets using Variation in Prices. Quarterly Journal of Economics 125(3), 877-921.

Einav, L., Finkelstein, A., Schrimpf, P., 2010b. Optimal Mandates and The Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market. Econometrica, 78(3), 1031-1092.

Einav, L., Finkelstein, A., Levin, J., 2010c. Beyond Testing: Empirical Models of Insurance Markets. Annual Review of Economics 2, 311-336.

Fang, H., Keane, M. 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. Multiple dimensions of private information: Evidence from the long-term care insurance market. American Economic Review, 96(4), 938–958.

Hamermesh, D., 1985. Expectations, Life Expectancy, and Economic Behavior. Quarterly Journal of Economics, 100(2), 389-408.

Hurd, M. and K. McGarry, 1995. Evaluations of the Subjective Probability of Survival in the Health and Retirement Survey. Journal of Human Resources, 30, 268-292.

Hurd, M. and K. McGarry, 2002. The Predictive Validity of Subjective Probabilities of Survival. Economic Journal, 112, 966-985.

Hurd, M., 2009. Subjective Probabilities in Household Surveys. Annual Review of Economics, 1, 543-562.

Rothschild, M., Stiglitz, J., 1976. Equilibrium in Competitive Insurance markets: An Essay on the Economics of Imperfect Information. Quarterly Journal of Economics 90, 630-649.

Slovic, P., 2000. The Perception of Risk. London: Earthscan.

Snowberg, E., and J. Wolfers, 2010. Explaining the Favorite–Long Shot Bias: Is it Risk-Love or Misperceptions? Journal of Political Economy, 118(4), 723-746.

Spinnewijn, J., 2010a. Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs. LSE mimeo.

Spinnewijn, J., 2010b. Insurance and Perceptions: How to Screen Optimists and Pessimists. LSE mimeo.

Sydnor, J., 2010. (Over)insuring Modest Risks. American Economic Journal: Applied Economics, 2(4), 177–99.

Tversky, A., Kahneman, D., 1974. Judgment under Uncertainty: Heuristics and Biases. Science, 185, 1124-1131.

Weinstein, N., 1980. Unrealistic Optimism about Future Life Events. Journal of Personality and Social Psychology 39(5), 806-820.

Weinstein, N., 1982. Unrealistic Optimism about Susceptibility to Health Problems. Journal of Behavioral Medicine 5(4), 441-460.

Weinstein, N., 1984. Why It Won't Happen to Me: Perceptions of Risk Factors and Susceptibility. Health Psychology 3 (5), 431-457.

Weitzman, M. 1974. Prices vs. Quantities. Review of Economic Studies 41(4), 477-491.

8 Tables

Table 1: Different measures of the cost of adverse selection as a function of the noise ratio.

Noise Ratio	**Cost of Adverse Selection**		
$\frac{cov(\varepsilon, \hat{v})}{cov(\varepsilon + r, \hat{v})}$	Actual Cost in \$ / indiv. Γ	Bias in Welfare Cost Γ/Γ^n	Actual Relative Cost $\Gamma/[\textit{Eff. Surplus}]$
0	9.6	1	.03
.01	9.8	1.03	.03
.10	12.5	1.30	.04
.25	18.7	1.95	.06
.50	38.5	4.03	.12
.75	67.7	7.70	.22
1	96.9	10.2	.31

Table 2: The welfare gain of government interventions as a function of the noise ratio.

Noise Ratio	**Government Interventions**	
$\frac{cov(\varepsilon, \hat{v})}{cov(\varepsilon + r, \hat{v})}$	Efficient Price Subsidy (in \$/ind.) $\Gamma - [0.3 \times q^* \times (p^c - p^*)]$	Universal Mandate (in \$/ind.) $\Gamma - \Phi$
0	-35.4	-8.8
.01	-36.2	-8.4
.10	-43.8	-4.3
.25	-62.2	6.7
.50	-125.6	45.7
.75	-96.4	52.4
1	-67.2	59.1

9 Appendix A: Proofs

Proof of Proposition 1

I assume that the random variables are draws from continuous distributions. The proof trivially generalizes for discrete distributions. Denote the density functions of $\hat{v}, v$ and ε by $f(\hat{v}), h(v)$ and $g(\varepsilon)$ respectively. Since by Bayes' law $g(\varepsilon|\hat{v}) = \frac{f(\hat{v}|\varepsilon)g(\varepsilon)}{f(\hat{v})}$, we can rewrite

$$\begin{aligned} g(\varepsilon|\hat{v} \geq p) &= \frac{\int_p g(\varepsilon|\hat{v})\, d\hat{v}}{\int_p f(\hat{v})\, d\hat{v}} \\ &= \frac{\int_p f(\hat{v}|\varepsilon)\, g(\varepsilon)\, d\hat{v}}{\int_p f(\hat{v})\, d\hat{v}} = \frac{\Pr(\hat{v} \geq p|\varepsilon)}{\Pr(\hat{v} \geq p)} g(\varepsilon), \end{aligned}$$

with $\int \frac{\Pr(\hat{v} \geq p|\varepsilon)}{\Pr(\hat{v} \geq p)} g(\varepsilon) = 1$. Moreover, since v and ε are independent, we have that $\Pr(\hat{v} \geq p|\varepsilon) = \int_{p-\varepsilon} h(v)\, dv$ is increasing in ε. Hence, the conditional distribution of $\varepsilon|\hat{v} \geq p$ first-order stochastically dominates the unconditional distribution of ε and thus

$$E(\varepsilon|\hat{v} \geq p) = \int \varepsilon g(\varepsilon) \frac{\Pr(\hat{v} \geq p|\varepsilon)}{\Pr(\hat{v} \geq p)} d\varepsilon \geq \int \varepsilon g(\varepsilon)\, d\varepsilon = E(\varepsilon) = 0.$$

Similarly, we find

$$E(\varepsilon|\hat{v} \leq p) = \int \varepsilon g(\varepsilon) \frac{\Pr(\hat{v} \leq p|\varepsilon)}{\Pr(\hat{v} \leq p)} d\varepsilon \leq \int \varepsilon g(\varepsilon)\, d\varepsilon = E(\varepsilon) = 0.$$

Proof of Proposition 2

By normality, we have

$$\begin{aligned} E(\hat{v}|\hat{v} \geq p) - E(v|\hat{v} \geq p) &= \mu_{\hat{v}} - \mu_v + \sigma_{\hat{v}} \frac{\phi\left(\frac{p-\mu_{\hat{v}}}{\sigma_{\hat{v}}}\right)}{1 - \Phi\left(\frac{p-\mu_{\hat{v}}}{\sigma_{\hat{v}}}\right)} - \sigma_v \rho \frac{\phi\left(\frac{p-\mu_{\hat{v}}}{\sigma_{\hat{v}}}\right)}{1 - \Phi\left(\frac{p-\mu_{\hat{v}}}{\sigma_{\hat{v}}}\right)} \\ &= [\sigma_{\hat{v}} - \sigma_v \rho] \frac{\phi\left(\frac{p-\mu_{\hat{v}}}{\sigma_{\hat{v}}}\right)}{1 - \Phi\left(\frac{p-\mu_{\hat{v}}}{\sigma_{\hat{v}}}\right)}. \end{aligned}$$

Hence, $E(\hat{v}|\hat{v} \geq p) \geq E(v|\hat{v} \geq p)$ iff $\sigma_{\hat{v}} \geq \sigma_v \rho$.

Proof of Proposition 3

This is an immediate application of Proposition 1 in Milgrom (1981). That is,

$$\int \varepsilon g(\varepsilon|\hat{v}_H)\, d\varepsilon \geq \int \varepsilon g(\varepsilon|\hat{v}_L)\, d\varepsilon \text{ for any } \hat{v}_H \geq v_L$$

iff

$$\frac{f(\hat{v}_H|\tilde{\varepsilon})}{f(\hat{v}_H|\varepsilon)} \geq \frac{f(\hat{v}_L|\tilde{\varepsilon})}{f(\hat{v}_L|\varepsilon)} \text{ for any } \tilde{\varepsilon} \geq \varepsilon.$$

Hence, the expected value of noise, conditional on the perceived value, is increasing in

the perceived value.□

Proof of Proposition 4

The perceived cost of adverse selection equals

$$\begin{aligned}\Gamma^n &= \int_{p^n}^{p^c} [p - MC(p)]\, dD(p) \\ &= \int_{p^n}^{p^c} \left[p - \frac{cov(\pi, \hat{v})}{var(\hat{v})}[p - \mu_v] - \mu_\pi\right] dD(p),\end{aligned}$$

where $p = MC(p)$ evaluated at $p = p^n$. Hence, the perceived cost of adverse selection is equal to the area between two proportional functions, relative to p^n. Now linearizing the demand function, (i.e., assuming that the density at each price level is the same and equal to $\bar{f}$), this is approximately equal to

$$\begin{aligned}\Gamma^n &\cong \left(1 - \frac{cov(\pi, \hat{v})}{var(\hat{v})}\right)[p^c - p^n]^2 \frac{\bar{f}}{2} \\ &= \frac{cov(r + \varepsilon, \hat{v})}{var(\hat{v})}[p^c - p^n]^2 \frac{\bar{f}}{2}.\end{aligned}$$

A similar argument allows to approximate the actual cost of adverse selection,

$$\begin{aligned}\Gamma &= \int_{p^*}^{p^c} [MV(p) - MC(p)]\, dD(p) \\ &= \int_{p^*}^{p^c} \left[\frac{cov(\pi + r, \hat{v})}{var(\hat{v})}[p - \mu_v] + \mu_v - \frac{cov(\pi, \hat{v})}{var(\hat{v})}[p - \mu_v] - \mu_\pi\right] dD(p) \\ &\cong \frac{cov(r, \hat{v})}{var(\hat{v})}[p^c - p^*]^2 \frac{\bar{f}}{2}.\end{aligned}$$

Hence, the ratio equals

$$\begin{aligned}\frac{\Gamma}{\Gamma^n} &\cong \frac{cov(r, \hat{v})}{cov(r + \varepsilon, \hat{v})} \frac{[p^c - p^*]^2}{[p^c - p^n]^2} \\ &= \frac{cov(r, \hat{v})}{cov(r + \varepsilon, \hat{v})}\left[1 + \frac{p^n - p^*}{p^c - p^n}\right]^2.\end{aligned}$$

Now we still want to substitute for the unobservable p^*. Notice that at $p = \mu_v$,

$$p - MC(p) = MV(p) - MC(p).$$

From the linearisation from before, we know that

$$\begin{aligned}p - MC(p) &\cong \frac{cov(r + \varepsilon, \hat{v})}{var(\hat{v})} \frac{1}{\bar{f}}[D(\mu_v) - D(p^n)] \\ MV(p) - MC(p) &\cong \frac{cov(r, \hat{v})}{var(\hat{v})} \frac{1}{\bar{f}}[D(\mu_v) - D(p^*)].\end{aligned}$$

Hence,

$$\frac{cov\left(r+\varepsilon,\hat{v}\right)}{var\left(\hat{v}\right)}\left[D\left(\mu_v\right)-D\left(p^n\right)\right]\cong\frac{cov\left(r,\hat{v}\right)}{var\left(\hat{v}\right)}\left[D\left(\mu_v\right)-D\left(p^*\right)\right]$$

and thus, linearizing the demand function,

$$cov\left(r+\varepsilon,\hat{v}\right)\left[\mu_v-p^n\right]\cong cov\left(r,\hat{v}\right)\left[\mu_v-p^*\right].$$

Rearranging, we find

$$\left[p^n-p^*\right]\cong\frac{cov\left(\varepsilon,\hat{v}\right)}{cov\left(r,\hat{v}\right)}\left[\mu_v-p^n\right].$$

Hence,

$$\begin{aligned}\frac{\Gamma}{\Gamma^n} &\cong \frac{cov\left(r,\hat{v}\right)}{cov\left(r+\varepsilon,\hat{v}\right)}\left[1+\frac{cov\left(\varepsilon,\hat{v}\right)}{cov\left(r,\hat{v}\right)}\frac{\mu_v-p^n}{p^c-p^n}\right]^2\\ &= \frac{\left[1+\frac{cov(\varepsilon,\hat{v})}{cov(r,\hat{v})}\frac{\mu_v-p^n}{p^c-p^n}\right]^2}{1+\frac{cov(\varepsilon,\hat{v})}{cov(r,\hat{v})}}.\end{aligned}$$

□

Proof of Result 5

The inverse demand function equals

$$D^{-1}\left(q\right)=\sigma_{\hat{v}}\Phi^{-1}\left(1-q\right)+\mu_{\hat{v}}.$$

Since $D^{-1}\left(0.5\right)=\mu_{\hat{v}}$, an increase in $\sigma_{\hat{v}}$ rotates the demand function clockwise around $p=\mu_{\hat{v}}$. For a given level of insurance coverage q, the marginal expected value of $x=r$ or $x=\pi$ equals

$$\begin{aligned}E\left(x|\hat{v}=D^{-1}\left(q\right)\right) &= E\left(x|\hat{v}=\sigma_{\hat{v}}\Phi^{-1}\left(1-q\right)+\mu_{\hat{v}}\right)\\ &= \frac{cov\left(x,\hat{v}\right)}{var\left(\hat{v}\right)}\sigma_{\hat{v}}\Phi^{-1}\left(1-q\right)+\mu_x\\ &= \rho_{\hat{v},x}\sigma_x\Phi^{-1}\left(1-q\right)+\mu_x.\end{aligned}$$

Hence, keeping $\rho_{\hat{v},x}$ and σ_x constant, an increase in $\sigma_{\hat{v}}$ does not affect the expected actual value and cost, conditional on the level of market coverage. The equilibrium price remains at $p^c=\mu_v$, where the demand curve and the average curve still intersect. This implies that the welfare benefit $\Gamma=\int_{p^*}^{p^c}\left[MV\left(p\right)-MC\left(p\right)\right]dD\left(p\right)$ from the efficient price subsidy remains unchanged. However, the efficient price is decreasing in $\sigma_{\hat{v}}$ and thus total cost of the price subsidy increases. The efficient price solves $MV\left(p^*\right)=MC\left(p^*\right)$. That is,

$$\frac{cov\left(\pi+r,\hat{v}\right)}{var\left(\hat{v}\right)}\left[p^*-\mu_{\hat{v}}\right]+\mu_\pi+\mu_r=\frac{cov\left(\pi,\hat{v}\right)}{var\left(\hat{v}\right)}\left[p^*-\mu_{\hat{v}}\right]+\mu_\pi$$

$\Rightarrow$

$$p^* = \mu_{\hat{v}} - \mu_r \frac{var(\hat{v})}{cov(r, \hat{v})}.$$

Hence,

$$p^c - p^* = \mu_r \frac{\sigma_{\hat{v}}}{\rho_{\hat{v},r}\sigma_r}.$$

A universal mandate would have the same welfare gain, while the cost of implementing a universal mandate does not depend on the demand function.□

Proof of Result 6

The correlation $\rho_{\varepsilon,\pi} = -\frac{1}{2}\frac{\sigma_\varepsilon}{\sigma_\pi}$ implies $cov(\pi, \varepsilon) = -\frac{1}{2}var(\varepsilon)$, while $\rho_{\varepsilon,r} = 0$ implies that $cov(r, \varepsilon) = 0$ and thus

$$\begin{aligned} var(\hat{v}) &= var(v) + var(\varepsilon) + 2cov(v, \varepsilon) \\ &= var(v). \end{aligned}$$

The demand function $D(p) = 1 - F_{\hat{v}}(p)$ is thus unaffected by σ_ε. Moreover, $\rho_{\varepsilon,r} = 0$ implies that $\frac{cov(r,\hat{v})}{var(\hat{v})} = \frac{cov(r,v)}{var(v)}$, such that the expected net-value at a price, $E(r|\hat{v}(\zeta) = p) \geq 0$, is unaffected by σ_ε as well. Finally, since $\frac{cov(\pi,\hat{v})}{\sqrt{var(\hat{v})}} = \frac{var(\pi) - \frac{1}{2}var(\varepsilon)}{\sqrt{var(\hat{v})}}$, the average cost,

$$AC(p) = \mu_\pi + \frac{cov(\pi, \hat{v})}{\sqrt{var(\hat{v})}} \frac{\phi\left(\frac{p-\mu_{\hat{v}}}{\sqrt{var(\hat{v})}}\right)}{1-\Phi\left(\frac{p-\mu_{\hat{v}}}{\sqrt{var(\hat{v})}}\right)},$$

increases when σ_ε decreases for any p. Hence, the equilibrium price $p^c = AC(p^c)$ increases. The welfare surplus, $\int_{p^c}^{\infty} E(r|\hat{v}(\zeta) = p)\, dF(p)$, decreases unambiguously.□

Proof of Result 7

The correlation $\rho_{\varepsilon,r} = -\frac{1}{2}\frac{\sigma_\varepsilon}{\sigma_r}$ implies $cov(r, \varepsilon) = -\frac{1}{2}var(\varepsilon)$, while $\rho_{\varepsilon,\pi} = 0$ implies that $cov(\pi, \varepsilon) = 0$ and thus $var(\hat{v}) = var(v)$. The demand function $D(p) = 1 - F_{\hat{v}}(p)$ is thus unaffected by σ_ε. Moreover, $\rho_{\varepsilon,\pi} = 0$ implies that $\frac{cov(\pi,\hat{v})}{\sqrt{var(\hat{v})}} = \frac{cov(\pi,v)}{\sqrt{var(v)}}$, such that the average cost $AC(p)$ is unaffected by σ_ε as well. Hence, the equilibrium price p^c remains the same. Finally, since $\frac{cov(r,\hat{v})}{\sqrt{var(\hat{v})}} = \frac{var(r) - \frac{1}{2}var(\varepsilon)}{\sqrt{var(\hat{v})}}$, the expected net-value at a price p,

$$E(r|\hat{v}(\zeta) \geq p) = \mu_r + \frac{cov(r, \hat{v})}{\sqrt{var(\hat{v})}} \frac{\phi\left(\frac{p-\mu_{\hat{v}}}{\sqrt{var(\hat{v})}}\right)}{1-\Phi\left(\frac{p-\mu_{\hat{v}}}{\sqrt{var(\hat{v})}}\right)},$$

is increasing in σ_ε. Hence, the welfare surplus, $\int_{p^c}^{\infty} E(r|\hat{v}(\zeta) = p)\, dF(p)$, decreases unambiguously.□

9.1 Appendix B: Discrete Distributions

The results analyzed in Section 2 extend for general distributions. I illustrate this by introducing two variations that change the perceived values relative to the actual values for any two types in a finite population $\vartheta = \{\zeta_1, \zeta_2, ..., \zeta_N\}$.

The first variation captures a reduction in the correlation between perceived and true values by making a type who has a higher actual value perceive her value as lower than another type.

Definition 1 *$\hat{v}$ is a confound of v if for some pairs of individuals characterized by ζ_x, ζ_y with $v(\zeta_x) < v(\zeta_y) : \hat{v}(\zeta_x) = v(\zeta_x) + \varepsilon$ and $\hat{v}(\zeta_y) = v(\zeta_y) - \varepsilon$ with $v(\zeta_x) + \varepsilon > v(\zeta_y) - \varepsilon$. For all other $\zeta : v(\zeta) = \hat{v}(\zeta)$.*

A natural example of a confound is when each type of a pair perceives to be the other type. This keeps the marginal distribution of the actual and perceived values identical, but reduces the correlation.

The second variation increases the spread of the perceived value relative to the true values.

Definition 2 *$\hat{v}$ is an exaggeration of v if for some pairs of individuals characterized by ζ_x, ζ_y with $v(\zeta_x) \geq v(\zeta_y) : \hat{v}(\zeta_x) = v(\zeta_x) + \varepsilon$ and $\hat{v}(\zeta_y) = v(\zeta_y) - \varepsilon$ for some $\varepsilon > 0$. For all other $\zeta : v(\zeta) = \hat{v}(\zeta)$.*

Notice that an exaggeration coincides with a mean-preserving spread when starting from $v_x = v_y$, in which case it corresponds to the introduction of random noise. Both confounds and exaggerations make types who overestimate their valuation to be overrepresented among the insured.

Proposition 9 *If the perceived values are the result of a sequence of exaggerations and confounds of the true values, the demand curve overestimates the value of the insured and underestimates the potential value for the uninsured.*

Proof. Assume $\hat{v}$ is a confound or exaggeration of v. Consider the set of individuals buying insurance at a price p, $\{\zeta \in \vartheta | \hat{v}(\zeta) \geq p\}$. If there is an individual $\zeta_y \in \vartheta$ buying insurance for whom $\hat{v}(\zeta_y) < v(\zeta_y)$, then there is also a individual $\zeta_x \in \vartheta$ buying insurance for whom $\hat{v}(\zeta_x) > v(\zeta_x)$, with $\hat{v}(\zeta_x) - v(\zeta_x) = v(\zeta_y) - \hat{v}(\zeta_y)$. However, the opposite is not true. That is, for any individual $\zeta_y \in \vartheta$ for whom $\hat{v}(\zeta_y) < v(\zeta_y)$, there is a price $p \in (\hat{v}(\zeta_y), \hat{v}(\zeta_x)]$ at which an individual $\zeta_x \in \vartheta$ for whom $\hat{v}(\zeta_x) > v(\zeta_x)$ buys while individual ζ_y does not. Hence, $E(v|\hat{v}(\zeta) \geq p) \geq E(\hat{v}|\hat{v}(\zeta) \geq p)$.□

We can adjust the notion of confounds and exaggerations for general discrete distributions such that the sign of the wedge between the demand and value curve only depends on the thickness of the market. For this to be the case, it is sufficient that all confounds and exaggerations are introduced around some value $\bar{v}$ in the following sense.

Definition 3 *$\hat{v}$ is a $\bar{v}$- centered confound of v if for some pairs of individuals characterized by ζ_x, ζ_y with $v(\zeta_x) \leq v(\zeta_y) : \hat{v}(\zeta_x) = v(\zeta_x) + \varepsilon$ and $\hat{v}(\zeta_y) = v(\zeta_y) - \varepsilon$ with $v(\zeta_x) + \varepsilon \geq \bar{v} \geq v(\zeta_y) - \varepsilon$. For all other $\zeta : v(\zeta) = \hat{v}(\zeta)$.*

Definition 4 *$\hat{v}$ is a $\bar{v}$- centered exaggeration of v if for some pairs of individuals characterized by ζ_x, ζ_y with $v(\zeta_x) \geq \bar{v} \geq v(\zeta_y) : \hat{v}(\zeta_x) = v(\zeta_x) + \varepsilon$ and $\hat{v}(\zeta_y) = v(\zeta_y) - \varepsilon$ for some $\varepsilon > 0$. For all other $\zeta : v(\zeta) = \hat{v}(\zeta)$.*

Proposition 10 *If the perceived value is the result of a sequence of $\bar{v}$- centered exaggerations and confounds of the true value, the demand function underestimates (overestimates) the value of insurance for the marginal buyer if the market is sufficiently thick (thin), i.e. $p < \bar{v}$ ($p > \bar{v}$).*

Proof. Assume $\hat{v}$ is a $\bar{v}$- centered confound or exaggeration of v. Consider the set of marginal buyers at a price p, $\{\zeta \in \vartheta | \hat{v}(\zeta) = p\}$. If p is above $\bar{v}$, for any marginal buyer $\hat{v}(\zeta) \geq v(\zeta)$. If p is below $\bar{v}$, for any marginal buyer $\hat{v}(\zeta) \leq v(\zeta)$. Hence, $E(v|\hat{v}(\zeta) = p) \geq E(\hat{v}|\hat{v}(\zeta) \geq p)$ for $p \leq \bar{v}$ and vice versa.$\square$

The Proposition implies that the demand and value curve intersect only once, at price $p = \bar{v}$. However, the difference between the two curves is not necessarily monotone in the price.