

NBER WORKING PAPER SERIES

MEASURING EX-ANTE WELFARE IN INSURANCE MARKETS

Nathaniel Hendren

Working Paper 24470

<http://www.nber.org/papers/w24470>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

March 2018, Revised October 2019

I am very grateful to Raj Chetty, David Cutler, Liran Einav, Amy Finkelstein, Ben Handel, Pat Kline, Tim Layton, Mark Shepard, Mike Whinston, seminar participants at the NBER Summer Institute and University of Texas, along with five anonymous referees and the editor, Adam Sziedl, for helpful comments and discussions. I also thank Kate Musen for helpful research assistance. Support from the National Science Foundation CAREER Grant is gratefully acknowledged. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Nathaniel Hendren. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Measuring Ex-Ante Welfare in Insurance Markets
Nathaniel Hendren
NBER Working Paper No. 24470
March 2018, Revised October 2019
JEL No. H0,I11,I3

ABSTRACT

The willingness to pay for insurance captures the value of insurance against only the risk that remains when choices are observed. This paper develops tools to measure the ex-ante expected utility impact of insurance subsidies and mandates when choices are observed after some insurable information is revealed. The approach retains the transparency of using reduced-form willingness to pay and cost curves, but it adds one additional sufficient statistic: the difference in marginal utilities between insured and uninsured. I provide an approach to estimate this statistic that uses only reduced-form willingness to pay and cost curves, combined with either a measure of risk aversion. I compare the approach to structural approaches that require fully specifying the choice environment and information sets of individuals. I apply the approach using existing willingness to pay and cost curve estimates from the low-income health insurance exchange in Massachusetts. Ex-ante optimal insurance prices are roughly 30% lower than prices that maximize market surplus. While mandates would increase deadweight loss, the results suggest they would actually increase ex-ante expected utility.

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

Measuring Ex-Ante Welfare in Insurance Markets

Nathaniel Hendren*

October, 2019

Abstract

The willingness to pay for insurance captures the value of insurance against only the risk that remains when choices are observed. This paper develops tools to measure the ex-ante expected utility impact of insurance subsidies and mandates when choices are observed after some insurable information is revealed. The approach retains the transparency of using reduced-form willingness to pay and cost curves, but it adds one additional sufficient statistic: the difference in marginal utilities between insured and uninsured. I provide an approach to estimate this statistic that uses only reduced-form willingness to pay and cost curves, combined with either a measure of risk aversion. I compare the approach to structural approaches that require fully specifying the choice environment and information sets of individuals. I apply the approach using existing willingness to pay and cost curve estimates from the low-income health insurance exchange in Massachusetts. Ex-ante optimal insurance prices are roughly 30% lower than prices that maximize market surplus. While mandates would increase deadweight loss, the results suggest they would actually increase ex-ante expected utility.

1 Introduction

Revealed preference theory is often used as a tool for measuring the welfare impact of government policies. Many recent applications use price variation to estimate the willingness to pay for insurance (Einav et al. (2010); Hackmann et al. (2015); Finkelstein et al. (2019); Panhans (2018)). Comparing willingness to pay to the costs individuals impose on insurers provides a measure of market surplus. This surplus potentially provides guidance on optimal

*Harvard University, nhendren@fas.harvard.edu. I am very grateful to Raj Chetty, David Cutler, Liran Einav, Amy Finkelstein, Ben Handel, Pat Kline, Tim Layton, Mark Shepard, Mike Whinston, seminar participants at the NBER Summer Institute and University of Texas, along with five anonymous referees and the editor, Adam Szedl, for helpful comments and discussions. I also thank Kate Musen for helpful research assistance. Support from the National Science Foundation CAREER Grant is gratefully acknowledged.

insurance subsidies and mandates (Feldman and Dowd (1982)). If individuals are not willing to pay the costs they impose on the insurer, then greater subsidies or mandates will lower market surplus. From this perspective, subsidies and mandates would reduce welfare and be socially undesirable.

Measures of willingness to pay are generally a gold standard input into welfare analysis. But, in insurance settings they can be misleading. Insurance obtains its value by insuring the realization of risk. Often, individuals make insurance choices after learning some information about their risk. It is well-known that this can lead to adverse selection. What is less appreciated is that observed willingness to pay will not capture value of insuring against this learned information.¹ As a result, welfare conclusions based on market surplus can vary with the information that individuals have when the economist happens to observe choices. Policies that maximize observed market surplus will not generally maximize canonical measures of expected utility.

To see this, consider the decision to buy health insurance coverage for next year. Suppose some people have learned they need to undergo a costly medical procedure next year. Their willingness to pay will include the value of covering this known cost plus the value of insuring other future unknown costs. Market surplus - defined as the difference between willingness to pay and costs - will equal the value of insuring their unknown costs. But, it will not include any insurance value from covering the known costly medical procedure. This risk has already been realized when willingness to pay is observed.

Now, consider an economist seeking to measure the welfare impact of extending health insurance coverage next year to everyone through a mandate or large subsidy. The market surplus or deadweight loss generated from the policy will depend on how much people have learned about their health costs at the time the economist happened to measure willingness to pay. Existing literature (and introspection) suggests that individuals know more about expected costs and events in the near future (e.g. Finkelstein et al. (2005); Hendren (2013, 2017); Cabral (2017)). If willingness to pay had been measured earlier, market surplus could be larger because it would include the value of insuring against the costly medical procedure. This occurs even though the economic allocation generated by a mandate does not vary depending on when the economist measures willingness to pay. In contrast, traditional notions of the expected utility impact of a mandate would not depend on when the economist happens to measure willingness to pay. Expected utility provides a consistent framework for identifying optimal insurance policies that depends on economic allocations.

The goal of this paper is to enable researchers to evaluate the impact of insurance market

¹This idea is related to Hirshleifer (1971), who shows that individuals may wish to insure against the realization of information.

policies on expected utility, where the expectation is taken prior to when insurance choices are made. Traditional methods to estimating ex-ante expected utility would estimate a structural model. Among other things, the model would specify what individuals know when choosing whether to buy an insurance plan. It would then be estimated using observed insurance choices along with data on the realized utility-relevant outcomes, such as health and consumption.² Intuitively, if one has a structural model and knows what information has been realized when individuals choose their insurance policies, one can infer the value of insuring the risk that has been revealed before making those choices. But, in practice it is especially difficult to observe individuals' information sets when they make choices. This is especially true in insurance markets that suffer from adverse selection driven by asymmetric information.

This paper develops a new approach to measure the expected utility impact of insurance market policies. The approach does not require specifying structural assumptions about individuals' information sets at the time of choice, nor does it require specifying a utility function or observing the distribution of utility-relevant outcomes in the economy. Instead, I exploit the information contained in reduced-form willingness to pay and cost curves as outlined by Einav et al. (2010). In this environment, I characterize the minimal additional sufficient statistics required to measure the expected utility of subsidies and mandates.

The first main result of the paper shows that one can measure ex-ante expected utility using one additional sufficient statistic: the difference in marginal utilities of income for those who do versus do not buy insurance. This measures how much individuals wish to move money to the state of the world in which they buy insurance. In the example above, it reflects the desire to insure the costly medical procedure. These individuals have a higher demand for insurance, and have a higher marginal utility of income.

In general, it is difficult to observe or measure differences in marginal utilities of income between those who do versus do not purchase insurance. The second result of the paper addresses this by providing a benchmark estimation method that uses only the reduced-form willingness to pay and cost curves combined with a measure of risk aversion. This additional risk aversion parameter can be assumed, or it can be inferred from the observed markup individuals are willing to pay for insurance, combined with the extent to which insurance reduces the variance in out of pocket expenditures.

This second main result follows two steps. First, building on the literature on optimal unemployment insurance (Baily (1978); Chetty (2006)), I use measures of consumption differences between insured and uninsured, combined with risk aversion, to measure differences in marginal utilities. Second, because consumption is seldom observed, I provide conditions

²For example, see Handel et al. (2015) or Section IV of Einav et al. (2016).

under which one can exploit the information in the reduced-form willingness to pay curve for insurance instead of consumption.

I apply the framework to study the optimal subsidies and mandates for low-income health insurance in Massachusetts. Finkelstein et al. (2019) use price discontinuities as a function of income to estimate willingness to pay and cost curves for those with incomes near 150% of the federal poverty level (FPL). Their results show that an unsubsidized private insurance market would unravel.³ Without subsidies, the market would not exist. I use the approach to ask what types of insurance subsidies or mandates individuals would want in the MA health insurance exchange if they had been asked prior to learning anything about whether they would actually be in the MA health insurance exchange and eligible for the subsidies in question. In this sense, it evaluates subsidies and mandates using an ex-ante notion of welfare that is both ex-ante expected utility and ex-post utilitarian social welfare.⁴

I use the framework to provide guidance on both budget neutral and non-budget neutral policies in this environment. Budget neutral policies ask whether the government should provide additional subsidies financed by increased prices/penalties for those not purchasing insurance – this is the canonical set of policies studied in Einav et al. (2010). To set the stage, traditional market surplus is maximized when insurance premiums are \$1,581 and 41% of those eligible for insurance choose to purchase. In contrast, I find that a 30% lower price of \$1,117 with 54% of the market insured maximizes ex-ante welfare. Those who are induced to purchase insurance from by lowering prices from \$1,581 to \$1,117 are not willing to pay the cost they impose on the insurance company. However, from behind a veil of ignorance they value the ability to purchase insurance at lower prices if they end up having a high demand for insurance.

One can also evaluate the welfare impact of a full insurance mandate (with sufficiently large penalties to generate full coverage). Such a policy would lower market surplus by \$45 – on average, those in the economy are not willing to pay the cost they impose on the insurance company. However, from behind a veil of ignorance individuals would be willing to pay \$169 to have a full insurance mandate. Indeed, the willingness to pay for a mandate remains positive for a wide range of plausible risk aversion parameters (e.g. with coefficients of relative risk aversion above 1.7). This illustrates how an ex-ante welfare perspective can lead to dramatically different normative conclusions about the desirability of commonly debated

³This unraveling is due to a combination of adverse selection and uncompensated care externalities.

⁴This welfare perspective provides a natural benchmark as it does not make a normative distinction between redistribution and insurance (e.g. Harsanyi (1978)). However, the approach does not require a utilitarian perspective, nor does it require measuring expected utility behind a complete veil of ignorance. One can conduct the analysis conditional on any observable subgroup, $X = x$ (e.g. old vs. young, male vs. female, with vs. without chronic health conditions, etc.). Doing so requires willingness to pay and cost curves for each subgroup.

insurance policies.

In practice, insurance subsidies in Massachusetts (as in many cases) were not paid by penalties targeting the uninsured in the market; rather they were paid by general government revenue and taxpayers more generally. To capture this, I estimate the marginal value of public funds (MVPF) of additional insurance subsidies. The MVPF for an additional insurance subsidy is the individual's willingness to pay for it divided by its net cost to the government (Hendren (2016)). Comparisons of MVPFs across policies correspond to statements about the value of redistributing across different subpopulations. For example, comparing the MVPF of insurance subsidies to low-income tax credits allows one to ask whether individuals at 150% would prefer additional insurance subsidies or prefer a tax subsidy.

The results suggest that when prices are set so that 30% of the market has insurance, individuals are willing to pay roughly 1.2 times the marginal cost they impose on the insurer to lower insurance prices. This implies an MVPF of 1.2, which is similar to the range of MVPF estimates for tax credits to low-income populations studied in Hendren and Sprung-Keyser (2019), which are around 0.9-1.3. However, those benefiting from the insurance subsidies have a higher marginal utility of income than those benefiting from a tax cut. From behind the veil of ignorance, the results suggest that individuals would be willing to pay 1.8 times the cost they impose on the insurer to lower insurance prices. Put differently, individuals would prefer from behind a veil of ignorance to spend money lowering insurance prices for those at 150% of FPL instead of providing them with a tax cut. In short, the ability to distinguish between ex-ante expected utility and market surplus provides different guidance on important questions of optimal insurance market policies.

Traditional approaches to measuring ex-ante expected utility would estimate a structural model. This would involve fully specifying not only a utility function but also the information set of individuals at the time they make insurance choices. The economic primitives estimated from the model would then provide an ex-ante measure of welfare. In contrast, the sufficient statistics approach developed here does not require researchers to know the exact utility function, nor does it require knowledge of individuals' information sets when they make insurance choices. Information sets can be particularly tough to specify in settings of adverse selection where even insurers have trouble worrying about the unobserved knowledge of the applicant pool. In addition, the approach developed here can be implemented using aggregate data from insurers or governments on the cost and fraction of the market purchasing insurance at different prices, as opposed to requiring observations on individual-level data that would be required to estimate the structural model.

To further understand the relationship to the structural approach, I develop a fully specified structural model with moral hazard and adverse selection that can fully match the

reduced form willingness to pay and cost curves in the Finkelstein et al. (2019) environment. The model builds upon the approach in Handel et al. (2015) but augments it with a moral hazard structure developed in Einav et al. (2013a). I use the model to verify that the approach developed here recovers the true ex-ante welfare quite well. However, the benchmark implementation relies on two key assumptions that may be violated in some applications. I use the structural environment to understand the impact of violating these assumptions, and validate proposed modifications to my approach that help recover ex-ante welfare when the key assumptions are violated.

First, using the demand curve to proxy for differences in consumption requires that there are no differences in liquidity or income between the insured and uninsured. While this is perhaps a natural assumption in the context subsidies to a given income level (i.e. 150% FPL in MA), it is quite restrictive in many other settings where income differences may be a key driver of willingness to pay for insurance. In these cases, I show that one can recover ex-ante welfare using my approach if one can observe the average consumption of the insured and uninsured in the market.

Second, the benchmark implementation requires that individuals have common coefficients of relative risk aversion. However, previous literature has highlighted a role of preference heterogeneity as an important driver of insurance demand. In this case, I show that the ideal measure of risk aversion is the “ex-ante” risk aversion that governs the markup that individuals would be willing to pay from behind a veil of ignorance for a product that helps lower their insurance prices. I show that the common structural models with preference heterogeneity raise the potential concern that the “ex-ante” risk aversion coefficient might differ from the average risk aversion of individuals expressed over coverage choices in the market. This means that, in practice, researchers will want to study the robustness of the results to a range of risk aversion parameters with the knowledge that the ideal risk aversion is the one that governs the hypothetical ex-ante willingness to pay for a financial product that helps reduce the price of insurance.

The ideas developed in this paper readily extend to other settings where individuals measure the value of insurance using principles of revealed preference. For example, often behavioral responses such as labor supply changes are used to measure the value of social insurance. The more individuals are willing to adjust their labor supply to become eligible for insurance, the more they value the insurance (e.g. Keane and Moffitt (1998); Gallen (2014); Dague (2014)). Such an approach captures the value of insurance against only the risk that is revealed after they adjust their behavior.⁵ Similarly, other papers infer willingness to

⁵Indeed, Gallen (2014) shows many individuals who respond have particularly costly health conditions. The measure of willingness to pay in Gallen (2014) will miss the value of insurance against those health

pay for social insurance from changes in consumption around a shock (e.g. Gruber (1997); Meyer and Mok (2013)). When information is revealed over time, the consumption change may vary depending on the time horizon used (Hendren (2017)). In the extreme, there may be no change around the event (e.g. smooth consumption around onset of disability or retirement). Consumption should change when information about the event is revealed, not when the event occurs. The methods in this paper can be applied to conceptualize which consumption difference is most appropriate for the desired notion of expected utility.

The rest of this paper proceeds as follows. Section 2 provides a stylized example that develops the intuition for the approach. Section 3 provides the general modeling framework and introduces the application from Finkelstein et al. (2019), who study health insurance subsidies for low-income adults in Massachusetts. Section 4 uses the model to define notions of ex-ante welfare and provides the general result that the ex-ante willingness to pay for insurance requires the difference in marginal utilities between insured and uninsured. Section 5 provides a benchmark method to estimate this difference in marginal utilities using willingness to pay and cost curves combined with a measure of risk aversion. Section 6 implements this approach using the estimates from MA. Section 7 develops a structural model and assesses the extent to which the proposed approach matches the structural notions of ex-ante welfare. Section 8 uses the model to assess the potential impact of violations of the implementation assumptions outlined in Section 5. Section 9 concludes.

2 Stylized Example

I begin with a stylized example to illustrate the distinction between market surplus and ex-ante expected utility and to summarize the paper’s main results. Suppose individuals have \$30 dollars but face a risk of losing \$ m dollars, where m is uniformly distributed between 0 and 10. Let $D^{Ex-ante}$ denote the willingness to pay or “demand” for a full insurance contract that is measured prior to individuals learning anything about their particular realization of m . This solves

$$u(30 - D^{Ex-ante}) = E[u(30 - m)] \tag{1}$$

where $E[u(30 - m)] = \frac{1}{10} \int_0^{10} u(30 - m) dm$ is the expected utility if uninsured.

Suppose individuals have a utility function with a constant coefficient of relative risk aversion of 3 (i.e. $u(c) = \frac{1}{1-\sigma} c^{1-\sigma}$ and $\sigma = 3$). This implies individuals are willing to pay $D^{Ex-ante} = 5.50$ for an insurance policy that fully compensates for their loss. The cost this policy would be $E[m] = 5$. Full insurance generates a market surplus of \$0.50.

conditions.

Figure 1: Example Willingness to Pay and Cost Curves

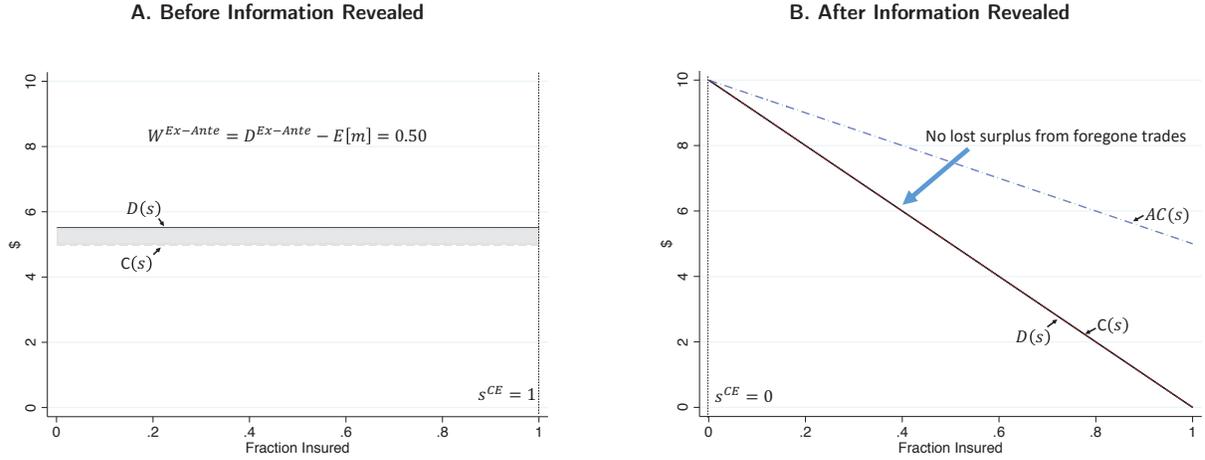


Figure 1A draws the demand and cost curves that would be revealed through random variation in prices in this environment, as formalized in Einav et al. (2010). The horizontal axis enumerates the population in descending order of their willingness to pay for insurance, indexed by $s \in [0, 1]$. The vertical axis reflects prices, costs, and willingness to pay in the market. Each individual is willing to pay \$5.50 for insurance, reflected in the horizontal demand curve of $D(s) = \$5.50$. In addition, each person imposes an expected cost of \$5 on the insurance company, which generates a flat cost curve of $C(s) = \$5$. If a competitive market were to open up in this setting, one would expect everyone to purchase insurance at a price of \$5, depicted by the vertical line at $s^{CE} = 1$. This allocation would generate $W^{Ex-Ante} = \$0.50$ of welfare, as reflected by the market surplus defined as the integral between demand and cost curve.

What happens if individuals learn about their costs before they choose whether to purchase insurance? For simplicity, consider the extreme case that individuals have fully learned their cost, m . Willingness to pay will equal individuals' known costs, $D(s) = m(s)$. Those who learn they will lose \$10 will be willing to pay \$10 for "insurance" against their loss; individuals who learn they will lose \$0 will be willing to pay nothing. The uniform distribution of risks generates a linear demand curve falling from \$10 at $s = 0$ to \$0 at $s = 1$. The cost imposed on the insurer by the type s , $C(s)$, will equal their willingness to pay of $D(s)$, as shown in Figure 1B.

If an insurer were to try to sell insurance, they would need to set prices to cover the average cost of those who purchase insurance. Let $AC(s) = E[C(S) | S \leq s]$ denote the

average cost of those with willingness to pay above $D(s)$.⁶ This average cost lies everywhere above the demand curve. Since no one is willing to pay the pooled cost of those with higher willingness to pay, the market would fully unravel. The unique competitive equilibrium would involve no one obtaining any insurance, $s^{CE} = 0$.

What is the welfare cost of this market unraveling? From a market surplus perspective, there is no welfare loss. There are no valuable foregone trades that can take place at the time insurance choices are made. This reflects an extreme case of a more general phenomenon identified in Hirshleifer (1971). The market demand curve does not capture the value of insurance against the portion of risk that has already been realized at the time insurance choices are made. This means that policies that maximize market surplus may not maximize expected utility if one measures expected utility prior to when all information about m is revealed to the individuals.

How can one recover measures of ex-ante welfare? The traditional approach to measuring $D^{Ex-Ante}$ and the value of other insurance market policies would require the econometrician to specify economic primitives, such as a utility function and an assumption about individuals' information sets at the time of choice. It would then also involve measuring the distribution of outcomes that enter the utility function, such as consumption, and use this information to infer the ex-ante value of insurance from the model. Intuitively, if one knows the utility function, u , and the cross-sectional distribution of consumption ($30 - m$ in the example above), then one can use this information to compute $D^{Ex-Ante}$ in equation (1). For recent implementations of this approach, see Handel et al. (2015), Section IV of Einav et al. (2016), or Finkelstein et al. (2016).

The goal of this paper is to measure the expected utility impact of insurance market policies, such as optimal subsidies and mandates, without knowledge of the full distribution of structural primitives in the economy (e.g. utilities, outcomes, and beliefs). Rather, the paper builds on the reduced form framework that uses price variation to identify demand and cost curves in the economy.

The core idea can be seen in the following example of a budget-neutral expansion of the market. To expand the size of the insurance market in a budget neutral way, one needs to subsidize insurance purchase and tax those who do not purchase insurance. These transfers between insured and uninsured do not affect market surplus. The market surplus from expanding the size of the insurance market from s to $s + ds$ is given by $D(s) - C(s)$. However, from an ex-ante perspective, these transfers affect welfare if the marginal utility of income is different for the insured versus uninsured.

⁶Throughout, we will let S denote the random variable corresponding to different types and s denote a number corresponding to the size of the insurance market.

The first main result shows that if individuals had been asked their willingness to pay to have a large insurance market prior to learning their risk type, they would have been willing to pay not just $D(s)$ but an additional amount $EA(s)$, where

$$EA(s) = s(1-s)(-D'(s)) \frac{E[u_c|Ins] - E[u_c|Unins]}{E[u_c]} \quad (2)$$

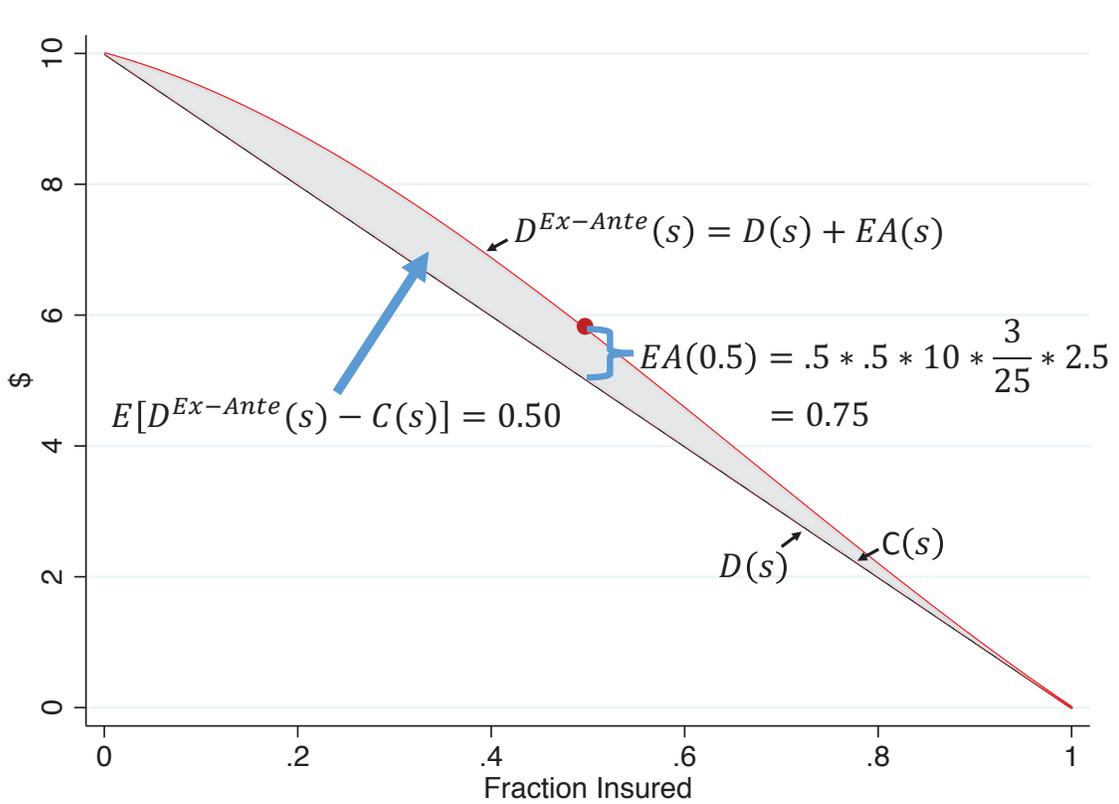
The first term, $-s(1-s)D'(s)$ characterizes the size of the transfer from uninsured to insured when expanding the size of the insurance market. The second term, $\frac{E[u_c|Ins] - E[u_c|Unins]}{E[u_c]}$, captures the value of this transfer using the difference in the marginal utilities of income between the insured and uninsured. If the insured have higher (lower) marginal utilities of income, then expanding the size of the insurance market by lowering the prices paid by the insured has ex-ante value beyond what is captured in traditional measures of market surplus.

Constructing $EA(s)$ in equation (2) requires knowledge of the difference in marginal utilities between insured and uninsured. Such differences are not directly observed. The second main result of the paper shows that if consumption levels are the only determinant of marginal utilities (as would be the case in canonical models of insurance like the model here or in more general models with additive separable utility over consumption), then one can approximate difference in marginal utilities using the difference in consumption between insured and uninsured, multiplied by a coefficient of risk aversion. However consumption is rarely observed in practice. Therefore, in a final step I show that if there are not systematic differences in income or liquidity between insured and uninsured, then one can actually proxy consumption differences using the demand curve. In the end, this yields conditions under which the difference in marginal utilities can be written as

$$\frac{E[u_c|Ins] - E[u_c|Unins]}{E[u_c]} \approx \gamma (D(s) - E[D(S)|S \geq s]) \quad (3)$$

where $D(s) - E[D(S)|S \geq s]$ is the difference in willingness to pay between the average uninsured individual and the marginal insured type and $\gamma = \frac{-u''}{u'}$ is the coefficient of absolute risk aversion. It is important to note that the assumptions needed to generate equation (3) are restrictive and stated formally in Section 5. Most notably, they require no income or liquidity differences between insured and uninsured and they assume no preference heterogeneity. While these assumptions are certainly not without loss of generality, the key advantage of this formula is that it provides a benchmark implementation to measure ex-ante expected utility with only the addition of a risk aversion coefficient. Section 8 provides a practical discussion of violations of these assumptions and what types of additional data or parameters can be useful in those cases to recover ex-ante welfare.

Figure 2: Recovering Ex-Ante Willingness to Pay



To illustrate how the formula for $EA(s)$ recovers ex-ante welfare, Figure 2 calculates $EA(s)$ for all values of $s \in [0, 1]$ using equations (2) and (3). The coefficient of risk aversion of 3 implies a coefficient of absolute risk aversion of $\gamma = \frac{3}{25}$, where 25 is the average consumption in the economy. At each value of s , $D^{Ex-Ante}(s)$ measures the impact on ex-ante expected utility of expanding the size of the insurance market from s to $s + ds$. For example, when 50% of the market is insured, the formula suggests that individuals are willing to pay an additional \$0.75 to expand the insurance market relative to what is revealed through the observed demand curve. The integral from $s = 0$ to $s = 1$ measures the ex-ante willingness to pay to fully insure the market (relative to having no one insured):

$$D^{Ex-Ante} = \int_0^1 D^{Ex-Ante}(s) = 5.50$$

Numerically integrating ex-ante demand curve in Figure 2 yields approximately \$5.50, which equals the integral under the demand curve in Figure 1A. The ex-ante demand curve recovers

the willingness to pay individuals would have for everyone to be insured ($s = 1$) if they were asked this willingness to pay prior to learning m .

The model in this section is highly stylized. There is no moral hazard, no preference heterogeneity, and the model assumed all information about costs, m , was revealed at the time of making the insurance decision. The next three sections extend these derivations to capture more realistic features of insurance markets encountered in common empirical applications and applies them to the study of health insurance subsidies to low-income adults in Massachusetts. The main result of Section 4 will be to show that equation (2) continues to be the key additional sufficient statistic required to construct the ex-ante willingness to pay for insurance. Section 5 will then establish conditions under which one can approximate this difference using the demand and cost curves combined with a measure of risk aversion, as in equation (3) above.

3 General Model

3.1 Setup

Individuals face evolving risk over their lives that is captured by the realization of a state variable, θ_t , that is realized in each period t . The realization of this state variable potentially impacts all aspects of an individual’s life, including their level of utility, their marginal utility of consumption or medical spending, their income, etc. Realized utility in period t is denoted $u_t(c, m; \theta_t)$, where c is non-medical consumption and m is medical consumption.⁷

For a sequence or realizations $\{\theta_t\}_t$ and choices $\{c_t(\theta_t), m_t(\theta_t)\}_t$, an individual’s realized utility over a lifetime that is given by $\sum_{t \geq 0} u_t(c_t(\theta_t), m_t(\theta_t); \theta_t)$. At any time t , individuals with state variable θ_t have beliefs about future states of the world that generate an expectation $E[\circ|\theta_t]$.⁸ The expected future utility of an individual from the perspective of what they know at time t is given by $E[\sum_{t' \geq t} u_{t'}(c_{t'}(\theta_{t'}), m_{t'}(\theta_{t'}); \theta_{t'})|\theta_t]$. Let $t = 0$ denote the time prior to any individuals learning that they differ from other individuals. This “ex-ante”

⁷This allows for but do not require standard discounting, e.g. $u_t(c, m; \theta_t) = \beta^t \tilde{u}(c, m; \theta_t)$. Note also that θ_t is not assumed to be uni-dimensional – it captures all elements of risk that are experienced by the individual at time t , such as income, shocks, health shocks, the quality of one’s friends, the extent to which one likes their job, and their satisfaction with referee comments about the generality of a model, etc. And, it includes individual’s information known at time t about the future evolution of these risks.

⁸The assumption that θ_t captures all aspects of an individuals’ knowledge at time t is without loss of generality because θ_t is not assumed to be uni-dimensional.

expected utility of individuals in the economy is given by

$$E \left[\sum_{t \geq 0} u_t(c_t, m_t; \theta_t) \right] = \sum_{t \geq 0} E[u_t(c_t, m_t; \theta_t)]$$

where E is taken with respect to the unconditional distribution of all possible realizations of sequences of $\{\theta_t\}_t$ in all periods. This is also equivalent to the utilitarian measure of social welfare (as shown in Harsanyi (1978)).

In each period t , individuals observe their realization of θ_t , which determines their utility function and net incomes in period t , $y_t(\theta_t)$. They then choose $c_t(\theta_t)$ and $m_t(\theta_t)$ subject to a budget constraint that depends upon whether they have purchased insurance to help cover their medical expenses, m . Insured individuals have an out of pocket cost⁹ of $x(m_\nu)$ for gross medical spending of m_ν , leading to the budget constraint in period $t = \nu$ for the insured of

$$c_\nu + x(m_\nu) + p_I \leq y_\nu(\theta_\nu) + \tau. \quad (4)$$

and

$$c_\nu + m_\nu + p_U \leq y_\nu(\theta_\nu) + \tau \quad (5)$$

for the uninsured. The values p_I and p_U are the prices paid by the insured and uninsured, respectively.¹⁰ The final term, τ , is a (hypothetical) transfer paid to individuals in the insurance market (i.e. $\theta_\nu \in \Theta_\nu$). This term that will be used below to define formal measures of willingness to pay. The choice of c and m depend upon whether the individual has chosen to be insured. I denote the choices of the insured and uninsured by $\{c_t^I(\theta_t), m_t^I(\theta_t)\}_t$ and $\{c_t^U(\theta_t), m_t^U(\theta_t)\}_t$, respectively.¹¹

For simplicity, the model does not explicitly allow a borrowing/savings decision, or other forms of insurance that individuals might own. But, these can be captured in the model by considering net income, $y(\theta_t)$, to be an endogenous choice made by individuals. To first order, the envelope theorem implies that changes in savings in response to changes in insurance subsidies will not affect their willingness to pay for those subsidies (Chetty (2006)).

⁹The model focuses on a health insurance contract that pays for medical spending, m . But, the results in Section 4 hold for a more general model in which the insurance contract pays an amount that depends on the state of the world, θ_ν , $x(m; \theta_\nu)$. This would allow the model to nest unemployment insurance settings where the contract pays out conditional on the state of the world. I leave a deeper analysis of this case for future work.

¹⁰I assume that the prices p_U and p_I are paid in period $t = \nu$ to allow the model to naturally nest a purely static environment, such as the environment in Einav et al. (2010). But, the model can readily be generalized to allow prices to be paid in other periods, such as $t = \mu$.

¹¹These choices may also depend upon p_I and p_U even conditional upon insurance purchase (e.g. through income effects). For brevity, this dependence is not included in the notation but is captured where applicable in the analysis, most notably in the construction of the derivative of the average cost curve below.

Therefore, the welfare conclusions derived in Section 4 will continue to hold if one were to allow y to be an endogenous choice. However, differential savings between the insured and uninsured would affect the implementation provided in Section 5, and I discuss how data on consumption can be preferred in settings where such differential savings/borrowing is important.

WTP for Insurance At some time $t = \mu$ (where $\mu \leq \nu$), individuals with $\theta_\mu \in \Theta_\mu$ have the opportunity to purchase insurance for the contract in period ν . In the example below, this notation will correspond to a focus on the 2011 subsidy schedule for those with incomes at 150% FPL in the Massachusetts health insurance exchange (i.e. $\nu = 2011$ and μ is the open enrollment period for the 2011 market in the fall of 2010).¹²

Individuals evaluate their decision to purchase insurance using their information set at time $t = \mu$, which means they evaluate utility using $E[\cdot | \theta_\mu]$. Let $d(p_U; \theta_\mu)$ denote the individual's marginal willingness to pay to be insured relative to uninsured. This solves:

$$E \left[u \left(y_\nu(\theta_\nu) - m^U(\theta_\nu) - p_U, m^U(\theta_\nu); \theta_\nu \right) | \theta_\mu \right] = E \left[u \left(y_\nu(\theta_\nu) - x \left(m^I(\theta_\nu) \right) - d(p_U, \theta_\mu) + p_U, m^I(\theta_\nu); \theta_\nu \right) | \theta_\mu \right] \quad (6)$$

which equates expected utility if uninsured (paying a price p_U) to the utility of the insured (paying a price $d(p_U, \theta_\mu) + p_U$). All individuals with $p_I - p_U \leq d(p_U, \theta_\mu)$ will choose to purchase insurance, whereas those with $d(p_U, \theta_\mu) < p_I - p_U$ will choose to remain uninsured and pay p_U .

For all individuals with $\theta_\mu \in \Theta_\mu$, let $Q(d(p_U, \theta_\mu))$ denote the quantile corresponding to a given level of willingness to pay, $d(p_U, \theta_\mu)$. Let $S(\theta_\mu)$ denote the share of the market that would purchase insurance if prices were set so that the type θ_μ is indifferent to purchasing insurance.¹³ And, let $D(p_U, s)$ denote the relative price required for a fraction s of the market to purchase insurance. This solves:

$$D(p_U, S(\theta_\mu)) = d(p_U, \theta_\mu) \quad \forall \theta_\mu \in \Theta_\mu$$

Throughout, I refer to $D(p_U, s)$ as the “demand curve” or “willingness to pay curve”. When viewed as a function of s , $D(p_U, s)$ corresponds to the curve in Einav et al. (2010) that would be revealed through random variation in prices, p_I and p_U . In practice, many empirical applications (e.g. Einav et al. (2010), Handel et al. (2015) Finkelstein et al. (2019)) impose

¹²Focusing on a single period insurance contract differs from the focus in Handel et al. (2015), who consider the welfare impact of reclassification risk, which results from repeated interaction in a market over time.

¹³Formally, this solves $S(\theta_\mu) = 1 - Q(d(p_U, \theta_\mu))$, where $Q(d(p_U, \theta_\mu))$ denote the quantile corresponding to a given level of willingness to pay, $d(p_U, \theta_\mu)$.

the empirical assumption that $D(p_U, s)$ does not depend on p_U . This can be micro-founded by assuming a CARA utility function.

Cost of Insured Population The average cost cost imposed on the insurer when a fraction s of the market owns insurance is given by

$$AC(s) = E [m_\nu^I(\theta_\nu) - x(m_\nu^I(\theta_\nu)) | S(\theta_\mu) \leq s] \quad (7)$$

so that $sAC(s)$ is the total cost of insuring a fraction s of the market.¹⁴ I assume individuals in the economy have rational expectations and that the realizations of $m_\nu^I(\theta_\nu)$ are idiosyncratic. This means that the realized distributions of $m_\nu^I(\theta_\nu)$ correspond to the cross-sectional distributions in the economy.¹⁵

The marginal cost curve is the derivative of total cost:¹⁶,

$$C(s) = \frac{d}{ds} sAC(s)$$

Market Clearing and the Government Budget Let the functions $p_I(s)$ and $p_U(s)$ denote the prices of insurance and remaining uninsured when a fraction s of the market owns insurance. By definition, these prices must be consistent with the definition of willingness to pay,

$$D(p_U(s), s) = p_I(s) - p_U(s) \quad (8)$$

Lastly, let $G(s)$ denote the total cost (net of premiums collected) to the insurer of insuring a fraction s of the market by setting prices $p_I(s)$ and $p_U(s)$:

$$G(s) = \underbrace{sAC(s)}_{\text{Cost of Insured}} - \underbrace{[sp_I(s) + (1-s)p_U(s)]}_{\text{Premiums Collected}} \quad (9)$$

$G(s)$ captures the net resource expenditures (e.g. government subsidies) for this health insurance market.

¹⁴The expectation conditioning on $S(\theta_\mu)$ is taken with respect to all $\theta_\mu \in \Theta_\mu$, as $S(\circ)$ is not defined outside of the set of individuals in the market.

¹⁵In contrast to structural approaches, I do not require that the econometrician knows the individual's information set, θ_μ , or the subjective belief distribution that generates $E[\circ|\theta_\mu]$.

¹⁶Imposing the assumption that price changes do not affect m , Einav et al. (2010) show that this can be represented as the difference between expenditures and out-of-pocket spending for the marginal purchasers of insurance:

$$C(s) = E [m_\nu^I(\theta_\nu) - x(m_\nu^I(\theta_\nu)) | S(\theta_\mu) = s]$$

3.2 Setting: Health Insurance Subsidies for Low-Income Adults

This section maps the model environment onto the health insurance marketplace for low-income adults in Massachusetts' subsidized insurance exchange, Commonwealth Care. The Massachusetts health reform of 2006 set up insurance exchanges, mandates, and health insurance exchange. This design of this insurance market later became a model for the Affordable Care Act. See Finkelstein et al. (2019) for a detailed discussion of this health insurance exchange.

Using administrative data from Massachusetts, Finkelstein et al. (2019) exploit discontinuities in the subsidy schedule to estimate willingness to pay and costs for low-income adults in the exchange. I focus on the baseline estimates from Finkelstein et al. (2019), which use the empirical discontinuities in 2011 to measure the willingness to pay, $D(s)$, and cost curve, $C(s)$, for those with incomes at 150% FPL.¹⁷¹⁸ In the language of the model in Section 3, the set of people eligible for the market, Θ_μ , corresponds to individuals with incomes at 150% FPL during the open enrollment period ($t = \mu$) in the fall of 2010.¹⁹ As is common in the empirical literature, their empirical model implies that the marginal willingness to pay for insurance does not depend on p_U . As noted above, this can be micro-founded in the model by assuming CARA utility, as will be done in the structural framework offered in Section 7. Therefore, when referring to the empirical demand curves in the Finkelstein et al. (2019) setting, I will often simply refer to these as $D(s)$ instead of $D(p_U, s)$.

Figure 3 presents the results for $D(s)$ and $C(s)$ plotted as a function of s . The patterns reveal that those with the highest willingness to pay (low values of s) are willing to pay more than their marginal cost for insurance, $D(s) > C(s)$. But those with the lowest willingness to pay are willing to pay less than the cost the cost of their insurance. The model captures $D(s) < C(s)$ by allowing part of the cost to be driven by the presence of moral hazard – part of the increase in cost from insurance is not valued by the individual because they do not have to pay for it.²⁰

If a private market were to try to sell this insurance without government subsidies, Finkelstein et al. (2019) show that their estimates imply that this market would fully unravel so that no one obtains insurance. Unraveling is the result of both adverse selection and uncompensated care externalities. Uncompensated care externalities arise because the total cost

¹⁷150% FPL corresponds to roughly \$16K in income for an individual with no children.

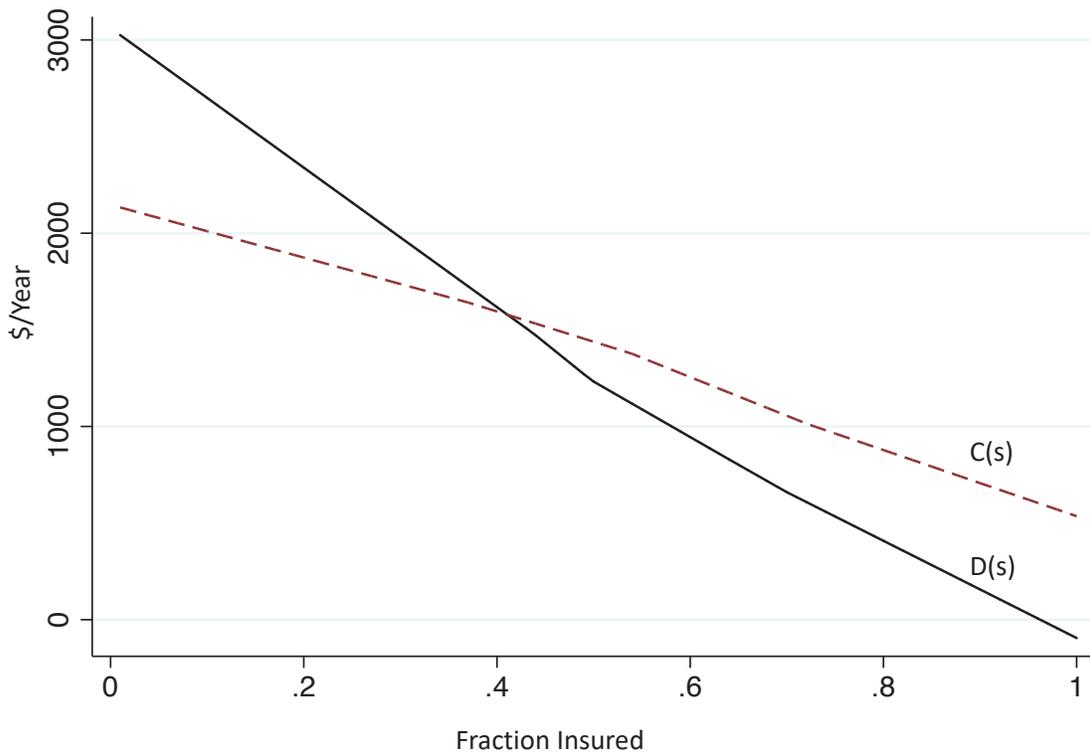
¹⁸Relative to the numbers presented in Finkelstein et al. (2019), these curves are scaled by a factor of 12 to translate the monthly premiums and costs into annual figures.

¹⁹Formally, they are individuals at 150% FPL in the fall of 2010 who reside in Massachusetts and who do not have access to employer-provided health insurance.

²⁰In contrast, this feature is not captured in the stylized model of Section 2 or in other previous models studying reclassification risk and notions of ex-ante expected utility (e.g. Handel et al. (2015); Einav et al. (2016))

to a private insurer would not only include the average resource cost of those insured, $C(s)$, but also the cost of care that would have otherwise been provided through uncompensated care programs. Because private insurance would have to pay these additional costs, it causes the average cost faced by the private insurer to lie everywhere above the demand curve.

Figure 3: Willingness to Pay and Cost for Health Insurance for Adults with Incomes at 150% of Federal Poverty Line in MA



Definition of “Ex-ante” The goal of the rest of the paper is to address normative questions such as: what is the welfare cost of this unraveling? Should the government raise penalties p_U to finance lower insurance prices, p_I ? What is the rationale for the government to further subsidize insurance prices, p_I ? Are these subsidies to those earning near 150% FPL a more (or less) efficient method of redistribution than tax credits, such as the earned income tax credit?

The aim is to answer these questions by measuring individuals expected utilities that are evaluated prior to deciding whether to purchase insurance in the fall of 2010. This means that

one can think of the approach as measuring expected utility behind a full “veil of ignorance” (e.g. prior to birth). For example, at birth one could conceptually ask individuals what subsidies they would like for insurance if they find themselves in the MA health insurance exchange in the fall of 2010 with incomes at 150% and deciding whether to purchase insurance for 2011? But, it is also equivalent to consider the same question at later points in an individuals’ life as long as they have not yet learned anything about their 2011 health risk. Formally, they can’t know anything about θ_ν or θ_μ .²¹ This means that one can think of the approach as measuring utility at birth or using a more traditional “veil of ignorance”. But, one can also use a later frame of reference as long as they have no knowledge about their particular risk realizations in the market.

The next section provides a general characterization of ex-ante measures of welfare that answer these questions. The subsequent section 5 provides a method of implementing these measures. Section 6 discusses the results for the MA health insurance context.

4 Measuring Ex-Ante Welfare

4.1 Budget Neutral Policies: “Ex-Ante” Demand Curve

I consider two classes of policies: Budget neutral and non-budget neutral policies. Budget neutral policies involve reductions in the price of insurance, p_I , financed by increases in the price/penalty of being uninsured, p_U . In a world where prices cover costs, $p_I(s)$ and $p_U(s)$ satisfy two equations:

$$D(p_U(s), s) = p_I(s) - p_U(s)$$

$$G(s) = 0$$

Traditional market surplus would evaluate the welfare impact of these policies using market surplus, $D(p_U, s) - C(s)$. The optimal size of the insurance market would equate demand to cost. Here, I use the modeling environment in Section 3 to measure the ex-ante welfare impact of such policies.

For any prices, p_I and p_U , and transfers, $\tau(\theta_t)$, let $v_t(p_I, p_U, \tau(\theta_t); \theta_t)$ be the realized utility of individual in period t :

$$v_t(p_I, p_U, \tau(\theta_t); \theta_t) = u(c_t(\theta_t), m_t(\theta_t); \theta_t)$$

²¹Formally, suppose there exists a period $t < \mu$ such that individuals do not know anything about their particular realization of θ_μ or θ_ν . Then, for any function $f(\theta_\nu)$ of θ_ν , it will necessarily be the case that $E[f(\theta_\nu) | \theta_\mu \in \Theta_\mu, \theta_\nu] = E[f(\theta_\nu) | \theta_\mu \in \Theta_\mu]$. Hence, the expected utility for those in the market is not dependent on θ_t .

where the chosen allocations of $c_t(\theta_t)$ and $m_t(\theta_t)$ depend upon the prices, p_I and p_U , and transfers, $\tau(\theta_t)$.²² The ex-ante expected utility of having prices p_I and p_U , along with transfers τ is given by:

$$V(p_I, p_U, \tau) = E \left[\sum_{t' \geq 0} v_{t'}(p_I, p_U, \tau; \theta_t) \right] \quad (10)$$

Let $W(s)$ denote the willingness to pay to have a fraction s of the market insured. Formally, I define this as the amount of income individuals would be willing to take in the uninsured world that equates their expected utility to its level if the fraction s is insured. This solves:

$$V(p_I(s), p_U(s), 0) = V(\infty, 0, W(s)) \quad (11)$$

where the LHS is the ex-ante expected utility of prices $p_I(s)$ and $p_U(s)$ and the RHS is the expected utility of simply giving a transfer in the world where everyone is insured.

The function $W(s)$ measures the ex-ante welfare impact of changes to prices of insurance in the market. Equation (11) shows that the size of the market that maximizes $W(s)$ is the size of the market maximizes expected utility, $V(p_I(s), p_U(s), 0)$.

Sufficient Statistic Representation of $W(s)$ In general, measuring $W(s)$ requires knowledge of the full underlying structural model embedded in equation (11). Traditional approaches to measuring $W(s)$ would follow a structural approach that specifies a utility function, u , budget constraints in all periods, beliefs about future outcomes, and one would estimate the model using data on observed data on c and m , combined with sufficient data variation to estimate all model components. In particular, one needs to separately estimate both utility and beliefs, which is often quite difficult and rests on assumptions that the econometrician perfectly observes individuals' information sets.

In contrast, the approach developed here is to characterize the sufficient statistics required for measuring $W(s)$. I exploit the envelope theorem to fully characterize the derivative of $W(s)$ at each s , which can then be used to assess the welfare impact of changes in the size of the insurance market. In particular, differentiating equation (11) with respect to s yields

$$W'(s) * \frac{\partial V(\infty, 0, S^{EA}(s))}{\partial \tau} = \frac{d}{ds} E \left[\sum_{t \geq 0} v_t(p_I(s), p_U(s), 0; \theta_t) \right] \quad (12)$$

$$= E \left[1 \{S(\theta_\mu) \geq s\} \frac{\partial u_\nu}{\partial c_\nu} (-p'_I(s)) + 1 \{S(\theta_\mu) < s\} \frac{\partial u_\nu}{\partial c_\nu} (-p'_U(s)) \mid \theta_\nu \in \Theta_\nu \right] \quad (13)$$

²²They also depend implicitly on whether individuals chose insurance at those prices.

The change in prices affect utility through the marginal utility of consumption in period ν . Behavioral responses from other arguments of the utility function drop out of this equation due to the envelope theorem. The key feature of equation 12 is that the price changes, $p'_I(s)$ and $p'_U(s)$, for the insured and uninsured are weighted by individual's marginal utilities of consumption, $\frac{\partial u_\nu}{\partial c_\nu}$. This means that transfers between insured and uninsured have value from an ex-ante perspective to the extent to which they help move resources from states of the world with low marginal utilities of income to states of the world with high marginal utilities of income.

Proposition 1 provides a characterization of $W'(s)$ for all s that illustrates how ex-ante welfare depends crucially on the difference in marginal utilities between the insured and uninsured.

Proposition 1. *For budget neutral policies satisfying $G'(s) = 0$, the marginal welfare impact of expanding the size of the insurance market from s to $s + ds$ is given by*

$$W'(s) = EA(s) + D(s) - C(s) + \delta_p(s) \quad (14)$$

where $EA(s)$ is the additional ex-ante value of expanding the size of the insurance market,

$$EA(s) = s(1-s) \left(-\frac{\partial D}{\partial s} \right) \beta(s) \quad (15)$$

and $\beta(s)$ is the percentage difference in marginal utilities of income for the insured relative to the uninsured,

$$\beta(s) = \frac{E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s \right]}{E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right]} \quad (16)$$

and $\delta_p(s)$ is a term that adjusts for the case when the marginal willingness to pay for insurance depends on p_U ,

$$\delta_p(s) = \frac{\partial D}{\partial p_U} \frac{-s}{1 + s \frac{\partial D}{\partial p_U}} \left(1 + (1-s) \beta(s) \left(1 + s \frac{\partial D}{\partial s} \right) \right)$$

where $\delta_p(s) = 0$ if $\frac{\partial D}{\partial p_U} = 0$, as is assumed in many empirical applications.

Proof. See Appendix A. □

Equation (14) shows that when $\frac{\partial D}{\partial p_U} = 0$ the marginal ex-ante willingness to pay for a larger insurance market is given by the sum of $D(s) + EA(s) - C(s)$. The term $D(s) - C(s)$ is traditional market surplus: expanding the size of the insurance market increases ex-ante welfare to the extent to which individuals are willing to pay more than their costs for insurance. $EA(s)$ captures the additional ex-ante value of expanding the size of the market

through its impact on insurance prices. Expanding the insurance market induces a transfer from uninsured to insured of size $(1 - s)(-sD'(s))$.²³ The transfer is valued according to the difference in marginal utilities between the insured and uninsured, $\beta(s)$.

Sign of $\beta(s)$ Canonical models of insurance predict that $\beta(s) > 0$. This is because those who choose not to purchase insurance expect to have out of pocket medical spending that is lower than the marginal price of insurance, which implies that the consumption of the insured are lower than those of the uninsured. Concavity of the utility function then implies that the marginal utilities of the insured are higher than the uninsured, so that $\beta(s) > 0$.

But, it is possible to have $\beta(s) < 0$. For example, if (a) moral hazard is sufficiently large and (b) differences in liquidity is a primary driver of insurance purchase, then it could be that the uninsured have a higher marginal utility of income than the insured, $\beta(s) < 0$. In this case, expanding the size of the insurance market will transfer resources from the liquidity constrained to those who are less constrained, which would suggest that $EA(s) < 0$.²⁴

4.2 Non-Budget Neutral Policies: The MVPF

In many cases, subsidies for the insured are paid by those not in the insurance market (i.e. $\theta_\mu \notin \Theta_\mu$), not through higher prices to the uninsured in the market, p_U . Massachusetts provides subsidies for insurance for low-income individuals that are financed by taxpayers more generally, not by the uninsured with incomes at 150% FPL. Health insurance subsidies are a method of redistribution to the low-income individuals who purchase health insurance. Are insurance subsidies to those earning near 150% FPL an efficient method of redistribution?

To answer this question, we follow the approach developed in Hendren (2016) and Hendren and Sprung-Keyser (2019) to construct the marginal value of public funds, MVPF, of increased subsidies that lower p_I . The MVPF is given by the ratio of individuals' willingness to pay for higher subsidies normalized by the net cost to the government of increasing

²³The original version of this manuscript defined $W(s)$ to be the willingness to pay out of the insured population as opposed to the average uninsured population when $s = 0$. As a result, the Taylor approximation in the earlier draft was performed around the average consumption of the insured as opposed to the average level of consumption in the population. The result of this is that the formula for $EA(s)$ included an additional $C(s) - D(s)$ term in earlier manuscripts. The revised definition of $W(s)$ now implies it is most appropriate to do the Taylor approximation around the average individual in the market. This leads to the formula for $EA(s)$ in equation (15) that does not include a $C(s) - D(s)$ term and is also identical to the formula provided in Section 2.

²⁴Relatedly, it is important to note that the model remains valid if the underlying heterogeneity generates advantageous as opposed to adverse selection. This could occur if S reflects a preference shock or liquidity shock such that those who prefer not to purchase insurance are also the latently higher risk population. In this case, the difference in marginal utilities between insured and uninsured remains the key sufficient statistic required for measuring ex-ante expected utility.

subsidies,

$$MVPF(s) = \frac{W'(s)}{G'(s)} = \frac{\text{Marginal WTP}}{\text{Marginal Cost}}$$

For every \$1 of net government spending, the policy delivers MVPF dollars of welfare to the beneficiaries in units of their own willingness to pay. Previous literature (e.g. Hendren and Sprung-Keyser (2019) and Hendren (2014, 2016)) shows how one can use MVPFs to make statements about the social welfare impact of hypothetical budget neutral policies. Given two policies, A and B, spending more on policy A financed by reduced spending on policy B increases social welfare if and only if one prefers giving $MVPF_A$ to policy A beneficiaries versus $MVPF_B$ to policy B beneficiaries.²⁵

Returning to the comparison of health insurance subsidies to tax cuts to those earning near 150% FPL, one can consider the MVPF of a reduction in p_I . The marginal cost has two components. First, lowering p_I by \$1 increases the cost by \$1 for each of the s enrollees. Second, there is an additional cost from those induced to purchase insurance by the lower prices. These enrollees pay $D(s) = p_I(s)$ but cost the insurer $C(s)$. So, the new enrollees impose a cost of $C(s) - D(s)$. A \$1 price reduction increases the size of the market by $\frac{1}{-D'(s)}$. Normalizing by the size of the market, s , the total cost of lowering premiums by \$1 is $1 + \frac{C(s)-D(s)}{s(-D'(s))}$.²⁶

For willingness to pay, an individual who has already learned that their value of θ_μ is such that $S(\theta_\mu) < s$ so that they decide to purchase insurance would be willing to pay \$1 to have premiums that are \$1 lower. So, if willingness to pay were calculated after learning θ_μ , the numerator would simply be 1 and the MVPF would be $\frac{1}{1 + \frac{C(s)-D(s)}{s(-D'(s))}}$.

Not everyone at 150% FPL purchases insurance. This means that one cannot readily compare the MPVF of health insurance subsidies to the MPVF of a tax credit that gives everyone at 150% FPL additional income. The two policies have different distributional incidence – those purchasing insurance might be sicker than the average population and have a higher (or lower) marginal utility of income. The ex-ante perspective provides a novel path forward to compare health insurance subsidies to a tax credit. From behind a veil of ignorance, one can ask whether individuals would want to have \$1 spent on tax cuts to those earning near 150% FPL or \$1 spent on health insurance subsidies to those earning near 150% FPL. Which policy increases welfare depends on which MVPF is higher,

²⁵When two policies have the same distributional incidence, they can provide a Pareto-based guidance on optimal policy: if the MVPF of lowering health insurance prices is higher than the MPVF of a tax cut to an identical population, then everyone's welfare can be increased by lowering health insurance prices financed by a reduction in tax subsidies.

²⁶More generally, if there are additional behavioral responses that affect the government budget (e.g. if insurance improves health and increases taxable income, or if the subsidies distort labor supply, etc.), these would also need to be incorporated into the marginal cost of lowering premiums.

where crucially one needs to use the WTP for health insurance subsidies from an ex-ante perspective prior to learning their health risk (i.e. prior to learning their particular value of $\theta_\mu \in \Theta_\mu$). Proposition 2 shows that one can construct exactly this measure by adding $(1 - s)\beta(s)$ to the willingness to pay in the MVPF for health insurance subsidies. This adjusts for the fact that those purchasing insurance might have a higher marginal utility of income as compared to the average population with $\theta_\mu \in \Theta_\mu$.

Proposition 2. *Suppose $p_V(s) = 0$. The MVPF of additional insurance market subsidies when a fraction s of the market is insured is given by*

$$MVPF(s) = \frac{W'(s)}{\text{Marginal Cost}} = \frac{1 + (1 - s)\beta(s)}{1 + \frac{C(s) - D(s)}{s(-D'(s))}} \quad (17)$$

where $\beta(s)$ is the percentage difference in marginal utilities of income for the insured relative to the uninsured given by equation (16).

Proof. See Appendix B □

From behind the veil of ignorance, \$1 for the insured is valued using the marginal utility of consumption for the insured, $E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right]$. In contrast, a tax cut to those with incomes near 150% FPL has a willingness to pay of \$1 for each person, $E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right]$. Therefore, the desired willingness to pay needs to be inflated by $E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] / E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right]$, which can be expressed as $1 + (1 - s)\beta(s)$.

Estimates of the MVPF for changes to the Earned Income Tax Credit (EITC) range between 0.9 and 1.3 (Hendren and Sprung-Keyser (2019)). To the extent to which this MVPF reflects the MVPF of a more targeted tax credit to those with incomes at 150% FPL, this provides a benchmark against which one can compare the MVPF for health insurance subsidies. If the MVPF in equation (17) exceeds this MVPF for the EITC, then reduced spending on the EITC that is used to finance greater health insurance subsidies for those with low incomes would increase ex-ante expected utility.

5 Implementation

The key additional parameter required to construct measures of ex-ante expected utility is the difference in marginal utilities between insured and uninsured, $\beta(s)$. The canonical approach to measuring ex-ante expected utility would estimate a structural model. In this sense, the results above show that the key parameter that drives the ability of structural

models to generate measures of ex-ante welfare is through their estimates of the difference in marginal utilities of income between the insured and uninsured, $\beta(s)$.

The aim of this section is to provide a transparent method for estimating ex-ante expected utility that does not require specification of the underlying market primitives that are often difficult to observe (e.g. information sets). I provide a set of assumptions that enable $\beta(s)$ to be written as a function of market demand curves combined with a measure of risk aversion. The required assumptions are not without loss of generality. To assess the quality of the fit and impact of violating these assumptions, I compare the estimates to a simulated implementation using a structural model in Section 8.

I begin by outlining the benchmark implementation using market demand and cost curves, combined with a measure of risk aversion.

5.1 Estimating $\beta(s)$ Using Market Demand and Cost Curves

To provide a method for estimating $\beta(s)$, I draw upon assumptions commonly used in the literature on optimal unemployment insurance. In particular, I begin by assuming that the marginal utility of income does not depend on θ_ν or m .

Assumption 1. *u_c depends only on the level of consumption, c , and does not vary across individuals, so that the marginal utility of consumption can be written as*

$$\frac{\partial u_\nu}{\partial c}(c, m; \theta_\nu) = f(c)$$

for all m and all θ_ν .

Assumption 1 implies that if one can observe the level of an individual's consumption, then one can infer their marginal utility of income. This is a common assumption imposed in the literature on optimal unemployment insurance (e.g. Baily (1978); Chetty (2006)), but this assumption is not without loss of generality. Most notably, it assumes away preference heterogeneity that is correlated with the marginal utility of income. I discuss violations of this assumption in Section 8.

Imposing Assumption 1 means $\beta(s)$ can be written as the coefficient of absolute risk aversion multiplied by the average difference in consumption between the insured and uninsured, as stated in Proposition 3.

Proposition 3. *Suppose Assumption 1 holds. Then, $\beta(s)$ can be written as*

$$\beta(s) \approx \gamma \Delta c \tag{18}$$

where “ \approx ” represents equality to a first order Taylor expansion of $f(c)$, $\Delta c = E[c(\theta_\nu) | S \geq s] - E[c(\theta_\nu) | S < s]$ is the difference in consumption between the uninsured and insured, and $\gamma = \frac{-f'(c)}{f(c)} = \frac{-\frac{\partial^2 u}{\partial c^2}}{\frac{\partial u}{\partial c}}$ is the coefficient of absolute risk aversion.

Proof. Recall

$$\beta(s) = \frac{E\left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s\right] - E\left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s\right]}{E\left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\nu \in \Theta_\nu\right]}$$

Begin with the denominator. Let c^* denote any value of consumption. Note to first order, $f(c) = f(c^*) + f'(c^*)(c - c^*)$. This implies

$$E\left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s\right] = f(c^*) + f'(c^*) E[c | S(\theta_\mu) < s]$$

and

$$E\left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s\right] = f(c^*) + f'(c^*) E[c | S(\theta_\mu) \geq s]$$

So,

$$E\left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s\right] - E\left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s\right] = f'(c^*) (E[c | S(\theta_\mu) < s] - E[c | S(\theta_\mu) \geq s])$$

Now, since c^* is as of yet left arbitrary, one can pick c^* to satisfy $E\left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\nu \in \Theta_\nu\right] = f(c^*)$ exactly, or one could choose c^* to be the population average consumption of those with $\theta_\nu \in \Theta_\nu$, which would satisfy $E\left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\nu \in \Theta_\nu\right] = f(c^*)$ exactly. In either case, the resulting expression yields

$$\frac{E\left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s\right] - E\left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s\right]}{E\left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\nu \in \Theta_\nu\right]} = \frac{f'(c^*)}{f(c^*)} (E[c | S(\theta_\mu) < s] - E[c | S(\theta_\mu) \geq s])$$

Multiplying each RHS term by -1 to give $\frac{-f'(c^*)}{f(c^*)}$ and $(E[c | S(\theta_\mu) \geq s] - E[c | S(\theta_\mu) < s])$ yields the result. □

If one observes consumption, then equation (18) provides a direct method for measuring the difference in marginal utilities. This is analogous to the methods used in the literature on optimal unemployment insurance (Baily (1978); Chetty (2006)).

In practice consumption is rarely observed. To provide an implementation method without consumption data, I make the additional assumption that incomes, $y_\nu(\theta)$ do not vary systematically between the insured and uninsured.

Assumption 2. *No differences in average liquidity/income between the insured and unin-*

sured,

$$E[y_\nu(\theta_\nu) | S \geq s] = E[y_\nu(\theta_\nu) | S \leq s] \quad \forall s$$

Assumption 2 is not without loss of generality, but is natural in contexts like the MA health insurance subsidies for those with incomes at exactly 150% FPL. But, in many contexts this assumption may be violated. Indeed, if the disposable income in period ν , $y_\nu(\theta_\nu)$, is the result of endogenous savings to whether individuals have purchased insurance then this could also generate a violation of Assumption 2. I return to a discussion of how one can use consumption data to relax Assumption 2 in Section 8.1 below. The key advantage of Assumption 2 is that it allows one to replace Δc in equation 18 with information contained in the market level demand curve, $D(s)$, as is shown in Proposition 4.

Proposition 4. *Suppose Assumptions 1 and 2 hold. Suppose moreover that the marginal willingness to pay for insurance does not depend on p_U , $\frac{\partial D}{\partial p_U} = 0$. Then, for any full insurance contract with $x(m_\nu) = 0$ ²⁷, the difference in marginal utilities is given by*

$$\beta(s) = \gamma \Delta D \tag{19}$$

where ΔD is the difference between the marginal price of insurance and the average willingness to pay for the uninsured population,

$$\Delta D = D(p_U(s), s) - E[D(p_U(s), S(\theta_\mu)) | S(\theta_\mu) \geq s]$$

Proof. See Appendix D. □

Proposition 4 provides a method of estimating ex-ante willingness to pay using the demand curve, combined with a measure of risk aversion. The estimate of risk aversion can either be imported from external settings, or it can be estimated internally using the relationship between the markup individuals are willing to pay and the reduction in consumption variance provided by the insurance, as discussed in Appendix C.

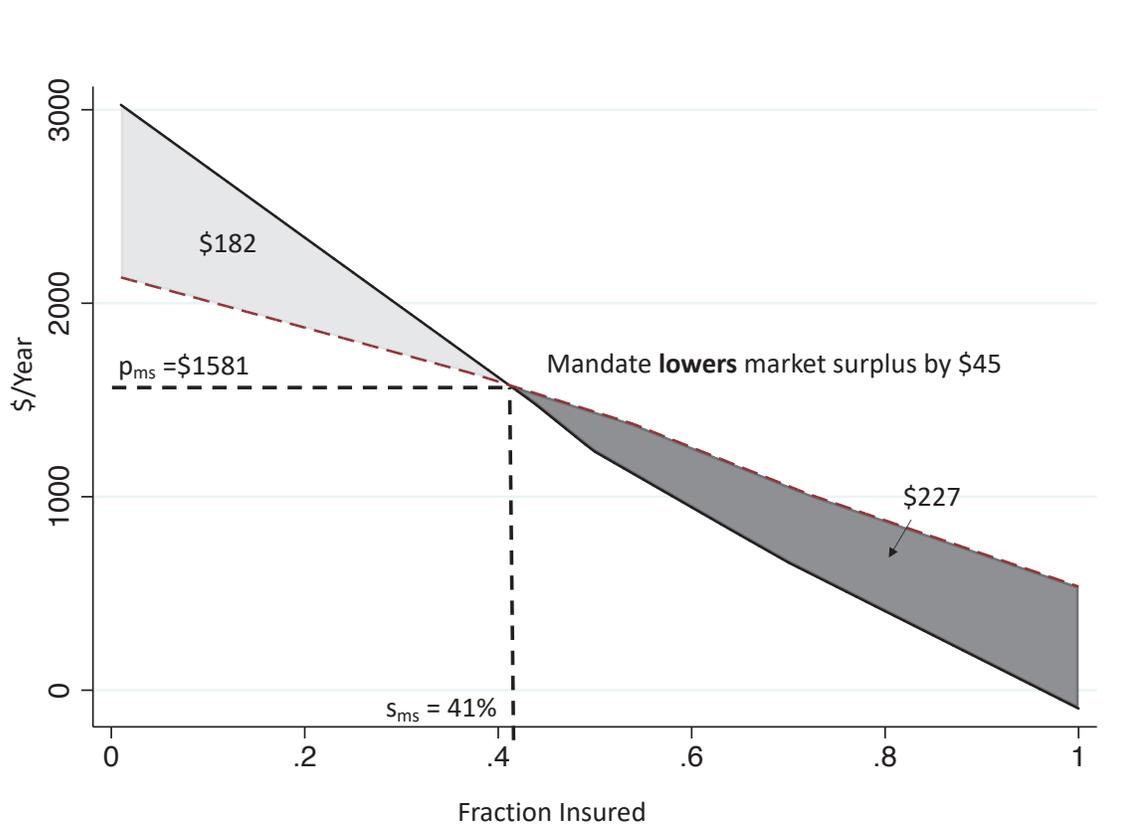
6 Application to MA Health Insurance Subsidies

This section measures the ex-ante welfare impacts of subsidies and mandates for low-income adults at 150% FPL in the MA health insurance exchange using the demand and cost curve

²⁷It is straightforward to extend the result for partial insurance contracts. The difference in marginal utilities is given by $\beta(s) = \gamma \Delta D + E[x(m_\nu) | S(\theta_\mu) < s]$, where $E[x(m_\nu) | S(\theta_\mu) < s]$ is the average out of pocket spending by the insured.

estimates from Finkelstein et al. (2019). I contrast these welfare conclusions to those generated from a more traditional analysis of market surplus.

Figure 4: Market Surplus Measures of Welfare



6.1 Budget neutral policies

I begin with the welfare impact of increasing subsidies for insurance if they are financed by increasing prices/penalties on the insured. Figure 4 shows the market surplus approach to this question by comparing $D(s)$ to $C(s)$. Market surplus is maximized when a fraction $s = 41\%$ of the market is insured. The marginal price for insurance would then be \$1581, equal to the marginal cost of this individual imposed on the insurer. The welfare gain from moving from $s = 0$ to $s = 41\%$ is given by the integral between $D(s)$ and $C(s)$. Individuals in the market are on aggregate willing to pay \$182 to have 41% of the market insured (instead of no one insured). One can also ask whether mandates ($s = 100\%$) are preferred to the full unraveling allocation of $s = 0$. For those with $D(s)$ below \$1581, they would on aggregate be willing to pay \$227 to avoid the mandate. This means that mandates would lower total

market surplus by \$45.

What insurance prices maximize ex-ante expected utility? To measure this, one requires an estimate of risk aversion. For the baseline case, I take a common estimate from the health insurance literature of $\gamma = 5x10^{-4}$ (e.g. similar to estimates in Handel et al. (2015)).²⁸ Table 1 also presents estimates for a range of alternative risk aversion coefficients discussed below.

Figure 5: Ex-Ante Welfare of Health Insurance for Low-Income Adults

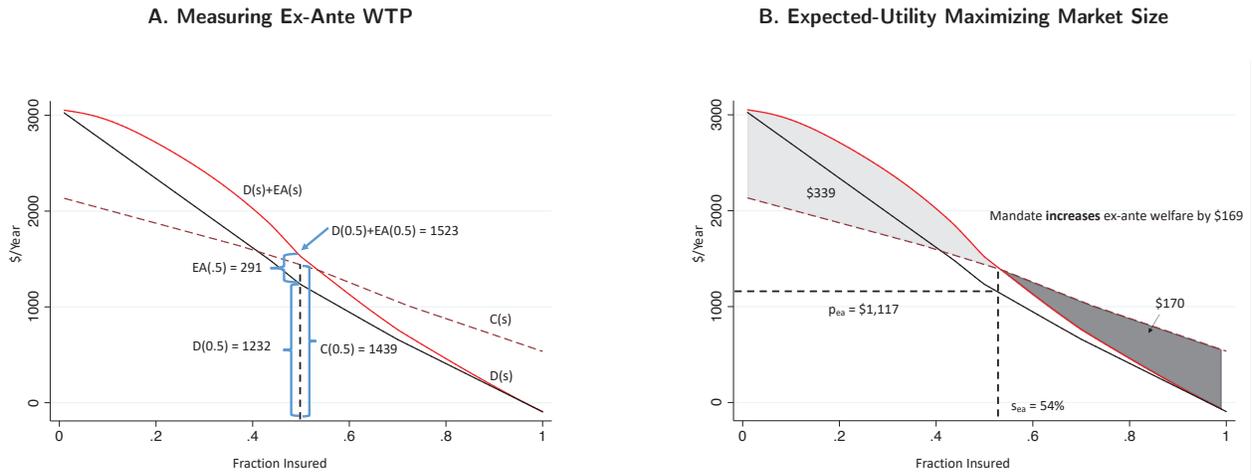


Figure 5 presents the ex-ante demand curve and welfare analyses using $\gamma = 5x10^{-4}$ and equation (19). Panel A illustrates the calculation of $EA(s)$ when 50% of the population owns insurance. The cost of the marginal enrollee is given by $C(0.5) = 1438$, willingness to pay is $D(0.5) = 1232$, and the slope of willingness to pay is $D'(0.5) = -3405$.²⁹ The average willingness to pay with those whose demand is below $D(0.5) = 1232$ is 548. Equation (15) implies that the ex-ante willingness to pay for a larger insurance market is

$$\begin{aligned}
 EA(s) &= (1-s)s(-D'(s))\gamma(D(s) - E[D(S)|S \geq s]) \\
 &= .5(0.5 * 3405)(5x10^{-4})(1232 - 548) \\
 &= 291
 \end{aligned}$$

²⁸As discussed in Appendix ??, Handel et al. (2015) estimates this for a relatively middle to high income population making choices over insurance plans. Under the natural assumption that absolute risk aversion decreases in consumption levels, this estimate is likely a lower bound on the size γ .

²⁹Finkelstein et al. (2019) estimate a piece-wise linear demand curve. To obtain smooth estimates of the slope of demand, I regress the estimates of $D(s)$ from Finkelstein et al. (2019) on a 10th order polynomial in s . The results are similar for other smoothed functions.

Individuals with median (0.5) levels of willingness to pay are willing to pay \$1,232 for insurance at the time the econometrician observes them in the market. But, prior to learning their willingness to pay for insurance, individuals are willing to pay to have lower insurance prices. From behind a veil of ignorance everyone would have been willing to pay \$2.91 to have the opportunity to purchase insurance at the prices that lead to 51% of the market insured instead of 50% of the market insured ($291 * (0.51 - 0.5)$).

Figure 5A also presents the ex-ante WTP curve, $D^{Ex-Ante}(s)$, for all values of s .³⁰ Expected utility is maximized when $W'(s) = 0$, or $D(s) + EA(s) = C(s)$. This occurs when 54% of the market owns insurance and the marginal price of insurance is \$1,117, as shown in Figure 5B. This contrasts with the market surplus-maximizing size of the market of 41%. The ex-ante optimal price is roughly 30% lower than the surplus-maximizing price of \$1,580.

What is the welfare gain from pricing insurance optimally relative to the full unraveling of the competitive market ($s = 0$)? Everyone would be willing to contribute \$340 per person if they could live in a world in which insurance prices set at $p = \$1117$ as opposed to having no one obtain insurance. This \$340 is much larger than the loss of market surplus of \$182 shown in Figure 4. After learning their willingness to pay for insurance, $D(s)$, individuals would only be willing to contribute an average of \$182 per person to set prices to maximize economic surplus. The ex-ante welfare perspective leads to different conclusions about optimal insurance prices and the welfare cost of adverse selection.

Mandates Is it better for the government to impose a binding mandate ($s = 1$) or allow the private market to unravel ($s = 0$)? As Figure 4 illustrates, mandates reduce market surplus by \$45 relative to $s = 0$. On aggregate, individuals in the market would be willing to pay \$45 per person to prevent a mandate.

In contrast, from an ex-ante expected utility perspective, the value of insuring the 54% of the market with the highest willingness to pay is \$339, and the welfare cost of insuring the remaining 46% of the market is \$170. Prior to learning their willingness to pay, individuals would pay an average of \$169 per person to have a mandate instead of having no insurance. Mandates increase ex-ante expected utility, but decrease market surplus.

Alternative Risk Aversion Values The results above rely on a coefficient of relative risk aversion of 5×10^{-4} . Table 1 presents estimates of the above results for alternative risk

³⁰It is perhaps surprising that $EA(s) < 0$ for low values of s . Mathematically, this is because for low values of s , $C(s) - D(s) < sD'(s)$. Economically, this means that expanding the size of the insurance market actually generates a Pareto improvement, as it can lower prices for both the insured and uninsured because the marginal cost of the new enrollees is sufficiently below their willingness to pay. As a result, market surplus actually over-states the welfare impact of expanding the insurance market for low values of s .

aversion measures. Columns (1) and (2) present the market surplus and baseline ex-ante welfare estimates for $\gamma = 5x10^{-4}$. Columns (3) and (4) consider alternative coefficients of absolute risk aversion of $1x10^{-4}$ and $10x10^{-4}$ and columns (5)-(10) consider alternative coefficients of relative risk aversion ranging from 1 to 10.³¹

The baseline specification of $\gamma = 5x10^{-4}$ is consistent with the estimates in Handel et al. (2015) but it implies a large coefficient of relative risk aversion of 8.2. Table 1 shows that a coefficient of relative risk aversion of 3 implies that the optimal insurance prices are \$1,351, which is 15% lower than the optimal price from a market surplus perspective of \$1581. Such prices would lead to 47% of the market insured, which is less than the 54% of the market that would be insured if prices were set to maximize ex-ante welfare in the baseline specification.

Although the precise optimal size of the market varies with the coefficient of risk aversion, the conclusion that mandates increase ex-ante welfare remains fairly robust across specifications. Mandates increase ex-ante expected utility as long as the coefficient of absolute risk aversion is above $1.05x10^{-4}$ and coefficients of relative risk aversion above 1.7. This means that for a range of plausible coefficients of risk aversion, an ex-ante welfare perspective leads to different normative conclusions about the optimal insurance subsidies and desirability of mandates.

Table 1: Alternative Risk Aversion Specifications

	Market Surplus	Baseline: CARA = $5x10^{-4}$	CARA Coefficient		CRRA Coefficient (150% FPL)					
			$1x10^{-4}$	$10x10^{-4}$	1	2	3	4	5	10
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Optimal market size	0.41	0.54	0.45	0.62	0.43	0.45	0.47	0.48	0.49	0.56
Optimal price of insurance	1581	1117	1432	888	1509	1432	1351	1311	1270	1060
WTP for optimal market size	182	339	210	525	199	217	235	254	274	379
WTP for mandate	-45	169	-2	383	-19	7	34	60	86	217
MVPF when 30% insured	1.28	1.79	1.38	2.29	1.34	1.41	1.47	1.53	1.59	1.90
MVPF when 90% insured	0.80	0.81	0.80	0.81	0.80	0.80	0.80	0.80	0.80	0.81

Notes: This table presents the welfare estimates under alternative risk aversion coefficients. Column (1) presents market surplus estimates and column (2) presents the baseline estimates that use a coefficient of absolute risk aversion of $5x10^{-4}$. Columns (3) and (4) use alternative coefficients of absolute risk aversion of $1x10^{-4}$ and $10x10^{-4}$. Columns (5)-(9) construct the ex-ante welfare measures using coefficients of relative risk aversion ranging from 1 to 5 and column (10) considers a coefficient of relative risk aversion of 10. The coefficient of relative risk aversion specifications assume a consumption level of 150% FPL for singles in 2011, which corresponds to $\$10,890 \times 150\% = \$16,335$.

6.2 Non-budget neutral policies

The insurance subsidies in Massachusetts are not paid by low-income individuals at 150% FPL choosing to forego insurance. Rather, they are ultimately paid by other taxpayers beyond those at 150% FPL. Here, I show how to estimate the marginal value of public funds of higher/lower subsidies. Recall, the MVPF equals the marginal willingness to pay of the

³¹To translate the coefficient of relative risk aversion into a coefficient of absolute risk aversion I multiply by $10,890 \times 1.5$, where 10,890 is the FPL for single adults.

beneficiaries for insurance subsidies per dollar of government expenditure. This is given by the formula:

$$MVPF(s) = \frac{1}{1 + \frac{C(s)-D(s)}{s(-D'(s))}} (1 + (1 - s) \gamma(D(s) - E[D(S) | S \geq s]))$$

where $\beta(s) = \gamma(D(s) - E[D(S) | S \geq s])$ is the difference in marginal utilities between the insured and uninsured.

Following Finkelstein et al. (2019), Figure 6 presents the MVPF for the case when 30% and 90% of the market have insurance. When 30% of the market is insured, annual costs are given by $C(0.3) = 1738$, willingness to pay is given by $D(0.3) = 1978$, and the slope of willingness to pay is given by $D'(0.3) = -3610$. The average willingness to pay for those with $s \geq 0.3$ is 853. Therefore, the MVPF is given by

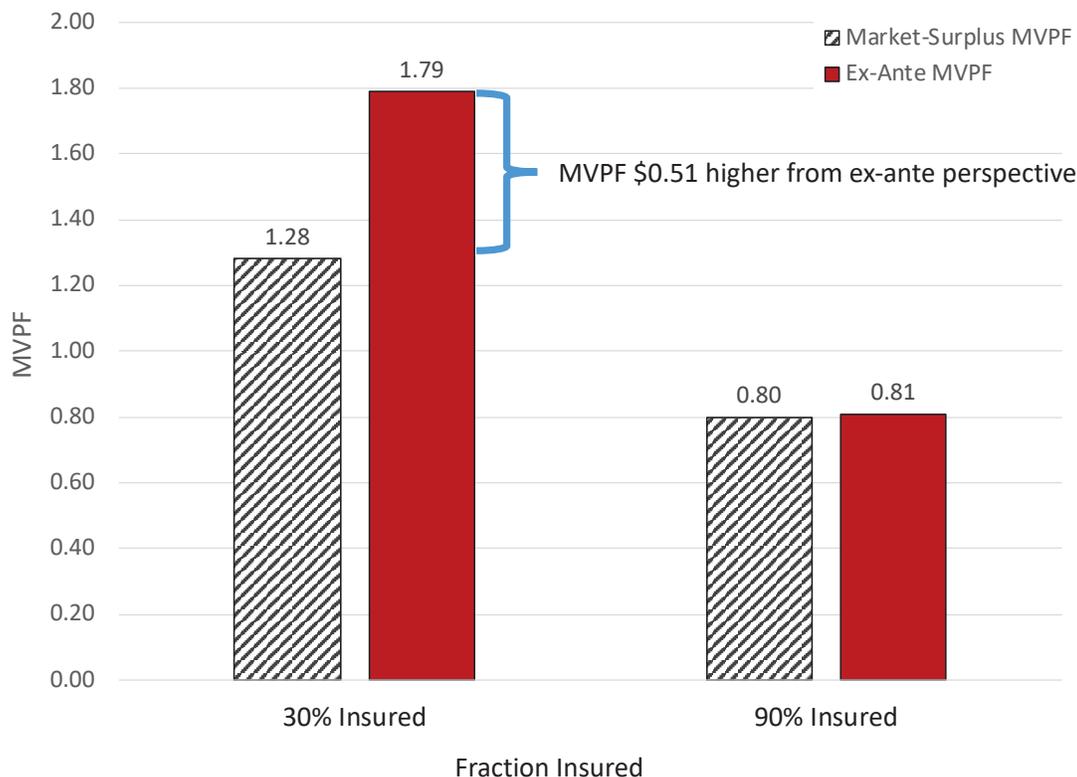
$$\begin{aligned} MVPF(0.3) &= \frac{1}{1 - \frac{1978-1738}{0.3*3654}} (1 + .7 * 5x10^{-4} * (1978 - 853)) \\ &= 1.282 * 1.394 \\ &= 1.79 \end{aligned}$$

The willingness to pay by those choosing to purchase insurance for additional subsidies is \$1.28 per dollar of government spending. Every \$1 of subsidy generates \$1.28 lower prices for the insured. This is greater than \$1 because the marginal types that are induced to enroll from lower prices have a lower cost of being insured, $D(0.3) > C(0.3)$.

Behind the veil of ignorance, these lower prices to the insured have additional value because the insured have a 40% higher marginal utility of income relative to the average person in this setting. Accounting for this, the results suggest that individuals would be willing to pay \$1.79 for every \$1 that the government spends to lower insurance prices for those with incomes at 150% FPL.

For comparison, the MVPF of low-income tax cuts, such as expansions of the Earned Income Tax Credit (EITC) have MVPFs ranging between 0.9-1.3 (Hendren and Sprung-Keyser (2019)). When only 30% of the market is insured, these results suggest expanded insurance subsidies financed by a budget-neutral reduction in EITC would increase ex-ante expected utility. Table 1 shows that with an alternative specification of a coefficient of relative risk aversion of 3, the MVPF would be 1.47. This remains above the general values found for the EITC. To the extent to which these EITC MVPF estimates generalize to a more targeted transfer to those at 150% FPL, it suggests that health insurance subsidies would increase ex-ante welfare more than tax subsidies when 30% of the market is insured.

Figure 6: MVPF for Health Insurance Subsidies for Low-Income Adults



In contrast, the MVPFs are lower when prices are more heavily subsidized so that more of the market has insurance. When $s = 0.9$, the willingness to pay of the marginal type is below her cost, $D(s) < C(s)$, so that $\frac{1}{1 + \frac{C(s) - D(s)}{s(-D'(s))}} = 0.8$. And, the distinction between market surplus and ex-ante expected utility is smaller. This is because the difference between $D(s^*)$ and $E[D(s) | s \geq s^*]$ is smaller when a larger fraction of the market has insurance. This means that the consumption of the insured relative to the average in the population population is smaller when a larger fraction of the market is insured. Comparing this to the MVPF for the EITC, this suggests that subsidies leading to 90% of the market being insured are too generous: reducing health insurance subsidies and using the resources to expand tax credits to those with incomes near 150% FPL would increase ex-ante expected utility.

7 Comparison to Structural Model

The most common approach to measuring ex-ante welfare estimates a structural model. If one knows the utility function, information sets, and distribution of outcomes, one can recover measures of ex-ante expected utility.

In this section, I fit the estimated willingness to pay and cost curves to a fully specified structural model. The model relies heavily on the setup in Einav et al. (2013b) to model both adverse selection and moral hazard.³² The model is parsimonious but flexible enough to perfectly match the estimated willingness to pay and cost curves in Finkelstein et al. (2019).

I then use the structural model to validate the sufficient statistics approach relative to a structural benchmark. In Section 8, I use extensions of the model to evaluate the impact of violating Assumptions 1 and 2 in the proposed implementation, and in particular explore the impact of preference heterogeneity.

7.1 Setup

The structural model follows the environment developed in Einav et al. (2013a). In period ν , individuals obtain a realization of θ_ν and choose c and m to maximize utility, which is given by:

$$u(c, m; \theta_\nu) = -\frac{1}{\gamma} e^{-\gamma[(m-\lambda(\theta_\nu))-\frac{1}{2w}(m-\lambda(\theta_\nu))^2+c(m)]} \quad (20)$$

where $c(m) = y - p_I$ if the individual is insured and $c(m) = y - m - p_U$ if the individual is uninsured (I assume a full insurance contrast). The realization of θ_ν affects utility through the function $\lambda(\theta_\nu)$, which determines an individual's demand for medical spending. To see this, note that the first order conditions for m imply $m^I = \lambda + w$ and $m^U = \lambda$. This means that $\lambda(\theta_\nu)$ is the baseline demand for medical spending by the uninsured who pay its full cost and w is the causal effect of insurance on the individual's medical spending. For simplicity, I assume the causal effect of insurance, w , income, y , and risk aversion, γ , are all constant across individuals. Allowing w to be heterogeneous does not affect the results (conditional on matching the demand and cost curves). Section 8 explores the robustness of the results when one allows y and γ to vary across individuals.

Plugging in the choices of individuals, the utility functions for type θ_ν are given by $-\frac{1}{\gamma} e^{-\gamma[y-\lambda(\theta_\nu)-p_U]}$ if they are uninsured. The realized utility for the insured is given by $-\frac{1}{\gamma} e^{-\gamma[\frac{1}{2}w+y-p_I]}$.

³²To my knowledge, no previous paper has evaluated the ex-ante welfare impact of insurance market policies using a structural approach that includes moral hazard. In the current setting, allowing for moral hazard is essential to match the fact that the demand curve lies below the cost curve for a broad range of the distribution, as shown in Figure 4.

7.2 Matching Demand and Cost Curves

At the time of deciding whether to purchase insurance, individuals have some knowledge, given by θ_μ , about their realization of $\lambda(\theta_\nu)$, which generates their marginal willingness to pay for insurance. As in Section 3, the population is ordered descending in their willingness to pay, given by $D(p_U, S(\theta_\mu))$ where $S(\theta_\mu) = s$ corresponds to individuals with willingness to pay at the $1 - s^{th}$ quantile of the willingness to pay distribution. When a fraction s purchases insurance, all those with θ_μ such that $S(\theta_\mu) < s$ will purchase insurance.

The structural model requires the researcher to specify individuals' beliefs about future realizations of θ_ν . To do so, I assume that the realizations of $\lambda(\theta_\nu)$ are normally distributed with mean $E[\lambda(\theta_\nu) | s] = C(s)$ and variance $\Sigma(s)$, where $C(s)$ is the marginal cost curve in the economy and $\Sigma(s)$ is a variance term. This variance will be set below to match the willingness to pay and cost curves. Both $C(s)$ and $\Sigma(s)$ are known to the individual at the time of insurance purchase, but they may vary for individuals with different levels of willingness to pay.

The utility function in equation (20) implies that the marginal willingness to pay for insurance solves:

$$D(p_U, s) = C(s) + \frac{w}{2} + \frac{\gamma \Sigma(s)}{2} \quad (21)$$

Individuals are willing to pay their expected costs, $C(s)$, plus half of the moral hazard induced spending, $\frac{w}{2}$, plus an additional term corresponding to the risk premium provided by insurance: to the extent to which the insurance reduces the variance of their consumption, $\Sigma(s)$, the CARA utility structure implies that they value this reduction according to the risk aversion parameter, γ , divided by 2. The CARA utility structure implies that willingness to pay does not depend on p_U . Therefore, as in earlier sections, I write $D(s)$ to refer to $D(p_U, s)$.

Matching WTP and Cost Curves Risk aversion is set to $\gamma = 5x10^{-4}$. To account for individuals having WTP below cost, I set w to be equal to mean net costs of 1,336, which ensures that $D(s) > 0$ for all s .³³ Given w , γ and $C(s)$, I set $\Sigma(s)$ to solve (21). The resulting model structure is able to perfectly match any observed willingness to pay and cost curves at each value of s .³⁴

³³This corresponds to roughly a 30% moral hazard affect on gross costs, roughly consistent with previous empirical findings.

³⁴Moreover, the model structure satisfies the baseline assumptions made in the sufficient statistics approach above. For example, the CARA utility structure implies that the marginal willingness to pay is not a function of p_U . The model also does not exhibit complementarities between consumption and medical spending, $u_{cm} = 0$, so that it satisfies Assumption 1 to first order.

7.3 Measuring Ex-Ante Expected utility

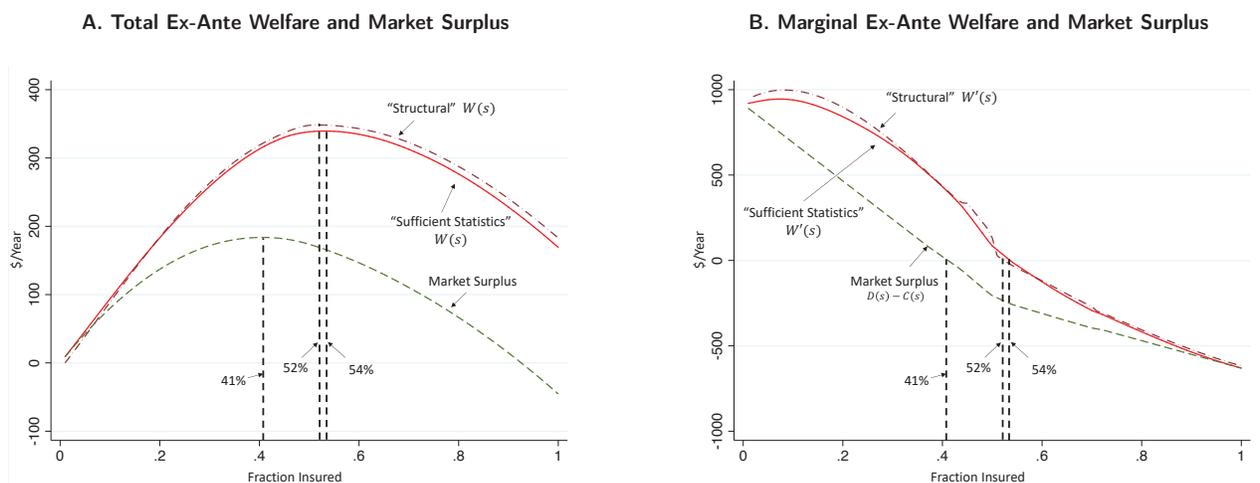
The structural model yields a straightforward solution for $W(s)$ in equation (11), as illustrated in the following proposition.

Proposition 5. *Under the modeling assumptions outlined above, $W(s)$ in equation (11) is given by*

$$W(s) = \frac{-1}{\gamma} \left[\log \left(se^{\gamma p_I(s)} e^{-\gamma \frac{w}{2}} + (1-s) e^{\gamma p_U(s)} \int_s^1 e^{\gamma(C(s)-w) + \frac{\gamma^2}{2} \Sigma(s)} d\bar{s} \right) - \log \left(E \left[e^{\gamma(C(S)-w) + \frac{\gamma^2}{2} \Sigma(S)} \right] \right) \right] \quad (22)$$

Proof. See Appendix F. □

Figure 7: Comparison to Structural Approach



7.4 Results

Figure 7A compares the value of $W(s)$ from the structural model computed using Equation (22) (shown in the dash-dot brown line), to the value of $W(s)$ computed using the sufficient statistics approach (in the solid red line). The latter is obtained by integrating $D(s) + EA(s) - C(s)$. The figure reveals that the sufficient statistics approach does a decent job of measuring the “true” ex-ante measure of welfare implied by the structural model.

Figure 7B compares the marginal willingness to pay measures, $W'(s)$. The again illustrates that the sufficient statistics approach approximates the structural model fairly well. The sufficient statistics correction clearly outperforms market surplus as a normative guide to ex-ante welfare. Ex-ante welfare is maximized in the structural model when $s = 52\%$

instead of 54% in the sufficient statistics implementation (compared with 41% for market surplus). The welfare gain from the optimal size of the insurance market is slightly larger in the structural approach (\$350 versus \$340), and the welfare gain from a full mandate is also slightly larger (\$180 versus \$170). Overall, the sufficient statistics approach proposed in this paper captures ex-ante welfare quite well.

8 Violations of Assumptions

The assumptions in Section 5 provide a benchmark method to estimate $\beta(s)$ that only requires knowledge of the market demand curve and risk aversion. However, these assumptions may be violated in practice. This section discusses these potential limitations and illustrates how to potentially overcome them with additional empirical estimates. Section 8.1 considers the case when Assumption 2 is violated. Section 8.2 considers the case when Assumption 1 is violated due to heterogeneity in risk preferences. Section 8.3 considers the case when Assumption 1 is violated due to correlations between health status and the marginal utility of income. Section 8.4 summarizes the results and provides a practical roadmap to researchers seeking to implement the sufficient statistics approach.

8.1 Income or Liquidity Differences

The benchmark implementation proxies for the difference in consumption between the insured and uninsured using the demand curve, $D(s) - E[D(S) | S \geq s]$. This would not be a good proxy if there were differences in liquidity or income between the insured and uninsured populations.³⁵ In the MA example, the insurance price variation is conditional on incomes of 150% FPL, which suggests this may not be a primary concern in this setting; but more generally it is likely that in many settings income is a key determinant of insurance demand.

This subsection uses the structural framework to explore the extent to which heterogeneity in income or liquidity can affect the proposed approach to measuring ex-ante welfare. If income variation is correlated with willingness to pay for insurance, the implementation using the market demand curve to capture $\beta(s)$ discussed in Section 5 will perform poorly. However, there is a straightforward way of correcting the estimates: instead of proxying consumption differences between the insured and uninsured, one could directly measure the

³⁵A more subtle violation of Assumption 2 arises when individuals can save across periods. In this case, those who spend money on insurance may be able to borrow or reduce savings, increasing their consumption to help cover the cost of the insurance. This would imply the baseline approach would over-state the ex-ante value of insurance. But, an approach that directly measures the difference in consumption - as suggested below - would correctly recover ex-ante welfare.

difference in consumption and avoid this potential bias.

To see this in the context of the model, suppose that individuals with different willingnesses to pay (generated by their knowledge at the time of choosing insurance, θ_μ) have different incomes in period ν , which I denote $y_\nu(\theta_\nu)$. To isolate variation in incomes that violates Assumption 2, I make the simplifying assumption that income does not vary conditional on θ_μ . Modifying the proof of Proposition 5 to allow for income heterogeneity, Appendix G shows that $W(s)$ solves

$$e^{-\gamma W(s)} \int_0^1 e^{\gamma(\bar{y}(\bar{s})+C(\bar{s})-w)+\frac{1}{2}\gamma^2\Sigma(\bar{s})} = e^{\gamma p_I(s)} e^{-\frac{1}{2}\gamma w} \int_0^s e^{-\gamma\bar{y}(\bar{s})} d\bar{s} + e^{\gamma p_U(s)} \int_s^1 e^{\gamma(\bar{y}(\bar{s})+C(\bar{s})-w)+\frac{1}{2}\gamma^2\Sigma(\bar{s})} d\bar{s}$$

To the extent to which individuals with different willingnesses to pay have different incomes, this will lead the $W(s)$ function to differ from the baseline case with no income variation.

I calibrate income heterogeneity by starting with a base income of \$16,335, which corresponds to 150% FPL for a single adult in 2011 in Massachusetts. I then consider two cases corresponding to whether those with a higher willingness to pay for insurance have (a) higher incomes or (b) lower incomes. For the higher incomes case, I assume that those with the highest level of willingness to pay have an average income that is \$1,382 higher than those with the lowest willingness to pay, where \$1382 is calibrated to be the mean health costs. In contrast, for the lower incomes case, I assume the opposite: those with the highest demand have incomes that are \$1382 lower than those with the highest willingness to pay. In both cases I assume average incomes are a linear function of s .

Figure 8A and 8B present the ex-ante willingness to pay curve, $W'(s)$, for these cases. Figure 8A reports the results when income is positively correlated with willingness to pay. Figure 8B presents the results when income is negatively correlated with willingness to pay. When the insured have higher incomes, their marginal utility is lower than is implied by the demand curve. The benchmark sufficient statistics implementation that does not account for this will over-state the true willingness to pay. Conversely, Figure 8B shows how the pattern reverses when the insured have higher incomes. In this case, the benchmark approach in Section 5 will under-state the ex-ante value of insurance because the insured have a higher marginal utility of income than is implied by the demand curve.

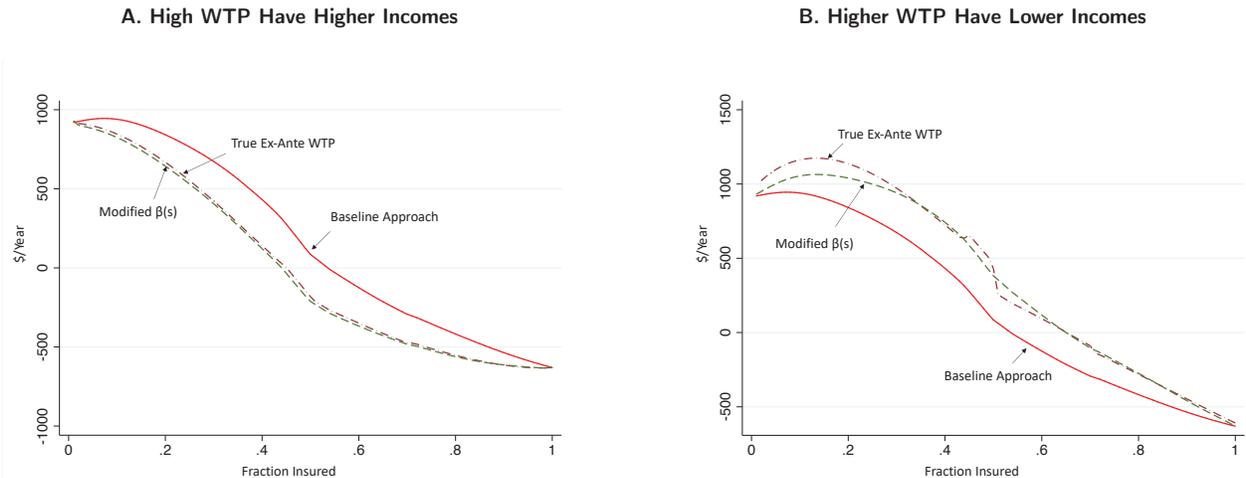
Consumption Data In the presence of income or liquidity heterogeneity, the demand curve no longer provides adequate guidance on the differences in consumption between the insured and uninsured. In principle, one can recover a measure of ex-ante welfare if the researcher observes consumption or disposable income data. To first order one can use a modified formula for $\beta(s)$ in the case when $x(m; \theta) = 0$ as

$$\beta(s) = \gamma (D(s) - E[D(S) | S \geq s] + E[y_\nu(\theta_\nu) | S \leq s] - E[y_\nu(\theta_\nu) | S > s])$$

where $E[y_\nu(\theta_\nu) | S \leq s] - E[y_\nu(\theta_\nu) | S > s]$ is the difference in disposable incomes between the insured and uninsured. More generally, if one observes consumption one can estimate $\beta(s) = \gamma \Delta c$ where Δc is the difference in consumption between the insured and uninsured.

The dashed line in Figures 8A and 8B shows the values of $W'(s)$ that use differences in consumption to measure the correct $\beta(s)$. As the figure reveals, one recovers a relatively accurate measure of ex-ante welfare if one can observe consumption data.³⁶

Figure 8: Income or Liquidity Differences



8.2 Heterogeneity in Risk Preferences

Most structural frameworks allow for heterogeneity in preferences, as this is often important to match observed variation in willingness to pay for insurance. However, preference heterogeneity – at least any heterogeneity that is correlated with the marginal utility of consumption – is ruled out by Assumption 1.

This subsection uses a generalized version of the structural framework in Section 7 to (a) document how preference heterogeneity can introduce bias in the benchmark implementation, (b) clarify why this bias arises, and (c) propose a modified approach that recovers ex-ante welfare. The main conclusion is that the benchmark implementation remains valid if one uses an “ex-ante” coefficient of risk aversion that corresponds to the markup individuals would be willing to pay to insure themselves against having to buy insurance. If preferences are stable, then the risk aversion observed from choices in the market can be used to measure

³⁶Absent consumption data, one can proxy for consumption using differences in income and assuming no differential savings between insured and uninsured.

the ex-ante risk aversion parameter. But, in the presence of preference heterogeneity there is in general no such guarantee that the risk aversion values observed in the market (from deductible or copay choices) correspond to the risk aversion that would govern the ex-ante demand for insurance against the risk of having to purchase insurance.

In principle, the observed risk aversion in the market (e.g. the desire to insure against the realization of θ_ν conditional on θ_μ , which is identified from deductible and copay choices in the market) could be higher or lower than ex-ante risk aversion governing the willingness to pay to insure against the need to buy insurance (i.e. against realizations of θ_μ). To see this, return to the modeling framework as in Section 7. Recall preferences are given by

$$u(c, m; \theta_\nu) = -\frac{1}{\gamma} e^{-\gamma[(m-\lambda(\theta_\nu))-\frac{1}{2w}(m-\lambda(\theta_\nu))^2+c(m)]}$$

where γ is a coefficient of absolute risk aversion. Section 7 assumed $\gamma = 5x10^{-4}$ for all individuals. Now, suppose $\gamma \sim N(\mu_\gamma, \sigma_\gamma^2)$, where $\mu_\gamma = 5x10^{-4}$ is the mean level of risk aversion and σ_γ is its population standard deviation. I consider two values of $\sigma_\gamma \in \{5x10^{-5}, 1x10^{-4}\}$. I refer to $\sigma = 5x10^{-4}$ as the “low” heterogeneity value and $\sigma = 1x10^{-4}$ as the “high” heterogeneity value.³⁷ I assume risk aversion is independent of the realization of costs.³⁸

When individuals have heterogeneity in γ , there are two reasons that individuals might have high willingness to pay: on the one hand, $D(s)$ might be driven by their expected costs (or higher variance of costs). On the other hand, $D(s)$ might be driven by high risk aversion. Therefore, for two individuals with the same willingness to pay, $D(s)$, it is no longer the case that it would be reasonable to assume they have the same expected costs that equals the marginal cost curve, $C(s)$. At each value of s , the cost curve reflects an average of individuals with different expected costs who have the same quantile of demand, $1 - s$, but arrive at that quantile through different levels of risk aversion. As a result, I fit a model with heterogeneity in individual’s expected costs that can differ from expected cost given $D(s)$. Appendix H develops this extension of the model in Section 7 and discusses the minimum distance estimator that is used to solve the model’s underlying cost heterogeneity parameters.

Figure 9 presents the results for the low heterogeneity case (Figure 9A) and high hetero-

³⁷This higher value of heterogeneity corresponds most closely to the estimates in Handel et al. (2015).

³⁸In contrast to this assumption, empirical applications of structural models of insurance often finding positive (Cohen and Einav (2007)) or negative (Handel et al. (2015)) correlations between risk and risk aversion. It is straightforward to establish that the conclusions developed here are similar for the case in which there is a correlation with costs. The presence of correlated risk aversion changes the extent to which the ex-ante risk aversion differs from the risk aversion distribution in the market. But conceptually, one continues to need to measure the ex-ante risk aversion coefficient that governs the ex-ante willingness to pay for a financial product that pays \$1 in the event one chooses to purchase insurance in the market.

generosity case (Figure 9B). The dot-dash brown line corresponds to the true ex-ante surplus $W'(s)$ that is implied by the structural framework with preference heterogeneity. The solid red line presents the benchmark implementation of $EA(s)$ using a coefficient of relative risk aversion of $\gamma = 5 \times 10^{-4}$, corresponding to the population average.

For both the high and low heterogeneity cases, the benchmark implementation over-states the ex-ante value of insurance. These differences are larger for the case with a high degree of preference heterogeneity, as illustrated by the difference between Figure 9A and Figure 9B.

Ex-ante vs. Ex-post Risk Aversion Why does the benchmark implementation over-state the ex-ante value of insurance implied by the structural model? The reason is that in the presence of preference heterogeneity the average risk aversion coefficient does not correspond to the “ex-ante” risk aversion that governs the willingness to pay for insurance against ex-ante risk. To see why, recall the utility function in the structural model has a CARA structure of the form $\frac{-1}{\gamma}e^{-\gamma c}$. This structure implies a marginal utility function of $e^{-\gamma c}$. Crucially, this is decreasing in γ . This implies that those who have higher risk aversion tend to have lower marginal utilities of consumption. Because those with higher risk aversion will also purchase insurance, it means that the difference in marginal utilities between the insured and uninsured will be smaller than if there was no heterogeneity in γ .

To see how this plays out in the structural model, one can calculate the implied risk aversion that individuals exhibit from behind a veil of ignorance in order to be consistent with the model’s true calculation of $EA(s)$. This is given by the dash-dot brown line in Figure 9).³⁹ The median risk aversion value is 3.9×10^{-4} for the low heterogeneity case and 6.7×10^{-5} for the high heterogeneity case. Because the CARA utility structure implies that the insured have lower *marginal* utilities of income, the implied ex-ante CARA coefficient is not equal to the mean risk aversion coefficient – it is instead much lower.

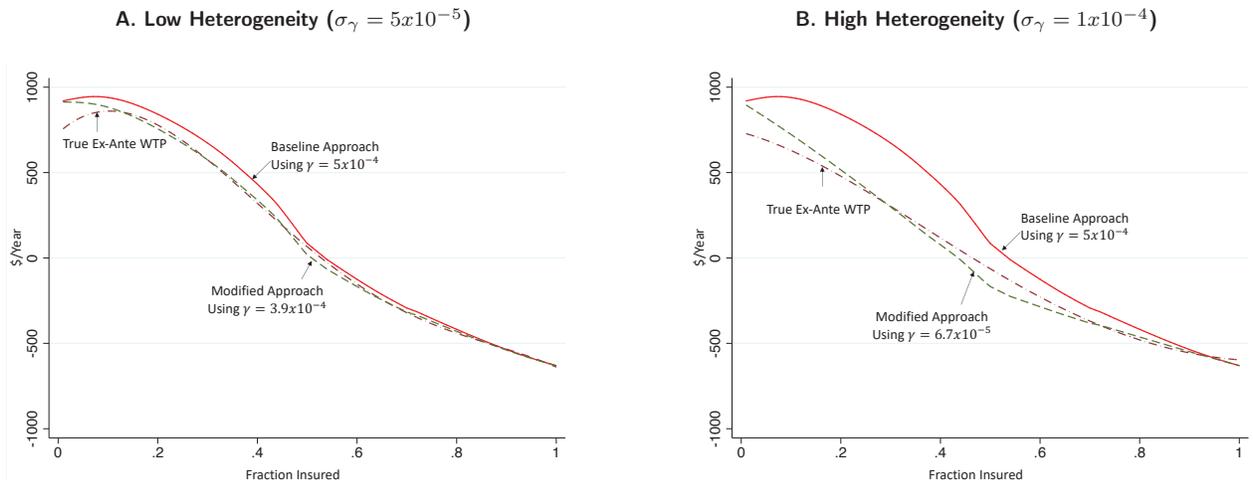
The fact that ex-ante risk aversion is lower than the mean risk aversion is not a general feature of all utility functions. Indeed, it is not even a feature of all utility functions with constant absolute risk aversion. For example, all utility functions of the form $-\frac{f(\gamma)}{\gamma}e^{-\gamma c}$ for any function $f(\circ)$ have a coefficient or absolute risk aversion of γ . As a result, there exists utility functions that perfectly match the demand and cost curves for any value of f . If f' is sufficiently large, the marginal utility of income will be *increasing* not decreasing with the level of risk aversion, not decreasing as is implied by $\frac{-1}{\gamma}e^{-\gamma c}$. This means that in general,

³⁹When prices are such that a fraction s of the market is insured, it is straightforward to show that is given by

$$\gamma^{ex-ante}(s) = \frac{-1}{s(1-s)D'(s)} \frac{EA(s)}{D(s) - E[D(S(\theta_\mu)) | S(\theta_\mu) \geq s]}$$

the relationship between ex-ante risk aversion – namely the risk aversion exhibited by those seeking to insure against the realization of θ_μ – can differ from the risk aversion exhibited by those in the insurance market – namely the risk aversion determining the copay and deductible choice to help insure against the realization of θ_ν given θ_μ .

Figure 9: Risk Aversion Heterogeneity



Correction Using “Ex-Ante” Risk Aversion One can recover a measure of ex-ante welfare through the reduced form approach if one uses the correct ex-ante risk aversion corresponding to the markups individuals would be willing to pay for insurance against the risk of having to buy insurance in the market (i.e. insurance against θ_μ , which may differ from insurance against θ_ν given θ_μ). To illustrate this, the green dashed lines in Figure 9 present the ex-ante surplus calculated using a coefficient of risk aversion of $3.9x10^{-4}$ for the low heterogeneity case and $6.7x10^{-5}$ for the high heterogeneity case. These risk aversion values correspond to the model’s implied risk aversion against the ex-ante realization of θ_μ . As illustrated in the figure, using these “ex-ante” risk aversion values generate measures of $W'(s)$ that are much closer to the true value of $W'(s)$.

In practice, this suggests that while the risk aversion that is revealed through insurance choices is a natural starting point for measuring ex-ante welfare, it may differ from the ex-ante coefficient of absolute risk aversion – especially in settings with significant preference heterogeneity. Ideally, one would like to elicit the risk aversion that determines the demand for an insurance product that would pay benefits in the even the individual found themselves in the insurance market and needing to purchase insurance. In practice, a feasible approach would consider the robustness of the results to a range of risk aversion parameters including

those beyond what is implied from copay/deductible choices that insure the risk of θ_ν given θ_μ to shed light on the willingness to pay for insuring ex-ante risk of the realization of θ_μ .

8.3 Marginal Utility and Health Status

In addition to assuming no heterogeneity in risk aversion, Assumption 1 also embeds the common assumption in the unemployment insurance literature that the marginal utility of income does not depend on health status, conditional on a given level of consumption. If the sick have higher marginal utilities of income and the sick buy insurance, then the ex-ante willingness to pay for insurance would be biased downwards: the difference in marginal utilities between insured and uninsured, $\beta(s)$, would be larger than what is implied by $\gamma\Delta c$. In contrast, if the sick have lower marginal utilities of income, then ex-ante willingness to pay for insurance would be biased upwards.

To address this, one can extend the estimation of $\beta(s)$ to include an additional term that captures the differences in marginal utilities arising from differences in health status. To provide some guidance on how this could be done in practice, Finkelstein et al. (2013) use a parametric utility model combined with self-reported health and well-being elicitation to infer that those who report being in poor health have a 10-25pp lower marginal utility of income, conditional on a given level of income. It is difficult to directly translate their health status results to the present setting because one does not know the relationship between subjective health reports and insurance purchase in the MA health insurance setting. But, to provide a back-of-the-envelope illustration of how this approach could proceed, suppose that 1/3 of the insured are sick and none of the uninsured are sick. Then, the insured would have a 3-8pp lower marginal utility of income than would be implied by their difference in consumption. Hence, one could consider adjusting $\beta(s)$ downward by 3-8pp. For comparison, the estimated value of $\beta(0.5)$ in the baseline implementation is 0.34. Hence, this back-of-the-envelope adjustment would lower $EA(s)$ by roughly 10-25%. While this exercise is highly stylized, it shows how one can incorporate future work estimating the complementarities between health status and the marginal utility of income.

8.4 Summary

In summary, the benchmark implementation approach measures ex-ante expected utility using only the market willingness to pay and cost curves, combined with a measure of risk aversion. Yet, the implementation assumptions that enable this implementation are strong. In practice, researchers will want to consider how these assumptions may affect their conclusions in their setting.

The robustness analysis presented here suggest proceeding as follows. First, if researchers suspect there are income or liquidity differences between the insured and uninsured, the benchmark implementation may be biased because the willingness to pay curve will no longer be a good proxy for the difference in consumption between insured and uninsured. This means that there would be a high value in obtaining consumption data that would allow one to construct the percentage difference in consumption between insured and uninsured. With data on consumption, one could estimate $\beta(s)$ using equation 18, $\beta(s) = \gamma\Delta c$.

Second, the ideal risk aversion parameter corresponds to an “ex-ante” measure of risk aversion: it determines an individual’s hypothetical willingness to pay for a financial product that pays in the event they end up in the market and need to buy insurance. In the presence of heterogeneity in risk aversion, it is natural to worry that the ex-ante measure of risk aversion differs from the risk aversion revealed through choices of copays or deductibles in the market, even if one measured the average risk aversion in the market. In practice, researchers will want to study the robustness of the results to a range of risk aversion parameters – akin to the exploration in Table 1 above - with the knowledge that the ideal risk aversion is the one that governs the hypothetical willingness to pay for a financial product that pays \$1 in the event they end up in the market and wish to buy insurance.

More broadly, the impact of preference heterogeneity on the benchmark approach highlights the value of future work that uses other methods to elicit the markup individuals would be willing to pay for a product that helps lower their future insurance prices. Such a markup would correspond directly to the percentage difference in marginal utilities between the insured and uninsured, $\beta(s)$, and could be directly used as an implementation of the general model outlined in Section 3 and 4, avoiding the strong implementation assumptions in Section 5. The unemployment insurance literature has developed a range of different approaches to measuring differences in marginal utilities of income (e.g. see Landais and Spinnewijn (2019)), which suggests future work could develop similar approaches tailored to the health insurance context.

9 Conclusion

This paper develops a set of tools to measure the impact of insurance market policies on ex-ante measures of welfare. Ex-ante measures of welfare differ from market surplus because they measure expected utility before individuals learn their willingness to pay for insurance. An ex-ante welfare perspective can lead to different normative conclusions as compared to a market surplus perspective. Policies that maximize ex-ante expected utility often involve lower insurance prices, a greater value of mandates, and a higher value of insurance subsidies.

Future work could measure the welfare consequences of contract distortions, such as the exclusion of high cost drugs for chronic conditions. It could also expand beyond the binary insurance decision considered here to consider menus of insurance contracts. Future work could also extend the results to normative frameworks that allow for behavioral biases (e.g. as in Spinnewijn (2017)), which have been shown to be important in health insurance settings.

Future work can also extend the ideas developed here to settings where prices are not observed. For example, many approaches use labor supply responses to infer the value of social insurance programs (e.g. Keane and Moffitt (1998); Gallen (2014); Dague (2014)). These approaches capture the value of insurance against only the risk that remains after choosing labor supply. Other approaches use changes in consumption around a shock to infer willingness to pay (e.g. Gruber (1997); Meyer and Mok (2013)). But consumption should change when information about the event is revealed, not when the event occurs. The approaches developed here could be extended to measure ex-ante expected utility in such settings.

Lastly, many macroeconomic welfare measures face similar conceptual issues. This includes the famous calculations of the welfare cost of business cycles in Lucas (2003). When consumption responds to information over time, the variance of consumption changes may under-state measures of ex-ante welfare. Future work could extend the tools in this paper to measure the ex-ante welfare cost of business cycles and other macroeconomic risk.

References

- Baily, M. N. (1978). Some aspects of optimal unemployment insurance. *Journal of Public Economics* 10(3), 379–402. 1, 5.1, 5.1
- Cabral, M. (2017). Claim timing and ex post adverse selection. *Review of Economic Studies* 84(1). 1
- Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics* 90(10), 1879–1901. 1, 3.1, 5.1, 5.1
- Cohen, A. and L. Einav (2007). Estimating risk preferences from deductible choice. *American economic review* 97(3), 745–788. 38
- Dague, L. (2014). The effect of Medicaid premiums on enrollment: A regression discontinuity approach. *Journal of Health Economics* 37, 1–12. 1, 9

- Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating welfare in insurance markets using variation in prices. *Quarterly Journal of Economics* 125(3). 1, 2, 10, 3.1, 16, C
- Einav, L., A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen (2013a). Selection on moral hazard in health insurance. *American Economic Review* 103(1), 178–219. 1, 7.1
- Einav, L., A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen (2013b). Selection on moral hazard in health insurance. *American Economic Review* 103(1), 178–219. 7
- Einav, L., A. Finkelstein, and H. Williams (2016). Paying on the margin for medical care: Evidence from breast cancer treatments. *American Economic Journal: Economic Policy* 8(1), 52–79. 2, 2, 20
- Feldman, R. D. and B. E. Dowd (1982). Simulation of a health insurance market with adverse selection. *Operations Research* 30(6), 1027–1042. 1
- Finkelstein, A., N. Hendren, and E. F. Luttmer (2016). The value of medicaid: Interpreting results from the oregon health insurance experiment. *NBER Working Paper*, <http://scholar.harvard.edu/hendren/publications/value-medicaid-interpreting-results-oregon-health-insurance-experiment>. 2
- Finkelstein, A., N. Hendren, and M. Shepard (2019). Subsidizing health insurance for low-income adults: Evidence from massachusetts. *American Economic Review* 109(4), 1530–67. 1, 3.1, 3.2, 18, 6, 29, 6.2, 7
- Finkelstein, A., E. F. Luttmer, and M. J. Notowidigdo (2013). What good is wealth without health: The effect of health on the marginal utility of consumption. *Journal of the European Economic Association* 11(s1), 221–258. 8.3
- Finkelstein, A., K. McGarry, and A. Sufi (2005). Dynamic inefficiencies in insurance markets: Evidence from long-term care insurance. *American Economic Review* 95(2), 224–228. 1
- Gallen, T. S. (2014). Measuring the value of Medicaid using TennCare. *Working Paper*. 1, 5, 9
- Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review* 87(1), 192–205. 1, 9
- Hackmann, M. B., J. T. Kolstad, and A. E. Kowalski (2015). Adverse selection and an individual mandate: When theory meets practice. *The American Economic Review* 105(3), 1030–1066. 1

- Handel, B., I. Hendel, and M. D. Whinston (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica* 83(4), 1261–1313. 2, 1, 2, 3.1, 12, 20, 6.1, 28, 6.1, 37, 38
- Harsanyi, J. C. (1978). Bayesian decision theory and utilitarian ethics. *The American Economic Review* 68(2), 223–228. 4, 3.1
- Hendren, N. (2013). Private information and insurance rejections. *Econometrica* 81(5), 1713–1762. 1
- Hendren, N. (2014). The inequality deflator: Interpersonal comparisons without a social welfare function. *NBER Working Paper*, <http://www.nber.org/papers/w20351>. 4.2
- Hendren, N. (2016). The policy elasticity. *Tax Policy and the Economy* 30. 1, 4.2
- Hendren, N. (2017). Knowledge of future job loss and implications for unemployment insurance. *American Economic Review* 107(7). 1
- Hendren, N. and B. D. Sprung-Keyser (2019). A unified welfare analysis of government policies. Technical report, National Bureau of Economic Research. 1, 4.2, 4.2, 6.2
- Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. *The American Economic Review* 61(4), 561–574. 1, 2
- Keane, M. and R. Moffitt (1998). A structural model of multiple welfare program participation and labor supply. *International economic review*, 553–589. 1, 9
- Landais, C. and J. Spinnewijn (2019). The value of unemployment insurance. *Working Paper*. 8.4
- Lucas, R. E. (2003). Macroeconomic priorities. *American economic review* 93(1), 1–14. 9
- Meyer, B. D. and W. K. Mok (2013). Disability, earnings, income and consumption. Technical report, National Bureau of Economic Research. 1, 9
- Panhans, M. (2018). Adverse selection in aca exchange markets: Evidence from colorado. *The American Economic Journal: Applied Economics*. 1
- Spinnewijn, J. (2017). Heterogeneity, demand for insurance, and adverse selection. *American Economic Journal: Economic Policy* 9(1), 308–43. 9

Online Appendix: Not For Publication

A Proof of Proposition 1

This Appendix walks through the proof of Proposition 1. To begin, the derivative, $\frac{\partial V(\infty, 0, W(s))}{\partial \tau}$, can be written as:

$$\frac{\partial V(\infty, 0, W(s))}{\partial \tau} = E \left[\frac{\partial u_\nu}{\partial c_\nu} \mid \theta_\mu \in \Theta_\mu \right]$$

where $\frac{\partial u_\nu}{\partial c_\nu}$ is evaluated using the choices uninsured individuals would make with incomes increased by $W(s)$ in all states of the world. Combining,

$$\begin{aligned} W'(s) &= \frac{sp'_I(s) E \left[\frac{-\partial u_\nu}{\partial c_\nu} \mid S(\theta_\mu) < s \right] + (1-s)p'_U(s) E \left[\frac{-\partial u_\nu}{\partial c_\nu} \mid S(\theta_\mu) \geq s \right]}{E \left[\frac{\partial u_\nu}{\partial c_\nu} \mid \theta_\mu \in \Theta_\mu \right]} \\ &= -(sp'_I(s) + (1-s)p'_U(s)) \frac{E \left[\frac{-\partial u_\nu}{\partial c_\nu} \mid S(\theta_\mu) < s \right]}{E \left[\frac{\partial u_\nu}{\partial c_\nu} \mid \theta_\mu \in \Theta_\mu \right]} + (1-s)p'_U(s) \frac{E \left[\frac{\partial u_\nu}{\partial c_\nu} \mid S(\theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c_\nu} \mid S(\theta_\mu) \geq s \right]}{E \left[\frac{\partial u_\nu}{\partial c_\nu} \mid \theta_\mu \in \Theta_\mu \right]} \end{aligned}$$

Now, turn to the resource constraint when $G(s) = 0$:

$$sp_I(s) + (1-s)p_U(s) = sAC(s)$$

Differentiating yields:

$$p_I(s) + s \frac{dp_I}{ds} + (1-s) \frac{dp_U}{ds} - p_U(s) = C(s)$$

or

$$sp'_I(s) + (1-s)p'_U(s) = C(s) - p_I(s) + p_U(s)$$

or

$$-(sp'_I(s) + (1-s)p'_U(s)) = D(p_U(s), s) - C(s)$$

Moreover, one can express $-p'_U(s)$ as follows. Differentiating the demand identity $p_I(s) - p_U(s) =$

$D(p_U(s), s)$ yields $p'_I(s) - p'_U(s) = \frac{\partial D}{\partial s} + \frac{\partial D}{\partial p_U} p'_U(s)$. So,

$$\begin{aligned} -s p'_I(s) - (1-s) p'_U(s) &= (p_I(s) - p_U(s)) - C(s) \\ -s \left[\frac{\partial D}{\partial s} + \frac{\partial D}{\partial p_U} p'_U(s) \right] - (1-s) p'_U(s) &= (p_I(s) - p_U(s)) - C(s) \\ -p'_U(s) \left[1 + s \left(\frac{\partial D}{\partial p_U} \right) \right] &= (p_I(s) - p_U(s)) - C(s) + s \frac{\partial D}{\partial s} \\ -p'_U(s) &= \frac{1}{1 + s \frac{\partial D}{\partial p_U}} \left[(p_I(s) - p_U(s)) - C(s) + s \frac{\partial D}{\partial s} \right] \end{aligned}$$

or

$$p'_U(s) = \frac{-1}{1 + s \frac{\partial D}{\partial p_U}} \left[D(p_U(s), s) - C(s) + s \frac{\partial D}{\partial s} \right]$$

Now, plugging back into the expression for $W'(s)$ yields

$$W'(s) = (D(p_U(s), s) - C(s)) \frac{E \left[\frac{-\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right]}{E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right]} - \frac{(1-s)}{1 + s \frac{\partial D}{\partial p_U}} \frac{E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s \right]}{E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right]} \left[D(p_U(s), s) - C(s) + s \frac{\partial D}{\partial s} \right]$$

or

$$\begin{aligned} W'(s) &= (D(p_U(s), s) - C(s)) \frac{E \left[\frac{-\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] - \frac{1-s}{1 + s \frac{\partial D}{\partial p_U}} \left(E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s \right] \right)}{E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right]} \\ &+ \frac{s(1-s)}{1 + s \frac{\partial D}{\partial p_U}} \left(-\frac{\partial D}{\partial s} \right) \frac{E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s \right]}{E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right]} \end{aligned}$$

And,

$$\begin{aligned} &E \left[\frac{-\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] - \frac{1-s}{1 + s \frac{\partial D}{\partial p_U}} \left(E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s \right] \right) \\ &= \frac{1 + s \frac{\partial D}{\partial p_U}}{1 + s \frac{\partial D}{\partial p_U}} E \left[\frac{-\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] - \frac{1-s}{1 + s \frac{\partial D}{\partial p_U}} \left(E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s \right] \right) \\ &= \frac{1}{1 + s \frac{\partial D}{\partial p_U}} \left(s \left(1 + \frac{\partial D}{\partial p_U} \right) E \left[\frac{-\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] + (1-s) E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s \right] \right) \\ &= \frac{1}{1 + s \frac{\partial D}{\partial p_U}} \left(s E \left[\frac{-\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] + (1-s) E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s \right] \right) + \frac{s \frac{\partial D}{\partial p_U}}{1 + s \frac{\partial D}{\partial p_U}} E \left[\frac{-\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] \\ &= \frac{1}{1 + s \frac{\partial D}{\partial p_U}} E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right] + \frac{s \frac{\partial D}{\partial p_U}}{1 + s \frac{\partial D}{\partial p_U}} E \left[\frac{-\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] \\ &= E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right] + E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right] \left(\frac{1}{1 + s \frac{\partial D}{\partial p_U}} - \frac{1 + s \frac{\partial D}{\partial p_U}}{1 + s \frac{\partial D}{\partial p_U}} \right) + \frac{s \frac{\partial D}{\partial p_U}}{1 + s \frac{\partial D}{\partial p_U}} E \left[\frac{-\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] \\ &= E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right] - \left(\frac{s \frac{\partial D}{\partial p_U}}{1 + s \frac{\partial D}{\partial p_U}} \right) \left(E \left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu \right] + E \left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s \right] \right) \end{aligned}$$

So,

$$\begin{aligned}
W'(s) &= (D(p_U(s), s) - C(s)) + s(1-s) \left(-\frac{\partial D}{\partial s}\right) \beta(s) - \frac{\left(s \frac{\partial D}{\partial p_U}\right) s(1-s)}{1+s \frac{\partial D}{\partial p_U}} \left(-\frac{\partial D}{\partial s}\right) \beta(s) - \left(\frac{s \frac{\partial D}{\partial p_U}}{1+s \frac{\partial D}{\partial p_U}}\right) \left(\frac{E\left[\frac{\partial u_V}{\partial c_V} | \theta_\mu \in \Theta_\mu\right] + E\left[\frac{\partial u_V}{\partial c_V} | S(\theta_\mu) < s\right]}{E\left[\frac{\partial u_V}{\partial c_V} | \theta_\mu \in \Theta_\mu\right]}\right) \\
&= (D(p_U(s), s) - C(s)) + s(1-s) \left(-\frac{\partial D}{\partial s}\right) \beta(s) + \frac{\partial D}{\partial p_U} \frac{-s}{1+s \frac{\partial D}{\partial p_U}} \left(1 + (1-s) \beta(s) \left(1 + s \frac{\partial D}{\partial s}\right)\right)
\end{aligned}$$

Lastly, note that $EA(s)$ is given by the second term in the equation above:

$$EA(s) = s(1-s) \left(-\frac{\partial D}{\partial s}\right) \beta(s)$$

and $\delta_p(s)$ is given by the third term:

$$\delta_p(s) = \frac{\partial D}{\partial p_U} \frac{-s}{1+s \frac{\partial D}{\partial p_U}} \left(1 + (1-s) \beta(s) \left(1 + s \frac{\partial D}{\partial s}\right)\right)$$

where clearly $\delta_p(s) = 0$ when the marginal willingness to pay does not depend on p_U , $\frac{\partial D}{\partial p_U} = 0$. Re-writing the expression for $W'(s)$, one arrives at $W'(s) = D(p_U(s), s) - C(s) + EA(s) + \delta_p(s)$.

B Proof of Proposition 2

Every dollar the government spends on additional subsidies imposes a cost on the government of $1 - \frac{C(s)-D(s)}{sD'(s)}$. And, from behind the veil of ignorance, the willingness to pay for these subsidies in units of the average income of those with $\theta_\mu \in \Theta_\mu$ is $\frac{E\left[\frac{-\partial u_V}{\partial c_V} | S(\theta_\mu) < s\right]}{E\left[\frac{-\partial u_V}{\partial c_V} | \theta_\mu \in \Theta_\mu\right]}$. Therefore, the MVPF is given by

$$MVPF(s) = \frac{E\left[\frac{-\partial u_V}{\partial c_V} | S(\theta_\mu) < s\right]}{E\left[\frac{-\partial u_V}{\partial c_V} | \theta_\mu \in \Theta_\mu\right]} \frac{1}{1 - \frac{C(s)-D(s)}{sD'(s)}}$$

Now,

$$\begin{aligned}
\frac{E\left[\frac{\partial u_V}{\partial c_V} | S(\theta_\mu) < s\right]}{E\left[\frac{\partial u_V}{\partial c_V} | \theta_\mu \in \Theta_\mu\right]} &= \frac{sE\left[\frac{\partial u_V}{\partial c_V} | S(\theta_\mu) < s\right] + (1-s)E\left[\frac{\partial u_V}{\partial c_V} | S(\theta_\mu) < s\right]}{E\left[\frac{\partial u_V}{\partial c_V} | \theta_\mu \in \Theta_\mu\right]} \\
&= \frac{sE\left[\frac{\partial u_V}{\partial c_V} | S(\theta_\mu) < s\right] + (1-s)E\left[\frac{\partial u_V}{\partial c_V} | S(\theta_\mu) \geq s\right]}{E\left[\frac{\partial u_V}{\partial c_V} | \theta_\mu \in \Theta_\mu\right]} + \frac{(1-s)\left(E\left[\frac{\partial u_V}{\partial c_V} | S(\theta_\mu) < s\right] - E\left[\frac{\partial u_V}{\partial c_V} | S(\theta_\mu) \geq s\right]\right)}{E\left[\frac{\partial u_V}{\partial c_V} | \theta_\mu \in \Theta_\mu\right]} \\
&= 1 + (1-s)\beta(s)
\end{aligned}$$

Hence, the MVPF is given by

$$MVPF(s) = \frac{1 + (1-s)\beta(s)}{1 + \frac{C(s)-D(s)}{s(-D'(s))}}$$

C Measuring Risk Aversion

In addition to the demand and cost curves in the Einav et al. (2010) framework, measuring ex-ante willingness requires an estimate of risk aversion, $\gamma(s)$. As discussed in Section 8.2, the relevant notion of risk aversion is the one that governs willingness to pay for insurance against ex-ante risk – namely the risk of having a risk type realization of θ_μ that leads you to buy insurance. In the absence of heterogeneity in risk aversion, one can principle infer the relevant risk aversion parameter within the demand and cost curve setup. Risk aversion is revealed by comparing individual’s willingness to pay for insurance to the reduction in variance of expenditures that is provided by the insurance product. For example, it is well-known that if preferences have a constant absolute risk aversion and the risk of medical expenditures is normally distributed (i.e. a “CARA-Normal” model), then the markup individual’s are willing to pay for insurance by is given by the variance reduction offered by the insurance multiplied by $\frac{\gamma(s)}{2}$.

More generally, one can consider a second-order Taylor approximation to equation (6) that characterizes willingness to pay, $D(\tilde{s})$. Again, I assume marginal willingness to pay does not depend on p_U . Let $p(\tilde{s}) = \frac{\partial x}{\partial m}$ denote the price of additional medical spending when insured. Under the additional assumption that $u_{mm} = 0$, then the it is straightforward to show⁴⁰ that the coefficient of absolute risk aversion is given by:

$$\gamma(s) = 2 \frac{D(s) - C(s) + (1 - p(s)) E[m^I(\theta_\nu) - m^U(\theta_\nu) | S(\theta_\mu) = s]}{V} \quad (23)$$

⁴⁰To see this, suppress notation w.r.t. θ_ν and condition all expectations on $S(\theta_\mu) = s$ for a given s . Let $(\bar{c}, \bar{m}, \bar{\theta})$ denote the average bundle of a type with $S(\theta_\mu) = s$. Taking a Taylor expansion to the utility function around this bundle in equation (6) yields

$$\begin{aligned} & u_c \left(E[y - x^I - D(s) - p_U - \bar{c}] \right) + \frac{1}{2} u_{cc} \left(E[y - x^I - D(s) - p_U - \bar{c}]^2 \right) + u_m E[m^I - \bar{m}] \\ &= u_c (E[y - p_U] - \bar{c}) + \frac{1}{2} u_{cc} (E[y - p_U - \bar{c}]^2) + u_m E[m^U - \bar{m}] \end{aligned}$$

or

$$\begin{aligned} & u_c \left((E[y - x^I - D(s) - p_U - \bar{c}]) - (E[y - p_U] - \bar{c}) \right) \\ &= \frac{-1}{2} u_{cc} \left[E[y - x^I - D(s) - p_U - \bar{c}]^2 - E[y - p_U - \bar{c}]^2 \right] + u_m E[m^U - m^I] \end{aligned}$$

or

$$\begin{aligned} D(s) - (E[m^I - x^I]) &= \frac{\gamma(s)}{2} V + \left(\frac{u_m}{u_c} - 1 \right) E[m^I - m^U] \\ \gamma(s) &= 2 \frac{D(s) - C(s) + \left(1 - \frac{u_m}{u_c} \right) E[m^I - m^U]}{V} \end{aligned}$$

where $C(s) = E[m^I(\theta_\nu) - x^I(\theta_\nu) | S(\theta_\mu) = s]$ is the cost to the insurer of enrolling the type s .

where $D(s) - C(s)$ is the markup individuals of type s are willing to pay above the cost they impose on the insurer, V is approximately the reduction in variance of consumption offered by the insurance:

$$V = E \left[(y - x^U - p_U - \bar{c})^2 | S(\theta_\mu) = s \right] - E \left[(y - x^I - D(\tilde{s}) - p_U - \bar{c})^2 | S(\theta_\mu) = s \right]$$

and $(1 - p(s)) E [m^I(\theta_\nu) - m^U(\theta_\nu) | S(\theta_\mu) = s]$ is a correction term to account for moral hazard. $E [m^I(\theta_\nu) - m^U(\theta_\nu) | S(\theta_\mu) = s]$ is the causal effect of insurance on medical spending to a type θ . If $p(s) < 1$, some of this additional cost that is imposed on the insurer will not be fully valued by the individual.

In this sense, one needs to observe two additional pieces of information in order to generate an internal measure of risk aversion, $\gamma(s)$: (1) the impact of insurance on medical spending for type s , $E [m^I(\theta_\nu) - m^U(\theta_\nu) | S(\theta_\mu) = s]$ and (2) the impact of insurance on the variance of consumption, $V(s)$. In this sense, one need not necessarily rely on an external measure of risk aversion, but can instead infer risk aversion from individuals revealed willingness to pay to reduce their variance in consumption.

D Proof of Proposition 4

Proposition 4 provides an expression for $\beta(s)$ that uses the demand curve to proxy for consumption. To arrive at this expression, let $D(\tilde{s}, p_U)$ denote the price that a type \tilde{s} is willing to pay for insurance when facing a price p_U of being uninsured. For a full insurance contract with $x(m) = 0$, this solves

$$E [u(y_\nu(\theta_\nu) - D(p_U, S(\theta_\mu)), m(\theta_\nu); \theta_\nu) | S(\theta_\mu) = s] = E [u(y_\nu(\theta_\nu) - m(\theta_\nu) - p_U, m(\theta_\nu); \theta_\nu) | S(\theta_\mu) = s]$$

So that

$$\frac{\partial}{\partial p_U} E [u(y_\nu(\theta_\nu) - m(\theta_\nu) - p_U, m(\theta_\nu); \theta_\nu) | S(\theta_\mu) = s] = \frac{\partial}{\partial p_U} E [u(y_\nu(\theta_\nu) - D(p_U, S(\theta_\mu)) - p_U, m(\theta_\nu); \theta_\nu) | S(\theta_\mu) = s]$$

So that the average marginal utility of income for the uninsured can be expressed as:

$$E \left[-\frac{\partial u}{\partial c}(y_\nu(\theta_\nu) - m(\theta_\nu) - p_U, m(\theta_\nu); \theta_\nu) | S(\theta_\mu) \geq s \right] = E \left[\left(\frac{\partial D}{\partial p_U} + 1 \right) \frac{\partial u}{\partial c}(y_\nu(\theta_\nu) - D(p_U, S(\theta_\mu)) - p_U, m(\theta_\nu); \theta_\nu) | S(\theta_\mu) \geq s \right]$$

The key advantage of this equation is that it relates the marginal utility of income for the uninsured (LHS) to the marginal utility of income these same people would experience if they instead chose to purchase insurance.

Now, imposing assumption 1 along with the assumption that $\frac{\partial D}{\partial p_U} = 0$ means that the

same Taylor expansion approach in Proposition 3 yields

$$\beta(s) = \gamma(\Delta y + \Delta D)$$

where

$$\Delta y = E[y_\nu(\theta_\nu) | S(\theta_\mu) \geq s] - E[y_\nu(\theta_\nu) | S(\theta_\mu) < s]$$

is the difference in income between insured and uninsured and

$$\Delta D = D(p_U(s), s) - E[D(p_U(s), S(\theta_\mu)) | S(\theta_\mu) \geq s]$$

is the difference between the price paid by the insured and the average willingness to pay for the uninsured.

E Insurance Versus Redistribution: Conditioning on $X = x$

The approach provided here can also be amended to facilitate welfare analysis after some observable information, X , has been revealed about θ . For example, perhaps one does not wish to incorporate the value of insurance to the extent to which it redistributes across those with different incomes or health conditions.

This appendix provides a brief sketch of how one can make adjustments to the baseline formula for $EA(s)$ by conditioning on the observable characteristics, $X = x$. To see how this can work, suppose prices, p_U and p_I , are charged uniformly to people with different values of X and that a fraction s of the market purchases insurance.⁴¹ Let s_x denote the fraction of the population with characteristics $X = x$ that are uninsured. Note that $s = E_X[s_x]$ is the total fraction of the market insured. Next, let $\beta(s, x)$ denote the difference in marginal utilities between the insured and uninsured given by a generalized version of equation (16):

$$\beta(s, x) = \frac{E\left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) < s, X = x\right] - E\left[\frac{\partial u_\nu}{\partial c_\nu} | S(\theta_\mu) \geq s, X = x\right]}{E\left[\frac{\partial u_\nu}{\partial c_\nu} | \theta_\mu \in \Theta_\mu, X = x\right]}$$

Now, note that the aggregate impact on p_U of expanding the size of the insurance market is still determined by the aggregate resource constraint. Therefore, the slope of the aggregate demand curve continues to determine how much prices change for a given group. This means that the ex-ante value of expanding the insurance market for those with characteristics $X = x$

⁴¹If prices, p_U and p_I , are charged differentially to those with different X characteristics, then one can simply conduct welfare analysis by conditioning on X everywhere in Proposition 1.

is given by

$$EA(s, x) = (1 - s_x) s \left(-\frac{\partial D}{\partial s} \right) \beta(s_x, x) \quad (24)$$

and aggregating across all values of X using equal weights on those with different X characteristics yields an ex-ante welfare value of $E_X [EA(s, X)]$. This approach aggregates welfare from behind a set of “veils of ignorance” – one for each value of X . In the limiting case where X incorporates all information about s , then there is no difference in marginal utilities across s conditional on X , $\beta(s_x, x) = 0$. Hence, there would be no additional ex-ante value to the insurance ($EA(s, x) = 0$). This is simply another way of saying that market surplus treats all sources of differences in demand as redistribution as opposed to having potential insurance value.

Analogous derivations show that the MVPF conditional on $X = x$ generalizes to

$$MVPF(s) = \frac{1 + (1 - s_x) \beta(s_x, x)}{1 + \frac{C(s) - D(s)}{s(-D'(s))}}$$

F Proof of Proposition 5

Note that the expected utility in the economy when a fraction s is insured, as specified in equation (10) is given by:

$$V(p_I(s), p_U(s), 0) = sE \left[-\frac{1}{\gamma} e^{-\gamma[y - p_I(s) + \frac{1}{2}w]} | S(\theta_\mu) < s \right] + (1 - s)E \left[-\frac{1}{\gamma} e^{-\gamma[y - p_U(s) - \lambda(\theta_\nu)]} | S(\theta_\mu) \geq s \right]$$

So, $S^{EA}(s)$ solves

$$-\frac{1}{\gamma} E \left[e^{-\gamma[y - \lambda(\theta_\nu) + W(s)]} \right] = sE \left[-\frac{1}{\gamma} e^{-\gamma[y - p_I(s) + \frac{1}{2}w]} | S(\theta_\mu) < s \right] + (1 - s)E \left[-\frac{1}{\gamma} e^{-\gamma[y - p_U(s) - \lambda(\theta_\nu)]} | S(\theta_\mu) \geq s \right]$$

or

$$E \left[e^{-\gamma[y - \lambda(\theta_\nu) + W(s)]} \right] = s e^{-\gamma[y - p_I(s) + \frac{1}{2}w]} + (1 - s) E \left[e^{-\gamma[y - p_U(s) - \lambda(\theta_\nu)]} | S(\theta_\mu) \geq s \right]$$

and,

$$\begin{aligned} e^{-\gamma W(s)} E \left[e^{\gamma \lambda(\theta_\nu)} \right] &= s e^{\gamma[p_I(s) - \frac{1}{2}w]} + (1 - s) E \left[e^{\gamma[p_U(s) + \lambda(\theta_\nu)]} | S(\theta_\mu) \geq s \right] \\ e^{-\gamma W(s)} \int_0^1 E \left[e^{\gamma \lambda(\theta_\nu)} | S(\theta_\mu) = \tilde{s} \right] d\tilde{s} &= s e^{\gamma p_I(s)} e^{-\gamma \frac{1}{2}w} + (1 - s) e^{\gamma p_U(s)} \int_s^1 E \left[e^{\gamma \lambda(\theta_\nu)} | S = \tilde{s} \right] d\tilde{s} \end{aligned}$$

Next, we note that $E \left[e^{\gamma \lambda(\theta_\nu)} | S(\theta_\mu) = \tilde{s} \right] = e^{-\gamma C(\tilde{s}) + \frac{\gamma^2}{2} \Sigma(\tilde{s})}$. So, $E_{\tilde{s}} \left[E \left[e^{\gamma \lambda(\theta_\nu)} | S(\theta_\mu) = \tilde{s} \right] \right] = \int_0^1 e^{-\gamma C(\tilde{s}) + \frac{\gamma^2}{2} \Sigma(\tilde{s})} d\tilde{s}$. Analogously, $\int_s^1 E \left[e^{\gamma \lambda(\theta_\nu)} | S = \tilde{s} \right] d\tilde{s} = \int_s^1 e^{-\gamma C(\tilde{s}) + \frac{\gamma^2}{2} \Sigma(\tilde{s})} d\tilde{s}$. Taking logs

then yields the result.

G Income Heterogeneity

To see this, note that with heterogeneity in $y(\theta_\nu)$ we have:

$$-\frac{1}{\gamma}E\left[e^{-\gamma[y(\theta_\nu)-\lambda(\theta_\nu)+W(s)]}\right] = sE\left[-\frac{1}{\gamma}e^{-\gamma[y(\theta_\nu)-p_I(s)+\frac{1}{2}w]}|S < s\right] + (1-s)E\left[-\frac{1}{\gamma}e^{-\gamma[y(\theta_\nu)-p_U(s)-\lambda(\theta_\nu)]}|S \geq s\right]$$

or

$$\begin{aligned} E\left[e^{-\gamma[y(\theta_\nu)-\lambda(\theta_\nu)+W(s)]}\right] &= sE\left[e^{-\gamma[y(\theta_\nu)-p_I(s)+\frac{1}{2}w]}|S < s\right] + (1-s)E\left[e^{-\gamma[y(\theta_\nu)-p_U(s)-\lambda(\theta_\nu)]}|S \geq s\right] \\ e^{-\gamma W(s)}E\left[e^{-\gamma[y(\theta_\nu)-\lambda(\theta_\nu)]}\right] &= se^{\gamma p_I(s)}e^{-\frac{1}{2}\gamma w}E\left[e^{-\gamma y(\theta_\nu)}|S < s\right] + (1-s)e^{\gamma p_U(s)}E\left[e^{-\gamma[y(\theta_\nu)-\lambda(\theta_\nu)]}|S \geq s\right] \end{aligned}$$

Now, conditioning first on $S = \bar{s}$ and then integrating over a uniform distribution of \bar{s} yields

$$e^{-\gamma W(s)}\int_0^1 E\left[e^{-\gamma[y(\theta_\nu)-\lambda(\theta_\nu)]}|S = \bar{s}\right]d\bar{s} = e^{\gamma p_I(s)}e^{-\frac{1}{2}\gamma w}\int_0^s E\left[e^{-\gamma y(\theta_\nu)}|S = \bar{s}\right]d\bar{s} + e^{\gamma p_U(s)}\int_s^1 E\left[e^{-\gamma[y(\theta_\nu)-\lambda(\theta_\nu)]}|S = \bar{s}\right]d\bar{s}$$

where

$$E\left[e^{-\gamma[y(\theta_\nu)-\lambda(\theta_\nu)]}|S = \bar{s}\right] = e^{\gamma(\bar{y}(\bar{s})+C(\bar{s})-w)+\frac{1}{2}\gamma^2\Sigma(\bar{s})}$$

So,

$$e^{-\gamma W(s)}\int_0^1 e^{\gamma(\bar{y}(\bar{s})+C(\bar{s})-w)+\frac{1}{2}\gamma^2\Sigma(\bar{s})}d\bar{s} = e^{\gamma p_I(s)}e^{-\frac{1}{2}\gamma w}\int_0^s e^{-\gamma\bar{y}(\bar{s})}d\bar{s} + e^{\gamma p_U(s)}\int_s^1 e^{\gamma(\bar{y}(\bar{s})+C(\bar{s})-w)+\frac{1}{2}\gamma^2\Sigma(\bar{s})}d\bar{s}$$

H Preference Heterogeneity

As noted in the text, there are two reasons that individuals might have high willingness to pay: on the one hand, WTP might be driven by their expected costs (or higher variance of costs). On the other hand, WTP might be driven by high risk aversion. Therefore, for two individuals with the same willingness to pay, i.e. $D(p_U, s)$, it is no longer the case that it would be reasonable to assume they have the same expected costs that equals the marginal cost curve, $C(s)$. Instead, at each level the cost curve reflects an average of individuals with different expected costs who have the same quantile of demand, $1-s$, but arrive at that quantile through different levels of risk aversion. As a result, I introduce heterogeneity in to the cost curve.

I proceed as follows. First, i enumerate all individuals who are in the insurance market. Let $\bar{\lambda}_i$ denote an uninsured individual's expected costs given what they know at the time of purchasing insurance and let V_i denote their subjective variance of their expected costs. I continue to assume moral hazard is constant (i.e. no variation in w). So, demand is given by

$$D_i = \bar{\lambda}_i + \frac{w}{2} + \frac{\gamma_i V_i}{2} \quad (25)$$

As noted above, demand could be high because of high expected costs, $\bar{\lambda}_i$, or because of high risk aversion, γ_i (or also high variance V_i). So, while previously in the subsection above we assumed that $\bar{\lambda}_i = C(s) - w$ for each person at the $1 - s^{th}$ quantile of demand, this assumption is no longer reasonable if there is heterogeneity in γ_i .

Instead, I proceed as follows. First, I draw a sample of individuals (I choose 10,000 individuals, but this is not essential). For each individual, I draw a random risk aversion coefficient as:

$$\gamma = \bar{\gamma} + \epsilon^\gamma$$

where $\bar{\gamma}$ is the mean level of risk aversion and $\epsilon^\gamma \sim N(0, \sigma_\gamma^2)$ is a normally distributed error term. I take $\bar{\gamma} = 5x10^{-4}$ and $\sigma_\gamma \in \{5x10^{-5}, 1x10^{-4}\}$. Note that in principle this can generate risk loving individuals with $\gamma < 0$. I therefore choose to censor the distribution to impose $\gamma > 0$, although this is not essential for the results.

Second, as before, fix the moral hazard component as $w = E[C(s)]$ so that average costs are doubled due to moral hazard (which corresponds to roughly a 30% increase in costs relative to their gross costs).

Third, I draw a realization of the expected cost and variance for each individual, $\bar{\lambda}_i$. I model the expected costs as a draw from the random distribution of $\{C(s), V(s)\}$ from the cost distribution and variance outcomes from the model with no heterogeneity. But, I then add an error term κ_i to the expected costs, so that

$$\bar{\lambda}_i = C(U) - w + \kappa_i$$

where $\kappa_i \sim N(0, \sigma_c^2)$ and $C(U)$ reflects the fact that we drew costs randomly from the 100 row dataset of the estimated cost curve.

I use a simple minimum distance estimator to estimate σ_c . For each individual, I construct D_i using equation (25) above. Sorting the data yields the predicted quantile s for each person, and a “demand curve” that would be implied by the structural model, $\hat{D}(s)$. I construct $\eta_D(s)$ to be the distance between the s^{th} quantile of the estimated demand and the model’s implied demand curve:

$$\eta_D^2(s) = \hat{D}(s) - D(s)$$

Next, I consider the cost curve. I take the sorted data and regress the costs of the insured population on each individuals quantile of the demand curve $1 - s$, or s for simplicity:

$$\bar{\lambda}_i + w = f(s)$$

where f is estimated using a local polynomial. The predicted \hat{f} yields the cost curve implied

by the model:

$$\hat{C}(s) = \hat{f}(s)$$

I then construct $\eta_C(s)$ to be the difference between the model's implied cost curve and the cost curve estimated in the data:

$$\eta_C^2 = \hat{C}(s) - C(s)$$

I then choose σ_c to minimize the squared loss, weighting each quantile equally:

$$L = \sum_s (\eta_C^2(s) + \eta_D^2(s))$$

This then generates the primitives of the model. I then use these to construct $W(s)$ using equation (11), being sure to correctly account for preference heterogeneity when taking expectations.