Heterogeneity, Demand for Insurance and Adverse Selection

Johannes Spinnewijn*
London School of Economics

April 18, 2012

Abstract

Recent empirical work finds that surprisingly little variation in the demand for insurance is explained by heterogeneity in risks. I distinguish between heterogeneity in risk preferences and risk perceptions underlying the unexplained variation. Heterogeneous risk perceptions induce a systematic difference between the revealed and actual value of insurance as a function of the insurance price. Using a sufficient statistics approach that accounts for this alternative source of heterogeneity, I find that the welfare conclusions regarding adversely selected markets are substantially different. The source of heterogeneity is also essential for the evaluation of different interventions intended to correct inefficiencies due to adverse selection like insurance subsidies and mandates, risk-adjusted pricing and information policies.

Keywords: Heterogeneity, Adverse Selection, Risk Perceptions, Welfare and Policy

JEL-codes: D60, D82, D83, G28

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 importance of adverse selection in insurance markets. An individual's risk type often plays a minor role in explaining his or her demand for

*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/).
I thank Pedro Bordalo, Gharad Bryan, Arthur Campbell, Raj Chetty, Jesse Edgerton, Erik Eyster, Amy Finkelstein, Philipp Kircher, Henrik Kleven, Botond Koszegi, Amanda Kowalski, David Laibson, Sendhil Mullainathan, Gerard Padró i Miquel, Matthew Rabin, Frans Spinnewyn and seminar participants at LSE-UCL, DIW Berlin, Lausanne, Max Planck Institute, Southampton, Zurich and the CESifo meetings for valuable discussions and comments. I would also like to thank Shantayne Chan for excellent research assistance.

insurance, which raises the important question what type of heterogeneity is actually driving the variation in insurance demand. Recent work attributes the unexplained variation to heterogeneity in preferences (Cohen and Einav 2007, Einav, Finkelstein and Cullen 2010a, Einav, Finkelstein and Schrimpf 2010b) and finds that the estimated welfare cost of inefficient pricing due to adverse selection is very small. The main reason is that the value of insurance for the uninsured is estimated to be small. Heterogeneity in preferences thus reduces the scope for policy interventions in insurance markets.

An alternative explanation why risks do not explain the demand for insurance is the discrepancy between perceived and actual risks. The formation of risk perceptions is inherently subjective and subject to biases and heuristics.[1] Risk perceptions are thus only a noisy measure of one's actual risk.[2] This also drives a wedge between the *actual* value of insurance and the value of insurance as *revealed* by an individual's demand. Recent empirical evidence identifies other behavioral and economic constraints causing a tenuous relation between choice and value in insurance markets (e.g., Abaluck and Gruber 2011, Handel 2011, Fang, Keane and Silverman 2008). To the extent that one cares about the actual value rather than the revealed value of insurance, the presence of these *non-welfarist* constraints - affecting the insurance demand, but not the insurance value - changes earlier welfare and policy conclusions.

This paper presents a simple model of insurance with heterogeneity in risk and preferences. The model introduces non-welfarist constraints through a noise term that distorts the insurance decision. This general framework is used to analyze how the different sources of heterogeneity underlying the insurance demand affect the welfare and policy analysis regarding adverse selection. The analysis extends the sufficient statistics approach by Einav et al. (2010a) and leads to two key insights. First, non-welfarist heterogeneity has an unambiguous impact on the estimated welfare cost of adverse selection due to a selection effect. Second, the effectiveness of all policy interventions used to tackle adverse selection depends on the source of heterogeneity underlying the demand for insurance. The paper also calibrates the model based on the empirical analysis in Einav et al. (2010a) and finds that both welfare and policy conclusions change substantially when accounting for non-welfarist heterogeneity.

At the heart of the analysis is a simple selection effect, which naturally applies in case of heterogeneous risk perceptions. Even when accurate on average, the insured individuals tend to overestimate, while the uninsured individuals tend to underestimate the value of insurance, regardless of the insurance 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] As a consequence, the demand curve

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

[2] For example, neighbors in a coastal area have very different perceptions about the risk of a natural disaster damaging their property, even though they face the same actual risk (Peacock et al. 2005).

[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.

overstates the surplus for the insured individuals and understates the potential surplus for the uninsured individuals. When taking the demand curve at face value, the evaluation of policy interventions which either target the insured or uninsured will be unambiguously biased in opposite directions. For example, the welfare gain of a universal mandate is unambiguously higher than the demand for insurance would suggest. The same selection mechanism tends to rotate the *value curve* in a counter-clockwise direction relative to the demand curve, where the value curve depicts the actual rather than revealed value of insurance for the marginally insured.[4] As consequence, the demand curve is more likely to underestimate the insurance value for individuals the lower their willingness to pay. For normally distributed heterogeneity, the rotation is counter-clockwise when the correlation between the perceived and actual risk is imperfect or the variance in perceived risks exceeds the variance in actual risks.

I use this systematic relation between the value and demand curve to extend the sufficient statistics approach by Einav et al. (2010a) for non-welfarist heterogeneity. One statistic is required in addition to the demand and cost curves, which are sufficient when the demand does reveal the actual insurance value. This statistic equals the share of the variation in insurance demand - left unexplained by heterogeneity in risks - that is driven by non-welfarist constraints (rather than by heterogeneous preferences). An advantage of the extended approach is that the welfare analysis can simply use existing empirical estimates of the demand and cost curves. However, additional data would be required to estimate the non-welfarist share. Building on the empirical analysis of employer-provided health insurance by Einav et al. (2010a), I find that the actual cost of adverse selection would be thirty percent higher when ten percent of the unexplained variation is driven by non-welfarist variation and four times as high when this share increases to fifty percent. While a precise empirical analysis of the heterogeneity underlying the demand curve is left for future work, back-of-the-envelope calculations using existing empirical evidence suggest a share of fifty percent to be plausible. The cost of adverse selection in this setting may thus be substantially higher than previously estimated and justify government interventions in this market.

I use the framework to analyze and calibrate the welfare impact of all relevant policies that are currently in place in insurance markets. I find that the presence of non-welfarist heterogeneity makes price policies less effective relative to insurance mandate. While price policies are constrained by individuals' perceived valuations, the welfare impact depends on the actual valuations. Subsidizing the insurance price to encourage the uninsured who underestimate the insurance value to buy insurance becomes very costly. Similarly, adjusting the insurance price for the buyer's particular risk type is only effective when individuals do perceive these risks accurately. The

[4] Johnson and Myatt (2006) analyze rotations of the demand curve when marketing and advertizing changes the distribution of the value of insurance. Here, the value curve is also a rotation of the demand curve, but the underlying distribution of perceived values is explicitly correlated with the distribution of actual values underlying the original demand curve.

calibrations show how non-welfarist heterogeneity reduces the net welfare gain from an efficient price subsidy and mitigates the efficiency gains from risk-adjusted pricing, as recently estimated by Bundorf, Levin and Mahoney (2012). The opposite is true for a universal mandate, which in addition can be implemented without any prior knowledge regarding the heterogeneity underlying the insurance demand. Finally, I evaluate the effect of policies that reduce the constraints distorting insurance choices. While relaxing constraints makes individuals better off at a given price, it also changes the selection of individuals buying insurance and thus the equilibrium price.[5] The framework with multi-dimensional heterogeneity allows to disentangle these two effects. I find that providing information to individuals about the expected risk they face individually always decreases welfare. In contrast, providing information about the variance of the risk increases welfare, since it induces those who previously underestimated (overestimated) the insurance value to become insured (uninsured), regardless of their expected cost to the insurance company.

1.1 Related Literature

Starting with the work by Chiappori and Salanié (1997, 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 reviewed in Cohen and Siegelman (2010), with some insurance markets being advantageously rather than adversely selected, inspired a new series of studies which estimate the heterogeneity in risk preferences jointly with the heterogeneity in risk types (Cohen and Einav, 2007; Einav et al. 2010a, 2010b). The estimated heterogeneity allows to move beyond testing for adverse selection and actually analyze the welfare cost of inefficient pricing. This cost is generally found to be small (see Einav, Finkelstein and Levin 2010c).

While attributing heterogeneity in insurance choices - unexplained by heterogeneity in risks - to heterogeneity in preferences is a natural first step and in line with the revealed preference paradigm, several papers have recently made the case that insurance behavior cannot be adequately explained with standard preferences and risk perceptions. Chiappori and Salanié (2012) emphasize the importance of understanding risk perceptions to analyze insurance behavior in future research. Cutler and Zeckhauser (2004) argue that distorted risk perceptions are one of the main reasons why some insurance markets do not work efficiently. Abaluck and Gruber (2011) identify important inconsistencies in the insurance choices of the elderly and document examples of insurance plans that offer better risk protection at a lower cost which are available, but not chosen. Fang et al. (2008) find that heterogeneity in cognitive ability is important

[5] Condon, Kling and Mullainathan (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.

(relative to risk aversion) in explaining the choice of elderly to buy Medigap. A number of related empirical papers analyze deviations from expected utility maximization in explaining insurance coverage and other choices under risk. For example, Barseghyan, Molinari, O'Donoghue and Teitelbaum (2011) find that a structural model with non-linear probability weighting explains the data better than a model with standard risk aversion looking at deductible choices in auto and house insurance. Other examples are Bruhin et al. (2010), Snowberg and Wolfers (2010) and Sydnor (2010). Notice that these papers restrict individuals who face the same actual risk to have the same risk perception. Most recently, the stability of an individual's risk preference across insurance domains has been challenged as well; Barseghyan, Prince and Teitelbaum (2011) reject the hypothesis of stable risk preferences across domains using a structural model. Einav, Finkelstein, Pascu and Cullen (2011) cannot reject the presence of a domain-general component, but also find that an individual's domain-specific risk plays a minor role in explaining insurance choices.

Accounting for non-welfarist heterogeneity when analyzing welfare and policy interventions in insurance markets seems the natural next step in light of the evidence above. The analysis in the paper relates to two recent trends in public economics; the first is the inclusion of non-standard decision makers in welfare analysis, the second is the expression of optimal policies in terms of sufficient statistics.[6] In a similar spirit, Chetty, Kroft and Looney (2009) extend the sufficient statistics approach to tax policy for tax salience and Spinnewijn (2010a) extends the sufficient statistics approach to unemployment policy for biased perceptions of employment prospects. Mullainathan, Schwartzstein and Congdon (2012) propose a unifying framework to examine the implications of behavioral biases for social insurance and optimal taxation. In contrast, the focus of this paper is on heterogeneity in behavioral tendencies and the implications for adverse selection. Sandroni and Squintani (2007, 2010) and Spinnewijn (2010b) also analyze heterogeneity in risk and risk perceptions, but focus on the characterization of the screening contracts offered in the equilibrium of Rotschild-Stiglitz type models and revisit whether an insurance mandate is Pareto-improving in the respective settings.

The remainder of the paper is as follows. Section 2 introduces a simple model of insurance demand and characterizes the difference between actual and revealed insurance values along the demand curve. Section 3 introduces heterogeneity in risk types and preferences to analyze and calibrate the cost of adverse selection depending on the role of non-welfarist heterogeneity, building on Einav et al. (2010a). Section 4 analyzes the effectiveness of different government interventions depending on the importance of non-welfarist heterogeneity. Section 5 discusses the empirical implementation and the robustness of the welfare and policy analysis. Section 6 concludes.

[6] See Congdon et al. 2011 and Chetty 2010 for recent discussions.

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 true value of insurance. The analysis deviates from the revealed preference paradigm and assumes that the variation in insurance decisions may be driven by heterogeneity in *non-welfarist* contraints, unrelated to the true value of insurance. These non-welfarist constraints relate to the notion of ancillary conditions, as introduced by Bernheim and Rangel (2009), which are features of the choice environment that may affect behavior, but not relevant to a social planner's choice. I assume that the social planner uses the true insurance value to evaluate welfare and refer to the policy maker who ignores non-welfarist heterogeneity as *naive*.[7]

I will mostly interpret the source of the non-welfarist heterogeneity as coming from differences between perceived and actual risks. Still, the analysis does apply more generally to heterogeneity in 'behavioral' constraints like inattention, cognitive inability or inertia, but also to heterogeneity in 'economic' constraints, like liquidity constraints or adjustment costs, which also restrict people's ability to buy insurance regardless of the value of insurance for those individuals.

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 interventions. The perceived value, however, refers to the value as perceived by this individual and determines his or her demand for insurance. The difference between the true and perceived value is driven by non–welfarist constraints, which are captured by a noise term ε,

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

and continuous distributions $F_{\hat{v}}, F_v$ and F_{ε}. For example, the noise term is positive when an individual overestimates the risk she is facing and negative when the individual underestimates that risk. I assume that the noise cancels out across the entire population. The true and perceived value are thus equal on average. However, since

[7] The difference between the optimal and naive welfare criterium thus relates to the difference between 'experienced utility' and 'decision utility' (Tversky and Kahneman 1979).

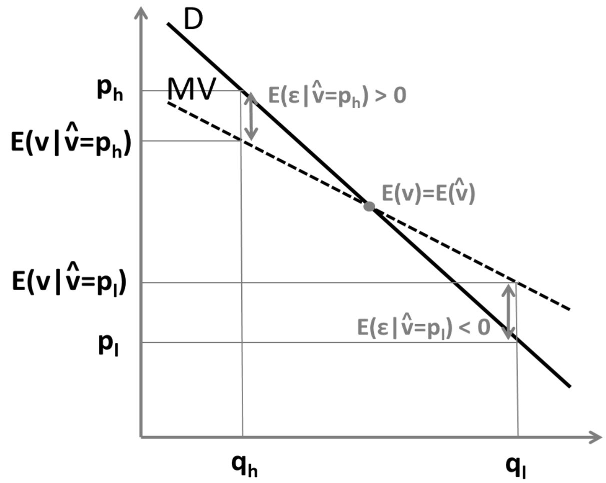

Figure 1: The Demand Curve and the Value Curve.

the demand for insurance depends only 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, the (average) true value for the marginal buyers is relevant, which I denote by $MV(p) \equiv E_\zeta(v|\hat{v} = p)$.[8] The central question is thus to what extent the true value co-varies 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, $cov(v, \hat{v})/var(\hat{v})$.

Graphically, one can construct the *value curve*, depicting the expected true value for the marginal buyers for any level of insurance coverage q, and compare this to the demand curve, depicting the perceived value $D^{-1}(q)$ for that level of insurance coverage, as shown in Figure 1. The mistake made by an *naive* policy maker who incorrectly assumes that the demand curve reveals the true value of insurance depends on the wedge between the two curves. I analyze the systematic nature of this difference along the demand curve.

2.2 Infra-marginal Policies: Robust Bias

I start by comparing the true and perceived insurance value for the infra-marginal individuals, as given by the area below the value and demand curve respectively. For

[8] 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 insured, the expected true value of insurance, $E_{\zeta}(v|\hat{v} \geq p)$, determines the actual 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, the expected true value of insurance, $E_{\zeta}(v|\hat{v} < p)$, determines the potential value of a universal mandate which forces all uninsured individuals to buy insurance.

Independence 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 insurance value for the insured and underestimates the insurance value for the uninsured,*

$$E_{\zeta}\left(\varepsilon|\hat{v} \geq p\right) \geq 0 \geq E_{\zeta}\left(\varepsilon|\hat{v} < p\right) \text{ for any } p.$$

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 though these characteristics cancel out over the entire population, they do not conditional on the decision to buy insurance. For example, optimistic beliefs discourage individuals from buying insurance, while pessimistic beliefs encourage individuals to buy insurance. Those buying insurance are thus 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 co-varies with the perceived value. For tractability, I only illustrate this here for normal distributions, but I extend this insight for more general distributions in Appendix. Denote the mean and variance of any variable x by μ_x and σ_x^2 and the correlation with any other variably y by $\rho_{x,y}$.

Proposition 2 *If the true and perceived value are normally distributed,*

$$E_\zeta\left(\varepsilon|\hat{v}\geq p\right)\geq 0\geq E_\zeta\left(\varepsilon|\hat{v}<p\right) \text{ for any } p \text{ if and only } \rho_{v,\hat{v}}\times\frac{\sigma_v}{\sigma_{\hat{v}}}\leq 1.$$

The condition is equivalent to $\rho_{v,\varepsilon}\geq -\frac{\sigma_\varepsilon}{\sigma_v}$. The proposition thus shows that the signs of the biases, as found in Proposition 1, remain the same as long as the correlation between the noise term and the true value is not too negative. The robust nature of the results seems confirmed when expressing the condition in terms of perceived and true value,

$$\frac{cov\left(v,\hat{v}\right)}{var\left(\hat{v}\right)}=\rho_{v,\hat{v}}\times\frac{\sigma_v}{\sigma_{\hat{v}}}\leq 1.$$

A naive policy maker will overestimate the insurance value for the insured and underestimate the insurance value for the uninsured when the true value changes less than one-for-one with the perceived value. A natural reason for this to be true is an imperfect correlation between the perceived and true value of insurance. For example, the assumption that the correlation between risk types and risk perceptions seems particularly strong. John C. Harsanyi (1968) observed that "by the very nature of subjective probabilities, even if two individuals have exactly the same information and are at exactly the same high level of intelligence, they may very well assign different subjective probabilities to the very same events." While rationality may restrict individuals to be Bayesian, it puts no restrictions on priors themselves, which are primitives of the model (Van Den Steen 2004). As long as learning is incomplete, the correlation $\rho_{v,\hat{v}}$ will be imperfect. An alternative interpretation of the non-welfarist heterogeneity leading to the same conclusion is the presence of some 'behavioral' individuals for whom the perceived value (or risk) is a random draw from the distribution of the true values (or risks), while for all other individuals the perceived value equals the true value. In this model, the correlation $\rho_{v,\hat{v}}$ equals $1-\alpha$, where α is the share of 'behavioral' individuals. Still, the estimated bias is also affected by the relative dispersion of the perceived and actual values. The bias would be reduced and potentially reversed if the perceived values are less dispersed than the actual values, for example when individuals underestimate the differences in their risk types. However, with imperfect correlation, the dispersion in perceived values should be sufficiently smaller than the dispersion in actual values to reverse the results.

2.3 Marginal Policies: Counter-clockwise Rotation

The results in the previous section apply to infra-marginal policies, affecting either all the insured or all the uninsured. To evaluate more targeted policies, like a small price subsidy, one needs to know the 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,*

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

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. The immediate policy implication is that a naive policy maker underestimates the value of an increase in insurance coverage more, the higher the share of insured individuals in the market. 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 demand curve and thus the naive policy maker overestimate the true value of additional insurance if and only if the market coverage is below one half.

The monotone likelihood ratio property is satisfied by a large class of distributions, including the normal distribution. With normal heterogeneity, the condition for the value curve to be a counter-clockwise rotation of the demand curve is $\rho_{v,\hat{v}} \times \dfrac{\sigma_v}{\sigma_{\hat{v}}} \leq 1$, exactly the same as in Proposition 2. Notice that 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.

3 Adverse Selection

I now introduce the cost of providing insurance and consider the supply of insurance contracts. Particular to insurance markets is that the cost of providing insurance to an individual depends on that individual's risk type. An individual's risk type thus influences both his or her demand for insurance, but also the cost to the insurer of providing insuranc. 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 not only determines the true value of insurance, but also the expected cost for the insurance company of providing insurance. In particular, I assume $c(\zeta) = \pi(\zeta)$. Like before, the perceived value equals the true value plus a noise term. The model thus captures heterogeneity in three different dimensions: risk types, risk preferences and non-welfarist constraints.

The setup is kept as simple as possible to keep the analysis insightful, clear and tractable. Notice that this exact setup arises 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. This suggests that 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. The nature of the results would not change if the value and cost function do not depend in an identical way on the individual's risk type $\pi(\zeta)$, neither if the value were not additive in the risk and preference type. Notice that the additivity is not restrictive without restrictions on the distribution of the heterogeneity in the different dimensions.[9]

3.2 Cost of Adverse Selection

The expected cost of an insurance contract depends on the types who decide to buy the contract. The average and marginal cost of providing a contract at price p equal

[9] The assumption of CARA preferences or additivity of the risk premium in the contract valuation is standard in the recent empirical insurance literature (see Einav et al. 2010c).

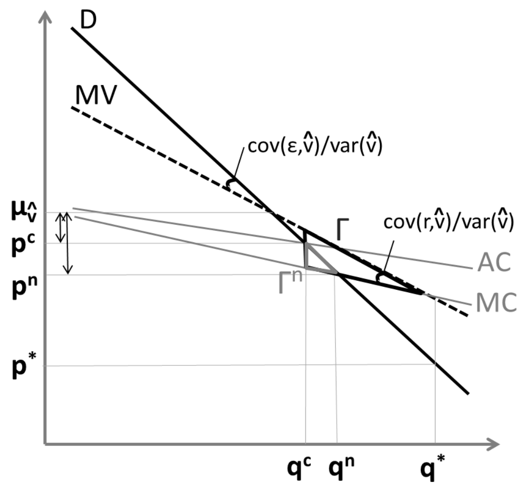

Figure 2: Adverse Selection: the naively estimated cost Γ^n vs. the actual cost Γ.

respectively,

$$AC\left(p\right)=E_{\zeta}\left(\pi|\hat{v}\geq p\right),\;MC\left(p\right)=E_{\zeta}\left(\pi|\hat{v}=p\right).$$

Adverse selection results when the marginal cost is an increasing function of the price. That is, the willingness to buy insurance is lower for lower risk types and they thus decide not to buy insurance at lower prices. Figure 2 plots the marginal and cost curve together with the demand curve. The marginal cost is decreasing with the share of insured individuals. The average cost function is thus decreasing as well, but at a slower rate, and lies above the marginal cost function. In advantageously selected markets, individuals with higher risk are less likely to buy insurance and the average cost function will be below rather than above the increasing marginal cost function. The less an individual's risk affects her insurance choice, the less the marginal cost will depend on the price.

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

$$AC\left(p^c\right)=p^c.$$

Graphically, this is the price for which the demand and average cost curve intersect. However, it is efficient for an individual to buy insurance as long as her valuation exceeds the cost of insurance. Hence, at the constrained efficient price p^*, the marginal cost of insurance equals the marginal actual value of insurance,[10]

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

[10]In the unconstrained efficient allocation, an individual buys insurance if and only if $r\geq 0$. Since individuals with the same perceived value cannot be separated, the constrained efficient allocation has individuals with perceived value $\hat{v}$ buying insurance if and only if $E_{\zeta}\left(r|\hat{v}\right)\geq 0$.

This price is given by the the intersection of the value curve and the marginal cost curve. When the market is adversely selected and the marginal cost is thus below the average cost ($MC(p) < AC(p)$), the competitive price is inefficiently high under the assumption that the demand curve reflects the value of insurance. When the demand curve underestimates the value of insurance ($p < MV(p)$), the inefficiency is further increased.

The total cost of adverse selection depends on the difference between the value and cost for the pool of inefficiently uninsured individuals with a perceived value between p^* and p^c,

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

Graphically, the cost equals the area between the value curve and the marginal cost curve from the competitive to the efficient level of insurance coverage, as shown in Figure 2. When the perceived and actual values coincide, the demand and cost curves are sufficient to determine the cost of adverse selection, as shown by Einav et al. (2010a). However, when the perceived and actual values do not coincide, the demand and cost curves are no longer sufficient. 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 Γ as he (1) misidentifies the pool of individuals who should be insured and (2) misestimates the welfare loss for the adversely uninsured. That is,

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

where $\Gamma^n = \int_{p^n}^{p^c} [p - MC(p)]\, dD(p)$ denotes the welfare cost as estimated by a naive policy maker. The difference between Γ and Γ^n depends on the share of insured individuals in the market ($MV(p) \gtrless p$) and the nature of selection ($AC(p) \gtrless MC(p)$). Figure 2 illustrates the difference between the actual and naively estimated inefficiency cost for an adversely selected market with high coverage. The inefficiency is higher than a naive policy maker thinks, both because the extent of underinsurance is worse ($p^* < p^n < p^c$) and the welfare loss of underinsurance at a given price is larger than expected ($p < MV(p)$).

3.3 Sufficient Statistics Formula

In order to derive a closed-form expression for the cost of adverse selection, I assume normal heterogeneity in all three dimensions. I put no restrictions on the covariance and use notation as before. Under normality, the expected value of any variable $z \in \{\pi, r, \varepsilon\}$, conditional on the perceived value, equals

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

The ratio $cov\left(z,\hat{v}\right)/var\left(\hat{v}\right)$ indicates how much the variable z moves with the price. The variation in demand can thus be attributed to the different sources of heterogeneity depending on the relative covariance of each component with the perceived value. Notice that if all terms are independent, the covariance of each term with the perceived value is equal to the variance of that term.

The misestimation by a naive policy maker crucially depends on the covariance ratio $cov\left(\varepsilon,\hat{v}\right)/cov\left(r,\hat{v}\right)$, capturing the extent to which the variation in demand is explained by noise rather than by preferences. Graphically, this ratio determines the position of the value curve between the demand curve and the marginal cost curve. This thus affects the wedge between the true surplus r and the perceived surplus of insurance $r+\varepsilon$,

$$\frac{E_{\zeta}\left(\varepsilon|\hat{v}=p\right)}{E_{\zeta}\left(r|\hat{v}=p\right)-\mu_{r}}=\frac{cov\left(\varepsilon,\hat{v}\right)}{cov\left(r,\hat{v}\right)},$$

and thus the misestimation of the welfare loss of the adversely uninsured, as in the earlier decomposition of Γ. In addition, the covariance ratio determines the difference between the price that is perceived to be efficient and the price that is actually efficient,

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

and thus the misidentification of the pool of inefficiently uninsured. By linearizing the demand curve through (p^n, q^n) and (p^c, q^c), we obtain the following approximate result.

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

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

The demand and cost curves allow estimating the cost of adverse selection in absence of non-welfarist heterogeneity Γ^n and the price ratio $\mathcal{P}=\dfrac{\mu_{\hat{v}}-p^{n}}{p^{c}-p^{n}}$. Hence, the covariance ratio $cov\left(\varepsilon,\hat{v}\right)/cov\left(r,\hat{v}\right)$ is the only additional sufficient statistic required to account for non-welfarist heterogeneity in the welfare analysis .The impact of the covariance ratio on the bias in the welfare cost estimation depends on the price ratio

$\mathcal{P} = \frac{\mu_{\hat{v}} - p^n}{p^c - p^n}$. This price ratio depends on the price difference $p^c - p^n$, which captures the nature of selection, and the price difference $\mu_{\hat{v}} - p^n$, which captures whether the pool of inefficiently selected over- or underestimate the value of insurance. Graphically, this depends on whether the inefficient pool is to the left or the right of the intersection between the demand and the value curve, as shown in Figure 2.

If the price ratio $\mathcal{P}$ is larger than one, the policy maker unambiguously underestimates the efficiency cost of selection. This is the case if the market is adversely selected, but coverage is large ($\mu_{\hat{v}} \geq p^c \geq p^n$) so that all adversely uninsured are underestimating the value of insurance on average.[11] This case arises in the empirical application. When exactly half of the market is covered ($p^c = \mu_{\hat{v}}$) so that $\mathcal{P}$ =1, the misestimation is approximately linear in the covariance ratio,

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

For higher market coverage ($\mu_{\hat{v}} > p^c$), the bias is larger and increases at a faster rate with the covariance ratio. For lower market coverage ($\mu_{\hat{v}} < p^c$), some of the adversely uninsured are overestimating rather than underestimating the value of insurance and the bias is thus smaller. If the market coverage is sufficiently low (e.g., $\mu_{\hat{v}} < p^n \leq p^c$), the policy maker will underestimate the inefficiency cost of selection.[12]

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 (2010a), henceforth EFC, illustrating their sufficient statistics approach. Based on the health insurance choices and medical insurance claims of the employees of Alcoa, a multinational producer of aluminium, EFC estimate the demand for insurance coverage and the associated cost of providing insurance.[13,14] They find that the marginal cost is increasing in the price, but the increase is small. The increase indicates the existence of adverse selection, but the small magnitude of the increase suggests that relatively little heterogeneity in insurance choices is explained by heterogeneity in risks. EFC 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

[11] The price ratio $\mathcal{P}$ is also larger than one if the market is advantageously selected, but coverage is small ($\mu_{\hat{v}} \leq p^c \leq p^n$).

[12] Not surprisingly, if the market is adversely selected ($p^c > p^n$), but coverage is very low ($p^n > \mu_{\hat{v}}$), the efficient price may be above the competitive price such that it becomes welfare improving to decrease rather than increase the level of market coverage.

[13] The price variation is argued to be exogenous, as business unit managers set the prices for a menu of different health insurance options, offered to all employees within their business unit.

[14] In particular, they consider a sample of 3,779 salaried employees, who chose one of the two modal health insurance choices, where one option provides more coverage at a higher price.

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 percent. Relative to the estimated surplus at efficient pricing, this suggests a welfare cost of only 3 percent.

I use the estimates in EFC to illustrate how welfare conclusions are affected when non-welfarist constraints affect insurance choices. This exercise does not require the data underlying the estimates in EFC, conditional on having an estimate of the covariance ratio $cov(\varepsilon, \hat{v})/cov(r, \hat{v})$. I apply the formula derived in Proposition 4, which was derived for a linear approximation of the demand curve under normal heterogeneity. EFC estimate a linear system which implies that the formula would be exact if the value curve is a rotation of the demand curve like in the case with normal heterogeneity.[15,16] Since the market is adversely selected and market coverage is large ($q^c > 0.5$), the bias in the estimation of the welfare cost increases as a function of $cov(\varepsilon, \hat{v})/cov(r+\varepsilon, \hat{v})$, as shown in Table 1. Using the earlier interpretation, I find that if 1 percent of the residual variation is explained by non-welfarist heterogeneity, the actual cost of adverse selection is 3 percent higher than estimated when using the demand function. If this share increases to 10 percent, the actual cost of adverse selection is already 31 percent higher. If half of the residual variation is explained by non-welfarist heterogeneity, the actual cost of adverse selection is more than 4 times higher than estimated based on the demand function. I find even a fifty percent share to be plausible based on back-of-the-envelope calculations using empirical evidence discussed in Section **??**. This would imply that rather than \$9.55 per employee per year, the cost of adverse selection would be \$38.4 per employee per year, corresponding to 25 percent of the surplus generated in this market at the efficient price.[17]

3.5 Discussion

The calibration suggests that the welfare cost of adverse selection is substantially higher in the presence of non-welfarist heterogeneity, potentially justifying the intervention of govenrments in insurance markets. While providing a precise estimate the importance of non-welfarist constraints is challenging and beyond the scope of this paper, exisiting empirical evidence suggests that the role of non-welfarist constraints may well be substantial.

To estimate demand and cost curves like in Einav et al. (2010a), exogenous price

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

[16] I have also evaluated the exact welfare cost when the demand components are normally distributed. The approach to calibrate the covariance matrix based on the estimates in EFC is the same as explained in the next subsections. The demand, value and cost curves are then calculated using this matrix. Table App1 in the web appendix shows that the welfare results are very similar for this system with normal heterogeneity. The final column shows the estimated bias based on the linear approximation in Proposition 4, suggesting that the linear approximation works very well when $cov(\varepsilon, \hat{v})/cov(\varepsilon + r, \hat{v})$ is small.

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

variation and data on insurance choices and claim rates are required. Additional data is required to disaggregate the revealed value of insurance into true value and constraints. One approach is to identify individuals for whom non-welfarist constraints do not bind. The demand elasticity estimated for these individuals could be used to uncover the value function associated with the observed demand. While similar in spirit to Chetty et al. (2009), the success of this approach depends entirely on the identification of unconstrained individuals. An alternative approach is to identify a non-welfarist constraint or variable which does affect the insurance decision, but is unrelated to the insurance value. This approach will provide a lowerbound for $cov(\varepsilon, \hat{v})/cov(r+\varepsilon, \hat{v})$, as other non-welfarist constraints may apply as well. One application of this approach is the evidence in Fang et al. (2008) 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.[18] In a similar spirit, the estimated relation between actual and perceived risks could be used to estimate how much the true and perceived insurance value co-vary. Surveyed risk perceptions are found to predict risk realizations, often better than any other set of covariates, but the estimated relation is generally very small.[19] For example, Finkelstein and McGarry (2006) find estimates smaller than 0.10 when estimating a linear probability model of nursing home use in the five years between 1995 and 2000 on the 1995 self-reported beliefs of this probability. Assuming that the perceived risk $\hat{\pi} = \pi + \varepsilon$, this would imply that $cov(\pi, \hat{\pi})/var(\hat{\pi}) = 0.10$. An increase in the perceived risk is thus associated with only a small increase in the actual risk.[20] When combined with the estimated relation between the insurance demand and the actual risk types, $cov(\pi, \hat{v})/var(\hat{v})$, this estimate can be used to recover the importance of risk perceptions underlying the demand for insurance.[21] Decomposing the covariances, we find

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

[18] Similarly, wealth, income and education are also often found to explain insurance choices. While these variables may be related to the true value of insurance, empirical evidence suggests that they are also strongly related to the mere quality of decisions under uncertainty (Choi et al. 2011).

[19] See Hurd (2009) for a recent overview of empirical work on the relation between surveyed risk perceptions and actual risks. For example, Hamermesh (1985) and Hurd and McGarry (1995, 2002) analyze subjective life expectations and survival probabilities.

[20] Clearly, the self-reported probability does not measure the demand-driving perceived probability $\hat{\pi}$ without error and measurement error attenuates the regression estimate of $cov(\pi, \hat{\pi})/var(\hat{\pi})$ towards 0. Kircher and Spinnewijn (2011) suggest an alternative approach using price variation to disentangle perceived risks from risk preferences. Another alternative to evaluate the impact of misperceived risks is to provide information about risks in a controlled experiment and analyze the effect on the demand for insurance and the associated costs.

[21] 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.

The approximation depends on the covariance between preferences and perceived or actual risks being small. After subtracting $\frac{cov(\pi,\hat{v})}{var(\hat{v})}$ from both sides, we find

$$\frac{cov\left(\pi,\hat{v}\right)}{var\left(\hat{v}\right)}\left[1/\frac{cov\left(\pi,\hat{\pi}\right)}{var\left(\hat{\pi}\right)}-1\right]\cong\frac{cov\left(\varepsilon,\hat{v}\right)}{var\left(\hat{v}\right)}.$$

The EFC analysis implies an estimate for $cov\left(\pi,\hat{v}\right)/var\left(\hat{v}\right)$ of about 1/3, which corresponds to the slope of the marginal cost curve relative to the demand curve. The approximation thus suggests that if $cov\left(\pi,\hat{\pi}\right)/var\left(\hat{\pi}\right)$ is smaller than 1/2, $cov\left(\varepsilon,\hat{v}\right)/var\left(\hat{v}\right)$ is greater than 1/3. Attributing the residual heterogeneity to preferences, we find that $cov\left(r,\hat{v}\right)/var\left(\hat{v}\right)$ is smaller than 1/3. Hence, this implies that the heterogeneity in risk perceptions explains more than 50 percent of the variation in demand that is left unexplained by the heterogeneity in actual risks. This back-of-the-envelope calculation thus suggests that our sufficient statistic $cov\left(\varepsilon,\hat{v}\right)/cov\left(\varepsilon+r,\hat{v}\right)$, used in Table 1 and 2, would exceed 0.5 when $cov\left(\pi,\hat{\pi}\right)/var\left(\hat{\pi}\right)$ is indeed smaller than 1/2.

While several papers have attributed the heterogeneity in insurance choices, left unexplained by heterogeneity in risks, to estimate heterogeneity in risk preferences (e.g., Cohen and Einav, 2007), only few papers use explicit measures of risk preferences to explain insurance choices (e.g., Cutler, Finkelstein and McGarry 2008). As discussed before, the empirical evidence across individuals (see Cohen and Siegelman, 2010) and across domains (Barseghyan et al., 2011, and Einav et al., 2011b) can attribute only a minor part of the variation in insurance demand to heterogeneous preferences. These findings are suggestive, but not sufficient to conclude that the link between choice and value is weak in insurance markets. Further empirical work is needed to provide more evidence on the role of both non-welfarist heterogeneity and preference heterogeneity.

Average Bias The analysis assumes that on average the demand function does reveal the actual value of insurance, i.e., $E_{\zeta}\left(\varepsilon\right)=0$. Regarding risk perceptions, various studies suggest that people may be 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). This causes a wedge between the actual and perceived value of insurance, as analyzed in Spinnewijn (2010a) and Mullanaithan et al. (2012), but does not affect the nature of the insights regarding the impact of heterogeneity, changing the wedge between the perceived and actual value along the demand curve. Still, the combination of both sources is relevant for welfare analysis. 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 the insurance value for the uninsured will be even larger and vice versa.

The welfare analysis can be easily extended for an average difference between the actual and revealed value of insurance, i.e., $E_{\zeta}\left(\varepsilon\right)\neq0$. In Proposition 4, only the price ratio $\mathcal{P} = \dfrac{\mu_{\hat{v}} - p^n}{p^c - p^n}$ should be adjusted to $\mathcal{P}^x \equiv \dfrac{x - p^n}{p^c - p^n}$, where x is determined by the intersection of the demand curve and the value curve, solving $F_{\hat{v}}(x) = F_v(x)$. Notice that $x \geq \mu_{\hat{v}}$ if and only if $\mu_v \geq \mu_{\hat{v}}$. Hence, in an adversely selected market with high coverage, the wedge Γ/Γ^n further increases if there is a pessimistic bias next to heterogeneity in perceptions. Similarly, changes in the symmetry of the distribution of the actual or perceived values would require the use of $\mathcal{P}^x$ rather than $\mathcal{P}$. Graphically, heterogeneity in perceptions induces a rotation of the value curve relative to the demand curve around $(p, q) = (\mu_{\hat{v}}, 0.5)$, while an average optimistic or pessimistic bias introduces a shift and thus changes the intersection of the demand and the value curve. Similarly, if liquidity constraints or inertia stop individuals from buying insurance, the value curve will be a rotation of the demand curve around $(p, q) = (\hat{v}_{\max}, 0)$. The demand curve would underestimate the actual value of insurance, but heterogeneity in liquidity constraints or inertia causes the bias to be particularly large for the insured relative ot the uninsured.

4 Policy Interventions

The cost of inefficient pricing due to adverse selection determines the welfare gain from correcting interventions. The analysis in the previous section suggests that the welfare gain may be substantially higher when accounting for non-welfarist heterogeneity. In this section, I analyze the welfare gain for the policy interventions that are currently in place in insurance markets and find that these policy interventions are differently affected by the nature of the heterogeneity driving the demand for insurance. To focus the analysis, I continue to assume normal heterogeneity and consider an adversely selected market with high coverage and non-welfarist constraints causing the value curve to be a counter-clockwise rotation of the demand curve, as discussed in the previous section.

4.1 Price vs. Quantity

The most common interventions in insurance markets are price subsidies and insurance mandates. The question whether health insurance should be subsidized or mandated still plays a central role in the policy debate in various countries. While in some circumstances price and quantity policies are equivalent (Weitzman 1974), this is no longer the case when perceived and actual values do not coincide. Price subsidies leave the choice to buy insurance to individuals. While the actual value of insurance determines the welfare impact of such a price policy, the perceived value determines how big the price incentives need to be. Encouraging the purchase of insurance through a price policy is more costly the less the value of insurance is appreciated. In contrast, a mandate forces an individual to buy insurance, regardless of her perceived value of the contract.[22]

To compare the effectiveness of the two types of policies in the presence and absence of non-welfarist heterogeneity, I consider an *efficient-price subsidy* and a *universal mandate*, following Einav et al. (2010a). An efficient-price subsidy reduces the price paid by the insured to the efficient price p^*. By inducing the pool of inefficiently uninsured individuals to buy insurance, the welfare gain from such subsidy equals Γ. The cost from such a subsidy equals $\Phi^S = \lambda q^* \left[p^c - p^*\right]$, where λ is the cost of public funds. A counter-clockwise rotation of the value curve due to non-welfarist heterogeneity unambiguously increases Γ, but also increases the cost of implementing the subsidy Φ^S by reducing the efficient price p^*. The change in the net welfare gain $\Gamma - \Phi^S$ is thus ambiguous. By forcing everyone to buy insurance, a universal mandate realizes the welfare gain Γ, but also entails a welfare cost $\Phi^M = \int^{p^*} \left[MC\left(p\right) - MV\left(p\right)\right] dD\left(p\right)$, since for individuals with perceived value below p^*, the expected surplus of insurance is negative. A counter-clockwise rotation of the value curve does not only increase the gain Γ, but also decreases the cost Φ^M. In line with Propositions 1 and 2, the presence of non-welfarist heterogeneity increases the insurance value for the uninsured and thus unambiguously increases the net welfare gain $\Gamma - \Phi^M$ from a universal mandate.[23]

Policy Result 1 *The presence of non-welfarist heterogeneity underlying the demand curve makes a universal mandate more desirable relative to an efficient-price policy.*

A naive policy maker underestimates the welfare gain Γ, but also overestimates the welfare cost Φ^M and thus will underestimate the value of a universal mandate. When intending to induce the efficient price, a naive policy maker would implement a subsidy equal to $p^c - p^n$ that is too small. Therefore, an additional advantage of the universal mandate is that the implementation requires no knowledge regarding the heterogeneity driving the demand for insurance.

Calibration EFC evaluate the welfare gains and losses from an efficient-price subsidy and a universal mandate based on the estimated demand and cost curves. Setting the cost of public funds λ equal to 0.3, EFC find that the welfare cost of the efficient price subsidy Φ^S equals \$45 per employee per year, almost five times as large as the welfare gain Γ. Table 2 shows how the implied estimates would change when the relative importance of non-welfarist heterogeneity underlying the estimated demand curve increases. The net loss from the efficient-price subsidy becomes even larger. Despite the increased social gain, the willingness to pay for insurance of the employees for whom insurance

[22] Notice that people's willingness to accept or vote for a mandate will depend on the perceived values.

[23] Notice that the results could be easily restated by considering an increase in the dispersion of perceived values causing a clockwise rotation of the demand curve, but keeping the value curve fixed. In case of independence, this is simply implied by an increase in the variance of the noise term. This would keep the level of efficient coverage fixed, but reduce the efficient price that induces that level of coverage.

is marginally efficient drops substantially. A larger subsidy is required to induce these employees to buy insurance. The net gain from a universal mandate unambiguously increases when non-welfarist constraints become more important. Columns (1) and (2) show that the increase in the net gain from a universal mandate dominates the change in the welfare gain from a price subsidy, in line with Policy Result 1. The calibration also illustrates that the source of heterogeneity may change the net welfare impact of a policy intervention and thus the decision to implement it or not. Without non-welfarist constraints, the estimates of EFC imply that a universal mandate decreases welfare by \$20 per employee. When more than 17 percent of the variation in demand, left unexplained by risks, is driven by non-welfarist constraints, the conclusion is reversed and a universal mandate becomes welfare increasing.

4.2 Information Policies

When choices are distorted by the presence of constraints, a natural government intervention is to alleviate these contraints. The provision of information, for example, can reduce information frictions and help individuals to improve the quality of insurance choices, as recently illustrated in the context of Medicare Part D by Kling et al. (2012). The issue with these interventions is that the pool of insured and thus the equilibrium price is affected. While an individual is always better off when unconstrained, if the more constrained individuals are more costly risk types, the intervention will increase the equilibrium price and reduce coverage in equilibrium. While the impact on welfare is ambiguous, the framework allows disentangling the two opposing effects precisely.

Consider two information policies; the first policy increases the correlation between the actual risk π and the *perceived* risk $\hat{\pi} \equiv \pi + \varepsilon$, the second policy increases the correlation between the actual net-value r and the *perceived* net-value $\hat{r} \equiv r + \varepsilon$.[24] The policies leave the aggregate demand for insurance as a function of the price unchanged, but change the selection of individuals buying insurance for a given price. The first policy induces 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 price 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 surplus is generated for those buying insurance, but less individuals buy insurance. Hence, the competitive surplus $S^c = E_\zeta \left(r | \hat{v} \geq p^c \right) \Pr \left(\hat{v} \geq p^c \right)$ is unambiguously lower.

Policy Result 2 *A policy that increases the correlation between the actual and perceived risk, ceteris paribus, unambiguously reduces the competitive surplus.*

[24] An alternative interpretation is that the information policy reduces the variance in the noise term. The noise term is independent of r, but negatively related to π in the first policy and vice versa in the second policy (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 the risk and preference term respectively.

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 competitive price remains unchanged as the expected cost of the individuals buying insurance is not affected. Hence, the competitive surplus unambiguously increases.

Policy Result 3 *A policy that increases the correlation between the actual and perceived net-value, ceteris paribus, unambiguously increases the competitive surplus.*

Better information induces people to make better decisions, but may increase the scope for adverse selection. The potential trade-off can be avoided 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 individuals buying insurance. Interestingly, if the policy is well designed, it will increase welfare, regardless of the exact nature and magnitude of the information frictions. In particular, for CARA-preferences and normally distributed risks, the net-value of insurance equals the risk premium, which depends on both the risk aversion and the variance of the risk. Providing individual-specific information about the variance of their risk increases welfare, while information about their expected risk decreases welfare. The trade-off is similar when other constraints drive a wedge between the perceived and actual value, but identifying policies that leave the equilibrium price unaffected may be more difficult. When switching costs prevent individuals from buying a new insurance contract, as considered by Handel (2010), a policy that reduces the switching costs will be welfare decreasing when the individuals facing higher switching costs face higher risks.

Calibration I use the empirical analysis in EFC to shed further light on the potential trade-off when eliminating non-welfarist constraints. In particular, I analyze how the welfare impact of a noise-reducing policy depends on the nature of the non-welfarist heterogeneity. Since these policies would change the selection of employees buying insurance contracts, the cost functions need to be recalibrated. I assume that all curves are linear as before, with the slopes depending on the covariance matrix of (π, r, ε).[25] I calibrate the covariance matrix under three different scenarios regarding the correlation between the noise term and the other demand components. I assume that an initial value for $cov\left(\varepsilon, \hat{v}\right)/cov\left(\varepsilon + r, \hat{v}\right)$ of 0.25, capturing the relative importance of non-welfarist heterogeneity underlying the demand function. This may well be conservative given the previous back-of-the-envelope calculations. Table 3 shows for each scenario how a reduction in the variance of the noise term σ_ε^2 changes equilibrium welfare S^c.

[25] The slope of the marginal cost curve and value curve equal $\frac{cov(\pi,\hat{v})}{var(\hat{v})}p'\left(q\right)$ and $\frac{cov(\pi+r,\hat{v})}{var(\hat{v})}p'\left(q\right)$ respectively.

Part of the welfare change is due to the changed cost of inefficient pricing Γ in the new equilibrium, which is also shown in the table.[26,27]

The first scenario assumes that the three demand components π, r and ε are independent. A reduction in σ_ε increases the net-value of the insured, but also their expected cost and thus the equilibrium price. While in theory the net impact is ambiguous, column (1a) in table 3 shows that in this case the first effect dominates. When the information policy eliminates all non-welfarist heterogeneity, the surplus in the new equilibrium is 4 percent higher. The cost of inefficient pricing has decreased from \$18.6 to \$14.2. The second and third scenario disentangle the importance of the two opposing effects. The second scenario assumes that the underlying dispersion in perceived and actual risks is the same (i.e., $var(\pi + \varepsilon) = var(\pi)$) such that the reduction in σ_ε is equivalent to increasing the correlation between the perceived and actual risks, like in Proposition 2. The information policy induces the more costly types to buy insurance and thus worsens the adverse selection. Welfare is lower in the new equilibrium and the cost of inefficient pricing has increased. With all non-welfarist heterogeneity eliminated, welfare decreases by 3 percent, while the cost of inefficient pricing has increased to \$26.5. Finally, the third scenario assumes that the underlying dispersion in perceived and actual net-value is the same (i.e., $var(r + \varepsilon) = var(r)$) such that a reduction in σ_ε is equivalent to increasing the correlation between the perceived and actual net-value, like in Proposition 3. The information policy improves the selection of individuals, without affecting the equilibrium price and thus welfare increases. With all non-welfarist heterogeneity eliminated, welfare increases by 12 percent, while the cost of inefficient pricing is halved. The welfare consequence of reducing non-welfarist constraints thus crucially depends on their specific nature.

4.3 Risk-Adjusted Pricing

The equilibrium of a competitive insurance market is inefficient when all the insured pay a uniform price for insurance, regardless of their risks. Adjusting the price to reflect an individual's risk could reduce adverse selection, but also introduces inequality between higher and lower risk types, with the higher risk types facing higher prices for the same insurance contract. While the equity argument has inspired more regulation of risk-adjustment pricing in recent times (e.g., the ban on gender discrimination in insurance pricing by the European Court of Justice), some recent work emphasizes the efficiency

[26] For each scenario, the variance $\sigma_{\hat{v}}^2$ is calibrated as follows. The estimated linear slope equals $-1/0.0007$, while the slope of the normal demand curve equals $p'(q) = \sigma_{\hat{v}} \left[\Phi^{-1}\right]'(1-q)$. Looking at the estimated demand for $q = 0.5$ and $q = 0.7$, in between which all observations in EFC are, we find

$$\frac{\Phi^{-1}(1-0.5) - \Phi^{-1}(1-0.7)}{0.5 - 0.7} = -2.5.$$

Hence, I set $\sigma_{\hat{v}} = 571.43 = \frac{1}{2.5 \times 0.0007}$.

[27] Table App3 in the web appendix shows that the results are again very similar when the demand components are normally distributed.

argument, showing that risk-adjusted pricing may substantially increase the net surplus generated in equilibrium (Bundorf, Levin and Mahoney 2011). However, the efficiency gain from adjusting premia to an individuals risk types crucially depends on these individuals perceiving their risk types accurately. Otherwise, the risk-adjustment may decrease rather than increase the net surplus generated in equilibrium.

Consider the adjusted insurance premium $p + \beta(\pi)$ for an individual with risk π, with the adjustment $\beta(\pi)$ weakly increasing in π and equal to 0 if $\pi = \mu_\pi$.[28] In general, the risk-adjustment can be only based on observable dimensions of the expected risk, but perfect risk-adjusted pricing is obtained when $\beta(\pi) = \pi - \mu_\pi$. An individual now buys insurance if and only if

$$\hat{v}(\zeta) \geq p + \beta(\pi(\zeta)) \Leftrightarrow \hat{v}^\beta(\zeta) \geq p,$$

where $\hat{v}^\beta(\zeta)$ denotes the perceived value of insurance net of the risk-adjustment. The cost for the insurer, net of the risk-adjustment, now equals

$$\begin{aligned} AC^\beta(p) &= E_\zeta\left(\pi - \beta(\pi) | \hat{v}^\beta \geq p\right). \\ MC^\beta(p) &= E_\zeta\left(\pi - \beta(\pi) | \hat{v}^\beta = p\right). \end{aligned}$$

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

The cost of adverse selection depends on the wedge between the competitive and the efficient price and the selection of individuals buying insurance at the competitive price. Adjusting the insurance price for risks affects both, but the selection effect crucially depends on the heterogeneity in risk perceptions. Pricing the risk changes the surplus generated for a given p,

$$E_\zeta\left(r | \hat{v}^\beta \geq p\right).$$

Intuitively, the insurance surplus will be higher the more risk preferences rather than any other variable drive the demand for insurance. Since the risk type π does not affect the net value of insurance, reducing the role that risk plays in the decision to buy insurance increases the equilibrium surplus. The issue here is that when perceived risks are different from true risks, adjusting the prices for the true risks does not reduce the impact of risk on insurance decisions as much. In fact, the impact of risk may even increase. For example, when the preference term is independently distributed, risk-adjusted pricing increases the surplus at a given equilibrium price only if

$$var(\pi + \varepsilon) \geq var(\pi + \varepsilon - \beta(\pi)).$$

For perfect risk-adjusted pricing, this simplifies to $\rho_{\varepsilon,\pi} \geq -\frac{1}{2}\frac{\sigma_\pi}{\sigma_\varepsilon}$. A negative correlation

[28] As the average risk is reflected in the competitive price, we can analyze risk-adjustments depending on an individual's risk relative to the average risk without loss of generality.

between the risk and noise term below this lower bound causes the introduction of risk-adjusted pricing to reduce the surplus.

While the surplus may be lower for a given price, risk-adjusted pricing will also lower the equilibrium price and thus increase equilibrium coverage, $\Pr_\zeta\left(\hat{v}^\beta \geq p^c\right)$. Pricing the risk (or part of the risk) mechanically reduces the difference between the average and marginal net-cost of providing insurance, 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. Pricing the risk also makes high risk types less likely to buy insurance and low risk types more likely to buy insurance. Both effects lower the average cost curve and thus the competitive price.[29] The Proposition considers two extreme cases to illustrate the opposing effects on the surplus in the competitive equilibrium, $S^c = E_\zeta\left(r|\hat{v}^\beta \geq p^c\right)\Pr_\zeta\left(\hat{v}^\beta \geq p^c\right)$.

Policy Result 4 *With accurate risk perceptions, $\hat{\pi} = \pi$, perfect risk-adjusting pricing unambiguously increases the equilibrium surplus. With no heterogeneity in risk perceptions, $\hat{\pi} = E(\pi)$, perfect risk-adjusted pricing unambiguously decreases the equilibrium surplus.*

When deciding whether or not to buy insurance, an individual does not internalize the cost she is imposing on the insurer. Perfect risk-adjusted pricing corrects this type of externality and induces an efficient decision if risk perceptions are accurate. If not, an individual does not accurately internalize the value of buying insurance for herself either. With no heterogeneity in perceived risks, $\hat{\pi} = E(\pi)$, this 'internality' exactly offsets the externality such that the introduction of risk-adjusted pricing creates the inefficiency that it is supposed to eliminate. The two considered cases are extreme, but make the policy implications very clear. By ignoring the heterogeneity in risk perceptions, a naive policy maker is likely to overestimate the efficiency gain realized by risk-adjusted pricing.

Calibration I build again on the empirical analysis in EFC to shed more light on the welfare impact of risk-adjusted pricing. I consider a linear risk-adjustment of the insurance premium $\beta(\pi) = \beta[\pi - \mu_\pi]$, where $\beta = 1$ implies perfect risk-adjusted pricing. Like for the information policies, I simulate the demand and cost curves and calculate the change in the competitive surplus S^c. I also report the cost of inefficient pricing Γ in the new equilibrium, driven by the wedge between p^c and p^*.

The first two columns (0a) and (0b) in Table 4 show the positive welfare impact of risk-adjusted pricing in the absence of non-welfarist heterogeneity like in EFC. Equilibrium welfare increases by up to 11 percent when the risk-adjustment is perfect, $\beta = 1$. The new equilibrium is first-best. At the new equilibrium price $p^c = p^*$, individuals

[29] The higher price for high-risk types induces more advantageous selection and may even lead to an increasing average cost curve. With perfect risk-adjusted pricing, the cost to the insurer is independent of whom is buying insurance. The average and marginal cost curve coincide.

buy insurance if and only if $r \geq 0$. The reduction in Γ due to the elimination of the inefficient wedge between the equilibrium and efficient price accounts for about one third of the welfare increase. These estimates are very similar to the estimates in Bundorf et al. (2011), analyzing the choice between HMO plans and PPO plans offered by 11 employers in the United States between 2004 and 2005. Bundorf et al. (2011) allow for private information about risks next to the observed risk scores, but assume accurate risk perceptions. They find a potential welfare increase of 2-11 percent from pricing the observable risk, where about one fourth is due to eliminating the wedge between the equilibrium and efficient price.

The remaining columns of Table 4 show how different the welfare conclusions are when the actual and perceived risks do not coincide. Like for the information policies, I assume an initial value for $cov(\varepsilon, \hat{v}) / cov(\varepsilon + r, \hat{v})$ of 0.25 and consider three different scenarios showing the importance of the negative correlation between the noise term and the actual risk $\rho_{\varepsilon,\pi}$. The first two scenario's are the same as for the information policies. In the first scenario, all components are independent and the welfare impact of risk-adjusted pricing is hardly affected (columns (1a) and (1b)). In the second scenario, the noise term is negatively correlated with the actual risk such that the dispersion in perceived and actual risks is the same, i.e., $var(\hat{\pi}) = var(\pi)$. Risk-adjusted pricing still increases welfare, but the increase is reduced to 7 percent for $\beta = 1$ (column 2a). The new third scenario increases the size of the negative correlation $\rho_{\varepsilon,\pi}$ reducing the variance in perceived risk to half of the variance in actual risk, i.e., $var(\hat{\pi}) = 0.5 var(\pi)$. This scenario illustrates that noisy risk perceptions may not only reduce but even reverse the positive welfare effect of risk-adjusted pricing. With little dispersion in the perceived relative to the actual risks, risks hardly affect the insurance choice. However, risk-adjusted pricing changes this, reducing the prices for the low risk types and thus inducing them to buy insurance, regardless of the net-value of insurance for these types. The opposite is true for the high risk types. The market thus becomes more advantageously selected. When the risk-adjustment is less than perfect, the advantageous selection initially offsets the adverse selection and thus increases welfare, as shown in column (3a). However, with perfect risk-adjustment, the inefficiently low price of insurance increases the cost Γ and welfare is reduced by 3 percent.[30]

5 Conclusion

What drives the heterogeneity in the demand for insurance? This difficult question has been central in a recent, but already prominent empirical literature. While a number of recent empirical studies suggest that what drives the selection into insurance con-

[30] Table App4 in the web appendix shows the equilibrium welfare and cost of adverse selection when the demand components are normally distributed. The results are very similar.

tracts is often unrelated to the actual value of these contracts, the studies analyzing the importance of adverse selection in insurance markets, have mostly evaluated potential government interventions under the assumption that individuals' choices reveal the actual value of insurance. This paper provides a simple framework to analyze the consequences of heterogeneity in the differences between 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 depend on the source of the heterogeneity underlying the demand for insurance. A calibration of the model illustrates that for plausible differences between the actual and perceived value of insurance, the policy conclusions are substantially different.

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.

Bundorf, K., Levin, J. and N. Mahoney, 2011. Pricing and Welfare in Health Plan Choice. American Economic Review. forthcoming

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, Jullien, B., Salanié, B., Salanié, F., 2006. Asymmetric Information in Insurance: General Testable Implications. Rand Journal of Economics 37, 783-798.

Chiappori, P., and B. Salanié, 2012. Asymmetric Information in Insurance Markets: Predictions and Tests. forthcoming in Handbook of Insurance, 2nd edition (G. Dionne, ed.).

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.

Einav, L., Finkelstein, A., Cullen, M., and I. Pascu 2011. How general are risk preferences? Choices under uncertainty in different domains. American Economic Review. forthcoming

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.

Kling, J., Mullainathan, S., Shafir, E., Vermeulen, L. and M. Wrobel. Comparison Friction: Experimental Evidence from Medicare Drug Plans. Quarterly Journal of Economics, 127(1), 199-235.

Mullainathan, S., Schwartzstein, J., and W. Congdon, 2012. A Reduced Form Approach to Behavioral Public Finance. forthcoming Annual Review of Economics.

Peacock, W., Brody, S., and W. Highfield, 2005. Hurricane Risk Perceptions among Florida's Single Family Homeowners. Landscape and Urban Planning, 73(2-3), 120-135.

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.

Sandroni, A., Squintani, F., 2007. Overconfidence, insurance and Paternalism. American Economic Review 97(5), 1994-2004.

Sandroni, A., Squintani, F., 2010. Overconfidence and Adverse Selection: The Case of Insurance. working paper

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.

Tables

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

Noise Ratio	Cost of Adverse Selection		
$\dfrac{cov\,(\varepsilon, \hat{v})}{cov\,(\varepsilon + r, \hat{v})}$	Γ (1)	Γ/S^* (2)	Γ/Γ^n (3)
0	9.5	.04	1
.01	9.8	.04	1.03
.10	12.4	.06	1.31
.25	18.6	.10	1.95
.50	38.4	.25	4.03
1	96.6	.62	10.1

Column (1) shows the actual cost of inefficient pricing due to adverse selection Γ expressed in \$ / indiv. Column (2) expresses this actual cost relative to the surplus S^* when the price is (constrained) efficient $p = p^*$. Column (3) expresses this actual cost relative to the estimated cost when ignoring non-welfarist noise, Γ^n. The first row corresponds to the welfare estimates in Einav, Finkelstein and Cullen (2010a), assuming the absence of non-welfarist heterogeneity. The covariance ratio $cov(\varepsilon, \hat{v})/cov(\varepsilon + r, \hat{v})$ captures the importance of non-welfarist heterogenity relative to preference heterogeneity in explaining insurance choices.

Table 2: The welfare gain of subsidies and mandates

Noise Ratio	Government Interventions	
$\dfrac{cov\,(\varepsilon, \hat{v})}{cov\,(\varepsilon + r, \hat{v})}$	Price Subsidy $\Gamma - \Phi^S$ (1)	Universal Mandate $\Gamma - \Phi^M$ (2)
0	-35.4	-19.8
.01	-35.7	-18.6
.10	-37.2	-8.1
.25	-41.1	9.3
.50	-125.7	38.4
1	-67.2	96.6

Column (1) shows the net welfare gain from the efficient-price subsidy closing the gap between the equilibrium price p^c and the efficient price p^*, with $\Phi^S = \lambda q^*[p^c - p^*]$. Column (2) shows the net welfare gain from a universal mandate obliging all individuals to buy insurance, where Φ^M denotes the welfare loss from mandating individuals with expected valuation below the expected marginal cost to buy insurance. The first row corresponds to the welfare estimates in Einav, Finkelstein and Cullen (2010a), assuming the absence of non-welfarist heterogeneity. The covariance ratio $cov(\varepsilon, \hat{v})/cov(\varepsilon + r, \hat{v})$ captures the importance of non-welfarist noise relative to welfarist noise in explaining insurance choices, conditional on risk.

Table 3: The Welfare Impact of Information Policies.

Noise Reduction	**Scenario 1**		**Scenario 2**		**Scenario 3**	
	Independence		$var\,(\pi+\varepsilon) = var\,(\pi)$		$var\,(r+\varepsilon) = var\,(r)$	
$\Delta\sigma^2_\varepsilon/\sigma^2_\varepsilon$	$\Delta S^c/S^c$ (1a)	Γ (1b)	$\Delta S^c/S^c$ (2a)	Γ (2b)	$\Delta S^c/S^c$ (3a)	Γ (3b)
0	0	18.6	0	18.6	0	18.6
.10	.00	18.1	$-.00$	19.2	.01	17.4
.25	.01	17.5	$-.01$	20.3	.03	15.7
.50	.02	16.4	$-.01$	22.2	.06	13.3
1	.04	14.2	$-.03$	26.5	.12	9.5

Columns (1a),(2a) and (3a) show the change in equilibrium welfare $S^c = E_\zeta(r|\hat{v} \geq p^c)\Pr(\hat{v} \geq p^c)$ when reducing the variance in noise under the three respective scenario's (relative to the case with no noise reduction). Columns (1b),(2b) and (3b) show the welfare cost in the new equilibrium due to the inefficient pricing Γ. Scenario 1 assumes independence between r, π and ε. Scenario 2 assumes that the variance in perceived risks equals the variance in actual risks. Scenario 3 assumes that the variance in perceived net-values equals the variance in actual net-values. The three scenario's start from an initial value for $cov(\varepsilon,\hat{v})/cov(\varepsilon+r,\hat{v})$ equal to .25. Notice that equilibrium welfare equals S^c= \$243 given this initial value. The demand, value and cost curves are linear with the slopes determined like with normal heterogeneity.

Table 4: The Welfare Impact of Risk-Adjusted Pricing.

Risk Adj.	**No Noise**		**Scenario 1**		**Scenario 2**		**Scenario 3**	
			Independ.		$var\,(\pi+\varepsilon) = var\,(\pi)$		$var\,(\pi+\varepsilon) = \frac{1}{2}var\,(\pi)$	
β	$\Delta S^c/S^c$ (0a)	Γ (0b)	$\Delta S^c/S^c$ (1a)	Γ (1b)	$\Delta S^c/S^c$ (2a)	Γ (2b)	$\Delta S^c/S^c$ (3a)	Γ (3b)
0	0	9.5	0	18.6	0	18.6	0	18.6
.10	.02	6.8	.02	15.0	.02	14.5	.03	13.2
.25	.05	3.6	.05	10.5	.05	9.9	.05	8.5
.50	.08	.8	.09	5.4	.08	5.7	.06	6.8
.75	.10	.1	.11	2.9	.08	4.7	.02	10.6
1	.11	0	.11	2.2	.07	6.4	$-.03$	19.8

Columns (0a),(1a),(2a) and (3a) show the change in equilibrium welfare $S^c = E_\zeta(r|\hat{v} \geq p^c)\Pr(\hat{v} \geq p^c)$ for positive linear shares of the risk-premium adjustment $\beta(\pi) = \beta[\pi - \mu_\pi]$ (relative to the case with no risk-adjustment, $\beta = 0$). Columns (0b),(1b),(2b) and (3b) show the welfare cost in the new equilibrium due to the inefficient pricing Γ. Scenario 1 assumes independence between r, π and ε. Scenario 2 assumes that the variance in perceived risks equals the variance in actual risks. Scenario 3 assumes that the variance in perceived net-values equals the variance in actual net-values. The three scenario's start from an initial value for $cov(\varepsilon,\hat{v})/cov(\varepsilon+r,\hat{v})$ equal to .25. Notice that equilibrium welfare equals S^c= \$243 given this initial value, while it equals S^c= \$272 without noise.

Appendix: Proofs

Proof of Proposition 1

I assume that the random variables are draws from continuous 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

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

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 $\sigma_{\hat{v}} \geq \sigma_v \rho$.$\square$

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

$$\begin{aligned}\Gamma^n &= \int_{p^n}^{p^c} [p - MC(p)]\, dD(p) \\ &= \int_{p^n}^{p^c} \left[p - \frac{cov(\pi, \hat{v})}{var(\hat{v})}[p - \mu_v] - \mu_\pi\right] dD(p) \\ &= \int_{p^n}^{p^c} \left(1 - \frac{cov(\pi, \hat{v})}{var(\hat{v})}\right)[p - p^n]\, dD(p),\end{aligned}$$

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,

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

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^*. By normality, we find that

$$\begin{aligned} p - MC(p) &= \frac{cov(r+\varepsilon,\hat{v})}{var(\hat{v})}[p - p^n], \\ MV(p) - MC(p) &= \frac{cov(r,\hat{v})}{var(\hat{v})}[p - p^*], \end{aligned}$$

since respectively $p^n = MC(p^n)$ and $MV(p^*) = MC(p^*)$. Moreover, notice that at $p = \mu_v$,

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

Hence,

$$\frac{cov(r+\varepsilon,\hat{v})}{var(\hat{v})}[\mu_v - p^n] = \frac{cov(r,\hat{v})}{var(\hat{v})}[\mu_v - p^*].$$

Rearranging, we find

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

Substituting this in the previous expression, we find

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

Proof of Policy Result 1

We consider the impact of a counter-clockwise rotation of the value curve, keeping the demand and cost functions unchanged (i.e., an increase in $\frac{cov(\varepsilon,\hat{v})}{cov(r,\hat{v})}$, keeping $cov(\pi,\hat{v})$ and $var(\hat{v})$ fixed). The counter-clockwise rotation increases $MV(p)$ for all $p \leq \mu_v$. It also increases the efficient price p^*, solving $MV(p) = MC(p)$ and thus

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

Hence,

$$\begin{aligned} p^* &= \mu_{\hat{v}} - \mu_r \frac{var(\hat{v})}{cov(r,\hat{v})} \\ &= \mu_r\left(\frac{cov(r,\hat{v}) - var(\hat{v})}{cov(r,\hat{v})}\right) + \mu_\pi \\ &= \mu_\pi - \mu_r\left(\frac{cov(\varepsilon,\hat{v})}{cov(r,\hat{v})} + \frac{cov(\pi,\hat{v})}{cov(r,\hat{v})}\right). \end{aligned}$$

In an adversely selected and thick market, $\mu_r > 0$ and $cov(\pi,\hat{v})$. Hence, an increase in $cov(\varepsilon,\hat{v})$ or decrease in $cov(r,\hat{v})$ decreases the efficient price p^*. The competitive price p^c, however, remains the same.

Hence, the cost of the efficient-price subsidy $\Phi^S = \lambda q^* [p^c - p^*]$ thus increases, since p^* decreases and $q^* = D(p^*)$ increases. The cost of the universal mandate $\Phi^M = \int_{-\infty}^{p^*} [MC(p) - MV(p)]\, dD(p)$ decreases, since $MC(p) \geq MV(p)$ for $p \leq p^*$. Since the gain from both policies is the same Γ, this proves the proposition.□

Proof of Policy Result 2

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, $E(r|\hat{v} = p) \geq 0$, 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 competitive price $p^c = AC(p^c)$ increases. The welfare surplus, $\int_{p^c}^{\infty} E(r|\hat{v} = p)\, dF(p)$, decreases unambiguously.□

Proof of Policy Result 3

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 competitive 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,

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

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

Proof of Policy Result 4

Consider first the case with accurate risk perceptions, $\hat{\pi} = \pi$. With perfect risk-adjusted pricing, $\beta(\pi) = \pi - \mu_\pi$, the average cost $E(\pi - \beta(\pi)|\pi + r \geq p + \beta(\pi)) = \mu_\pi$, independent of the price. Hence, $p^c = \mu_\pi$. An individual thus buys insurance if and only if

$$\pi + r \geq p^c + \beta(\pi) \Leftrightarrow r \geq 0.$$

This is the first-best. Hence, perfect risk-adjusted pricing improves welfare in an adversely selected market. Consider now the case with no heterogeneity in risk perceptions, $\hat{\pi} = E(\pi)$. Without risk-adjusted pricing, the average cost $E(\pi | \mu_\pi + r \geq p) = \mu_\pi$, independent of the price. Hence, $p^c = \mu_\pi$. An individual thus buys insurance if and only if

$$\mu_\pi + r \geq p^c \Leftrightarrow r \geq 0.$$

This is the first-best. However, with perfect risk-adjusted pricing, the competitive price still equals $p^c = \mu_\pi$. However, an individual buys insurance if and only $r \geq \pi$, which is inefficient.$\square$