

Measuring Ex-Ante Welfare in Insurance Markets

Nathaniel Hendren*

February 2020

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 percentage difference in marginal utilities between insured and uninsured. I provide an approach to estimate this additional statistic that uses only the reduced-form willingness to pay curve, combined with 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 observed market surplus. While mandates reduce market surplus, 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);

*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 Sziedl, for helpful comments and discussions. I also thank Kate Musen and Peter Ruhm for helpful research assistance. Support from the National Science Foundation CAREER Grant is gratefully acknowledged.

Panhans (2018)). Comparing willingness to pay to the costs individuals impose on insurers provides a traditional measure of market surplus. This surplus potentially provides guidance on optimal 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 ex-ante 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 - measured as the difference between observed willingness to pay and costs in the market - 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)). This means that if willingness to pay had been measured earlier, market surplus would be larger. This is because it would include the value of insuring against the costly medical procedure. While ex-ante market surplus would be larger if it were measured earlier, the economic allocation generated by a mandate does not vary depending on when the economist measures willingness to pay. This means that the ex-ante expected utility impact of a mandate would not depend on when the economist happens to measure

¹This idea dates to Hirshleifer (1971), who shows that individuals may wish to insure against the realization of information that is revealed prior to making choices.

willingness to pay. While the average willingness to pay tends to decline with the amount of information revealed at the time of making insurance market choices, expected utility should not change. Ex-ante expected utility provides a consistent welfare framework to study optimal insurance policies that does not depend upon how much information individuals know at the time they choose to purchase insurance.

The goal of this paper is to enable researchers to evaluate the impact of insurance market policies on ex-ante welfare, defined as ex-ante expected utility.² Traditional methods to estimating ex-ante welfare 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.³ 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 asymmetric information.

This paper develops a new approach to measure the ex-ante welfare 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 defined in Einav et al. (2010). In this environment, I characterize the minimal additional sufficient statistics required to measure the ex-ante welfare impact of subsidies and mandates.

The first main result shows that one can measure ex-ante welfare using one additional sufficient statistic: the percentage 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 ex-ante 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 issue by providing a benchmark estimation method that uses only the reduced-

²Throughout the paper, I adopt the common assumption in health insurance models that there is no aggregate risk and rational expectations. This means that the ex-ante risk distribution corresponds to the realized cross-sectional distribution. As a result, ex-ante welfare is also equivalent to measuring (ex-post) utilitarian welfare.

³For example, see Handel et al. (2015) or Section IV of Einav et al. (2016).

form willingness to pay curve 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 approximate differences in marginal utilities using measures of consumption differences between insured and uninsured combined with risk aversion. Second, since consumption is seldom observed, I provide conditions under which one can exploit the information in the reduced-form willingness to pay curve for insurance instead of using consumption data. When these conditions hold, a high willingness to pay for insurance signals a greater desire for money to help cover medical expenses. This enables the information in the willingness to pay curve to substitute for the consumption difference between the insured and uninsured.

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.⁴ Without subsidies, the market would not exist. I use my approach to ask what types of insurance subsidies or mandates individuals would want from an ex-ante perspective – prior to learning anything about their risk.

I evaluate the welfare impact of both budget neutral and non-budget neutral policies. Budget neutral insurance subsidies are financed by increased prices or 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. From behind a veil of ignorance, individuals value the ability to purchase insurance at lower prices if they end up having a high demand for insurance.

What about a mandate? The sum of willingness to pay across individuals when they are observed in the market is less than the cost they impose on the insurance company. Lowering relative prices enough to yield full coverage would lower the traditional market surplus measure by \$45. However, from behind a veil of ignorance mandates increase ex-ante welfare: individuals would be willing to pay \$169 to have a full insurance mandate. The ex-ante value of 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

⁴This unraveling is due to a combination of adverse selection and uncompensated care externalities.

an ex-ante welfare perspective can lead to very different normative conclusions about the desirability of commonly debated insurance policies.

As in many settings, insurance subsidies in Massachusetts were financed by general government revenue, not by imposing penalties on the uninsured who were eligible for the subsidies. 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 provide a method to assess the relative efficiency of the subsidies as a method of redistribution. 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 credit.

The results suggest the ex-ante and traditional MVPF can differ significantly. For example, starting with low subsidies such that 30% of the market is insured, the MVPF of increasing subsidies is 1.2 if one uses observed willingness to pay. Individuals are willing to pay roughly 1.2 times the marginal cost they impose on the insurer to lower insurance prices. This 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. Yet, from behind the veil of ignorance, individuals would be willing to pay 1.8 times the cost they impose on the insurer to lower insurance prices. Ex-ante, individuals prefer that the government spend money lowering insurance prices for those at 150% of FPL instead of providing them with a tax credit. In summary, ex-ante measures of welfare can lead to different conclusions than those based on observed willingness to pay and traditional measures of market surplus.

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 individual-level data.

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 MA setting. The model builds upon the

approach in Handel et al. (2015) but augments it with a moral hazard structure developed in Einav et al. (2013). 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 to 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 (e.g. the example above where subsidies are provided to those at 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 if one can measure the difference in average consumption between the insured and uninsured.

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, the risk premium individuals are willing to pay to insure against ex-ante risk may differ from the risk premium they are willing to pay to insure against risk that remains at the time they choose to buy insurance. This can generate bias in the benchmark implementation. This means in practice researchers will want to study the robustness of the results to a range of risk aversion parameters. But more generally, the first main result provided in Proposition 1 continues to hold even in the presence of preference heterogeneity. This provides a potential roadmap for future work to develop methods to measure the markup individuals are willing to pay for insurance against risk that is realized prior to making insurance coverage choices.

In the broader context of existing literature, 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 (2015); Dague (2014)). However, this approach only captures the value of insurance against the risk that remains after adjusting their labor supply. Similarly, other papers infer willingness to pay for social insurance from changes in consumption around a shock (e.g. Gruber (1997); Meyer and Mok (2019)). 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 useful to devise strategies to recover ex-ante measures of welfare in these settings.

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. 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 percentage 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 curve combined with a measure of risk aversion. Section 6 implements this approach to study optimal health insurance subsidies for low-income adults in Massachusetts using the estimates from Finkelstein et al. (2019). Section 7 develops a structural model to compare the validity of the proposed approach to the model’s measure of ex-ante welfare and also uses the model to study the impact of violations of the implementation assumptions outlined in Section 5. Section 8 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. I adopt a rational expectations framework with no aggregate risk. This means that the realized cross-sectional distribution in the economy corresponds to the ex-ante distribution of risk. 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

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

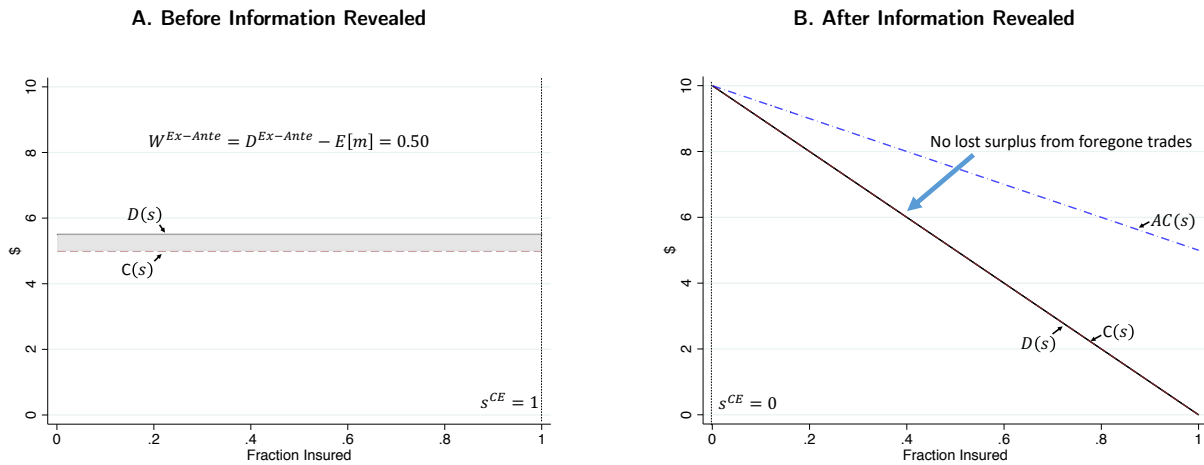
where $E[u(30 - m)] = \frac{1}{10} \int_0^{10} u(30 - m) dm$ is the expected utility if uninsured.

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 of this policy would be $E[m] = 5$. Full insurance generates a market surplus of \$0.50.

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.

Figure 1: Example Willingness to Pay and Cost Curves



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 $S \sim Uniform[0, 1]$ be a uniform random variable and define the average cost of those with willingness to pay above $D(s)$ by $AC(s) = E[C(S) | S \leq s]$. 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. 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. (2019).

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. I use these curves to calculate the sufficient statistics necessary to measure the utility impact insurance market policies.

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.

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 what is measured when making choices to purchase insurance to purchase

insurance in the market, $D(s)$ 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.⁵ 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 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 percentage 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, then one can approximate this difference in marginal utilities using the difference in consumption between insured and uninsured, multiplied by a coefficient of risk aversion.

Consumption is rarely observed in practice. However, notice that in the model those with high willingness to pay for insurance are those with lower consumption. Therefore, in a final step, I provide conditions under which information in the demand curve can proxy for the consumption difference. This leads to the formula:

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


where $\gamma(s) = \frac{-u''}{u'}$ is the coefficient of absolute risk aversion for those indifferent to purchasing insurance and $D(s) - E[D(S)|S \geq s]$ is the difference in willingness to pay between the average uninsured individual and the marginal insured type. This latter term captures the difference in average consumption between the insured and uninsured. When 50% of the market own insurance, this difference is \$2.50: the insured pay \$5 for insurance and the uninsured pay \$0 but on average experience a \$2.5 loss, generating a difference in consumption of \$2.5 on average.

The assumptions needed to generate equation (3) are stated formally in Section 5. Most

⁵The term $D'(s)$ captures how changes in the size of the market translate into changes in the relative price of insurance, $p_I - p_U$. This is weighted by $s(1-s)$ to account for the fact that the size of price increase for the insured is inversely proportional to $1-s$ (a high $1-s$ means insurance prices decline rapidly when the market expands because more people pay p_U). Similarly, the price decrease for the uninsured is inversely proportional to s . As shown in Appendix A, these two forces imply that the size of the transfer is $s(1-s)(-D'(s))$.

notably, they require no income or liquidity differences between insured and uninsured and they assume no heterogeneity in risk aversion. While not without loss of generality, the formula provides a benchmark implementation to measure ex-ante expected utility with only the addition of a risk aversion coefficient. Section 7.4 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):

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

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 statistic using the demand and cost curves combined with a measure of risk aversion, as in equation (3) above.

3 General Model

This section develops a general model environment that can be applied to realistic empirical applications such as the health insurance exchange for low-income adults in Massachusetts studied in Finkelstein et al. (2019).

3.1 Setup

Individuals face evolving risk over periods $t = 1, \dots, T$ that is captured by the realization of a random variable or “shock”, Θ'_t , whose realizations are in \mathbb{R}^N . The shock can be multi-dimensional ($N > 1$) so that it captures all aspects of an individual’s life (level of utility, marginal utility of medical spending, information about future values of $\Theta'_{t'}$ for $t' > t$, etc.). Shocks may be correlated over time. I define $\Theta_t = \{\Theta'_1, \dots, \Theta'_t\}$ to be the history of shocks up to period t . I let θ_t denote particular realizations of Θ_t . For notational brevity, I abstract from the distinction between the random variable and its realizations and generally refer to the variable by its realizations, θ_t .⁶ As in Section 2, I assume rational expectations and no aggregate risk so that the distribution of θ_t is equal to the cross-sectional population distribution of realizations of θ_t .⁷ I let period $t = 0$ denote the ex-ante period before anyone

⁶More formally, if $G_{\theta_t}(\theta_s)$ describes the distribution of Θ_s given a particular realization of $\Theta_t = \theta_t$, I will use the notation $E[f(\theta_s)|\theta_t]$ to denote $\int f(\theta_s) dG_{\theta_t}(\theta_s) d\theta_s$.

⁷Note that for any $t > t'$, only the more recent realization governs beliefs, $E[f(\theta_s)|\theta_t, \theta_{t'}] = E[f(\theta_s)|\theta_t]$. For simplicity, I refer to realizations of θ_t as a “realization of a shock”, even though the first $t - 1$ components $\theta'_1, \dots, \theta'_{t-1}$ will be known to the individual in time t .

learns any shocks, so that individuals have identical beliefs that correspond to the population distributions of θ_t .

Realized utility in period t , $u_t(c, m; \theta_t)$, is a function of the history of shocks up to period t , θ_t , and choices of non-medical consumption, c , and medical consumption, m .⁸ For each θ_t and t , I assume u_t is twice-continuously differentiable and strictly concave in c and m and strictly increasing⁹ in c but not necessarily increasing in m (so that fully insured individuals may choose finite m). In each period t , individuals learn their realization of θ_t and then choose c and m subject to a constraint $g_t(c, m; \theta_t) \leq 0$. I assume $g_t(c, m; \theta_t)$ is twice continuously differentiable and weakly convex in both c and m . I assume g_t is strictly increasing in c and weakly increasing in m (which allows for full insurance contracts in period t).¹⁰ The budget constraint can depend on the full history of shocks, θ_t . For the model in the main text, I assume that the budget constraint in period t is independent of consumption and medical spending choices in other periods (i.e. no savings). In Appendix, A.2, I discuss a generalization of the model that allows for savings and show that Proposition 1 continues to hold.¹¹

To this standard environment, I add the ability in period $t = \mu$ for individuals with $\theta_\mu \in M$ to decide whether or not to purchase insurance that affects their budget constraint in a single period $t = \nu > \mu$. In the example in Section 6, individuals with incomes near 150% of the Federal Poverty Level (FPL) will have the opportunity in the fall of 2010 to purchase insurance in the Massachusetts subsidized health insurance exchange for the 2011 calendar year. In the notation, $v = 2011$, μ is the open enrollment period in the fall of 2010, and M is the set of individuals whose incomes are near 150% FPL. Individuals who chose to purchase insurance for period ν have the budget constraint

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

where $x(m_\nu)$ is the out of pocket costs required for gross medical spending of m_ν , $y_\nu(\theta_\nu)$ is the individual's income, and p_I is a price paid by the insured. I assume $x(m_\nu)$ is continuously differentiable, weakly increasing, and weakly convex in m_ν . Individuals who chose to be

⁸Note the utility function allows for discounting, e.g. $u_t(c, m; \theta_t) = \beta^t \tilde{u}(c, m; \theta_t)$.

⁹To ensure existence of willingness to pay below, I assume $\lim_{c \rightarrow \infty} u_t(c, m; \theta_t) = \infty$ for each m .

¹⁰I assume that the maximization program is bounded so that exists values \bar{c}, \bar{m} such that (a) for any \tilde{m} , $g_t(c, \tilde{m}; \theta_t) > 0$ for all $c \geq \bar{c}$ and (b) for any \tilde{c} $u_t(\tilde{c}, m; \theta_t)$ is decreasing in m for all $m > \bar{m}$. This ensures choices c and m lie below \bar{c} and \bar{m} for all t and θ_t .

¹¹While Proposition 1 continues to hold in the presence of more general constraints that allow for savings, the baseline implementation that uses the demand curve to proxy for differences in consumption between the insured and uninsured in equation (3) does not necessarily hold. Instead, one can recover ex-ante welfare using information on consumption of the insured and uninsured (as discussed in Proposition 3).

uninsured have a budget constraint

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

where p_U is the price of being uninsured.

After observing θ_ν , individuals choose (c, m) to maximize their utility, leading to indirect utility functions for those who chose to be insured and uninsured:

$$v^I(p_I, \theta_\nu) = \max_{c, m} u(c, m; \theta_\nu) \quad s.t. \quad (4)$$

$$v^U(p_U, \theta_\nu) = \max_{c, m} u(c, m; \theta_\nu) \quad s.t. \quad (5)$$

and optimal allocations for the insured and uninsured given by $(c_\nu^I(p_I, \theta_\nu), m_\nu^I(p_I, \theta_\nu))$ and $(c_\nu^U(p_U, \theta_\nu), m_\nu^U(p_U, \theta_\nu))$.¹²

For any individual who is eligible to purchase insurance, $\theta_\mu \in M$, and any price of being uninsured, p_U , I define their relative willingness to pay for insurance, $d(p_U, \theta_\mu)$, as the solution to

$$E[v^U(p_U, \theta_\nu) | \theta_\mu] = E[v^I(d(p_U, \theta_\mu) + p_U, \theta_\nu) | \theta_\mu] \quad (6)$$

which equates expected utility of the uninsured to the insured who pay a relative price $d(p_U, \theta_\mu)$.¹³ Note that the model assumptions imply that d is continuously differentiable in p_U .¹⁴

While much of the results will allow for a general utility function, it will be useful to also discuss the case where there are no income effects in period $t = \nu$ on choices of m and the level of willingness to pay, d .

Definition. The utility function, $u_\nu(c, m; \theta_\nu)$ is said to have **no income effects** if there exists positive constants a and b and a function $\kappa(m, \theta_\nu)$ such that

$$u_\nu(c, m; \theta_\nu) = -ae^{-b[c + \kappa(m, \theta_\nu)]}$$

for all θ_ν .

A utility function that satisfies no income effects implies that $m_\nu^I(p_I, \theta_\nu)$, $m_\nu^U(p_U, \theta_\nu)$,

¹²These choices are guaranteed to exist, to be unique, and to be continuously differentiable in p_U and p_I because x is continuously differentiable, increasing, and convex in m , and the utility function is strictly concave and twice differentiable in m .

¹³Note that $d(p_U, \theta_\mu)$ is well-defined because utility is strictly increasing in c .

¹⁴The envelope theorem (Milgrom and Segal (2002)) implies $\frac{\partial v^U}{\partial p_U} = -\frac{\partial u_\nu}{\partial c}(c_\nu^U(p_U, \theta_\nu), m_\nu^U(p_U, \theta_\nu))$ so that this follows from the fact that choices c_ν^U and m_ν^U are differentiable in p_U and utility is twice continuously differentiable in (c, m) . Similarly, this holds for v^I as well.

and $d(p_U, \theta_\mu)$ do not depend on the amount of income individuals have in period $t = \nu$, and therefore these functions do not depend on p_I or p_U .¹⁵ Assuming no income effects is restrictive, but it enables the environment to nest results in existing literature and provide more precise characterizations of some results below.

3.2 Aggregating to Market Willingness to Pay and Cost Curves

To aggregate the model to market-level willingness to pay and cost curves, I impose a smoothness condition that requires the population distribution of willingness to pay, $d(p_U, \theta_\mu)$, to be continuously distributed with positive mass throughout its support. Formally, let $\zeta(\Delta, p_U) = Pr \{d(p_U, \theta_\mu) > \Delta\}$ denote the fraction of the market that purchases insurance when prices are p_U and $\Delta = p_I - p_U$. The assumption that $d(p_U, \theta_\mu)$ is continuously distributed means that $\zeta(\Delta, p_U)$ is continuously differentiable in (Δ, p_U) . The assumption that d has positive mass throughout its support means that $\frac{\partial \zeta}{\partial \Delta} < 0$ for each (Δ, p_U) .¹⁶ The smoothness condition is an implicit assumption on the utility function and smoothness of the shock distribution (i.e. distribution of θ_μ). Section 7 provides a class of utility functions and type distributions where this smoothness condition is satisfied. In Appendix I, I use the structural model in Section 7 to consider the case when this assumption is violated because of discontinuous demand curves.¹⁷

Market Willingness to Pay Curve Let $D(p_U, s)$ denote the marginal price of insurance required for a fraction $s \in [0, 1]$ of the market to purchase insurance. This is given by the solution to $\zeta(D(p_U, s), p_U) = s$.¹⁸ The assumption that $\zeta(\Delta, p_U)$ is continuously differentiable in Δ means that D is continuously differentiable by the implicit function theorem, and the fact that $\zeta(\Delta, p_U)$ is strictly decreasing in Δ implies D is strictly decreasing in Δ .

¹⁵Appendix A provides proofs of these statements. The fact that m does not depend on transfers follows from quasi-linearity between c and $\kappa(m, \theta_\nu)$; the fact that d does not depend on transfers follows from the fact that $\frac{\partial u}{\partial c} = \gamma u$ (where $\gamma = ab$ is the coefficient of absolute risk aversion). Because a and b are constant, this specification also rules out income effects for ex-ante willingness to pay for insurance ($W(s)$ defined below).

¹⁶Moreover, $Pr \{d(p_U, \theta_\mu) > \Delta\} = Pr \{d(p_U, \theta_\mu) \geq \Delta\}$ so that the precise specification of inequalities does not affect the market size, $\zeta(\Delta, p_U)$.

¹⁷I simulate an econometrician who estimates a continuously differentiable approximation to a discontinuous demand curve. The simulations suggest that applying my proposed method to continuously differentiable approximations to the demand curve provide relatively accurate approximations to the true measures of ex-ante welfare ($W(s)$ defined below), even when aggregating across the points of discontinuity. But, as one would expect, estimates of the marginal willingness to pay (estimates of $W'(s)$ below) can be biased near the point of discontinuity in the demand curve.

¹⁸Note that the solution exists and is unique because ζ is strictly decreasing and continuous in Δ and there exists prices such that no one purchases insurance and everyone purchases insurance (e.g. $\zeta(x, p_U) = 1$ as $x \rightarrow -\infty$ and $\zeta(x, p_U) = 0$ as $x \rightarrow \infty$).

For any type θ_μ , I let $S(p_U, \theta_\mu) = \zeta(d(p_U, \theta_\mu), p_U)$ denote the fraction of the market that is insured when prices are such that type θ_μ is indifferent to purchasing insurance. Because $d(p_U, \theta_\mu)$ is continuously distributed, its quantiles are uniquely defined and uniformly distributed over $[0, 1]$ with no point masses. Because the quantiles of willingness to pay correspond to $1 - S(p_U, \theta_\mu)$, this means that $S(p_U, \theta_\mu) \sim \text{Uniform}[0, 1]$ for each p_U . If the utility function has no income effects, I write $D(s)$ in place of $D(p_U, s)$ and $S(\theta_\mu)$ in place of $S(p_U, \theta_\mu)$.

Cost of Insured Population For any s , let $AC(p_U, s)$ denote the average cost when a fraction s of the market owns insurance.

$$AC(p_U, s) = E [m_\nu^I (D(p_U, s) + p_U, \theta_\nu) - x(m_\nu^I (D(p_U, s) + p_U, \theta_\nu)) | d(p_U, \theta_\mu) \geq D(p_U, s)] \quad (7)$$

where the expectation is taken with respect to all $\theta_\mu \in M$ that choose to purchase insurance when prices are such that a fraction s of the market is insured, $d(p_U, \theta_\mu) \geq D(p_U, s)$.¹⁹ Insurer costs are given by the average difference between an individual's medical costs, m_ν^I , and the portion of these costs they pay out-of-pocket, $x(m_\nu^I)$. Since $d(p_U, \theta_\mu)$ is continuously distributed and continuously differentiable in p_U , $AC(p_U, s)$ is differentiable in (p_U, s) .²⁰

4 Measuring Ex-Ante Welfare

This section derives methods to measure the ex-ante welfare impact (from the perspective of $t = 0$) of policies that change the prices p_I and p_U in the market for insurance that affects the budget constraint in period $t = \nu$. I define ex-ante welfare formally as follows.

¹⁹This is equivalent to the set of all θ_μ such that $S(p_U, \theta_\mu) \leq s$.

²⁰This follows because the fact that d is continuously distributed in the population means the quantiles $1 - s$ of willingness to pay can be represented by a uniform distribution. Total costs when s of the market is insured is given by

$$sAC(p_U, s) = \int_0^s E [m_\nu^I (D(p_U, s) + p_U, \theta_\nu) - x(m_\nu^I (D(p_U, s) + p_U, \theta_\nu)) | S(p_U, \theta_\mu) = s] ds$$

So, the derivative of total costs is given by

$$\frac{d}{ds} [sAC(p_U, s)] = E [m_\nu^I - x(m_\nu^I) | S(p_U, \theta_\mu) = s] + sE \left[\frac{dm_\nu^I}{ds} (1 - x'(m_\nu^I)) | S(p_U, \theta_\mu) \leq s \right]$$

where m_ν^I is evaluated at $m_\nu^I (D(p_U, s) + p_U, \theta_\nu)$. Note that $\frac{dm_\nu^I}{ds} = \frac{\partial D}{\partial s} \frac{\partial m_\nu^I}{\partial p_I}$ exists because m_ν^I is continuously differentiable in p_I . Finally, if there are no income effects, then $\frac{dm_\nu^I}{ds} = 0$ so that the derivative of total cost is the cost of the marginal enrollees, as in Einav et al. (2010).

For any sequence of shocks realized over an individual's lifetime, $\{\theta_t\}_{t=1}^T$, let $c_t^*(p_I, p_U, \theta_t)$ and $m_t^*(p_I, p_U, \theta_t)$ denote individuals' choices of consumption and medical spending in each period, t . For example, in period $t = \nu$ the optimal allocation of medical spending is²¹

$$m_\nu^*(p_I, p_U, \theta_\nu) = 1 \{d(p_U, \theta_\mu) > p_I - p_U\} m_\nu^I(p_I, \theta_\nu) + 1 \{d(p_U, \theta_\mu) \leq p_I - p_U\} m_\nu^U(p_U, \theta_\nu)$$

In other periods $t \neq \nu$, these choices will not depend on p_I and p_U because there is no scope for savings behavior in the baseline model. Here, I maintain p_I and p_U in the notation for choices in other periods and discuss the extension to allow for savings in Appendix A.2.

Let $v_t(p_I, p_U, \theta_t)$ denote the utility level in period t that is attained by an individual with realization θ_t . This solves

$$v_t(p_I, p_U; \theta_t) = u_t(c_t^*(p_I, p_U, \theta_t), m_t^*(p_I, p_U, \theta_t); \theta_t)$$

The ex-ante expected utility of having prices p_I and p_U is given by:

$$V(p_I, p_U) = E[\sum_{t=1}^T v_t(p_I, p_U; \theta_t)] \quad (8)$$

where the expectation is taken at $t = 0$. The absence of aggregate risk and the assumption of rational expectations means that the expectation is taken with respect to the population distribution of sequences of realizations of $\{\theta_t\}_{t=1}^T$.²² This means that V is not only ex-ante welfare but also corresponds to (ex-post) utilitarian welfare.

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

I consider two classes of policies that change the fraction of the market that is insured in period $t = \nu$: 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 of being uninsured, p_U , charged to those in the market $\theta_\mu \in M$. In a world where prices cover costs,

²¹Note this is well-defined because θ_μ is contained in θ_ν so that $d(p_U, \theta_\mu)$ is implied by p_U and θ_ν .

²²Ex-ante utility measures preferences for prices p_I and p_U at time $t = 0$ before individuals learn anything about themselves. But, it is equivalent measuring ex-ante utility at any period t' such that individuals do not yet have any particular knowledge about their particular values of $v_\nu(p_I, p_U; \theta_\nu)$. To see this, note that if for some $t < \nu$ we have $E[v_\nu(p_I, p_U; \theta_\nu) | \theta_t] = E[v_\nu(p_I, p_U; \theta_\nu)]$, for all realizations of θ_t , then $E[v_\nu(p_I, p_U; \theta_\nu)]$ will correspond to $E[v_\nu(p_I, p_U; \theta_\nu) | \theta_t]$ where the expectation is evaluated at time t . In the MA example, this means that measuring ex-ante welfare corresponds to measuring the subsidies or mandates that individuals would desire to have in the MA health insurance exchange if one asked them prior to learning anything about their particular utility-relevant risks.

$p_I(s)$ and $p_U(s)$ satisfy two equations:

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

$$sAC(p_U(s), s) = [sp_I(s) + (1 - s)p_U(s)]$$

where the first equation requires the fraction insured, s , be consistent with the fraction that wish to purchase at prices $p_I(s) - p_U(s)$ and the second equation requires that the total cost of insurance equals the sum of premiums collected. Recall that $AC(p_U, s)$ is differentiable in p_U and s . Therefore, I define the marginal cost of expanding the insurance market through budget neutral price changes as the derivative of total costs,

$$C(s) = \frac{d}{ds} [sAC(p_U(s), s)]$$

This derivative includes the impact of changes to the composition of who is insured and any income effects from the price changes on medical spending, m_v^I .

I define $W(s)$ to be the ex-ante equivalent-variation measure of willingness to pay to have a fraction s of the market insured. Formally, this is the amount of income that makes individuals indifferent between being uninsured with additional income $W(s)$ and having prices of insurance given by $p_I(s)$ and $p_U(s)$. This solves:

$$V(p_I(s), p_U(s)) = V(\infty, -W(s)) \tag{9}$$

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 giving a transfer of size $W(s)$ in the world where no one is insured.²³ Equation (9) implies that the size of the market that maximizes $W(s)$ is the size of the market maximizes ex-ante welfare, $V(p_I(s), p_U(s))$.

Sufficient Statistic Representation of $W(s)$ Traditional approaches to measuring $W(s)$ would follow a structural approach that specifies a utility function, u , budget constraints, and beliefs. It would then estimate the model using individual-level data on c and m , combined with sufficient identifying variation to estimate all model components. In particular, this approach would attempt to separately estimate both utility and beliefs, which is often quite difficult and usually rests on the assumption that the econometrician perfectly observes individuals' information sets.

In contrast, I exploit the envelope theorem to characterize the derivative of $W(s)$ at each

²³I write $p_I = \infty$ but formally the RHS of equation (9) can be written as $V(\bar{c}, -W(s))$ where \bar{c} is defined in Footnote 10.

s , and use this to derive the minimal sufficient statistics required to measure the welfare impact of changes in the size of the insurance market.²⁴ Differentiating equation (9) with respect to s yields

$$\begin{aligned} -W'(s) * \frac{\partial V}{\partial p_U}(\infty, -W(s)) &= \frac{d}{ds} E \left[\sum_{t \geq 1} v_t(p_I(s), p_U(s); \theta_t) \right] \\ &= E \left[1 \{S(p_U(s), \theta_\mu) < s\} \frac{\partial u_\nu}{\partial c}(-p'_I(s)) + 1 \{S(p_U(s), \theta_\mu) \geq s\} \frac{\partial u_\nu}{\partial c}(-p'_U(s)) \mid \theta_\mu \in M \right] \end{aligned} \quad (10)$$

where $1 \{S(p_U, \theta_\mu) < s\}$ is an indicator for the event that an individual of type θ_μ purchases insurance when a fraction s of the market is insured and the uninsured pay p_U .²⁵ The second line invokes the envelope theorem: because prices p_U and p_I only affect the budget constraint in period $t = \nu$, the impact of expanding s only affects ex-ante utility through its mechanical effect on consumption in period ν as if individuals do not change their choices.²⁶ The key insight in equation 10 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, $\frac{\partial u_\nu}{\partial c}$. 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 percentage difference in marginal utilities of consumption between the insured and uninsured.

Proposition 1. *For budget neutral policies, the marginal welfare impact of expanding the size of the insurance market from s to $s + ds$ is given by*

$$W'(s) = D(p_U(s), s) - C(s) + EA(s) + \delta_p(s) \quad (11)$$

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

$$EA(s) = s(1-s) \left(-\frac{\partial D}{\partial s} \right) \beta(s) \quad (12)$$

²⁴Because of the convexity of the constraints and differentiability of the utility function, the environment satisfies the assumptions outlined in Milgrom and Segal (2002) so that the envelope theorem holds. Appendix A.2 shows this explicitly for the more general case that allows for savings.

²⁵The derivatives $\frac{\partial u_\nu}{\partial c}$ exist by assumption and the differentiability of the prices $p_I(s)$ and $p_U(s)$ follows from the differentiability of demand and total cost as a function of s .

²⁶Appendix A.2 shows that equation (10) continues to hold even in a more general model that allows for endogenous savings decisions. Because of the envelope theorem, endogenous savings responses to changes in prices do not have first order impacts on V .

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

$$\beta(s) = \frac{E\left[\frac{\partial u_v}{\partial c} | S(p_U(s), \theta_\mu) < s\right] - E\left[\frac{\partial u_v}{\partial c} | S(p_U(s), \theta_\mu) \geq s\right]}{E\left[\frac{\partial u_v}{\partial c} | \theta_\mu \in M\right]} \quad (13)$$

and $\delta_p(s)$ is an adjustment for income effects. $\delta_p(s)$ is continuously differentiable in s and $\delta_p(0) = 0$. If the utility function satisfies no income effects, then $\delta_p(s) = 0$ for all s .

Proof. See Appendix A. □

Proposition 1 shows that the marginal welfare impact of expanding the size of the insurance market is given by the sum of four terms. The first two terms, $D(p_U(s), s) - C(s)$, in equation (11) correspond to 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. As in Section 2, expanding the insurance market induces a transfer from uninsured to insured of size $(1-s)\left(-s\frac{\partial D}{\partial s}\right)$. The transfer is valued according to the difference in marginal utilities between the insured and uninsured, $\beta(s)$. The final term, $\delta_p(s)$, is an adjustment for the presence of income effects that leads to differences between equivalent variation and consumer surplus measures of welfare. This adjustment is equal to zero whenever the utility function satisfies no income effects. More generally, $\delta_p(0) = 0$ and $\delta_p(s)$ is continuously differentiable in s .²⁷

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, in the presence of advantageous selection whereby the insured are healthier than the uninsured, then it could be that the uninsured have a higher marginal utility of income than the insured, $\beta(s) < 0$. Or, if liquidity effects were a driver of insurance purchase so that the uninsured choose to forego both medical spending and the purchase of insurance, it could be that the uninsured have a higher marginal utility of income. In these cases, expanding the size of the insurance market

²⁷In Section 7.4.2 I consider a utility function that does not satisfy the no income effects assumption and the approximations shown in Figure 8 reveal that ignoring $\delta_p(s)$ provides a good measure of $W'(s)$. Moreover, the utility function in Section 2 did not satisfy the no income effects assumption, but nonetheless recovered ex-ante utility quite well as shown in the numerical example in Figure 2.

will transfer resources from the liquidity constrained to those who are less constrained, which would suggest that $EA(s) < 0$.

4.2 Non-Budget Neutral Policies: The MVPF

In many cases including the example in Section 6, insurance subsidies are redistributive: they are financed by those not in the insurance market (i.e. $\theta_\mu \notin M$), as opposed to being financed by higher prices to the uninsured in the market, p_U . This section asks whether health insurance subsidies are an efficient method of redistribution.

To do so, I construct the marginal value of public funds (MVPF) for lower insurance prices. The MVPF equals the ratio of individuals' willingness to pay for higher subsidies normalized by the net cost to the government of increasing subsidies,

$$MVPF = \frac{\text{Marginal WTP}}{\text{Marginal Govt Cost}}$$

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. Hendren and Sprung-Keyser (2019) provide a library of MVPF estimates that allow one to compare across non-budget neutral policies. This means that one can assess whether the health insurance subsidies are an efficient method of redistribution by comparing the MVPF of lower health insurance prices to the MVPF of other policies that spend resources on low-income populations, such as the Earned Income Tax Credit (EITC).

I construct the MVPF for lower insurance prices in an environment where the uninsured pay no prices, $p_U = 0$. This means that the size of the insurance market, s , is given by the solution to $p_I(s) = D(0, s)$. The cost function, $C(s)$, is then given by the derivative of total costs as the size of the market expands, $C(s) = \frac{d}{ds} [sAC(0, s)]$. And, the net cost to the government is given by the difference between total costs and premiums collected, $sAC(0, s) - sp_I(s)$. Differentiating, this yields the marginal cost to the government of expanding the market:

$$\begin{aligned} \text{Marginal Govt Cost} &= \frac{d}{ds} [sAC(s) - sp_I(s)] \\ &= -s \frac{\partial D}{\partial s}(0, s) + C(s) - D(0, s) \end{aligned}$$

The marginal cost to the government has two components. First, there is the mechanical cost of the lower prices $-sp'_I(s) = -s \frac{\partial D}{\partial s}(0, s)$. In addition, there is a cost (or benefit) from those induced to purchase insurance. Insuring these individuals increases costs by $C(s)$,

from which we subtract the prices they pay, $p_I(s) = D(0, s)$.

To measure WTP, one would traditionally use the observed average willingness to pay for those in the market. A fraction s of the market receives a price change of $-p'_I(s)$ so that the marginal WTP is $-sp'_I(s)$. This would imply an MVPF of $\frac{-s \frac{\partial D}{\partial s}(0, s)}{-s \frac{\partial D}{\partial s}(0, s) + C(s) - D(0, s)} = \frac{1}{1 + \frac{C(s) - D(0, s)}{s \left(-\frac{\partial D}{\partial s}(0, s)\right)}}$. The MVPF exceeds 1 to the extent to which the marginal enrollees are willing to pay more than their marginal cost.²⁸

But, this traditional approach ignores the ex-ante value individuals obtain from having lower insurance prices. To measure ex-ante marginal WTP, let $\tilde{W}(\tilde{s}, s)$ denote the ex-ante willingness to pay for having a fraction \tilde{s} of the market insured as opposed to having a fraction s of the market insured. This is given by the solution to $V(p_I(\tilde{s}), 0) = V(p_I(s) - \tilde{W}(\tilde{s}, s), -\tilde{W}(\tilde{s}, s))$ so that individuals are indifferent to having \tilde{s} insured and having s insured but receiving a transfer of $\tilde{W}(\tilde{s}, s)$ regardless of whether they are insured.²⁹ The ex-ante marginal WTP is then given by $\frac{\partial \tilde{W}}{\partial \tilde{s}}|_{\tilde{s}=s}$. Proposition 2 shows that one can measure ex-ante WTP by adding $(1 - s)\beta(s)$ into the numerator.

Proposition 2. *If $\frac{C(s) - D(0, s)}{s \left(-\frac{\partial D}{\partial s}(0, s)\right)} > -1$, the MVPF of non-budget neutral reduction in p_I combined with $p_U = 0$ when a fraction s of the market is insured is given by*

$$MVPF(s) \equiv \frac{\frac{\partial \tilde{W}}{\partial \tilde{s}}|_{\tilde{s}=s}}{\frac{d}{ds} [sp_I(s) - sAC(s)]} = \frac{1 + (1 - s)\beta(s)}{1 + \frac{C(s) - D(0, s)}{s \left(-\frac{\partial D}{\partial s}(0, s)\right)}} \quad (14)$$

where $\beta(s)$ is the percentage difference in marginal utilities of income for the insured relative to the uninsured given by equation (13). If $\frac{C(s) - D(0, s)}{s \left(-\frac{\partial D}{\partial s}(0, s)\right)} \leq -1$, the MVPF is infinite, as lower insurance prices generate a Pareto improvement.

Proof. See Appendix B □

Proposition 2 shows that the ex-ante marginal WTP differs from the observed willingness to pay by a factor of $1 + \beta(s)(1 - s)$. The $\beta(s)(1 - s)$ captures the additional markup that individuals are willing to pay for the ability to purchase insurance at lower prices.

²⁸When $\frac{C(s) - D(0, s)}{s \left(-\frac{\partial D}{\partial s}(0, s)\right)} < -1$, expanding the size of the market actually leads to lower prices for the insured, which makes everyone better off. This case when benefits are positive and costs are negative is defined as an infinite MVPF in Hendren and Sprung-Keyser (2019). It implies lower insurance prices would generate a Pareto improvement.

²⁹Note that $W(s)$ is well-defined because utility is strictly increasing in consumption.

5 Implementation

The key additional parameter required to construct measures of ex-ante welfare is the percentage difference in marginal utilities between insured and uninsured, $\beta(s)$. This section provides a method for estimating $\beta(s)$ using the market demand curve, $D(p_U, s)$, combined with a measure of risk aversion. This provides a benchmark method for measuring ex-ante welfare without needing to specify a utility function or the information sets of individuals in the economy. The implementation 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 those from a fully-specified structural model in Section 7.

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 consumption depends only on c , not m and θ_ν .

Assumption 1. *The marginal utility of consumption, u_c , depends only on the level of an individual's consumption, c , so that there exists a function f such that*

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

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 7.

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

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

$$\beta(s) \approx \gamma \Delta c(s) \tag{15}$$

where “ \approx ” represents equality to a first order Taylor approximation to $f(c)$,

$$\Delta c(s) = E[c_\nu(p_I(s), p_U(s), \theta_\nu) | S(p_U(s), \theta_\mu) \geq s] - E[c_\nu(p_I(s), p_U(s), \theta_\nu) | S(p_U(s), \theta_\mu) < s]$$

is the difference in consumption between the uninsured and insured, and $\gamma = \frac{-f'(E[c])}{f(E[c])}$ is the coefficient of absolute risk aversion ($-\frac{\partial^2 u_\nu}{\partial c^2}$) evaluated at the population average level of consumption, $E[c]$.

Proof. See Appendix C. □

Using consumption differences combined with a measure of risk aversion is analogous to the methods used to measure the welfare impact of unemployment insurance (Baily (1978); Chetty (2006)). But, in practice consumption is rarely observed. To provide an implementation method without consumption data, I make the additional assumption that incomes do not vary systematically between the uninsured and insured.

Assumption 2. *No differences in average liquidity/income between the insured and uninsured,*

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

Assumption 2 is not without loss of generality. I return to a discussion of how one can use consumption data to relax Assumption 2 in Section 7.4.1 below.³⁰ The key advantage of Assumption 2 is that it allows the difference in market demand between the insured and uninsured to proxy for their differences in consumption, as in the stylized model in Section 2.

Proposition 4. *Suppose Assumptions 1 and 2 hold, suppose the utility function has no income effects. Then, the percentage difference in marginal utilities is given by*

$$\beta(s) \approx \gamma(\Delta D(s) + \Delta x(s)) \tag{16}$$

where “ \approx ” represents equality to a first order Taylor approximation to $f(c)$, $\Delta D(s)$ and $\Delta x(s)$ are given by:

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

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

and $\gamma = \frac{-f'(E[c])}{f(E[c])}$ is the coefficient of absolute risk aversion. Moreover, for full insurance contracts (i.e. $x(m) = 0$), the Taylor approximation to $\beta(s)$ is given by

$$\beta(s) \approx \gamma \Delta D(s).$$

³⁰Assumption 2 will be a natural assumption 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.

Proof. See Appendix C. □

As in Section 2, the difference between the marginal willingness to pay and the average willingness to pay for the uninsured types can proxy for their difference in consumption. For the general case when $x(m) \neq 0$, one also needs to adjust for any difference in out of pocket spending between the insured and uninsured that would occur in a world where everyone made the choice of m as if they were insured.

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 D. The next section takes this approach to the data.

6 Application to MA Health Insurance Subsidies

I apply the approach to study the optimal health insurance subsidies and mandates in the subsidized insurance marketplace for low-income adults in Massachusetts, Commonwealth Care. Developed as part of the 2006 Massachusetts health insurance reform, it later became a model for the health insurance exchanges for low-income adults constructed in the Affordable Care Act. The marketplace provides subsidies to low-income individuals who are not eligible for Medicaid and do not have access to employer-provided insurance. Finkelstein et al. (2019) provide further details. Importantly for the modeling purposes, these contracts involved virtually no cost sharing, $x(m) = 0$.

Using administrative data from Massachusetts, Finkelstein et al. (2019) exploit discontinuities in the health insurance subsidy schedule to estimate willingness to pay and cost curves. I focus here on the baseline estimates from Finkelstein et al. (2019), which use the empirical discontinuities in 2011 to measure $D(s)$ and $C(s)$, for those with incomes at 150% of the Federal Poverty Level (FPL).³¹ In the language of Section 3, this means that the set of people eligible for the market, M , corresponds to individuals with incomes at 150% FPL, μ is the open enrollment period in the fall of 2010, and $v = 2011$.

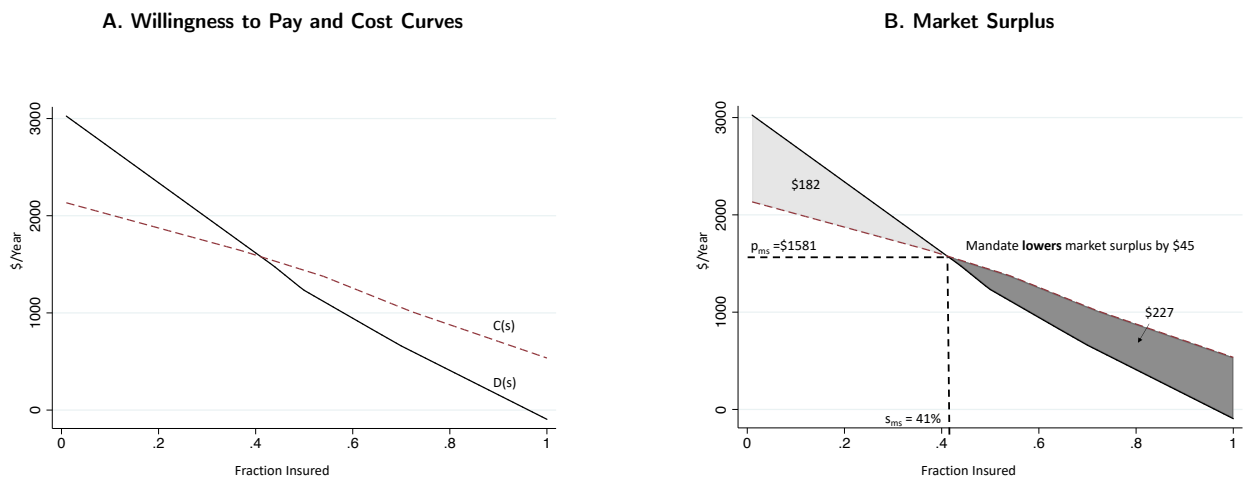
Figure 3A presents the results for $D(s)$ and $C(s)$ in Finkelstein et al. (2019) plotted as a function of s .³² The patterns reveal that those with the highest willingness to pay (low

³¹150% FPL corresponds to roughly \$16K in income for an individual with no children.

³²The estimates from Finkelstein et al. (2019) correspond to those that assume the government is the payer of uncompensated care, and they are scaled by a factor of 12 to correspond to yearly values.

values of s) are willing to pay more than their marginal cost for insurance, $D(s) > C(s)$.³³ But those with the lowest willingness to pay are willing to pay less than the cost the cost of their insurance. The model in Section (3) captures $D(s) < C(s)$ by allowing part of the cost to be driven by moral hazard: when $x'(m) < 1$, insured individuals may choose medical services that they don't value at their full resource cost.³⁴

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



Finkelstein et al. (2019) show that this market would fully unravel without subsidies, so that no one obtains insurance ($s = 0$). This unraveling is the result of both adverse selection and uncompensated care externalities. Uncompensated care externalities arise in this environment because the total cost 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. Private insurance would have to pay these additional costs, which leads the average cost faced by the private insurer to lie everywhere above the demand curve. This generates a full unraveling of the market in the absence of subsidies.

³³Finkelstein et al. (2019) do not report the joint sampling distribution for $D(s)$ and $C(s)$, which prevents a formal construction of standard errors on the estimates I provide below. But, they test whether the negative slopes for $D(s)$ and $C(s)$ are statistically significant. They find high t-stats of 6.1 (47.3/7.7) at 150% FPL for the cost curve and 13.2 (1735/131) for the demand curve. Equation 3 shows that testing whether $EA(s)$ is positive is equivalent to testing whether $D(s)$ slopes downward. Therefore, the appropriate t-stat for testing $EA(s) > 0$ is 13.2.

³⁴In contrast, this moral hazard response is assumed not to exist in the stylized model of Section 2 and in other previous models studying reclassification risk and notions of ex-ante expected utility (e.g. Handel et al. (2015); Einav et al. (2016))

The goal of this section is to evaluate the ex-ante welfare impacts of subsidies and mandates in this market and compare the conclusions a more traditional analysis of market surplus. I begin with the welfare impact of increasing subsidies for insurance if they are financed by increasing prices/penalties on the uninsured.

6.1 Budget neutral policies

Figure 3B shows that market surplus is maximized when $s = 41\%$ of the market is insured and the marginal price for insurance is \$1581. The market surplus from insuring 41% of the market is \$182. But, expanding coverage beyond this would lower market surplus since those with $D(s) < 1581$ are not willing to pay the cost they impose on insurers. On net, mandates would lower total market surplus by \$45.

What insurance prices maximize ex-ante welfare? 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 $\gamma = 5 \times 10^{-4}$ (e.g. similar to estimates in Handel et al. (2015)).³⁵ Table 1 presents estimates for a range of alternative risk aversion coefficients.

Figure 4: Ex-Ante Welfare of Health Insurance for Low-Income Adults

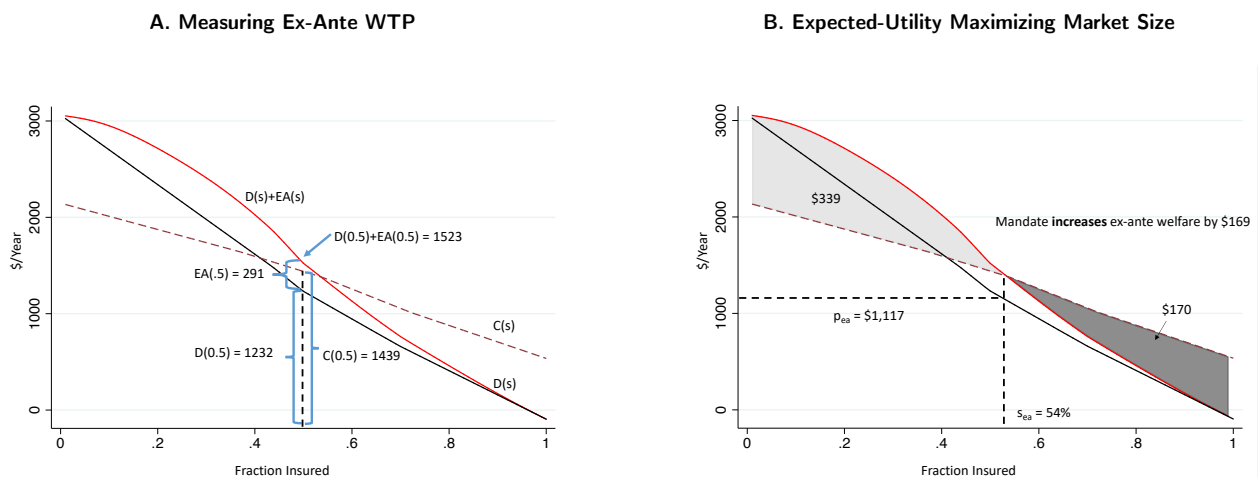


Figure 4 presents the ex-ante demand curve, $D(s) + EA(s)$, using equation (16). Panel A illustrates the calculation of $EA(s)$ when 50% of the population owns insurance. The cost

³⁵Handel et al. (2015) estimate this risk aversion coefficient 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 .

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 of the marginal enrollee is $D'(0.5) = -3405$.³⁶ The average $D(s)$ for those with $s > 0.5$ is 548. Equation (16) implies that the ex-ante willingness to pay for a larger insurance market is

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

Individuals with median (0.5) levels of $D(s)$ are willing to pay \$1,232 for insurance at the time the econometrician observes them in the market. But, from behind a veil of ignorance before knowing $D(s)$, 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 ($291 * (0.51 - 0.5)$).

Ex-ante welfare 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 4B. 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.

The ex-ante welfare gain from insuring $s = 54\%$ of the market is large. 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. The ex-ante welfare cost of insuring the remaining 46% of the market is \$170. This means that imposing a mandate would generate lower ex-ante welfare than the optimal price of \$1117, but individuals would prefer a mandate relative to no insurance. 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 Table 1 presents estimates of the above results for alternative risk 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

³⁶Finkelstein 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.

alternative coefficients of relative risk aversion ranging from 1 to 10.³⁷

The baseline specification of $\gamma = 5 \times 10^{-4}$ 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 1.05×10^{-4} – or, equivalently, 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 Coefficient		CRRR Coefficient (150% FPL)					
		CARA = 5×10^{-4}	1×10^{-4}	10×10^{-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 5×10^{-4} . Columns (3) and (4) use alternative coefficients of absolute risk aversion of 1×10^{-4} and 10×10^{-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, but by taxpayers at other income levels. The subsidies are a method of redistribution. To compare the welfare impact of these subsidies to other forms of redistribution such as tax credits, I construct the MVPF as described in Section 4.2. This is given by the formula:

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

³⁷To 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.

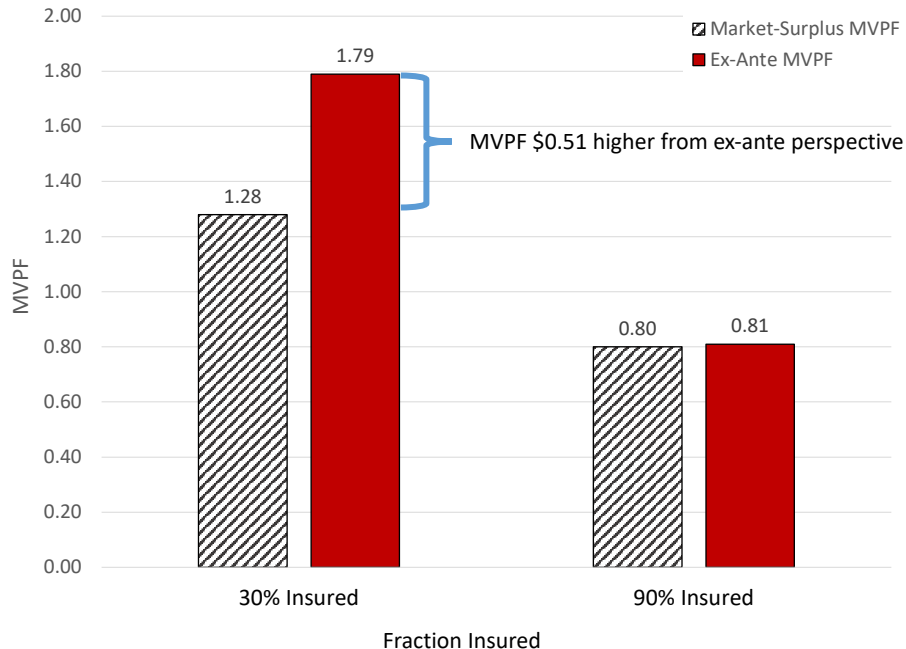
where $\beta(s) = \gamma(D(s) - E[D(S) | S \geq s])$ is the difference in marginal utilities between the insured and uninsured.

Figure 5 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

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

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, $D(0.3) > C(0.3)$. Using observed WTP would imply an MVPF of 1.28, as shown in the left bar in Figure 5. Behind the veil of ignorance, individuals are willing to pay a 39.4% markup to have the ability to purchase insurance at lower prices. This means that from an ex-ante perspective, individuals would be willing to pay \$1.79 for every \$1 of government spending on insurance subsidies.

Figure 5: MVPF for Health Insurance Subsidies for Low-Income Adults



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)). This suggests expanded insurance subsidies financed by a budget-neutral reduction in EITC would increase ex-ante welfare when $s = 0.3$.³⁸

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 ex-ante value of having marginally lower premiums is smaller because insurance premiums are already low ($D(s)$ is similar to $E[D(S) | S \geq s]$ when s is close to 1). 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 welfare.

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 is parsimonious but flexible enough to perfectly match the estimated willingness to pay and cost curves in Finkelstein et al. (2019). I first use the model to validate the sufficient statistics approach relative to this structural benchmark. Then, I use extensions of the model to evaluate violations of Assumptions 1 and 2.

7.1 Setup

The structural model follows the environment developed in Einav et al. (2013) that allows for both adverse selection and moral hazard.³⁹ In period ν , individuals obtain a realization of θ_ν and choose c and m to maximize utility, which is given by:

$$u(c, m; \theta_\nu) = -\frac{1}{\gamma} e^{-\gamma[(m-\lambda(\theta_\nu)) - \frac{1}{2w}(m-\lambda(\theta_\nu))^2 + c(m)]} \quad (17)$$

³⁸Table 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.

³⁹To my knowledge, no previous paper has evaluated the ex-ante welfare impact of insurance market policies using a structural approach that includes moral hazard. 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.

where $c(m) = y - p_I$ if the individual is insured and $c(m) = y - m - p_U$ if the individual is uninsured (recall $x(m) = 0$ for the MA setting). 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.⁴⁰

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 $-\frac{1}{\gamma}e^{-\gamma[\frac{1}{2}w+y-p_I]}$.

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 $D(S(\theta_\mu))$ 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 utility function in equation (17) satisfies the no income effects condition so that willingness to pay does not depend on p_U .

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 (17) implies that the marginal willingness to pay for insurance solves:

$$D(s) = C(s) + \frac{w}{2} + \frac{\gamma\Sigma(s)}{2} \quad (18)$$

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,

⁴⁰I assume w , y , and γ are constant across individuals. Allowing w to be heterogeneous does not affect the results (conditional on matching the demand and cost curves). Section 7.4 explores the robustness of the results when one allows y and γ to vary across individuals.

$\Sigma(s)$, the CARA utility structure implies that they value this reduction according to the risk aversion parameter, γ , divided by 2. The model matches heterogeneity in the markup ($D(s) - C(s)$) that individuals are willing to pay through heterogeneity in individuals' belief variances, $\Sigma(s)$.

I parameterize risk aversion to $\gamma = 5 \times 10^{-4}$. I set w to be equal to mean net costs of 1,336, which corresponds to roughly a 30% moral hazard effect on gross medical spending (roughly consistent with previous empirical findings). Given w , γ and $C(s)$, I set $\Sigma(s)$ to be the unique value of $\Sigma(s)$ that solves equation (18). This means that the model parameters perfectly match the reduced form cost curve, $C(s)$, and willingness to pay curve, $D(s)$, at each value of s . Moreover, the model structure satisfies the baseline implementation assumptions above.⁴¹

Finally, Proposition 5 shows that the structural model yields a solution for $W(s)$ given by equation (9).

Proposition 5. *Under the modeling assumptions outlined above, $W(s)$ in equation (9) is given by*

$$W(s) = \frac{-1}{\gamma} \left[\log \left(s e^{\gamma p_I(s)} e^{-\gamma \frac{w}{2}} + (1-s) e^{\gamma p_U(s)} \int_s^1 e^{\gamma(C(\bar{s})-w) + \frac{\gamma^2}{2} \Sigma(\bar{s})} d\bar{s} \right) - \log \left(\int_0^1 \left[e^{\gamma(C(\bar{s})-w) + \frac{\gamma^2}{2} \Sigma(\bar{s})} \right] d\bar{s} \right) \right] \quad (19)$$

Proof. See Appendix F. □

7.3 Results


Figure 6 compares the value of $W'(s)$ from the structural model computed using Equation (19) (shown in the dash-dot brown line), to the value of $W'(s)$ computed using the sufficient statistics approach (in the solid red line).⁴² 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. 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

⁴¹It is straightforward to show that the utility function satisfies the no income effects condition and Assumption 2. Moreover, the utility function does not exhibit complementarities between consumption and medical spending, $u_{cm} = 0$, so that it satisfies Assumption 1 to first order.

⁴² $W'(s)$ is computed using a numerical derivative. As shown in an earlier draft and available upon request, $W(s)$ is estimated to be smooth. Therefore, the results are robust to the methods used to compute this derivative.

a full mandate is also slightly larger (\$180 versus \$170). Overall, the approach proposed in this paper closely mirrors the ex-ante welfare as measured in the structural model.

Figure 6: Comparison to Structural Approach



7.4 Violations of Assumptions

While Assumptions 1 and 2 provide a benchmark method to measure $\beta(s)$, they are restrictive. Here, I use the structural model to explore the impact of violating these assumptions and discuss potential additional data elements that can recover $\beta(s)$ when these assumptions do not hold. Section 7.4.1 focuses on the impact of income or liquidity differences between the insured and uninsured, and Section 7.4.2 studies the implications of heterogeneity in risk aversion. In addition, Appendix J discusses the implications of complementarities between the marginal utility of consumption and health status.

7.4.1 Income or Liquidity Differences

If the insured have different income or liquidity than the uninsured, then their difference in willingness to pay for insurance will no longer proxy for their difference consumption.⁴³ In the MA example, the demand for insurance subsidies is estimated conditional on incomes of 150% FPL, which suggests this may not be a primary concern; but more generally income or liquidity may be a key determinant of insurance demand. In this case, I show how one can recover ex-ante measures of willingness to pay if one can directly observe the difference in consumption or income between the insured and uninsured, as suggested in Proposition 3.

To see this, suppose that individuals with different willingnesses to pay have different incomes in period ν . For example, let $\bar{y}(s)$ to be the income of an individual who is indifferent to purchasing insurance when prices are such that a fraction s of the market owns insurance. Appendix G modifies Proposition 5 to show that $W(s)$ solves

$$e^{-\gamma W(s)} \int_0^1 e^{\gamma(\bar{y}(\bar{s})+C(\bar{s})-w)+\frac{1}{2}\gamma^2\Sigma(\bar{s})} d\bar{s} = e^{\gamma p_I(s)} e^{-\frac{1}{2}\gamma w} \int_0^s e^{-\gamma\bar{y}(\bar{s})} d\bar{s} + e^{\gamma p_U(s)} \int_s^1 e^{\gamma(\bar{y}(\bar{s})+C(\bar{s})-w)+\frac{1}{2}\gamma^2\Sigma(\bar{s})} d\bar{s} \quad (20)$$

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 \$1,382 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 \$1,382 lower than those with the highest willingness to pay. In both cases I assume average incomes are a linear function of s .


Figure 7A and 7B present the ex-ante willingness to pay curve, $W'(s)$, for these cases. When the insured have higher incomes, Figure 7A shows that the benchmark implementation over-states the true ex-ante willingness to pay: it is optimal for just 44% of the market to be insured in contrast to the 54% implied in the baseline implementation. Conversely, Figure 7B shows how the pattern reverses when the insured have higher incomes. In this case, the optimal size of the market involves 65% of the market being insured.

In the presence of liquidity or income differences between the insured and uninsured,

⁴³A 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 potentially 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.

consumption data provides a path to accurate measurement of $W'(s)$. The “modified $\beta(s)$ ” curve shows how using consumption or income data to measure Δc as in Proposition 3 recovers the true $W'(s)$.⁴⁴ When the insured have higher incomes, using consumption data to measure the ex-ante willingness to pay leads to a predicted optimal size of the market of 45%, close to the true optimal size of the market of 44%. Similarly, when the insured have lower incomes, using consumption data to measure ex-ante willingness to pay implies an optimal size of the market of 65%, very close to the true optimum implied by the structural model.

Figure 7: Income or Liquidity Differences



7.4.2 Heterogeneity in Risk Preferences

A classic problem in economics to separate out risk preferences from beliefs. Individuals may be willing to pay a higher markup for insurance either because they have higher risk aversion, γ , or more uncertainty in their outcomes, Σ . To be consistent with Assumption 1, the specification in Section 7.2 ruled out heterogeneity in γ and assumed differences in willingness to pay was due to differences in $\Sigma(s)$ (see equation (18)). Here, I assess the potential bias that arises when there is heterogeneity in γ . I provide intuition for the potential direction of this bias and make suggestions for how future work can overcome these potential biases.

I consider two specifications that parameterize heterogeneity in γ as $\gamma \sim N(\mu_\gamma, \sigma_\gamma^2)$. In

⁴⁴I use the formula $\beta(s) = \gamma(D(s) - E[D(S)|S \geq s] + E[y_\nu(\theta_\nu)|S \leq s] - E[y_\nu(\theta_\nu)|S > s])$ 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.

both specifications, mean risk aversion is $\mu_\gamma = 5x10^{-4}$, but they differ in their population standard deviations, σ_γ . The first “low heterogeneity” specification calibrates $\sigma_\gamma = 5x10^{-5}$ and the second “high heterogeneity” specification calibrates $\sigma_\gamma = 1x10^{-4}$ (close to the estimates in Handel et al. (2015)). I fit the model using a minimum distance estimator discussed in Appendix H.

Figure 8 presents the results for the low heterogeneity case (Figure 8A), high heterogeneity case (Figure 8B). The dot-dash brown line corresponds to the true ex-ante marginal 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 homogeneous coefficient of relative risk aversion of $\gamma = 5x10^{-4}$, which corresponds to the population average in both specifications.

For both the high and low heterogeneity cases, the benchmark implementation overstates 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 8A and Figure 8B. Ex-ante welfare is maximized when 53% of the market owns insurance in the low-heterogeneity specification and 46% in the high-heterogeneity specification. These contrast with the optimal ex-ante size of the market of 54% that is implied by the baseline implementation and the 41% that would maximize market surplus.

The true ex-ante welfare is lower because the marginal utility of income is lower for the insured relative to uninsured than is implied in a model with fixed risk aversion. To see this, note that the marginal utility of consumption for CARA utility is $e^{-\gamma C}$ where C is the net consumption (e.g. $C = (m - \lambda(\theta_\nu)) - \frac{1}{2w}(m - \lambda(\theta_\nu))^2 + c(m)$). With heterogeneity in γ , those with high willingness to pay for insurance are more likely to have a higher γ .⁴⁵ Because $e^{-\gamma C}$ is then declining in γ , this means the insured have a lower marginal utility of income relative to the uninsured, even conditional on their level of C .⁴⁶


Stepping back, the core issue is that individuals may have different risk preferences over realizations of θ_μ than their risk preferences over realizations of θ_ν given θ_μ (the latter is what is identified from insurance choices). The benchmark implementation above assumes

⁴⁵Note that the model with heterogeneity in γ continues to match the willingness to pay curve, $D(s)$, and cost curve, $E[\lambda(\theta_\nu)|s] = C(s)$, at each value of s . This means that the difference in average willingness to pay and costs between the insured and uninsured is held constant as one changes the amount of preference heterogeneity, σ_γ .

⁴⁶While the marginal utility of income declines in γ in a benchmark CARA specification, $u(C; \gamma) = \frac{-1}{\gamma}e^{-\gamma C}$, it is also possible to have CARA specifications in which the marginal utility of income increases in γ . To see this, let \bar{C} denote the upper bound of consumption, C , in the economy and let $u(C; \gamma) = -\frac{e^{\gamma \bar{C}}}{\gamma}e^{-\gamma C}$. This utility function also exhibits constant absolute risk aversion: $\frac{-u''(C)}{u'(C)} = \frac{\gamma e^{\gamma \bar{C}} e^{-\gamma C}}{e^{\gamma \bar{C}} e^{-\gamma C}} = \gamma$. But, $u'(C; \gamma) = e^{\gamma(\bar{C}-C)}$ in this case is increasing in γ . With this specification, the marginal utility of the insured would be higher than the marginal utility of the uninsured, conditional on consumption. This would lead the baseline approach to under-state ex-ante willingness to pay.

risk aversion over these risks are the same. As these examples highlight, this need not be the case. However, Proposition 1 continues to hold even with heterogeneity in γ . This provides a potential roadmap for future work: one can seek to directly estimate the percentage difference in marginal utility of incomes between those who do versus do not purchase insurance.

Figure 8: Risk Aversion Heterogeneity



To see this, the dashed blue line in Figures 8A and 8B present estimates of ex-ante welfare using the formula $EA(s) = s(1-s)D'(s)\beta(s)$, where $\beta(s)$ is taken to be the “true” percentage difference in marginal utilities of income between the insured and uninsured. The dashed blue line approximates the true ex-ante welfare quite well in both specifications. The optimal size of the market implied by this implementation is 47% and 52% in the high and low heterogeneity specifications, respectively, which correspond closely to the true optimal size of the market of 46% and 53%. This shows that if one could observe the needed additional sufficient statistic, $\beta(s)$, one could recover ex-ante welfare even in the presence of preference heterogeneity.

Recent and ongoing work develops a range of strategies for estimating percentage differences in marginal utilities of income. For example, in the unemployment insurance context, Hendren (2017) provides three potential approaches that use consumption data, ex-ante realization of knowledge about future job loss, and spousal responses to infer the differences in marginal utilities. Landais and Spinnewijn (2019) also discuss a novel method using consumption responses to income shocks to infer percentage differences in marginal utilities of income. Future work could develop similar approaches tailored to the health insurance context and assess whether ex-ante risk preferences differ systematically from those governing

willingness to pay in observed insurance markets.

8 Conclusion

Traditional market surplus does not capture the value of insurance against risk that has been revealed at the time individuals choose insurance. In contrast, ex-ante expected utility provides a consistent welfare framework to study optimal insurance policies. Ex-ante measures of welfare differ from traditional market surplus because they measure expected utility before individuals learn their willingness to pay for insurance.

This paper develops a set of tools to measure the ex-ante welfare impact of insurance market policies and applies the approach to existing estimates of willingness to pay and cost curves for low-income health insurance in Massachusetts. Applying the model to the study of the Massachusetts health insurance exchange, the results show that an ex-ante welfare perspective can lead to different normative conclusions. Policies that maximize ex-ante welfare 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. One 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 (2015); 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 (2019)). 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.
- Cabral, M. (2017). Claim timing and ex post adverse selection. *Review of Economic Studies* 84(1).
- Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics* 90(10), 1879–1901.
- Dague, L. (2014). The effect of Medicaid premiums on enrollment: A regression discontinuity approach. *Journal of Health Economics* 37, 1–12.
- Einav, L., A. Finkelstein, and M. R. Cullen (2010). Estimating welfare in insurance markets using variation in prices. *Quarterly Journal of Economics* 125(3).
- Einav, L., A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen (2013). Selection on moral hazard in health insurance. *American Economic Review* 103(1), 178–219.
- 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.
- Feldman, R. D. and B. E. Dowd (1982). Simulation of a health insurance market with adverse selection. *Operations Research* 30(6), 1027–1042.
- Finkelstein, A., N. Hendren, and E. F. Luttmer (2019). The value of Medicaid: Interpreting results from the Oregon health insurance experiment. *Journal of Political Economy* 127(6).
- 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.
- 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.
- 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.

- Gallen, T. S. (2015). Using participant behavior to measure the value of social programs: The case of Medicaid. *Working Paper*.
- Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review* 87(1), 192–205.
- Hackmann, M. B., J. T. Kolstad, and A. E. Kowalski (2015). Adverse selection and an individual mandate: When theory meets practice. *American Economic Review* 105(3), 1030–1066.
- Handel, B., I. Hendel, and M. D. Whinston (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica* 83(4), 1261–1313.
- Hendren, N. (2013). Private information and insurance rejections. *Econometrica* 81(5), 1713–1762.
- Hendren, N. (2016). The policy elasticity. *Tax Policy and the Economy* 30.
- Hendren, N. (2017). Knowledge of future job loss and implications for unemployment insurance. *American Economic Review* 107(7).
- Hendren, N. and B. Sprung-Keyser (2019). A unified welfare analysis of government policies.
- Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. *American Economic Review* 61(4), 561–574.
- Keane, M. and R. Moffitt (1998). A structural model of multiple welfare program participation and labor supply. *International Economic Review*, 553–589.
- Landais, C. and J. Spinnewijn (2019). The value of unemployment insurance. *Working Paper*.
- Lucas, R. E. (2003). Macroeconomic priorities. *American Economic Review* 93(1), 1–14.
- Meyer, B. D. and W. K. Mok (2019). Disability, earnings, income and consumption. *Journal of Public Economics* 171, 51 – 69.
- Milgrom, P. and I. Segal (2002). Envelope theorems for arbitrary choice sets. *Econometrica* 70(2), 583–601.
- Panhans, M. (2018). Adverse selection in ACA exchange markets: Evidence from Colorado. *The American Economic Journal: Applied Economics*.

Spinnewijn, J. (2017). Heterogeneity, demand for insurance, and adverse selection. *American Economic Journal: Economic Policy* 9(1), 308–43.

Online Appendix: Not For Publication

A Proof of Proposition 1

This Appendix walks through the proof of Proposition 1. I begin with the general case, then consider no income effects.

A.1 Proof of Theorem

To begin, note that prices solve the equations:

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

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

As noted in the text, $sAC(s) = \int_0^s [m_\nu^I - x(m_\nu^I) | S(\theta_\nu) = s] ds$ is differentiable in s because $D(p_U, s)$ is differentiable in p_U and s and individuals choice of m^I is differentiable in prices. Differentiating, we have

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

or

$$sp_I'(s) + (1-s)p_U'(s) = C(s) - p_I(s) + p_U(s)$$

or

$$-(sp_I'(s) + (1-s)p_U'(s)) = D(p_U(s), s) - C(s)$$

Intuitively, the average change in prices (weighted by market shares) is equal to the marginal market surplus, given by the difference between willingness to pay and cost (if there is no marginal surplus, then the price decreases for the uninsured must fully pay for the price decreases for the insured).

Next, note that one can express $-p_U'(s)$ as follows. Differentiating the demand identity

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

yields

$$p_I'(s) - p_U'(s) = \frac{\partial D}{\partial s}(p_U(s), s) + \frac{\partial D}{\partial p_U}(p_U(s), s) p_U'(s)$$

So,

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

or

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

which provides an expression for the change in prices for the uninsured in the budget neutral policy.

Now, recall that marginal ex-ante welfare, $W'(s)$, solves:

$$\begin{aligned}
W'(s) * \frac{-\partial V(\infty, -W(s))}{\partial p_U} &= \frac{d}{ds} E \left[\sum_{t \geq 0} v_t(p_I(s), p_U(s); \theta_t) \right] \\
&= E \left[1 \{S(p_U(s), \theta_\mu) < s\} \frac{\partial u_\nu}{\partial c} (-p'_I(s)) + 1 \{S(p_U(s), \theta_\mu) \geq s\} \frac{\partial u_\nu}{\partial c} (-p'_U(s)) \mid \theta_\mu \in M \right] \\
&= sp'_I(s) E \left[\frac{-\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) < s \right] + (1-s)p'_U(s) E \left[\frac{-\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) \geq s \right] \\
&= -(sp'_I(s) + (1-s)p'_U(s)) E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) < s \right] + \dots \\
&\quad \dots + (1-s)p'_U(s) \left[E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) \geq s \right] \right]
\end{aligned}$$

Note the term $-(sp'_I(s) + (1-s)p'_U(s))$ is equal to $D(p_U(s), s) - C(s)$, so that

$$W'(s) * \frac{-\partial V(\infty, -W(s))}{\partial p_U} = (D(p_U(s), s) - C(s)) E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) < s \right] + (1-s)p'_U(s) \left[E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) \geq s \right] \right]$$

Now, plugging in $p'_U(s)$ from above yields:

$$\begin{aligned}
W'(s) * \frac{-\partial V(\infty, -W(s))}{\partial p_U} &= (D(p_U(s), s) - C(s)) E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) < s \right] \\
&\quad + \frac{-(1-s)}{1 + s \frac{\partial D}{\partial p_U}} \left[D(p_U(s), s) - C(s) + s \frac{\partial D}{\partial s} \right] \left[E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) \geq s \right] \right]
\end{aligned}$$

which can be re-written as

$$\begin{aligned}
W'(s) \frac{-\partial V(\infty, -W(s))}{\partial p_U} &= (D(p_U(s), s) - C(s)) \left[E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) < s \right] - \frac{1-s}{1 + s \frac{\partial D}{\partial p_U}} \left(E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) \geq s \right] \right) \right] \\
&\quad + \frac{s(1-s)}{1 + s \frac{\partial D}{\partial p_U}} \left(-\frac{\partial D}{\partial s} \right) \left[E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) < s \right] - E \left[\frac{\partial u_\nu}{\partial c} \mid S(p_U(s), \theta_\mu) \geq s \right] \right]
\end{aligned}$$

Next, focus on the term multiplying $D(p_U(s), s) - C(s)$ to see how it can be reduced:

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

where $E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right]$ denotes the average marginal utility of income for those in the market when a fraction s is insured:

$$E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right] = s E \left[\frac{\partial u_\nu}{\partial c} | S(p_U(s), \theta_\mu) < s \right] + (1-s) E \left[\frac{\partial u_\nu}{\partial c} | S(p_U(s), \theta_\mu) \geq s \right]$$

So, returning to $W'(s) \frac{-\partial V(\infty, -W(s))}{\partial p_U}$, we have

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

which implies

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

where the last term is equal to zero if $\frac{\partial D}{\partial p_U} = 0$.

Next, note that the second term above can be written as

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

so that

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

Now, consider $\frac{\partial V(\infty, W(s))}{\partial p_U}$. The envelope theorem implies:

$$\frac{-\partial V(\infty, -W(s))}{\partial p_U} = E \left[\frac{\partial u_\nu(\infty, -W(s))}{\partial c} | \theta_\mu \in M \right]$$

where $\frac{\partial u_\nu(\infty, -W(s))}{\partial c} = \frac{\partial u_\nu(c^U(\infty, -W(s)), m_\nu^U(\infty, -W(s), \theta_\nu))}{\partial c}$ is the marginal utility of consumption evaluated at the optimal choices when $p_I = \infty$ and $p_U = -W(s)$. This is not to be confused with the average marginal utility of consumption that occurs when prices are set according to $p_I(s)$ and $p_U(s)$ so that a fraction s of the market has insurance, which I denote by $E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right]$. But, I note below that if the utility function has no income effects, these marginal utilities of income will be the same so that $E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right] = E \left[\frac{\partial u_\nu}{\partial c}(\infty, -W(s)) | \theta_\mu \in M \right]$. Intuitively, when the marginal utility of income only depends on the level of utility, these must be identical because the transfer $W(s)$ equates the utilities. Combining terms,

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

And now replacing the denominator in the first two terms, $E \left[\frac{\partial u_\nu}{\partial c}(\infty, -W(s)) | \theta_\mu \in M \right]$,

with $E \left[\frac{\partial u_{\nu}}{\partial c} | \theta_{\mu} \in M \right]$ and collecting the difference in a fifth and sixth term yields:

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

So, now let

$$EA(s) = s(1-s) \left(-\frac{\partial D}{\partial s} \right) \left[\frac{E \left[\frac{\partial u_{\nu}}{\partial c} | S(p_U(s), \theta_{\mu}) < s \right] - E \left[\frac{\partial u_{\nu}}{\partial c} | S(p_U(s), \theta_{\mu}) \geq s \right]}{E \left[\frac{\partial u_{\nu}}{\partial c} | \theta_{\mu} \in M \right]} \right]$$

and

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

Which can be re-written as

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

With these definitions,

$$W'(s) = D(p_U(s), s) - C(s) + EA(s) + \delta_p(s)$$

Now, to see that $\delta_p(0) = 0$, note that the only term of question is the third term in the definition of $\delta_p(s)$ (since all others are multiplied by s). Here, note that when $s = 0$, it must be the case that $W(0) = 0$, so that the equilibrium allocations when $s = 0$ are equal to the equilibrium allocations when no one has insurance and pays $p_U = -W(0) = 0$. Hence, when $s = 0$ we have $E \left[\frac{\partial u_{\nu}}{\partial c} | \theta_{\mu} \in M \right] = E \left[\frac{\partial u_{\nu}}{\partial c} (\infty, -W(s)) | \theta_{\mu} \in M \right]$.

Hence, $\delta_p(0) = 0$.

A.2 General Case with Savings

This section discusses the generalization of the model in Section 3 to allow for savings behavior. The key insight is that in the presence of savings, the model continues to nest the assumptions required in the general framework of Milgrom and Segal (2002) that enables one to invoke the envelope theorem. This means that the marginal welfare impact of changing prices p_I and/or p_U only depends on the size of the price change multiplied by the marginal utilities of consumption. This means that the results in Proposition 1 and 2 would continue to hold. Moreover, assuming the marginal utility function continues to satisfy $\frac{\partial u_\nu}{\partial c} = f(c)$ for some function c , the results in Proposition 3 will continue to hold even when individuals respond to changes in prices, p_I and p_U , by increasing/decreasing savings or other choices in other periods.⁴⁷

The main modification relative to the baseline model is as follows. For any potential sequence of realizations, $\{\theta_t\}_{t=1}^T$, individuals make two classes of choices. First, they choose their state-contingent plan of consumption and medical spending $\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T \in \mathbb{R}^{\sum_{t=1}^T tN}$ for all periods. Second, in period $t = \mu$ they realize $\theta_\mu \in \mathbb{R}^{\mu N}$ decide whether to purchase health insurance for period $t = \nu$, characterized by an indicator $I_\nu(\theta_\mu) \in \{0, 1\}$. I assume individuals make these choices, $\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T$ and $I_\nu(\theta_\mu)$, subject to K constraints:

$$g^k \left(\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T, I_\nu(\theta_\mu); p_I, p_U \right) \leq 0 \quad \forall k = 1, \dots, K$$

where $g^k(c; p_I, p_U) : \mathbb{R}^{\sum_{t=1}^T tN} \times \mathbb{R}^{\mu N} \rightarrow \mathbb{R}$ for each $k = 1, \dots, K$. I assume each g^k is twice continuously differentiable in all c and m arguments. Note in particular that by choosing appropriate functions g^k one can allow for savings.⁴⁸

Individuals choose consumption, medical spending, along with the decision to purchase

⁴⁷It is perhaps worthwhile to note that I do need to assume that providing insurance in period $t = \nu$ does not induce externalities on other individuals or insurers in other periods $t \neq \nu$.

⁴⁸For example, perfect credit markets with no interest rate and no insurance in any other periods aside from $t = \nu$ would correspond to a single constraint,

$$g \left(\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T, 0; p_I, p_U \right) = \sum_{t \neq \nu} [c_t(\theta_t) - m_t(\theta_t)] + c_\nu(\theta_\nu) - I_\nu(\theta_\mu) m_\nu(\theta_\nu) - (1 - I_\nu(\theta_\mu)) x(m_\nu(\theta_\nu)) \leq \sum_t y_t(\theta_t)$$

insurance in period μ for coverage in period ν to maximize ex-ante expected utility:

$$V(p_I, p_U) = \max_{\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T, I_\nu(\theta_\mu)} \sum_{t=1}^T E[u_t(c_t(\theta_t), m_t(\theta_t); \theta_t)]$$

$$s.t. \quad g^k \left(\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T, I_\nu(\theta_\mu); p_I, p_U \right) \leq 0 \quad \forall k \quad (21)$$

Given the differentiability assumptions, the optimal choices maximize the Lagrangian,

$$\mathcal{L}(p_I, p_U) = E \left[\sum_{t=1}^T u_t(c_t(\theta_t), m_t(\theta_t); \theta_t) \right] - \sum_{k=1}^K \lambda_k g^k \left(\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T, I_\nu(\theta_\mu); p_I, p_U \right)$$

where λ_k are the multipliers on each constraint $k = 1, \dots, K$.

The constraints depend upon the prices p_I and p_U that individuals pay in period $t = \nu$ if they are insured or uninsured. I assume that in the absence of behavioral changes to consumption in other periods or changes in medical spending, the impact of a price increase is to reduce the amount of consumption individuals can consume in period ν . Formally, I assume

$$\frac{\partial g^k}{\partial p_I} = \frac{\partial g^k \left(\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T, 1; p_I, p_U \right)}{\partial c_t(\theta_t)}$$

so that increasing consumption has the same impact on the constraints as increasing p_I in the event of choosing $I_\nu(\theta_\nu) = 1$. Similarly, for the uninsured, I assume

$$\frac{\partial g^k}{\partial p_U} = \frac{\partial g^k \left(\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T, 0; p_I, p_U \right)}{\partial c_t(\theta_t)}$$

so that increasing consumption has the same impact on the constraints as increasing p_U in the event of choosing $I_\nu(\theta_\nu) = 0$.

Notice that the maximization program in equation (21) meets the conditions in Milgrom and Segal (2002) so that the envelope theorem holds. Therefore,

$$\frac{\partial V}{\partial p_I} = -E \left[\sum_k \lambda_k \frac{\partial g^k \left(\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T, 1; p_I, p_U \right)}{\partial p_I} \right]$$

and

$$\frac{\partial V}{\partial p_U} = -E \left[\sum_k \lambda_k \frac{\partial g^k \left(\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T, 0; p_I, p_U \right)}{\partial p_U} \right]$$

Now, note that the choice of consumption satisfies

$$\frac{\partial u_t}{\partial c} (c_t(\theta_t), m_t(\theta_t); \theta_t) = \sum_k \lambda_k \frac{\partial g^k \left(\{c_t(\theta_t), m_t(\theta_t)\}_{t=1}^T, I_\nu(\theta_\mu); p_I, p_U \right)}{\partial c_t(\theta_t)}$$

So,

$$\frac{\partial V}{\partial p_I} = -E \left[1 \{I_\nu(\theta_\mu) = 1\} \frac{\partial u_t}{\partial c} \right] = \Pr \{I_\nu(\theta_\mu) = 0\} E \left[\frac{\partial u_t}{\partial c} | I_\nu(\theta_\mu) = 0 \right]$$

and

$$\frac{\partial V}{\partial p_U} = -E \left[1 \{I_\nu(\theta_\mu) = 0\} \frac{\partial u_t}{\partial c} \right] = \Pr \{I_\nu(\theta_\mu) = 1\} E \left[\frac{\partial u_t}{\partial c} | I_\nu(\theta_\mu) = 1 \right]$$

so that the welfare impact of a marginal change in the price of insurance p_I (uninsurance p_U) is given by the average marginal utility of consumption for the insured (uninsured) multiplied by the fraction of insured (uninsured).

Because the envelope theorem holds, the formula for ex-ante welfare remains identical to the case with no savings,

$$\begin{aligned} W'(s) * \frac{-\partial V(\infty, -W(s))}{\partial p_U} &= \frac{d}{ds} E \left[\sum_{t \geq 0} v_t(p_I(s), p_U(s); \theta_t) \right] \\ &= E \left[1 \{S(p_U(s), \theta_\mu) < s\} \frac{\partial u_\nu}{\partial c}(-p'_I(s)) + 1 \{S(p_U(s), \theta_\mu) \geq s\} \frac{\partial u_\nu}{\partial c}(-p'_U(s)) | \theta_\mu \in M \right] \end{aligned}$$

and it is straightforward to show that the remainder of the proof of Proposition 1 continues to hold. Moreover, analogous to the discussion for Proposition 1 above, it is straightforward to verify that the formulas for Proposition 2 continue to hold: the welfare impact of lower p_I depend on the size of the price change multiplied by the marginal utility of consumption. Finally, to see that the results in Proposition 3 continue to hold, note that the requirements for this proposition to hold only rely on a Taylor expansion of the utility function and do not require assumptions on the budget constraints. Thus, it continues to be the case that

$$\frac{E \left[\frac{\partial u_t}{\partial c} | I_\nu(\theta_\mu) = 1 \right] - E \left[\frac{\partial u_t}{\partial c} | I_\nu(\theta_\mu) = 0 \right]}{E \left[\frac{\partial u_t}{\partial c} \right]} \approx \gamma(s) \Delta c(s)$$

However, it is no longer the case that the willingness to pay curve continues to proxy for the difference in consumption between insured and uninsured. Intuitively, some of the consumption impact of a high willingness to pay can be spread onto consumption in other periods, so that willingness to pay will not adequately capture the difference in consumption between insured and uninsured.

A.3 No income effects

Now, suppose there exists positive constants a and b and a function $\kappa(m, \theta_\nu)$ such that

$$u_\nu(c, m; \theta_\nu) = -ae^{-b[c+\kappa(m, \theta_\nu)]}$$

Implication #1: m does not respond to prices p_I and p_U .

To begin, consider the choice of $m(\theta_\nu)$. For an uninsured individual, this will solve

$$1 = \frac{\partial \kappa}{\partial m}(m^U(\theta_\nu))$$

and for an insured individual

$$x'(m^I(\theta_\nu)) = \frac{\partial \kappa}{\partial m}(m^I(\theta_\nu))$$

so that the choice of m^I and m^U will not depend on prices p_I and p_U (or τ). Note that these equations yield a unique solution for $m^I(\theta_\nu)$ because of the assumption that (a) utility is concave in m so that $\frac{\partial \kappa}{\partial m}$ is decreasing in m and (b) x' is weakly increasing in m .

Implication #2: demand, $d(p_U, \theta_\mu)$ does not depend on p_U .

Next, note that the marginal utility of income can be written as a linear function of the level of utility:

$$\begin{aligned} \frac{\partial u_\nu}{\partial c}(c, m; \theta_\nu) &= bae^{-b[c+\kappa(m, \theta_\nu)]} \\ &= -bu_\nu(c, m, \theta_\nu) \end{aligned}$$

for all (c, m, θ_ν) .

Note that $d(p_U, \theta_\mu)$ solves

$$E[u_\nu(y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | \theta_\mu] = E[u_\nu(y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - d(p_U, \theta_\mu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | \theta_\mu]$$

Differentiating the LHS with respect to p_U yields

$$\frac{d}{dp_U} E[u_\nu(y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | \theta_\mu] = E\left[-\frac{\partial u_\nu}{\partial c}(y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | \theta_\mu\right]$$

and the RHS yields:

$$\frac{d}{dp_U} E[u_\nu(y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - d(p_U, \theta_\mu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | \theta_\mu] = E\left[-\left(1 + \frac{\partial d}{\partial p_U}\right) \frac{\partial u_\nu}{\partial c}(y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - d(p_U, \theta_\mu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | \theta_\mu\right]$$

where one can ignore responses of m^U and m^I to p_U due to the envelope theorem (and, under

the functional form of the utility function m^U doesn't respond to p_U regardless). Combining, we have

$$E \left[-\frac{\partial u_\nu}{\partial c} (y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | \theta_\mu \right] = E \left[-\left(1 + \frac{\partial d}{\partial p_U}\right) \frac{\partial u_\nu}{\partial c} (y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - d(p_U, \theta_\mu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | \theta_\mu \right]$$

Note, note that

$$E \left[-\frac{\partial u_\nu}{\partial c} (y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) \right] = -bE \left[u_\nu (y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | \theta_\mu \right]$$

and

$$E \left[-\frac{\partial u_\nu}{\partial c} (y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - d(p_U, \theta_\mu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) \right] = -bE \left[u_\nu (y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - d(p_U, \theta_\mu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | \theta_\mu \right]$$

So,

$$E \left[b u_\nu (y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | \theta_\mu \right] = E \left[\left(1 + \frac{\partial d}{\partial p_U}\right) b u_\nu (y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - d(p_U, \theta_\mu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | \theta_\mu \right]$$

or

$$\begin{aligned} E \left[u_\nu (y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | \theta_\mu \right] &= E \left[u_\nu (y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - d(p_U, \theta_\mu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | \theta_\mu \right] \\ &\quad + \frac{\partial d}{\partial p_U} E \left[u_\nu (y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - d(p_U, \theta_\mu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | \theta_\mu \right] \end{aligned}$$

where one can pull $\frac{\partial d}{\partial p_U}$ out of the expectation because it does not vary conditional on θ_μ .

Now, note that by definition of $d(p_U, \theta_\mu)$, the levels of utility are equated. Hence,

$$\frac{\partial d}{\partial p_U} E \left[u_\nu (y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - d(p_U, \theta_\mu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | \theta_\mu \right] = 0$$

or

$$\frac{\partial d}{\partial p_U} = 0$$

Implication #3: $\delta_p(s) = 0$.

Recall that the average marginal utility of income in the market is defined as

$$E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right] = sE \left[\frac{\partial u_\nu}{\partial c} (c_\nu^I(\theta_\nu), m^I(\theta_\nu), \theta_\nu) | S(\theta_\mu) < s \right] + (1-s)E \left[\frac{\partial u_\nu}{\partial c} (c_\nu^U(\theta_\nu), m^U(\theta_\nu), \theta_\nu) | S(\theta_\mu) \geq s \right]$$

Now, note that $\frac{\partial u_\nu}{\partial c} = -b u_\nu$ so that this implies

$$E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right] = -b \left[sE \left[u_\nu (c_\nu^I(\theta_\nu), m^I(\theta_\nu), \theta_\nu) | S(\theta_\mu) < s \right] + (1-s)E \left[u_\nu (c_\nu^U(\theta_\nu), m^U(\theta_\nu), \theta_\nu) | S(\theta_\mu) \geq s \right] \right]$$

Now, turning to $E \left[\frac{\partial u_\nu}{\partial c} (\infty, -W(s)) | \theta_\mu \in M \right]$, note that

$$E \left[\frac{\partial u_\nu}{\partial c} (\infty, -W(s)) | \theta_\mu \in M \right] = -bE [u_\nu (\infty, -W(s)) | \theta_\mu \in M]$$

Finally, note that by definition of $W(s)$ we have

$$E [u_\nu (\infty, -W(s)) | \theta_\mu \in M] = sE [u_\nu (c_\nu^I(\theta_\nu), m^I(\theta_\nu), \theta_\nu) | S(\theta_\mu) < s] + (1-s)E [u_\nu (c_\nu^U(\theta_\nu), m^U(\theta_\nu), \theta_\nu) | S(\theta_\mu) \geq s]$$

so that the utility levels are equated. Hence, $E \left[\frac{\partial u_\nu}{\partial c} (\infty, -W(s)) | \theta_\mu \in M \right] = E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right]$.

Now to see that $\delta_p(s) = 0$, recall

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

Note that when there are no income effects, $\frac{\partial D}{\partial p_U} = 0$ which means the first two terms are equal to zero. In addition, $E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right] - E \left[\frac{\partial u_\nu}{\partial c} (\infty, -W(s)) | \theta_\mu \in M \right]$ so that the last two terms equal zero.

B Proof of Proposition 2

The MVPF is defined as the ratio of marginal willingness to pay from lower insurance prices to the marginal cost to the government of providing lower insurance prices.

$$MVPF = \frac{\text{Marginal WTP}}{\text{Marginal Cost}}$$

where one can either measure willingness to pay behind a veil of ignorance (prior to learning θ_μ) or the marginal willingness to pay after learning θ_μ .

Ex-Post MVPF After learning θ_μ , the aggregate marginal willingness to pay for a larger insurance market is

$$\begin{aligned} \text{Marginal WTP} &= -sp'_I(s) \\ &= -s \frac{\partial D}{\partial s} (0, s) \end{aligned}$$

By the envelope theorem, this is equal to the fraction of people receiving the lower prices, s , multiplied by the marginal price decline, $-p'_I(s)$. The second equality follows from the fact that $p_U(s) = 0$ so that $p'_I(s) = \frac{\partial D}{\partial s}$.

Since government expenditure is $sAC(s) - sp_I(s)$, the marginal cost of this expansion is given by

$$\begin{aligned} \text{Marginal Cost} &= -sp'_I(s) - p_I(s) + \frac{d}{ds} [sAC(s)] \\ &= -sp'_I(s) - p_I(s) + C(s) \\ &= -s \frac{\partial D}{\partial s}(0, s) - D(0, s) + C(s) \end{aligned}$$

So, for any value of s where $\frac{\partial D}{\partial s} \neq 0$, the MVPF using ex-post willingness to pay is given by

$$MVPF^{Ex-Post}(s) = \frac{1}{1 + \frac{C(s) - D(0, s)}{-s \frac{\partial D}{\partial s}(0, s)}}$$

and is by definition infinite if $\frac{C(s) - D(0, s)}{-s \frac{\partial D}{\partial s}(0, s)} < -1$.

Ex-Ante MVPF Now, consider the MVPF where marginal willingness to pay is measured from behind the veil of ignorance. This is given by

$$MVPF^{Ex-Ante}(s) = \frac{\frac{\partial \tilde{W}}{\partial \tilde{s}}|_{\tilde{s}=s}}{-s \frac{\partial D}{\partial s}(0, s) - D(0, s) + C(s)}$$

where $\tilde{W}(\tilde{s}, s)$ is given by the solution to:

$$V(p_I(\tilde{s}), 0) = V\left(p_I(s) - \tilde{W}(\tilde{s}, s), -\tilde{W}(\tilde{s}, s)\right)$$

where $\tilde{W}(s, s) = 0$ by construction.

Differentiating with respect to \tilde{s} and evaluating at $\tilde{s} = s$ yields:

$$p'_I(s) \frac{\partial V}{\partial p_I}(p_I(s), 0) = \frac{\partial \tilde{W}}{\partial \tilde{s}}|_{\tilde{s}=s} \left[s \frac{-\partial V}{\partial p_I}(p_I(s), 0) + (1-s) \frac{-\partial V}{\partial p_U}(p_I(s), 0) \right]$$

So that

$$\frac{\partial \tilde{W}}{\partial \tilde{s}}|_{\tilde{s}=s} = \frac{p'_I(s) \frac{\partial V}{\partial p_I}(p_I(s), 0)}{s \frac{-\partial V}{\partial p_I}(p_I(s), 0) + (1-s) \frac{-\partial V}{\partial p_U}(p_I(s), 0)}$$

Now, note that the envelope theorem implies that the marginal welfare impact of higher

prices is given by the average marginal utility of consumption of the insured:

$$-\frac{\partial V}{\partial p_I}(p_I(s), 0) = sE \left[-\frac{\partial u_\nu}{\partial c} | d(p_U(s), \theta_\mu) > p_I(s) \right]$$

and the average ex-ante welfare impact of a marginal transfer to both the uninsured and insured state of the world is given by the average marginal utility of consumption for those in the market, $\theta_\mu \in M$:

$$s \frac{-\partial V}{\partial p_I}(p_I(s), 0) + (1-s) \frac{-\partial V}{\partial p_U}(p_I(s), 0) = E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right]$$

Hence,

$$\begin{aligned} \frac{\partial \tilde{W}}{\partial \tilde{s}} \Big|_{\tilde{s}=s} &= -sp'_I(s) \frac{E \left[\frac{\partial u_\nu}{\partial c} | d(p_U(s), \theta_\mu) > p_I(s) \right]}{E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right]} \\ &= -sp'_I(s) \frac{E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right] + E \left[\frac{\partial u_\nu}{\partial c} | d(p_U(s), \theta_\mu) > p_I(s) \right] - E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right]}{E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right]} \\ &= -sp'_I(s) \frac{E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right] + (1-s) \left(E \left[\frac{\partial u_\nu}{\partial c} | d(p_U(s), \theta_\mu) > p_I(s) \right] - E \left[\frac{\partial u_\nu}{\partial c} | d(p_U(s), \theta_\mu) \leq p_I(s) \right] \right)}{E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right]} \\ &= -sp'_I(s) [1 + \beta(s)(1-s)] \end{aligned}$$

where $\beta(s) = \frac{E \left[\frac{\partial u_\nu}{\partial c} | d(p_U(s), \theta_\mu) > p_I(s) \right] - E \left[\frac{\partial u_\nu}{\partial c} | d(p_U(s), \theta_\mu) \leq p_I(s) \right]}{E \left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M \right]}$. Hence, whenever $\frac{\partial D}{\partial s}(0, s) \neq 0$, one has:

$$MVPF^{Ex-Ante}(s) = \frac{1 + \beta(s)(1-s)}{1 + \frac{C(s) - D(0, s)}{s \frac{\partial D}{\partial s}(0, s)}}$$

C Proof of Propositions 3 and 4

This appendix provides proofs for Proposition 3 and 4.

Proposition 3 Recall

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

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

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

and

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

So,

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

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

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

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

Proposition 4 Proposition 4 provides an expression for $\beta(s)$ that uses the demand curve to proxy for consumption. To arrive at this expression, equation (6) implies that for any p_U we have the equation:

$$E \left[u_\nu(y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | S(\theta_\mu) = s \right] = E \left[u_\nu(y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - D(p_U, S(\theta_\mu)) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | S(\theta_\mu) = s \right]$$

where now I write $S(\theta_\mu)$ in place of $S(p_U(s), \theta_\mu)$ because of the assumption of no income effects. Next, taking a derivative with respect to p_U yields:

$$\frac{\partial}{\partial p_U} E \left[u_\nu(y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | S(\theta_\mu) = s \right] = \frac{\partial}{\partial p_U} E \left[u_\nu(y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - D(p_U, S(\theta_\mu)) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | S(\theta_\mu) = s \right]$$

And, imposing the no income effects assumption so that $\frac{\partial m^I}{\partial p_I} = \frac{\partial m^U}{\partial p_U} = 0$ and $\frac{\partial D}{\partial p_U} = 0$, this yields:

$$E \left[-\frac{\partial u_\nu}{\partial c} (y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | S(\theta_\mu) = s \right] = E \left[-\frac{\partial u_\nu}{\partial c} (y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - D(S(\theta_\mu)) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | S(\theta_\mu) = s \right]$$

Now, integrating over all uninsured types, the average average marginal utility of income for the uninsured can be expressed as:

$$E \left[\frac{\partial u_\nu}{\partial c} (y_\nu(\theta_\nu) - m_\nu^U(\theta_\nu) - p_U, m_\nu^U(\theta_\nu); \theta_\nu) | S(\theta_\mu) \geq s \right] = E \left[\frac{\partial u_\nu}{\partial c} (y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - D(S(\theta_\mu)) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | S(\theta_\mu) \geq s \right]$$

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 and taking a Taylor expansion yields,

$$\begin{aligned} E\left[\frac{\partial u_\nu}{\partial c}(y_\nu(\theta_\nu) - m_\nu^I(\theta_\nu) - p_U, m_\nu^I(\theta_\nu); \theta_\nu) | S(\theta_\mu) \geq s\right] &= E[f(y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - D(S(\theta_\mu)) - p_U) | S(\theta_\mu) \geq s] \\ &\approx f(c^*) + f'(c^*) \left(E[y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - D(S(\theta_\mu)) - p_U | S(\theta_\mu) \geq s] - c^*\right) \end{aligned}$$

A similar expression holds for the insured so that

$$\begin{aligned} &E\left[\frac{\partial u_\nu}{\partial c} | S(\theta_\mu) < s\right] - E\left[\frac{\partial u_\nu}{\partial c} | S(\theta_\mu) \geq s\right] \\ &\approx f'(c^*) \left[E[y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - D(s) - p_U | S(\theta_\mu) < s] - E[y_\nu(\theta_\nu) - x(m_\nu^I(\theta_\nu)) - D(S(\theta_\mu)) - p_U | S(\theta_\mu) \geq s]\right] \\ &= -f'(c^*) [\Delta y + \Delta x + \Delta D] \end{aligned}$$

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,

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

is the difference in out of pocket spending between the insured and uninsured using the choice of m that the uninsured would make if they were insured, $m_\nu^I(\theta_\nu)$, and

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

is the difference between the willingness to pay of the marginal type (i.e. $p_I - p_U$), and the average willingness to pay for the uninsured.⁴⁹

Setting c^* so that $E\left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M\right] = f(c^*)$ yields

$$\begin{aligned} \beta(s) &= \frac{\left(E\left[\frac{\partial u_\nu}{\partial c} | S(\theta_\mu) < s\right] - E\left[\frac{\partial u_\nu}{\partial c} | S(\theta_\mu) \geq s\right]\right)}{E\left[\frac{\partial u_\nu}{\partial c} | \theta_\mu \in M\right]} \\ &= \frac{-f'(c^*) [\Delta y + \Delta x + \Delta D]}{f(c^*)} \\ &= \gamma (\Delta y + \Delta x + \Delta D) \end{aligned}$$

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

⁴⁹Note that the first term, $D(s)$, is not $E[D(S(\theta_\mu)) | S(\theta_\mu) < s]$ because all insured individuals pay $D(s)$ when a fraction s of the market is insured.

D 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 in 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 assumptions that the utility function satisfies no income effects and that $\frac{\partial^2 u}{\partial m^2} = \frac{\partial^2 u}{\partial m \partial c} = 0$, then it is straightforward to show that the coefficient of absolute risk aversion is given by:

$$\gamma(s) = 2 \frac{D(s) - C(s) + (1 - p(s)) E [m^I(\theta_\nu) - m^U(\theta_\nu) | S(\theta_\mu) = s]}{V(s)} \quad (22)$$

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 the reduction in variance of consumption offered by the insurance which is approximated by:

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

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 .⁵⁰ If $p(s) < 1$, some of this additional cost that is imposed on the insurer will not be fully valued by the individual.

⁵⁰To see this, suppress all notation w.r.t. θ_ν for brevity. Let $(\bar{c}, \bar{m}, \bar{\theta})$ denote the average bundle of a type with $S(\theta_\mu) = s$. Note that under the no income effects condition (and aggregating across all θ_μ such that $S(\theta_\mu) = s$), equation (6) can be re-written by substituting in the optimized utility function

$$E [u(y - D(s) - x(m^I) - p_U, m^I) | S(\theta_\mu) = s] = E [u(y - m^U - p_U, m^U) | S(\theta_\mu) = s]$$

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.

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

The no income effects assumption allows the LHS to continue to depend on the medical spending chosen by the insured, m^I (no income effects ensures that m^I is unchanged when individuals pay $D(s)$).

Taking a second-order Taylor expansion and imposing $\frac{\partial^2 u}{\partial m \partial c} = 0$ and $\frac{\partial^2 u}{\partial m^2} = 0$ yields

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

Re-writing,

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

Or,

$$\left(E \left[m^U - x(m^I) - D(s) | S(\theta_\mu) = s \right] \right) = \frac{1}{2} \frac{u_{cc}}{u_c} V(s) + \frac{u_m}{u_c} E \left[m^U - m^I | S(\theta_\mu) = s \right]$$

or

$$D(s) + E \left[x(m^I) | S(\theta_\mu) = s \right] - E \left[m^U | S(\theta_\mu) = s \right] = \frac{-1}{2} \frac{u_{cc}}{u_c} V(s) + \frac{u_m}{u_c} E \left[m^I - m^U | S(\theta_\mu) = s \right]$$

and noting that $C(s) = E \left[m^I - x(m^I) | S(\theta_\mu) = s \right]$ yields

$$D(s) - C(s) + E \left[m^I - m^U | S(\theta_\mu) = s \right] = \frac{-1}{2} \frac{u_{cc}}{u_c} V(s) + \frac{u_m}{u_c} E \left[m^I - m^U | S(\theta_\mu) = s \right]$$

or

$$D(s) - C(s) = \frac{\gamma(s)}{2} V(s) - \left(1 - \frac{\partial x}{\partial m} \right) E \left[m^I - m^U | S(\theta_\mu) = s \right]$$

where $\frac{\partial x}{\partial m}$ is the price paid by an insured individual with average medical spending, $E \left[m^I | S(\theta_\mu) = s \right]$. So,

$$\frac{\gamma(s)}{2} V(s) = D(s) - C(s) + \left(1 - \frac{\partial x}{\partial m} \right) E \left[m^I - m^U | S(\theta_\mu) = s \right]$$

or

$$\gamma(s) = 2 \frac{D(s) - C(s) + \left(1 - \frac{\partial x}{\partial m} \right) E \left[m^I - m^U | S(\theta_\mu) = s \right]}{V(s)}$$

concluding the proof.

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.⁵¹ 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 (13):

$$\beta(s, x) = \frac{E\left[\frac{\partial u_V}{\partial c} | S(\theta_\mu) < s, X = x\right] - E\left[\frac{\partial u_V}{\partial c} | S(\theta_\mu) \geq s, X = x\right]}{E\left[\frac{\partial u_V}{\partial c} | \theta_\mu \in M, X = x\right]}$$

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 $X = x$ is given by

$$EA(s, x) = (1 - s_x) s \left(-\frac{\partial D}{\partial s} \right) \beta(s_x, x) \quad (23)$$

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

Recall this proposition states:

Proposition. *Under the modeling assumptions outlined above, $W(s)$ in equation (9) is given*

⁵¹If 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.

by

$$W(s) = \frac{-1}{\gamma} \left[\log \left(s e^{\gamma p_I(s)} e^{-\gamma \frac{w}{2}} + (1-s) e^{\gamma p_U(s)} \int_s^1 e^{\gamma(C(\tilde{s})-w) + \frac{\gamma^2}{2} \Sigma(\tilde{s})} d\tilde{s} \right) - \log \left(\int_0^1 \left[e^{\gamma(C(\tilde{s})-w) + \frac{\gamma^2}{2} \Sigma(\tilde{s})} \right] d\tilde{s} \right) \right]$$

Because of the absence of savings behavior, one can focus exclusively on the period $t = \nu$ utility. Note that the expected utility in the economy when a fraction s is insured, as specified in equation (8) is given by:

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

So, $W(s)$ solves

$$-\frac{1}{\gamma} E \left[e^{-\gamma[y - \lambda(\theta_\nu) + W(s)]} \right] = s E \left[-\frac{1}{\gamma} e^{-\gamma[y - p_I(s) + \frac{1}{2}w]} | S(\theta_\mu) < s \right] + (1-s) E \left[-\frac{1}{\gamma} e^{-\gamma[y - p_U(s) - \lambda(\theta_\nu)]} | S(\theta_\mu) \geq s \right] \quad (24)$$

For fixed risk aversion, this implies:

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

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

This Appendix provides an expression for $W(s)$ in the presence of heterogeneity in $y(\theta_\nu)$. To do so, note that with heterogeneity in income, equation (24) becomes:

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

or

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

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

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

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

This section discusses the solution to the model in the presence of heterogeneity in risk aversion. I begin by discussing the estimation of the model for the high and low heterogeneity specifications where $\gamma \sim N(\mu_\gamma, \sigma_\gamma^2)$. 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 even conditional on their willingness to pay quantile, s .

I proceed as follows. First, let 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 Σ_i denote their subjective variance of their expected costs (e.g. for an individual i with knowledge θ_μ at time $t = \mu$, their subjective variance is $Var(\lambda(\theta_\nu) | \theta_\mu)$). I continue to assume moral hazard is constant (i.e. no variation in w). So, demand is given by

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

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 Σ_i). So, while previously in subsection 7.4.2 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 $\bar{\gamma} = 5 \times 10^{-4}$ and $\sigma_\gamma \in \{5 \times 10^{-5}, 1 \times 10^{-4}\}$. 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 from the distribution of $\{C(s), V(s)\}$ estimated in the model with no heterogeneity. To account for the presence of cost variation conditional on s , I next add an error term κ_i to the expected costs, so that

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

where κ_i is drawn from a distribution $N(0, \sigma_c^2)$ and $C(U_i)$ is drawn by taking a draw from a uniform distribution, U_i .

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. Given these estimates, I then calculate $W(s)$ using equation (9).⁵²

I Discontinuous Willingness to Pay

This section uses the structural model developed in Section 7 to illustrate the potential biases that can arise when demand is discontinuous. The results show that estimates of $EA(s)$ can be biased in regions of s near a discontinuity in the market willingness to pay curve, but perform generally quite well away from the discontinuity. And, aggregating across s to form $W(s)$ generally approximates the true $W(s)$ quite well.

To illustrate, I apply the model to $\tilde{D}(s)$ that is given by

$$\tilde{D}(s) = D(s) + 0.25 * 1\{s < 0.5\} E[D(s)]$$

so that willingness to pay is $0.25 * E[D(s)] = \$375$ higher for those with below-median willingness to pay. This introduces a discontinuity in modify the baseline specification by adding a term $1\{s < 0.5\} * E[D(s)]$ to all that $D(s)$.

Given $\tilde{D}(s)$, I estimate the baseline structural model in Section 7 by fitting the model to $\tilde{D}(s)$. And, I also construct a 10th order polynomial approximation to the $\tilde{D}(s)$ curve to simulate what would happen if an econometrician attempted to empirically approximate this curve and apply the methods proposed in this paper to measure ex-ante welfare. The results are shown in Figure A1.

⁵²Because γ is now heterogeneous in the population, income (y) does not drop out of the estimation. For simplicity, I assume $y = \$10,000$. The general result that higher levels of risk aversion leads to lower ex-ante welfare does not depend on this choice for income.

Figure A1: Discontinuous Willingness to Pay




Figure A1A plots the value of true $W(s)$ implied by the structural model for all s alongside the benchmark approximation to $W(s)$ and market surplus. This shows that overall, the aggregated approximation of $W(s)$ formed using the baseline approach performs quite well even when one aggregates over the point of discontinuity at $s = 0.5$. Figure A1B compares the marginal ex-ante willingness to pay from the structural model, $W'(s)$, to the baseline approximation and to market surplus. Here, the true $W(s)$ is not differentiable near $s = 0.5$ so that polynomial approximations can be biased in a region near that value. But, away from the point of discontinuity, the approximations to $W'(s)$ perform similarly well as in the baseline specification.

J 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 measures 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 directly incorporate estimates of the complementarities between health status and the marginal utility of income to accurately recover $\beta(s)$ even when the marginal utility of income varies with health status, conditional on the level of consumption.