Estimating the Value of Public Insurance Using Complementary Private Insurance*

Marika Cabral and Mark R. Cullen †

November 7, 2016

Abstract

The welfare associated with public insurance is often difficult to quantify. Relative to private insurance, a fundamental difficulty is that public insurance is typically compulsory, so the demand for coverage is unobserved and thus cannot be used to analyze welfare. However, in many public insurance settings, individuals can purchase private insurance to supplement their public coverage. In this paper, we outline an approach to use data and variation from private complementary insurance to quantify welfare associated with several counterfactuals related to compulsory public insurance. Using administrative data from one large firm on employee long-term disability insurance, we then apply this approach empirically to quantify the value of disability insurance among this population. We use premium variation among the employer-provided disability policies to quantify the surplus that would be generated by increasing the replacement rate of disability insurance for our sample population—a counterfactual that is within the set of insurance contracts observed in this setting. In addition, we estimate a lower bound on the surplus generated by public disability insurance in this context. Our findings suggest that public disability insurance generates substantial surplus for this population, and there may be gains to increasing the generosity of coverage in this context.

*We thank David Beheshti and Frank Martin-Buck for their excellent research assistance. For providing helpful comments, we thank Michael Geruso, Amanda Kowalski, Neale Mahoney, Mark Shepard, and seminar participants at Carnegie Mellon Heinz College, Indiana University Kelley School of Business, University of Texas at Austin, the 2016 Chicago Booth Whistler Junior Health Economics Summit, the 2016 IFN Stockholm conference, the 2016 NBER Summer Institute Public Economics Meeting, and the 2016 STATA Texas Empirical Micro Conference. The Alcoa data were provided as part of an ongoing service and research agreement between Alcoa, Inc. and academic institutions to perform jointly agreed-upon ongoing and ad hoc research projects on workers' health, injury, disability, and health care, and research using this data is supported by a grant from the National Institute on Aging (Disease, Disability and Death in an Aging Workforce, NIH/NIA, 1 R01 AG026291). Mark R. Cullen serves as Senior Medical Advisor for Alcoa, Inc.

†Cabral: University of Texas at Austin Department of Economics and NBER, marika.cabral@austin.utexas.edu. Cullen: Stanford School of Medicine and NBER.

Social insurance programs are ubiquitous and cover many of the largest risks individuals face. To determine the welfare generated by public insurance, it is important to quantify both the benefits and the costs of this insurance. While the costs of public insurance are relatively straightforward to calculate, the benefits of public insurance are often difficult to quantify. A fundamental difficulty is that public insurance is typically compulsory, so the demand for this non-market good is unobserved and thus cannot be used to infer the value individuals place on coverage. While a relatively large literature has worked to quantify welfare in private insurance markets, this difficulty may explain why there have been far fewer studies investigating the welfare associated with public insurance.

Several recent studies investigate welfare within private insurance markets (e.g., Hackmann, Kolstad and Kowalski (2015), Bundorf, Levin and Mahoney (2012), Einav, Finkelstein and Cullen (2010)). Although recent studies investigating welfare in private insurance settings use a range of methods, the commonality among these studies is that they use price variation to identify the demand for insurance and thus the value individuals place on coverage.[1] Perhaps because of the unique challenges that arise in the setting of compulsory public insurance, a recent literature on the welfare generated by compulsory public insurance has evolved independently of the literature on welfare in private insurance settings. A handful of recent studies have analyzed welfare within public insurance settings, such as Medicaid (Finkelstein, Hendren and Luttmer (2015)), unemployment insurance (e.g., Chetty (2008), Gruber (1997)), and disability insurance (Low and Pistaferri (2015), Chandra and Samwick (2005)). While these studies employ a range of approaches across these settings, all of these studies confront the fact that demand for compulsory public coverage is unobserved by making several critical assumptions on the nature of individual utility (or marginal utility), uncertainty, and heterogeneity to estimate welfare, often employing data on consumption or assets.[2] Importantly, these studies make several assumptions that can meaningfully affect the resulting welfare analysis and are often difficult to empirically validate.

In this paper, we propose a complementary approach to analyzing the welfare associated with public insurance that leverages the fact that individuals can often purchase supplemental private insurance to top-up the benefits of compulsory public insurance. The basic idea behind this approach is simple: The existence of complementary private insurance allows us to extend standard willingness-to-pay approaches used in private insurance settings and apply these to welfare questions associated with compulsory public insurance. Relative to other approaches to value compulsory public insurance, the approach outlined in this paper requires minimal assumptions and can be implemented using commonly available data and variation.

While the existence of parallel private and public insurance is a necessary condition to implement this approach, it turns out that this is a common feature of many social insurance programs in the United States and abroad.[3] Some of the largest social insurance programs in the United States have this feature including Medicare (private Medigap insurance), Social Security disability insurance (private long-term disability insurance), and Social Security retirement benefits (private annuities). The model of complementary public and private insurance is also very common outside of the United States, particularly in the context of universal

[1] See Einav, Finkelstein and Levin (2010) for a review of the literature on welfare analysis in private insurance markets and a discussion of the various empirical methods employed in this literature.

[2] See Chetty and Finkelstein (2013) for a review of the recent literature on social insurance. In addition, see Finkelstein, Hendren and Luttmer (2015) for a discussion of some of the trade-offs between modeling assumptions and data requirements in conducting welfare analysis in compulsory public insurance settings.

[3] Despite that fact that private insurance is often available to top-up compulsory public insurance, most of the prior literature on welfare analysis in compulsory public insurance settings abstracts from opportunities individuals have to purchase private complementary coverage. A notable exception is Chetty and Saez (2010) who characterize the impact of endogenous private insurance on the welfare associated with government intervention under a range of different modeling assumptions.

public health insurance. Many countries that provide universal public health insurance also allow individuals to purchase complementary private health insurance coverage to top-up the incomplete public health insurance benefits including France, the Netherlands, Canada, Denmark, Japan, Switzerland, New Zealand, Italy, England, Norway, and Sweden. [4]

Intuitively, in settings in which individuals can purchase private complementary insurance, coverage decisions individuals make in the private supplemental insurance market can inform us about individuals' willingness-to-pay for extending the generosity of the compulsory coverage to include the benefits of supplemental insurance, and, in some settings, these decisions can inform us about individuals' underlying valuation of the inframarginal compulsory public coverage. We formalize this intuition by developing a theoretical framework which illustrates that data and variation from the market for supplemental insurance can be used to quantify welfare associated with several policy-relevant counterfactuals related to compulsory public insurance. We then apply this framework to the context of disability insurance using administrative data on enrollment and claims from one large employer that provides its employees long-term disability insurance that supplements the wage replacement benefits of public disability insurance.

The paper begins by outlining a framework that can be used to evaluate welfare in the setting of complementary public and private insurance. The modeling approach builds upon prior work by Einav and Finkelstein (2011) and Einav, Finkelstein and Cullen (2010). The key empirical inputs into the welfare analysis are the demand and cost curves associated with the private supplemental insurance market. We describe how these demand and costs curves can be used to evaluate a wide range of counterfactuals related to compulsory public insurance, including marginal counterfactuals (related to the incremental coverage sold in the existing market for supplemental insurance) and inframarginal counterfactuals (related to the inframarginal coverage provided by the compulsory public insurance).

For instance, consider the welfare associated with a particular marginal counterfactual: an extension of the generosity of compulsory public insurance to include the coverage provided by private supplemental insurance. The existence of a private supplemental insurance market provides the opportunity to study the welfare from extending the generosity of public insurance to incorporate the coverage provided by supplemental insurance because we effectively observe a market for the extension of this public coverage. Thus, we illustrate how a straightforward extension of the Einav and Finkelstein (2011) framework can be used to investigate the welfare associated with an expansion of the generosity of compulsory public insurance, where we incorporate key features that can matter in this setting such as externalities induced by the supplemental insurance and crowd-out of the private supplemental insurance market.

Beyond estimating the welfare generated by a marginal extension of public insurance, one might also be interested in broader counterfactuals that involve quantifying the welfare associated with insurance coverage for which we cannot directly estimate demand. Within the setting of social insurance, a particular broader counterfactual of interest is quantifying the total welfare generated by the social insurance program relative to the absence of insurance for this risk. In the recent literature on welfare analysis within private insurance markets, one common approach to this type of inframarginal welfare analysis is to specify a structural model of decision-making, use empirical variation in premiums and other product attributes to estimate the primitives of the model, and then use the fitted model to investigate broader welfare questions of interest. While this basic structural approach could be applied in settings with complementary private and public insurance, a main disadvantage of this approach is that it generally requires a lot of assumptions on the form of utility,

[4] In some of these countries, private supplemental insurance covers the cost-sharing associated with services partially covered by the compulsory public insurance; in other countries, private supplemental insurance covers complementary services not covered by the compulsory public insurance (e.g., drugs, dental, out-of-network doctors/hospitals). For a detailed overview of parallel public and private health insurance internationally, see Thomson et al. (2013).

the distribution of risk individuals face, and the degree of heterogeneity across individuals.

We outline an alternative, complementary approach to study the total welfare generated by compulsory public insurance, which relies on the demand and cost curves associated with supplemental insurance (the same objects required to implement the marginal welfare analysis described above). Looking toward the empirical application to disability insurance, we focus attention on the case where public and private insurance provide coverage that is linear in the insured loss. We then derive a lower bound on the total surplus generated by the inframarginal compulsory public insurance coverage in terms of the demand and cost curves that can be estimated within the supplemental insurance market. This derivation relies on minimal assumptions on the nature of individual preferences, and we provide a simple and intuitive sufficient condition applicable when preferences are represented by a univariate utility function: utility exhibits non-increasing absolute risk aversion. Note that this property is satisfied by most common utility functions used by empiricists, including CARA, CRRA, and a broader class of empirically relevant HARA utility functions. Importantly, while this approach requires some minimal restrictions on the nature of individual preferences, this bounding approach does not require specifying a particular utility function, specifying the distribution of risk individuals face, or restricting heterogeneity across individuals. Thus, relative to a more structural approach, the key advantage of this approach is that it requires many fewer assumptions, while the disadvantage is that it delivers a bound rather than a precise welfare number.

Using administrative data from one large firm on employee long-term disability (LTD) insurance, we then apply this approach empirically to quantify the value of the wage replacement benefits of disability insurance among this population. Disability insurance is a particularly important setting in which to understand the welfare associated with insurance, as the threat of a career-ending disability is one of the largest financial risks many individuals face, and the public Social Security Disability Insurance (SSDI) program is one of the largest social insurance programs in the United States. In 2014, approximately 9 million disabled workers received SSDI benefits, and the total SSDI benefits paid exceeded $140 billion.[5] Despite the large size of this social insurance program, there has been very little research quantifying the welfare provided by disability insurance.[6,7] In addition to the coverage available to workers through SSDI, 34% of US workers have the opportunity to purchase private supplemental LTD coverage through their employer to top-up the benefits of SSDI.[8] A typical empirical challenge to investigating the welfare associated with public SSDI coverage (relative to other social insurance programs) is that it is a national program with little variation in coverage across workers. Thus, disability insurance is a natural context to apply the framework described above as it allows us to use variation within employer-provided private supplemental LTD insurance to overcome this key empirical challenge and analyze the welfare associated with disability insurance.

[5]Source: `https://www.ssa.gov/policy/docs/statcomps/supplement/2015/highlights.pdf`.

[6]While there has been limited research on the consumption-smoothing benefits of disability insurance, there are a small number of recent related studies. A few recent papers document consumption changes in response to disability onset or major health shocks: Meyer and Mok (2013) document changes in consumption and income that follow a change in self-reported disability status using PSID data; Kostol and Mogstad (2015) document changes in income and consumption among denied and approved disability insurance applicants in Norway; Dobkin et al. (2016) document consumption changes following major hospitalizations, providing indirect evidence on individuals' exposure to disability-induced earnings losses. In addition, a few papers have investigated the ex ante welfare provided by disability insurance using calibrated life-cycle models (see Low and Pistaferri (2015), Chandra and Samwick (2005)). In another recent study, Autor, Kostol and Mogstad (2015) employ a structural model along with earnings and consumption data to estimate Norwegian disability insurance applicants' implied willingness-to-pay to be approved for disability insurance. In contrast to these few related papers which rely on consumption data and/or calibrated life-cycle models to infer the value of disability insurance, the approach in the present paper employs data and variation from the market for supplemental disability coverage to directly investigate the willingness-to-pay for disability insurance and the welfare associated with this coverage.

[7]While there has been relatively little research on the welfare associated with disability insurance, there has been extensive research on the causes and consequences of the vast growth in disability insurance rolls over time. See Autor and Duggan (2006) and Liebman (2015) for a review of this literature.

[8]These statistics are reported in Table 16 of the U.S. Department of Labor's National Compensation Survey, Employee Benefits in the U.S., March 2015 publication at: `http://www.bls.gov/ncs/ebs/benefits/2015/ownership/civilian/table16a.pdf`.

We apply the welfare framework by leveraging premium variation and the subsequent decisions employees make in the context of LTD insurance at one large firm, Alcoa, Inc. There are several nice features of the data and environment for this empirical application. First, the firm offers employees three vertically differentiated plans with wage replacement rates of 50%, 60%, and 70% in the event of disability. Second, the basic 50% replacement rate plan is free for all employees. This is convenient as no one opts out of the plans, and thus claims data is available for all employees. Third, there is variation over time in the incremental premium for the highest generosity plan. This variation allows us to trace out the relative demand curves for an incremental 10% replacement rate, starting from a baseline 60% replacement rate. Fourth, the administrative data include information on disability claims in addition to disability insurance enrollment, so the demand estimates can be paired with cost data to evaluate welfare. Lastly, the firm's disability plans explicitly require workers to apply to SSDI and, if approved, SSDI benefits crowd-out LTD benefits dollar for dollar. In this way, the LTD plans at this firm top-up the compulsory public insurance, providing a natural interpretation to our estimates.

Using premium variation among the LTD policies available to employees, we find that the demand for supplemental disability coverage is price-sensitive. Our baseline estimates indicate that if the premium for the most generous LTD plan increases by 0.1% of annual earnings, enrollment would decline by 7 percentage points. This is precisely estimated and robust to controlling for time trends and individual fixed-effects. Based on our estimates, the implied mean willingness-to-pay for the incremental 10% replacement rate (starting from the baseline of 60% wage replacement) is roughly 0.3% of annual earnings, or $202 annually for a worker earning $65,000 a year (roughly the mean annual earnings in the population). When scaling this by the costs of providing this incremental coverage, we find that on average individuals value this incremental coverage at more than two and a half times the cost of providing this incremental coverage, indicating that the incremental disability coverage moving from a 60% to a 70% replacement rate is associated with substantial welfare. While these estimates suggest that the benefits of disability coverage on the margin far outweigh the costs for this population, to evaluate the welfare associated with an extension of compulsory public coverage one would need to account for any potential crowd-out of private supplemental insurance. In the empirical context, crowd-out of the private LTD insurance available within the firm reduces the value of an extension of compulsory coverage substantially.

In addition to evaluating the welfare associated with a marginal extension of disability coverage, we also use the estimates to obtain a lower bound on the surplus generated by the inframarginal public compulsory disability insurance relative to the absence of insurance for this risk. Because the estimates indicate that individuals highly value disability coverage on the margin, we obtain a meaningful lower bound on the value of the inframarginal coverage. Based on these estimates, the total surplus generated by public insurance is at least 0.7% of annual earnings, or $440 annually for a worker with annual earnings of $65,000. It is important to emphasize that these estimates come from a particular population that is not representative of a broader set of workers. However, based on these estimates, workers in this population seem to highly value the inframarginal disability coverage offered by public disability insurance, and, if it were not for the fact that this firm offers fairly priced supplemental coverage, there would be significant welfare gains from extending the generosity of public coverage for these workers. We demonstrate that these results are robust to a wide range of alternative specifications.

While the empirical results are specific to the population we examine, our analysis also highlights some general strengths of this approach for evaluating the welfare generated by compulsory public insurance. First, this method requires straightforward variation and limited data to be implemented empirically. Implementing this approach simply requires sufficient variation in the price of supplemental coverage to estimate

the demand curve for supplemental coverage and standard data on prices, insurance enrollment, and costs. In contrast, other approaches to valuing public insurance often require isolating random (or like-random) assignment of benefits within the compulsory public insurance program and often require data on consumption or assets (data that is typically difficult to obtain and can be subject to measurement error). Second, this approach does not require us to specify many of the underlying primitives of individuals' decisions. The fact that this approach requires so few assumptions is in stark contrast to other approaches of valuing compulsory public coverage which rely on fully specifying individual utility (or marginal utility), obtaining data on all the inputs in the utility (or marginal utility), specifying the uncertainty individuals face, and/or assumptions on the nature of individuals' optimization. Third, this approach is flexible enough to be applied in a broad range of settings. It is straightforward to apply this approach to any setting with complementary public and private insurance to estimate the welfare associated with extending the generosity of public insurance to include the benefits provided by complementary private insurance. In settings with linear coverage, this approach can be used to go beyond marginal welfare questions and bound the welfare generated by the inframarginal compulsory public coverage under minimal assumptions. Of course, this approach is not without limitations. A primary limitation of this approach is that it can only be applied in settings in which individuals are permitted to buy supplemental insurance coverage. While there are many public insurance settings in which individuals can buy complementary coverage (e.g., disability insurance, universal health insurance settings such as Medicare, etc.), there are some public insurance settings in which no complementary coverage is typically available (e.g., unemployment insurance, means-tested public insurance such as Medicaid), and in these settings this approach cannot be applied.

The remainder of the paper proceeds as follows. Section 1 describes the framework for evaluating welfare. Section 2 presents background on the empirical application and describes the data. Section 3 describes the empirical strategy, and Section 4 presents the empirical results and implied welfare analysis. Lastly, Section 5 concludes.

1 Framework

1.1 Model

Setup and Notation The setup of our framework draws heavily upon prior work by Einav and Finkelstein (2011) and Einav, Finkelstein and Cullen (2010). Suppose there is a population of heterogeneous individuals, indexed by φ. Let $G(\varphi)$ represent the distribution of the population. An important aspect of this approach is that heterogeneity across individuals is unrestricted. Importantly, this means φ may be multi-dimensional to include variation in preferences and risk, and there are no restrictions on G.

Individuals each face some downside disposable income risk. Looking forward to our empirical application, let us consider linear insurance products where the generosity of insurance is indexed by the fraction of the risk covered by the insurance.[9] Specifically, consider a supplemental insurance product that provides δ generosity coverage on top of baseline public compulsory insurance coverage of generosity α, so that an individual who holds both the baseline and supplemental insurance has total coverage of generosity $\alpha + \delta$. Let $\pi(\theta, \gamma|\varphi)$ represent an individual's willingness-to-pay for coverage of generosity θ relative to the outside option of coverage of generosity γ. In this notation, we can represent the willingness-to-pay for supplemental coverage ($\pi(\alpha + \delta, \alpha|\varphi)$) and the willingness-to-pay for the compulsory baseline coverage ($\pi(\alpha, 0|\varphi)$). Let $c(\theta, \gamma|\varphi)$ represent the expected cost borne by insurance providers for an individual with coverage of gen-

[9]While the framework is described in terms of linear contracts, all of the welfare analysis within the "marginal counterfactuals" can easily be extended to non-linear insurance settings as discussed further below.

erosity θ relative to coverage of generosity γ. In this notation, we can describe the expected cost associated with an individual buying supplemental coverage as the sum of two components,

$$c(\alpha+\delta,\alpha|\varphi) = c^S(\alpha+\delta,\alpha|\varphi) + c^P(\alpha+\delta,\alpha|\varphi), \tag{1}$$

where $c^S(\alpha+\delta,\alpha|\varphi)$ is the expected cost to the supplemental insurer for providing this incremental coverage of generosity δ to the individual, and $c^P(\alpha+\delta,\alpha|\varphi)$ is the expected external cost associated with the individual having this incremental coverage which are borne by the primary insurer.[10] In the special case where the supplemental coverage is not associated with an externality on the primary insurer (with linear contracts, this is equivalent to the case of no moral hazard), then there is no external cost for the primary insurer, $c^P(\alpha+\delta,\alpha|\varphi)=0$, and thus the total expected cost is simply the expected cost to the supplemental insurer, $c(\alpha+\delta,\alpha|\varphi) = c^S(\alpha+\delta,\alpha|\varphi)$.

Supplemental Insurance Market: Demand, Costs, and Equilibrium Suppose each individual makes a discrete choice whether to buy supplemental insurance of generosity δ for (relative) price p or go with the outside option of the basic compulsory insurance of generosity α. Define the demand curve for this supplemental coverage as follows:

$$D(p|\alpha+\delta,\alpha) = \int \mathbb{1}\big(\pi(\alpha+\delta,\alpha|\varphi)\geq p\big)dG(\varphi), \tag{2}$$

where we assume the underlying primitives in the population are such that demand is decreasing, continuous, and differentiable. Throughout we ignore income effects associated with changes in the price of supplemental insurance.[11] As will become clear, the theoretical framework explicitly accommodates income effects associated with a change in the generosity of the inframarginal coverage.

Define the social marginal cost and social average cost curve for an insurance product that provides δ generosity coverage on top of a baseline coverage of generosity α as follows:

$$MC(p|\alpha+\delta,\alpha) = E[c(\alpha+\delta,\alpha|\varphi)|\pi(\alpha+\delta,\alpha|\varphi)=p] \tag{3}$$

$$AC(p|\alpha+\delta,\alpha) = \int c(\alpha+\delta,\alpha|\varphi)\mathbb{1}\big(\pi(\alpha+\delta,\alpha|\varphi)\geq p\big)dG(\varphi) \tag{4}$$

where the cost curves $MC(p|\alpha+\delta,\alpha)$ and $AC(p|\alpha+\delta,\alpha)$ represent the costs inclusive of those incurred by both the supplemental insurer and the primary insurer. We can also define the marginal and average cost curves for the supplemental insurer, who is only responsible for the incremental coverage:

$$MC^S(p|\alpha+\delta,\alpha) = E[c^S(\alpha+\delta,\alpha|\varphi)|\pi(\alpha+\delta,\alpha|\varphi)=p] \tag{5}$$

$$AC^S(p|\alpha+\delta,\alpha) = \int c^S(\alpha+\delta,\alpha|\varphi)\mathbb{1}\big(\pi(\alpha+\delta,\alpha|\varphi)\geq p\big)dG(\varphi) \tag{6}$$

Consider the benchmark case of perfect competition, where we define the equilibrium price P^{CE}, such that supplemental insurers break even on average,

$$P^{CE}(\alpha+\delta,\alpha) = min\{p : p = AC^S(p|\alpha+\delta,\alpha)\}. \tag{7}$$

[10] Throughout the discussion, the cost of insuring an individual is treated as invariant to the insurer's identity. The model could easily be extended to incorporate systematic cost differences across insurers.

[11] Ignoring income effects allows us to use estimates of the Marshallian demand curve for supplemental insurance in the welfare analysis that follows. Abstracting from income effects may be reasonable in settings where the variation in premiums for supplemental insurance is small relative to income (as is the case in our empirical application).

1.2 Welfare Measures

Next, we apply the framework above to describe welfare associated with various counterfactuals. First, we consider welfare associated with "marginal counterfactuals", or counterfactuals related to the incremental coverage sold in the existing market for supplemental insurance. Second, we consider welfare associated with "inframarginal counterfactuals", or broader counterfactuals associated with the inframarginal baseline compulsory insurance.

1.2.1 Marginal Counterfactuals: Welfare Associated with Extending the Generosity of Compulsory Public Insurance

A basic policy-relevant parameter of interest in social insurance settings is the value of extending the generosity of compulsory public insurance. In a setting in which the only insurance available is the baseline compulsory insurance of generosity α, the welfare from extending the generosity of the compulsory insurance coverage can be described using the notation above as,

$$\int \Big(\frac{\partial}{\partial \theta} \big(\pi(\theta, \gamma | \varphi) - c(\theta, \gamma | \varphi) \big) \Big|_{\theta, \gamma = \alpha} \Big) dG(\varphi). \tag{8}$$

In many settings, individuals have access to complementary private insurance. In these settings, there is a direct empirical analog to the objects in the expression above. Notice that the surplus generated by supplemental insurance for an individual, $\pi(\alpha + \delta, \alpha | \varphi) - c(\alpha + \delta, \alpha | \varphi)$, is simply a discretized version of $\frac{\partial}{\partial \theta} \big(\pi(\theta, \gamma | \varphi) - c(\theta, \gamma | \varphi) \big) \Big|_{\gamma, \theta = \alpha}$. In other words, relative to a benchmark of only compulsory insurance of generosity α, the social surplus associated with increasing the generosity of the compulsory coverage by δ for an individual is simply the social surplus associated with that individual purchasing supplemental insurance, $\pi(\alpha + \delta, \alpha | \varphi) - c(\alpha + \delta, \alpha | \varphi)$.

Note that equation 8 measures the welfare associated with a marginal extension of compulsory coverage relative to a benchmark scenario where only compulsory coverage exists (i.e., a world with no private supplemental insurance market). While this is precisely the object of interest in a setting where supplemental insurance is not available or is prohibited by law, we may be interested in measuring the effect of extending the generosity of compulsory coverage relative to a benchmark scenario in which individuals are able to purchase private supplemental insurance themselves. This alternative benchmark scenario is of particular interest because a necessary condition for empirically implementing this approach is the existence of a private supplemental insurance market.

Let $MeanWTP_Extn(p | \alpha + \delta, \alpha)$ and $Welfare_Extn(p | \alpha + \delta, \alpha)$ represent the mean willingness-to-pay and per-capita welfare, respectively, associated with extending compulsory coverage to include the supplemental coverage relative to a benchmark scenario where private supplemental coverage is available for price p:

$$\begin{aligned} MeanWTP_Extn(p | \alpha + \delta, \alpha) &= \int \Big(\pi(\alpha + \delta, \alpha | \varphi) \mathbb{1}\big(\pi(\alpha + \delta, \alpha | \varphi) < p \big) + p \mathbb{1}\big(\pi(\alpha + \delta, \alpha | \varphi) \geq p \big) \Big) dG(\varphi) \\ Welfare_Extn(p | \alpha + \delta, \alpha) &= \int \big(\pi(\alpha + \delta, \alpha | \varphi) - c(\alpha + \delta, \alpha | \varphi) \big) \mathbb{1}\big(\pi(\alpha + \delta, \alpha | \varphi) < p \big) dG(\varphi). \end{aligned} \tag{9}$$

In words, the above expression for welfare tells us that the value of extending compulsory public insurance comes from aggregating individual values among people who would not have otherwise bought the private supplemental coverage.[12]

[12] While the discussion here focuses on describing the mean willingness-to-pay and welfare associated with an extension, this setup

There are a few important points to highlight about the above expressions. First, these expressions nest the case in which no supplemental insurance is available (where effectively the price for supplemental insurance is infinite),

$$\begin{aligned} MeanWTP_Extn(\infty|\alpha+\delta,\alpha) &= \int \pi(\alpha+\delta,\alpha|\varphi)dG(\varphi) \\ Welfare_Extn(\infty|\alpha+\delta,\alpha) &= \int \big(\pi(\alpha+\delta,\alpha|\varphi)-c(\alpha+\delta,\alpha|\varphi)\big)dG(\varphi). \end{aligned} \tag{10}$$

Second, these expressions provide an intuitive way to characterize how crowd-out affects the value of extending the generosity of public coverage. To see this, note that we can decompose these expressions as follows:

$$\begin{aligned} MeanWTP_Extn(p|\alpha+\delta,\alpha) &= MeanWTP_Extn(\infty|\alpha+\delta,\alpha) - MeanWTP_PrivateSupp(p|\alpha+\delta,\alpha) \\ Welfare_Extn(p|\alpha+\delta,\alpha) &= Welfare_Extn(\infty|\alpha+\delta,\alpha) - Welfare_PrivateSupp(p|\alpha+\delta,\alpha), \end{aligned} \tag{11}$$

where, in each expression, the first term represents the value of the extension if there were no private market available to crowd-out, and the second term represents crowd-out of the voluntary private market for supplemental insurance that would have existed in the absence of the compulsory public insurance extension:

$$\begin{aligned} MeanWTP_PrivateSupp(p|\alpha+\delta,\alpha) &= \int \big(\pi(\alpha+\delta,\alpha|\varphi)-p\big)\mathbb{1}\big(\pi(\alpha+\delta,\alpha|\varphi)\geq p\big)dG(\varphi) \\ Welfare_PrivateSupp(p|\alpha+\delta,\alpha) &= \int \big(\pi(\alpha+\delta,\alpha|\varphi)-c(\alpha+\delta,\alpha|\varphi)\big)\mathbb{1}\big(\pi(\alpha+\delta,\alpha|\varphi)\geq p\big)dG(\varphi). \end{aligned} \tag{12}$$

Graphical Illustration We build intuition further through a graphical example in the spirit of Einav and Finkelstein (2011). For the moment, let us abstract from moral hazard meaning that the social and private cost curves are one and the same. Figure 1 Panel (A) plots the demand, marginal cost, and average cost curves where the horizontal axis represents the fraction with supplemental insurance, and the vertical axis is measured in dollars. For the purposes of this example, let us consider a private supplemental insurance market that is characterized by competitive average-cost pricing, and let us assume the market is adversely selected (as depicted in the figure by the downward sloping cost curves).

Consider an extension of the compulsory public coverage to include the benefits of the supplemental insurance coverage. It is straightforward to measure the objects of interest in the context of this example. Relative to a benchmark where no supplemental coverage is available, the mean willingness-to-pay and the per-capita welfare generated by this extension are: $MeanWTP_Extn(\infty|\alpha+\delta,\alpha) =$ Area ABCD and $Welfare_Extn(\infty|\alpha+\delta,\alpha) =$ Area ABCD $-$ Area EFCD. Because a necessary condition to empirically implement this estimation is the existence of a private supplemental insurance market, it may be of particular interest to consider measures of welfare that take into account crowd-out of the private market for supplemental coverage that would otherwise exist. Relative to a benchmark of a competitive private supplemental insurance market, the welfare generated by extending compulsory coverage is: $Welfare_Extn(P^{CE}|\alpha+\delta,\alpha) =$ AreaHGI $-$ Area GFB $= Welfare_Extn(\infty|\alpha+\delta,\alpha) -$ Area AHIE, where this second expression makes clear that the term (Area AHIE) represents crowd-out of the private supplemental market that otherwise would have existed. The mean willingness-to-pay for the extension relative to a competitive private supplemental insurance market is $MeanWTP_Extn(P^{CE}|\alpha+\delta,\alpha) =$ Area ABCD $-$ Area AHK, where Area AHK rep-

offers a natural way to characterize the effect of an extension on several other measures of interest including the government's budget, total use of resources, and consumer and producer surplus.

resents the consumer valuation of the private supplemental market that is crowded-out by the insurance extension.

There are several additional counterfactual scenarios one can consider within this graphical example. For instance, fixing the level of compulsory public coverage, the welfare associated with allowing a private supplemental insurance market relative to banning supplemental coverage is represented by the trapezoid AHIE. Welfare under the first-best efficient allocation of supplemental insurance relative to a world with no supplemental insurance is represented by the triangle AGE. Relative to the first-best efficient allocation, a private competitive market for supplemental coverage is associated with welfare losses due to adverse selection represented by the triangle HGI.

It is also straightforward to generalize this graphical example along several dimensions. First, one can calculate analogous welfare measures relative to any equilibrium in the private supplemental market, regardless of whether the equilibrium represents perfect competition. Figure 1 Panel (B) depicts a generic equilibrium (P^*, Q^*), and the welfare measures above are equally applicable with a simple re-labeling of points. Second, the welfare measures are easy to translate to the case with moral hazard. Figure 1 Panel (C) illustrates this point. With moral hazard, the supplemental insurer does not internalize the full cost of the incremental insurance it provides. Thus, a supplemental insurer's cost curves are shifted downward relative to the social cost curves associated with the incremental coverage, where the vertical distance between these curves represents the externality associated with the supplemental coverage. While the competitive equilibrium price and quantity are different than in the situation with no moral hazard, the points have been re-labeled to emphasize that the analogous welfare measures are represented as in the discussion above.

1.2.2 Inframarginal Counterfactuals: Total Welfare Generated by Compulsory Public Insurance

Beyond estimating the welfare associated with marginal counterfactuals, economists are sometimes interested in estimating welfare associated with broader counterfactuals which involve changes outside the empirical context or contracts outside of those for which we can directly estimate demand. In public insurance settings, a broader counterfactual of particular interest is the welfare generated by the inframarginal compulsory public insurance relative to the absence of insurance for this risk.

In the recent literature on welfare analysis within private insurance markets, one common approach to inframarginal welfare analysis is to specify a structural model of decision-making, use empirical variation in premiums or other product attributes to estimate the primitives of the model, and then use the fitted model to investigate broader welfare questions of interest. This basic approach could be applied to investigate the welfare within compulsory public insurance settings as well, using variation within the private market for supplemental insurance. While this type of analysis would allow for a broad range of counterfactuals, one main disadvantage of this more structural approach is that the results are typically quite sensitive to the assumptions the analysis relies on in terms of the form of utility, the distribution of risk individuals face, and the degree of heterogeneity across individuals. Below, we outline an alternative, complementary approach to investigating the total surplus generated by compulsory public insurance relative to the absence of insurance for this risk using data and variation from supplemental private insurance. Relative to a more structural approach, the advantage of this approach is that it requires fewer assumptions; the disadvantage of this approach is that it delivers a bound rather than a precise welfare number.

Once again, consider a setting in which compulsory public insurance provides coverage α and individuals have the opportunity to buy supplemental coverage of generosity δ for price p. In the interest of investigating counterfactuals requiring minimal assumptions, our primary focus within the inframarginal counterfactual analysis will be to investigate the mean willingness-to-pay and welfare provided by the base-

line compulsory coverage relative to a world with no insurance for this risk. Note that this welfare measure does not necessarily represent the value of compulsory public insurance relative to a world without this compulsory insurance, as it could be the case that alternative private insurance or alternative public programs providing coverage for this risk would exist in the absence of compulsory public insurance.[13] Because nothing about the observed supplemental insurance market can reveal what private or public coverage might exist in the absence of the compulsory public insurance, we avoid making arbitrary assumptions about such a counterfactual and focus instead on the surplus generated by the inframarignal compulsory public insurance relative to a world with no insurance for this risk:

$$MeanWTP_Baseline(\varnothing, p|\alpha+\delta,\alpha) = \int \Big\{\pi(\alpha,0|\varphi)\mathbb{1}\big(\pi(\alpha+\delta,\alpha|\varphi)<p\big) + \big(\pi(\alpha+\delta,0|\varphi)-p\big)\mathbb{1}\big(\pi(\alpha+\delta,\alpha|\varphi)\geq p\big)\Big\}dG(\varphi), \quad (13)$$

$$Welfare_Baseline(\varnothing, p|\alpha+\delta,\alpha) = \int \Big\{\big(\pi(\alpha,0|\varphi)-c(\alpha,0|\varphi)\big)\mathbb{1}\big(\pi(\alpha+\delta,\alpha|\varphi)<p\big) + \big(\pi(\alpha+\delta,0|\varphi)-c(\alpha+\delta,0|\varphi)\big)\mathbb{1}\big(\pi(\alpha+\delta,\alpha|\varphi)\geq p\big)\Big\}dG(\varphi).$$

The first term in each of the expressions above captures the value of compulsory public coverage among individuals who do not buy supplemental insurance, and the second term in each expression captures the value of both the compulsory public coverage and the opportunity to buy private supplemental coverage among individuals who do purchase supplemental insurance at price p.[14] Note that the key challenge to estimating the expressions above is the fact that the willingness-to-pay for the baseline coverage, $\pi(\alpha,0|\varphi)$ (and the willingness-to-pay for the combination of the baseline and supplemental coverage $\pi(\alpha+\delta,0|\varphi)$), is unobserved. To be able to empirically use this measure of welfare, we need to have a way to connect what we can estimate in the data to these objects in the above expressions. The following proposition helps us do this:

Proposition 1. Suppose $\frac{\partial}{\partial\theta}\big(\frac{\pi(\theta,\gamma|\varphi)}{\theta-\gamma}\big)\leq 0$, and $\frac{\partial}{\partial\gamma}\big(\frac{\pi(\theta,\gamma|\varphi)}{\theta-\gamma}\big)\leq 0$. Then,

$$\text{[i]}\quad \pi(\alpha,0|\varphi)\geq\frac{\alpha}{\delta}\pi(\alpha+\delta,\alpha|\varphi), \quad \text{and} \quad \text{[ii]}\quad \pi(\alpha+\delta,0|\varphi)\geq\frac{\alpha+\delta}{\delta}\pi(\alpha+\delta,\alpha|\varphi) \quad (14)$$

Proof. Under the assumptions, we get the following:

$$\frac{\pi(\theta,\gamma|\varphi)}{\theta-\gamma}\geq\frac{\pi(\theta+\epsilon,\gamma|\varphi)}{\theta+\epsilon-\gamma}\geq\frac{\pi(\theta+\epsilon,\gamma+\mu|\varphi)}{\theta+\epsilon-(\gamma+\mu)} \quad (15)$$

where the first inequality follows from $\frac{\partial}{\partial\theta}\big(\frac{\pi(\theta,\gamma|\varphi)}{\theta-\gamma}\big)\leq 0$ and the second from $\frac{\partial}{\partial\gamma}\big(\frac{\pi(\theta,\gamma|\varphi)}{\theta-\gamma}\big)\leq 0$. This holds for all $\mu,\epsilon\geq 0$. Evaluating this at $\mu=\theta-\gamma$, $\theta=\alpha$, $\gamma=0$, and $\epsilon=\delta$ gives result [i] and result [ii]. □

Proposition 1 allows us to get a lower bound on welfare provided by the baseline coverage that can be applied empirically. To see this, note that the left-hand-side of inequalities [i] and [ii] enter equation 13, while

[13]The welfare measure we focus on may characterize the welfare associated with compulsory public insurance more generally in insurance settings in which there is a large degree of private information, in which we might expect complete (or nearly complete) unraveling of private insurance in the absence of public coverage. For instance, this is arguably the relevant case for disability insurance and long-term care insurance, as Hendren (2013) argues that the large degree of asymmetric information in these settings can explain the almost complete unraveling of private, non-group insurance for these risks.

[14]Since we aim to measure welfare relative to a world in which no insurance exists for this risk, it is important to have this welfare measure account for both the direct value of public coverage and the opportunity to buy supplemental coverage (as the existence of the supplemental market is an indirect consequence of the public coverage).

the right-hand-side of these inequalities is a linear transformation of the willingness-to-pay for supplemental coverage, $\pi(\alpha+\delta,\alpha|\varphi)$, an object whose distribution one can estimate with sufficient data and variation from the private supplemental insurance market. Applying Proposition 1, we get the following natural corollary:

Corollary 1. Suppose $\frac{\partial}{\partial\theta}\big(\frac{\pi(\theta,\gamma|\varphi)}{\theta-\gamma}\big)\leq 0$, and $\frac{\partial}{\partial\gamma}\big(\frac{\pi(\theta,\gamma|\varphi)}{\theta-\gamma}\big)\leq 0$, $\forall\varphi$. Then we obtain the following lower bound on the mean willingness-to-pay and welfare generated by the compulsory baseline coverage relative to a world without insurance,

$$MeanWTP_Baseline(\varnothing,p|\alpha+\delta,\alpha)\geq\int\frac{\alpha}{\delta}\pi(\alpha+\delta,\alpha|\varphi)\mathrm{I}\big(\pi(\alpha+\delta,\alpha|\varphi)<p\big) + \big(\frac{\alpha+\delta}{\delta}\pi(\alpha+\delta,\alpha|\varphi)-p\big)\mathrm{I}\big(\pi(\alpha+\delta,\alpha|\varphi)\geq p\big)dG(\varphi) \quad (16)$$

$$Welfare_Baseline(\varnothing,p|\alpha+\delta,\alpha)\geq\int\big(\frac{\alpha}{\delta}\pi(\alpha+\delta,\alpha|\varphi)-c(\alpha,0|\varphi)\big)\mathrm{I}\big(\pi(\alpha+\delta,\alpha|\varphi)<p\big) + \big(\frac{\alpha+\delta}{\delta}\pi(\alpha+\delta,\alpha|\varphi)-c(\alpha+\delta,0|\varphi)\big)\mathrm{I}\big(\pi(\alpha+\delta,\alpha|\varphi)\geq p\big)dG(\varphi).$$

Proof. Result follows directly from Proposition 1. □

The assumptions within Proposition 1 have a simple intuitive basis and are related to concepts in the prior theoretical literature related to decisions under uncertainty.[15] Both of these assumptions rely on the basic notion that an individual's willingness-to-pay per unit of insurance is declining in the amount of insurance he/she already has. The first assumption, $\frac{\partial}{\partial\theta}\big(\frac{\pi(\theta,\gamma)}{\theta-\gamma}\big)\leq 0$, says the following: Holding the generosity of the baseline coverage fixed, the willingness-to-pay per unit of supplemental insurance coverage is declining in the amount of supplemental coverage. In words, the second assumption, $\frac{\partial}{\partial\gamma}\big(\frac{\pi(\theta,\gamma)}{\theta-\gamma}\big)\leq 0$, says: Holding the generosity of the total coverage fixed, the willingness-to-pay per unit of supplemental insurance coverage is declining in the generosity of the baseline coverage. Under these assumptions, Proposition 1 gives us two intuitive inequalities (14 above): [i] the per unit of coverage willingness-to-pay for the baseline coverage is at least as large as the per unit of coverage willingness-to-pay for supplemental coverage, and [ii] the per unit of coverage willingess-to-pay for the combination of the baseline and supplemental coverage is at least as large as the per unit of coverage willingness-to-pay for only the supplemental coverage.

In the case that preferences can be represented by a univariate utility function, we obtain a simple sufficient condition for the assumptions underlying Proposition 1:

Proposition 2. Suppose an individual's utility can be represented by the increasing, univariate function $u(c)$, so that $\pi(\theta,\gamma)$ is defined as:

$$E[u(w+(1-\theta)x-\pi(\theta,\gamma))]=E[u(w+(1-\gamma)x)], \quad (17)$$

where the expectation is taken over $x\leq 0$, representing the uncertain losses the individual faces. Additionally, suppose the individual is risk averse and his/her utility exhibits (weakly) decreasing absolute risk

[15] Perhaps because it is so intuitive that an individual's willingness-to-pay to avoid taking on risk should scale with the size of the risk, prior theoretical papers have worked to derive sufficient conditions on utility functions such that the underlying preferences satisfy similar properties to the properties we study here. See Eeckhoudt and Gollier (2001) and Menezes and Hanson (1970) for examples. While much of the related theoretical literature focuses on risks with both positive and negative realizations (as is applicable in a finance setting), we show that in the context of an insurance problem (where all realizations of risk are non-positive), we obtain a simple sufficient condition for the assumptions within Proposition 1 when preferences can be represented by a univariate utility function (see Proposition 2).

aversion. Then,

$$\frac{\partial}{\partial\theta}\left(\frac{\pi(\theta,\gamma)}{\theta-\gamma}\right)\leq 0, \text{ and } \frac{\partial}{\partial\gamma}\left(\frac{\pi(\theta,\gamma)}{\theta-\gamma}\right)\leq 0, \tag{18}$$

and thus we obtain the results of Proposition 1.

Proof. See Appendix A. □

The above proposition tells us that if utility exhibits weakly decreasing absolute risk aversion, we can apply a simple method to bound the welfare generated by compulsory public insurance using information from the associated private supplemental insurance market. There are two important things to note about this. First, weakly decreasing absolute risk aversion is satisfied by most of the common utility functions used by empiricists, including Constant Absolute Risk Aversion (CARA), Constant Relative Risk Aversion (CRRA), and a broader class of empirically relevant Hyperbolic Absolute Risk Aversion (HARA) utility functions.[16] Second, this bound does not require specifying a particular utility function or restricting heterogeneity across individuals. In other words, each individual may have his/her own distinct utility function that varies arbitrarily throughout the population, so long as each individual's preferences exhibit weakly decreasing absolute risk aversion (or more generally, satisfy the conditions within Proposition 1), the lower bound on welfare described above is applicable.[17]

It is worth emphasizing that this approach has both advantages and disadvantages when compared to a more structural analysis of inframarginal counterfactuals. The key advantage of the bounding approach above is that it places no restrictions on: (i) utility beyond the minimal assumptions within the propositions above, (ii) the nature of the risk each individual faces, and (iii) heterogeneity across individuals. In contrast, other studies that evaluate broader welfare counterfactuals typically make assumptions on all of these objects, and these assumptions can substantially affect the results (and are often fundamentally untestable). Thus, the main advantage of this bounding approach is that it requires much less restrictive assumptions, while a primary disadvantage is that it provides a bound rather than a precise welfare number. The following numerical example illustrates this trade-off more concretely.

Numerical Example Consider an individual who faces some potential loss and has compulsory public insurance that covers 50% of this risk. Further, suppose the individual has the option to purchase private supplemental insurance to cover an additional 10% of the risk on top of the public insurance. Following the definitions above, let $\pi(0.5,0)$ represent the individual's willingness-to-pay for compulsory public insurance relative to no insurance, and let $\pi(0.6,0.5)$ represent the individual's willingness-to-pay for supplemental insurance to top-up the compulsory public insurance. Suppose the individual's preferences are represented by a univariate utility function $u(c)$, and the individual is risk averse with weakly decreasing absolute risk aversion. Based on Propositions 1 and 2, we get a lower bound on $\pi(0.5,0)$ for any given observed $\pi(0.6,0.5)$. The black solid line in Figure 2 plots this lower bound for a range of potential values of $\pi(0.6,0.5)$.

This bound is applicable regardless of the distribution of uncertainty the individual faces and regardless of what form the utility takes. To contrast this bound with a more structural approach, Table 1 and Figure

[16]These assumptions hold for a wide range of utility functions used by empiricists, including a broad class of utility functions within the HARA family, which includes as special cases CRRA and CARA. The general version of the HARA family is characterized by a utility function of the following form: $u(c)=\xi(\nu+\frac{c}{\gamma})^{1-\gamma}$. The absolute risk aversion for this family is then equal to $r_A=(\nu+\frac{c}{\gamma})^{-1}$, which is clearly decreasing in c for all parameter values for which the individual is risk averse over the entire range of possible consumption values. Note the values for which the individual is risk averse are those that satisfy: $\xi(1-\gamma)\gamma^{-1}(\nu+\frac{c}{\gamma})^{-\gamma-1}>0, \quad \forall c\geq 0$.

[17]It is important to highlight that the sufficient condition in Proposition 2 applies when utility is defined over a single argument. Most of the theory surrounding decisions under uncertainty is developed within the setting of univariate utility, and Proposition 2 follows in this spirit. In situations in which consumers face more complicated optimization problems with utility defined over multiple goods, as is the case in insurance models with moral hazard, one would want to rely on the broader assumptions within Proposition 1 to obtain the lower bound on the welfare provided by the inframarginal coverage.

2 display the implied value of $\pi(0.5, 0)$ for a range of potential observed values of $\pi(0.6, 0.5)$ under various assumptions. Specifically, for each set of assumptions and each value of the willingness-to-pay for supplemental coverage ($\pi(0.6, 0.5)$), Table 1 reports the calibrated utility parameters to match this willingness-to-pay for supplemental insurance and the implied willingness-to-pay for the inframarginal coverage, $\pi(0.5, 0)$, based on these calibrated parameters. The overwhelming fact that jumps out from this table and figure is that the precise implied valuation for public coverage is extremely sensitive to the underlying assumptions on the form of the utility, the benchmark income, and the distribution of uncertainty. To see this, consider the case when $\pi(0.6, 0.5) = \$150$. Comparing Table 1 columns (3) and (5), we see that the implied valuation of public coverage is 41% higher under the CRRA specification than under the CARA specification, with the same assumed risk and benchmark income. Comparing Table 1 columns (3), (7), (11), and (17), we see that the implied valuation of public coverage varies greatly with small changes in the assumed risk distribution, even conditional on specifying the same form of utility and benchmark income. A comparison of columns (3) and (15) reveals that the implied valuation of public coverage is very sensitive to the benchmark wage; in this example, when $\pi(0.6, 0.5) = \$150$, the implied value of public coverage based on CRRA utility increases 35% when the benchmark wage is cut by \$10k. In summary, this numerical example highlights the central trade-off between these complementary approaches to quantifying welfare. The bounding approach outlined above is much more robust than a more structural approach in terms of perturbations in the underlying utility, risk distribution, liquidity, and heterogeneity across individuals; however, this robustness comes at the cost of obtaining a lower bound on welfare rather than a precise welfare number.

Graphical Illustration We illustrate this approach graphically to demonstrate how the demand and cost curves in the supplemental market can be used to calculate a lower bound on the total welfare generated by compulsory public coverage. To do this, it is helpful to define a few additional objects. Let the expected willingness-to-pay for coverage of generosity θ (relative to no insurance) among individuals marginal to buying supplemental insurance at price p be represented by:

$$MTV(p|\alpha + \delta, \alpha; \theta) = E[\pi(\theta, 0|\varphi)|\pi(\alpha + \delta, \alpha|\varphi) = p]. \tag{19}$$

Analogously, let the expected cost of providing θ generosity coverage (relative to no insurance) to individuals marginal to buying supplemental insurance at price p be defined as follows:

$$MTC(p|\alpha + \delta, \alpha; \theta) = E[c(\theta, 0|\varphi)|\pi(\alpha + \delta, \alpha|\varphi) = p]. \tag{20}$$

In the discussion that follows, we will refer to the objects defined above as the *marginal total value* (MTV) and *marginal total cost* (MTC) curves, respectively.[18]

In Figure 3, the horizontal axis represents the fraction with supplemental insurance, and the vertical axis is measured in dollars. Suppose the assumptions of Proposition 1 hold. In addition, suppose there are sufficient data and variation in the supplemental insurance market to estimate the inverse demand curve for supplemental insurance ($D^{-1}(p|\alpha + \delta, \alpha)$), the marginal cost curve for supplemental insurance ($MC(p|\alpha + \delta, \alpha)$), and the marginal total cost curves defined above ($MTC(p|\alpha + \delta, \alpha; \alpha)$ and $MTC(p|\alpha + \delta, \alpha; \alpha + \delta)$).

[18]Note that marginal here refers to individuals on the margin of being willing to purchase the supplemental coverage (not the inframarginal coverage). Thus, these curves have no direct analog to the demand curve or cost curves associated with a hypothetical market for the inframarginal coverage, because the ordering of individuals' willingness-to-pay for supplemental coverage may depart from the ordering of individuals' willingness-to-pay for the inframarginal coverage. The inframarginal counterfactuals (defined in Corollary 1) rely on the MTC curve and the lower bound on the MTV curve, not the demand and cost curves associated with a hypothetical market for inframarginal coverage. Of course, it is easy to imagine other counterfactuals where one would need the demand and cost curves for inframarginal coverage (instead of or in addition to the MTV and MTC curves). Other counterfactuals that require the demand and cost curves for the inframarginal coverage would necessitate more assumptions and structure than what we impose here.

Further, suppose there is no moral hazard so the competitive equilibrium in the supplemental market is then represented by (P^{CE}, Q^{CE}), where the average cost curve intersects the inverse demand curve.[19]

Based on Proposition 1, the observed demand curve can be used to bound both $MTV(p|\alpha+\delta, \alpha; \alpha)$ and $MTV(p|\alpha+\delta, \alpha; \alpha+\delta)$, as displayed in the figure under the assumptions described above. It is then straightforward to use the figure to describe the lower bound on mean willingness-to-pay associated with compulsory coverage given that a competitive supplemental market exists: $MeanWTP_Baseline(\varnothing, P^{CE}|\alpha + \delta, \alpha) \geq$ Area ABCD + Area EFGH, where the first term and second term, respectively, give us the mean willingness-to-pay for compulsory coverage among those who do and who do not purchase supplemental coverage. The analogous lower bound on welfare is represented as: $Welfare_Baseline(\varnothing, P^{CE}|\alpha + \delta, \alpha) \geq$ Area ABIJ + Area ERL − Area RKF. These curves can also be used to evaluate welfare under other counterfactual scenarios. For example, the bounds on the mean willingness-to-pay and welfare generated by compulsory public coverage in a world in which supplemental insurance is not available can be represented by: $MeanWTP_Baseline(\varnothing, \infty|\alpha + \delta, \alpha) \geq$ Area MFGO and $Welfare_Baseline(\varnothing, \infty|\alpha + \delta, \alpha) \geq$ Area MRN − Area RKF.

1.3 Comparison to Alternative Methods

It is worth highlighting a few strengths of this willingness-to-pay approach to analyzing welfare associated with compulsory public insurance relative to other approaches. First, this method requires straightforward variation and limited data to be implemented empirically. To estimate all of the inputs of the marginal and inframarginal welfare analysis discussed above, you simply need sufficient variation in the price of supplemental coverage to estimate the demand curve for supplemental coverage and standard data on prices, insurance enrollment, and costs. In contrast, other approaches to valuing compulsory public insurance often require isolating random (or like-random) assignment within the compulsory public insurance program (which by virtue of the program being compulsory is often difficult to do). In addition, alternative approaches often require data on consumption or assets, data which is often not available and can suffer from considerable measurement error when it is available.

Second, an advantage of this approach is that it allows us to remain largely agnostic as to the underlying primitives of individuals' decisions (similar to Einav and Finkelstein (2011)). It is possible to investigate a wide range of marginal welfare questions (welfare analysis in the spirit of the sufficient statistics literature on social insurance (e.g., Chetty, 2008)) without making restrictions on the form of individual utility or heterogeneity across individuals, as evidenced in the framework above by the fact that there are no restrictions on φ or the distribution G. In settings with linear insurance coverage, going beyond marginal welfare questions and bounding the welfare provided by the compulsory inframarginal coverage requires only minimal assumptions on individual preferences and no restrictions on heterogeneity across individuals. The fact that this approach requires so few assumptions is in stark contrast to other approaches of valuing compulsory public coverage which rely on fully specifying utility (or marginal utility), specifying the nature of uncertainty individuals face, restricting heterogeneity across individuals, obtaining data on all the inputs in the utility (or marginal utility), and/or assumptions on the nature of the individual's optimization.

Third, this approach is flexible enough to be applied in a broad range of settings. It is straightforward to show that this approach can accommodate features that may be present in the market for supplemental insurance, such as imperfect competition, administrative costs, multiple contract options, and pricing regu-

[19] The assumption of no moral hazard in this graphical example is purely for convenience, not necessity. In the case of moral hazard, one would need to appropriately adjust the cost curves for inframarignal coverage ($MTC(p|\alpha+\delta, \alpha; \alpha)$ and $MTC(p|\alpha+\delta, \alpha; \alpha+\delta)$) relative to the cost curve for supplemental insurance ($MC(p|\alpha+\delta, \alpha)$) . We discuss accounting for moral hazard in more detail within the context of the empirical application.

lations. Importantly, while the framework above is discussed in the context of linear insurance (which is the relevant case in the empirical application that follows), the general approach to valuing a marginal extension of compulsory coverage to include benefits covered by supplemental insurance can be applied to any setting with complementary private insurance provided sufficient data and variation exist, even settings with highly non-linear coverage. For example, if compulsory public insurance includes a deductible and supplemental insurance fills in this deductible, an individual's willingness-to-pay for supplemental insurance reveals his valuation of extending public insurance to fill in this deductible. Or alternatively, in the setting of health insurance, if private supplemental insurance improves upon basic compulsory public insurance by providing access to a broader group of health care providers, an individual's willingness-to-pay for supplemental insurance reveals his/her valuation of an extension of the compulsory public insurance coverage to include this broader group of health care providers. To the extent that the available supplemental insurance fills the obvious gaps in the compulsory public insurance coverage in a given setting, expanding compulsory public coverage to include the benefits of supplemental insurance may be one of the most policy-relevant marginal expansions to consider.

Fourth, this approach can be used to measure welfare in wide range of policy-relevant counterfactuals. As discussed above, this approach can be used to value the extension of compulsory coverage, taking into account any crowd-out of the private market or externalities induced by the incremental coverage. This approach can also be used to value alternative government interventions such as banning supplemental insurance and taxing or subsidizing supplemental insurance. In settings with linear coverage, this approach can be used to go beyond marginal welfare questions and bound the welfare generated by the inframarginal compulsory public coverage under minimal assumptions discussed above.

Of course, an important limitation of this approach is that it can be applied only in settings in which individuals are able to purchase complementary private insurance to top-up the compulsory public coverage. While there are many public insurance settings in which individuals can buy complementary coverage (e.g., disability insurance, universal health insurance settings such as Medicare, etc.), there are some public insurance settings in which no complementary coverage is typically available (e.g., unemployment insurance, means-tested public insurance such as Medicaid), and in these settings this approach does not provide an alternative to existing approaches to valuing public insurance.

It is also important to note is that welfare analysis following this approach is subject to the same limitations of welfare analysis in private insurance markets more generally. For instance, this approach relies on the notion that individuals' decisions to buy or forgo insurance can be used within a revealed preference framework to do welfare analysis. While this is a common approach to evaluating welfare in private insurance settings, there are potential criticisms and drawbacks to relying on revealed preference as individuals may suffer from behavior biases or may face informational frictions. Additionally, as in typical analyses within private insurance markets, this approach relies on estimating demand, which requires isolating price variation that is plausibly exogenous and often requires making parametric assumptions when there is not sufficient price variation to estimate demand non-parametrically. Despite the non-trivial challenges associated with using revealed preference for welfare analysis, researchers conducting welfare analysis in private insurance settings typically opt to employ demand estimates and rely on revealed preference (as opposed to employing methods similar to those used in the prior compulsory public insurance literature). In this way, the framework described above provides a constructive illustration of how to extend this revealed preference style analysis often employed in settings with private insurance to investigate questions related to compulsory public insurance.

2 Empirical Application: Background and Data Description

The remainder of the paper applies the framework discussed in the prior section to evaluate welfare in the context of public and private disability insurance. Specifically, we leverage price variation in the private supplemental long-term disability (LTD) policies offered by one large employer to quantify the value of disability insurance among workers in this population. While the empirical setting is interesting in its own right, the empirical application also highlights some advantages of applying this willingness-to-pay approach to valuing compulsory public insurance. This section begins by briefly describing some relevant background on public and private disability coverage. We then proceed to discuss the data and identifying variation.

2.1 Brief Background on Public and Private Disability Coverage

In the United States, the public disability insurance program, Social Security Disability Insurance (SSDI), is one of the largest social insurance programs. In 2014, approximately 9 million disabled workers received SSDI benefits, and the total SSDI benefits paid exceeded $140 billion.[20]. Eligible workers may apply for public disability insurance benefits in the event of a disability. The applicant then goes through an evaluation of his case that includes an assessment of the severity of his disability and an assessment of his residual work capacity.[21] If an individual's SSDI case is approved, the individual receives a monthly benefit that is a function of his prior earnings history. In particular, the Social Security Administration (SSA) calculates the individual's Average Indexed Monthly Earnings (AIME), which is defined as the average inflation-adjusted monthly earnings in the individual's 35 highest earning years of employment (or all years of employment if the beneficiary has worked fewer than 35 years).[22] The monthly benefit for the beneficiary, the Primary Insurance Amount (PIA), is then determined by a piece-wise linear function of the AIME. Since SSDI pays a monthly benefit for each month the individual is out of the labor force due to disability, SSDI insurance provides coverage that is linear in the risk (lost wages due to time out of work with disability), where the effective replacement rate for individual i in the event that he becomes disabled in terms of his current annual earnings, w_i, is,

$$\gamma_i = \frac{\text{Annual SSDI Benefit}_i}{w_i}, \tag{21}$$

where Annual SSDI Benefit$_i = 12 \times PIA(AIME_i)$. Among US workers, the mean pre-tax replacement rate for public disability insurance through SSDI is roughly 45%.[23]

Some SSDI beneficiaries have access to additional benefits beyond the wage-replacement benefits described above. For example, after two years of SSDI coverage, beneficiaries receive health insurance through Medicare. Additionally, in some circumstances, SSDI provides additional financial benefits to beneficiary dependents (i.e., young children or elderly spouses). Throughout the following discussion, we focus on the welfare generated by the primary wage-replacement benefits of SSDI abstracting from these additional features.

Many employers offer LTD insurance to supplement the wage replacement benefits of SSDI. Approximately 34% of all workers have the option to buy LTD coverage through their employer to top-up the coverage from SSDI.[24] While on average 34% of all workers have access to LTD coverage, the percent of workers

[20]Source: `https://www.ssa.gov/policy/docs/statcomps/supplement/2015/highlights.pdf`.

[21]The following SSA publication contains a basic description of the evaluation criteria: `http://www.ssa.gov/pubs/EN-05-10029.pdf`.

[22]A calculator on the SSA website describes the disability benefit formula: `https://www.ssa.gov/planners/retire/AnypiaApplet.html#&sb=-1`.

[23]This statistic is based on authors' calculations using employees in the CPS, as described in Table 3. Note this replacement rate is calculated before considering the tax treatment of wages and SSDI benefits.

[24]These statistics are reported in Table 16 of the U.S. Department of Labor's National Compensation Survey, Employee Benefits in the

with access to LTD is much higher among particular subgroups: 44% of full-time workers, 53% of managerial and professional workers, and 57% high income workers (workers in the top 25% of earners). The median replacement rate of employer-provided LTD coverage is 60%, with only 1% of plans providing coverage in excess of 67% and 24% of plans providing coverage less than 60%.[25]

2.2 Data Description and Identifying Variation

In this paper, we use data from one large employer that offers LTD coverage. The data comes from Alcoa, Inc., a large, multinational manufacturing firm that annually employed roughly 48,000 employees within the United States residing within 24 different states during our sample period. Approximately half of the the firm's employees are offered LTD benefits. The data include information on LTD benefit menus, enrollment decisions, wages, employment status, basic demographics, and LTD claim payments.

As is typical with employer-provided LTD policies, the policies offered by this employer guarantee a stated replacement rate in terms of annual earnings, with the explicit stipulation that LTD claimants must apply for SSDI, and, if approved, SSDI benefits will crowd-out LTD benefits dollar-for-dollar.[26] The LTD policies at this firm insure relatively long duration disability spells, as workers cannot place an LTD claim until he/she has been disabled for at least six months.[27] In this way, LTD policies at this firm largely operate as a means to top-up the replacement rate of the public program. For the remainder of the empirical application, we make the simplifying assumption that the LTD plans provide pure top-up insurance for SSDI.[28]

2.2.1 Estimation Setting

For the empirical application, we focus on employees offered the most common LTD benefit menu containing three vertically differentiated plans with replacement rates of 50%, 60%, and 70%, which we denote as Plan L, Plan M, and Plan H, respectively. We focus on the years in which this menu is consistently available, 2003-2006. Table 2 describes the menu of LTD policies offered during our sample period, displaying the plan premiums by year both in levels and relative to the plan of adjacent quality. In this setting, both benefits and premiums are described in terms of an individual's annual earnings. To ease interpretation, Table 2 also displays the premiums scaled by $60,000, which is roughly the median annual earnings in the estimation sample. From this table, one can see that there are several nice properties of the data for this empirical application. First, the basic 50% replacement rate plan is free for all employees. This is convenient as no one opts out of the plans, and thus claims data are available for all employees in the sample. Second, there is variation over time in the incremental premium for the highest generosity coverage. In particular, the premium for Plan H (relative to Plan M) falls by nearly a third between 2003 and 2004, providing variation to estimate the demand for this plan.

In this setting with vertically differentiated plans, the demand for a plan is determined by the premium relative to the premiums for the adjacent plans, provided that no plans are dominated in the observed range of premiums. Consistent with this intuition, the bottom panel of Table 2 illustrates that the decline in the

U.S., March 2015 publication at: `http://www.bls.gov/ncs/ebs/benefits/2015/ownership/civilian/table16a.pdf`.

[25] These statistics are reported in Table 30 of the U.S. Department of Labor's National Compensation Survey, Employee Benefits in the U.S., March 2015 publication at: `http://www.bls.gov/ncs/ebs/benefits/2015/ownership/civilian/table30a.pdf`.

[26] This crowd-out is in terms of pre-tax monetary benefits (not post-tax benefits), which is an important consideration when incorporating the differential tax treatment of LTD benefits and SSDI benefits. We discuss this further in the robustness analysis in Section 4.3.

[27] This required period of disability before applying for benefits is often referred to as the "elimination period". Autor, Duggan and Gruber (2012) analyze data from a large LTD insurer covering policies offered by nearly ten thousand unique employers, and they report that nearly two-thirds of beneficiaries have LTD insurance with a shorter elimination period than the firm we study, with the most common elimination period being 90 days.

[28] As will become clear, this simplification is helpful in interpreting the inframarginal welfare analysis; however, the simplification is not necessary to have a meaningful interpretation of the marginal welfare analysis.

Plan H premium between 2003 and 2004 is associated with a sharp increase in Plan H enrollment and a sharp decline in Plan M enrollment, while Plan L enrollment is largely unaffected. The primary estimation focuses on the demand and costs associated with the incremental 10% replacement rate moving from Plan M to Plan H. That is, the estimation focuses on estimating the demand for a 10 percentage point increase in the replacement rate starting from a baseline replacement rate of 60%. Although there is a small amount of variation in the Plan M premium that could be used to investigate the demand for Plan M (relative to Plan L), there are two key reasons why we focus on the Plan H demand for our baseline estimation. First, from inspecting Table 2, one can see that the relative premium for Plan H (compared to Plan M) varies more than five times as much as the relative premium for Plan M (compared to Plan L) varies. Because most of the premium variation is in the relative premium for Plan H (compared to Plan M), we obtain more precise estimates by focusing on the demand for Plan H coverage. Second, under the assumptions used to derive the bound in Section 1, focusing on the demand for Plan H provides a more conservative lower bound on the welfare provided by public insurance. Section 4 discusses this further within the context of the estimates.

2.2.2 Measuring Variables for Analysis

The primary data needed for the analysis are cost data, premium data, and coverage data. Both premiums and costs in this setting are measured relative to annual earnings. While measuring premiums and coverage are straightforward, there are a few subtle issues that arise when measuring costs in this setting. The ideal cost data would follow all realized disability spells through their entire duration to calculate the average cost measures used in the welfare analysis. Because the disability spells we focus on originated relatively recently (our sample period is 2003-2006), many of the disability spells are still unfinished, and thus, this ideal information following each spell through its entire duration is not available.[29] In practice, we have data on LTD claims paid through 2011, which means that all disability spells originating during the period of coverage (2003-2006) are observed through at least the first five years of the spell. Although complete data is not available for censored spells, we can construct proxies for the total realized costs using the data that is available. In our baseline analysis, we use a conservative measure of costs where we calculate the maximum potential costs by extending those spells truncated at five years to their maximum possible duration and imputing the lost wages.[30] Because our aim is to estimate a lower bound on the welfare provided by disability insurance, this conservative measure is appropriate. However, we also repeat the welfare analysis with alternative (less conservative) cost measures and display the results in Appendix D for comparison. The primary cost measure used in the welfare analysis is the mean present discounted value of the realized costs relative to annual earnings associated with the incremental 10 percentage point increase in the replacement rate. The baseline cost measure is constructed using an interest rate of 4%, and additional analysis presented in the Appendix D illustrates that the results are robust to a range of alternative interest rates.

One aim of the welfare analysis is to estimate a lower bound on the wage replacement benefits of public disability insurance. In this analysis, we focus on bounding the surplus generated by compulsory disability insurance coverage of generosity equal to the mean SSDI replacement rate in the estimation sample, where we calculate an individual's SSDI replacement rate using the SSA formula replacing the AIME with the monthly wage rate according to the firm's human resources records. We note this is an imperfect calculation because the data does not contain each worker's full work history which would be needed to precisely calculate the

[29] Notice that this is not just a data availability issue. Many of the disability spells originating during the period of coverage (2003-2006) are currently still unfinished. Thus, even with the most up-to-date data available to claims administrators, there would still be a censoring issue.

[30] The maximum possible duration of benefits is set by the LTD insurance policies and depends on the age at disability. For those disabled before age 60, benefits may be paid until the individual turns age 65. For those disabled after age 60, the benefits may continue for a maximum of five years.

AIME following the SSA formula.

2.2.3 Summary Statistics and Identifying Variation

The estimation sample is restricted to salaried employees who are actively employed in the previous calendar year, who are subject to the relevant LTD benefit menu, and who have an SSDI replacement rate less than 50%.[31] This final restriction ensures that for all the employees in the sample, the LTD plans offered by the firm all provide coverage that goes beyond SSDI's coverage. To ensure the workers in the estimation sample are comparable over time, we drop workers at job sites that open or close during the period of analysis (2003-2006) and workers at job sites that experience very large fluctuations in employment during the period of analysis (job sites at which the percent change in employment from the minimum to the maximum employment year exceeded 75%).[32]

Figure 4 describes the public and top-up private disability coverage available to employees at the firm. The horizontal axis displays the monthly earnings, and the right vertical axis displays the associated pre-tax monthly potential disability benefit. The solid line displays the SSDI public disability benefit formula for the year 2003 (the first year of our sample), while the remaining three lines display the benefit formula for those who have one of the firm's three available top-up private LTD plans in addition to SSDI public disability insurance. The estimation sample focuses on individuals with an SSDI replacement rate less than 50% (the Plan L replacement rate). In terms of this figure, these are the individuals whose monthly earnings lie to the right of the intersection of the "with Plan L top-up" line and the "SSDI" line. This level of monthly earnings based on the 2003 benefit formula is roughly $2,020 which is indicated in the figure by the vertical reference line. In addition to displaying the public and private potential disability benefits, this figure also displays a histogram of the monthly earnings for all salaried employees.[33] As can be seen in this figure, the vast majority of salaried employees at the firm have high enough earnings that their implied SSDI replacement rate lies below the 50% threshold used to define the estimation sample.

Table 3 compares employees in the estimation sample (column 3) to the entire employee population at the firm (column 1) and to all salaried employees at the firm (column 2). From inspecting Table 3, one can see that the mean employee tenure is approximately 13 years. The majority of employees are male, and the median age is 45. The estimation sample contains roughly 23% of the overall employee population. Those in the estimation sample differ from the wider employee population at the firm along a couple of dimensions: those in the estimation sample are salaried employees (as opposed to hourly employees) and earn higher wages on average. The median wage is roughly $60,000 in the estimation sample, while the median wage is approximately $36,000 across the entire firm. Along several other dimensions such race, age, sex, and company tenure, the estimation sample and the overall employee population look quite similar.

For comparison, Table 3 columns (4) through (7) report analogous summary statistics for employees in the Current Population Survey (CPS), pooling data across 2004-2007. Comparing columns (1) and (4), we see that employees in the firm and a nationally representative sample of employees in the CPS appear broadly similar in terms of mean age and mean annual earnings, but quite different in terms of gender composition. The difference in gender composition can be explained by industry, as the employees at the firm (column 1) and manufacturing employees in the CPS (column 5) look very similar along this dimension. Examining the table further, we can see that the salaried employees at the firm (column 2) look similar in terms of

[31] The restriction that workers be "actively employed" means that the workers were employed continuously at the firm and were not on leave, paid or unpaid, during the previous calendar year.

[32] Because the identification strategy leverages over time variation, we make this restriction to ensure the sample is relatively homogenous year-to-year.

[33] In this figure, monthly earnings are censored at ($200,000/12).

observable characteristics to manufacturing employees in the CPS with white collar occupations (column 6).[34] Lastly, comparing the estimation sample (column 3) to the sample in the CPS that meets an analogous selection criteria (column 7), one can see that they look broadly similar on observables such as age, gender, race, and annual earnings. Overall, it is important to emphasize that neither the firm's overall employee population nor the estimation sample are representative of employees in the United States as a whole, so one should exercise appropriate caution in interpreting the estimates. That being said, along the lines of observable attributes, the employees in the estimation sample look similar to manufacturing, white collar employees more broadly.

Table 4 displays some summary statistics by year for the estimation sample. Since the identifying variation stems from a sharp change in premiums between 2003 and 2004, it is important to demonstrate that the estimation sample is comparable across the years in all other ways. Column (1) summarizes demographic information for employees in the sample during 2003 (the year in which the relative price for Plan H was the highest), while columns (2) through (4) display summary statistics for the subsequent three years, during which the relative price for Plan H was lower and quite stable. Across these columns, the table summarizes observable employee information by year, along with the overall mean predicted probability of the worker claiming LTD based on the worker's observable characteristics (denoted by "predicted LTD claims"). This predicted probability is simply the fitted value from regressing an indicator for having a disability claim on all the observable attributes summarized in Table 4. Inspecting this table, we see the sample looks very similar across the four years in terms of age, job tenure, fraction male, earnings, and predicted LTD claims. In particular, there is no evidence of a problematic break in observables coincident with the Plan H price decrease between 2003 and 2004. We also investigate this in a slightly different way, relating worker characteristics to the continuous (relative) price for Plan H including a time trend. The coefficients and p-values associated with these regressions, reported in columns (5) and (6), collectively illustrate that the identifying variation is unrelated to the observable characteristics. We can also reject that all of these observables are jointly related to the premium variation (p-value 0.33). Overall, the evidence from Table 4 illustrates that the identifying premium variation is unrelated to observable characteristics. In addition to this evidence supporting the identification assumption, we also demonstrate that the welfare analysis is robust to the inclusion/exclusion of covariates in the demand estimation (e.g., a time trend, demographic controls, individual fixed effects, etc.).

3 Empirical Application: Welfare Framework and Estimation Strategy

3.1 Adapting Welfare Framework to Empirical Context

Below we discuss how the framework outlined in Section 1 is adapted for the purpose of the empirical application. As discussed in the prior section, we focus on estimating the demand and costs associated with a 10 percentage point increase in the replacement rate starting from a baseline of 60% wage replacement in the event of disability (moving from Plan M to Plan H among the plans offered by the firm). Thus, the empirical application departs from the initial basic framework in Section 1 in that there is a wedge between the willingness-to-pay we estimate (moving from a 60% to 70% replacement rate) and the replacement rate of compulsory public insurance (on average 34% in our estimation sample). We adapt the welfare measures and interpretation to account for this difference as described below.

[34] The definition of white collar occupations used here encompasses individuals who report professional or managerial occupations.

3.1.1 Marginal Welfare Analysis

First, we consider the welfare associated with the incremental coverage provided by supplemental insurance, the coverage for which there is variation to directly estimate demand. Specifically, we consider a supplemental insurance product that offers an additional 10% replacement rate on top of a compulsory baseline replacement rate of 60%. We use the variation in this context to obtain estimates of the mean willingness-to-pay and welfare in two different scenarios: (i) an extension of the compulsory baseline coverage to include the incremental coverage if individuals would otherwise have no access to the incremental coverage ($MeanWTP_Extn(\infty|0.7, 0.6)$, $Welfare_Extn(\infty|0.7, 0.6)$), and (ii) an extension of the compulsory baseline coverage to include the incremental coverage relative to a competitive market for the incremental coverage ($MeanWTP_Extn(P^{CE}|0.7, 0.6)$, $Welfare_Extn(P^{CE}|0.7, 0.6)$).

These measures tell us the value associated with extending compulsory disability coverage, starting from a baseline of compulsory coverage with a 60% replacement rate. Note that these measures do not directly measure the value of extending compulsory public coverage which has a much lower replacement rate than the baseline coverage in the empirical context (on average 34% among individuals in the estimation sample). However, under the assumption used to calculate the inframarginal welfare bound (Assumption 1 below), the value of the incremental coverage, $MeanWTP_Extn(\infty|\alpha + 0.1, \alpha)$ and $Welfare_Extn(\infty|\alpha + 0.1, \alpha)$, is decreasing in the baseline coverage α. Thus, under the assumption used to derive the inframarginal welfare bound, we can interpret the marginal welfare measures above as a lower bound on the welfare obtained by a 10 percentage point increase in the replacement rate, starting from the lower replacement rate of public disability coverage.

Following the same logic, note that there is a testable implication of the assumption needed for the inframarginal welfare analysis (Assumption 1 below): the mean willingness-to-pay for incremental coverage $MeanWTP_Extn(\infty|\alpha + 0.1, \alpha)$ should be decreasing in the baseline coverage, α. As described further in Section 4, we provide suggestive evidence on this by comparing our baseline demand estimates for Plan H coverage (relative to Plan M coverage) to those obtained for Plan M coverage (relative to Plan L coverage), leveraging the limited available variation in the relative price for Plan M.

3.1.2 Inframarginal Welfare Analysis

We next discuss welfare within the context of broader counterfactuals which involve changes beyond the incremental coverage observed in the data. As discussed in Section 1, in order to bound welfare generated by the inframarginal compulsory disability coverage, we make the following assumption:

Assumption 1. The willingness-to-pay for coverage of generosity $(\theta - \gamma)$ on top of a baseline coverage of γ, $\pi(\theta, \gamma)$, satisfies the following properties: $\frac{\partial}{\partial \theta}\left(\frac{\pi(\theta,\gamma)}{\theta-\gamma}\right) \leq 0$ and $\frac{\partial}{\partial \gamma}\left(\frac{\pi(\theta,\gamma)}{\theta-\gamma}\right) \leq 0$.

Section 1 describes the intuition behind this assumption and describes a simple sufficient condition for this assumption (i.e., utility satisfies weakly decreasing absolute risk aversion). Further, suggestive evidence on the plausibility of this assumption in the empirical context is discussed in Section 4.

Maintaining this assumption, we consider the value of the inframarginal compulsory disability coverage relative to a setting in which no insurance is available for this risk. First, let us consider compulsory baseline insurance which provides a 60% replacement rate. Under Assumption 1, we calculate a lower bound on mean willingness-to-pay and welfare associated with the baseline coverage in two scenarios: (i) individuals have compulsory baseline coverage of generosity 60% with no opportunity to buy supplemental insurance ($MeanWTP_Baseline(\varnothing, \infty|0.7, 0.6)$, $Welfare_Baseline(\varnothing, \infty|0.7, 0.6)$) and (ii) individuals have compulsory baseline coverage of generosity 60% with the option to buy private supplemental insurance that pro-

vides an additional 10% replacement rate in a competitive market ($MeanWTP_Baseline(\varnothing, P^{CE}|0.7, 0.6)$, $Welfare_Baseline(\varnothing, P^{CE}|0.7, 0.6)$). Note this latter measure which accounts for the opportunity to buy supplemental coverage is a tighter lower bound than the former measure, provided that the competitive private supplemental market generates positive surplus (as is the empirically relevant case in this context).[35]

Beyond considering welfare of the baseline coverage in the empirical context, we consider welfare provided by public disability coverage. Specifically, we calculate a lower bound on welfare generated by compulsory insurance that provides coverage of generosity equal to the average replacement rate of public disability coverage, α_{Public} (approximately 34% in our sample). Under Assumption 1, $\pi(\alpha + \delta, \alpha)$ is decreasing in α and we obtain a lower bound on the welfare associated with compulsory disability insurance of generosity α_{Public} relative to no insurance coverage:

$$MeanWTP_Public(\varnothing, \infty|0.7, 0.6; \alpha_{Public}) = \int \pi(\alpha_{Public}, 0|\varphi)dG(\varphi) \geq \int \frac{\alpha_{Public}}{0.1}\pi(0.7, 0.6|\varphi)dG(\varphi) \quad (22)$$

$$\begin{aligned} Welfare_Public(\varnothing, \infty|0.7, 0.6; \alpha_{Public}) &= \int \big(\pi(\alpha_{Public}, 0|\varphi) - c(\alpha_{Public}, 0|\varphi)\big)dG(\varphi) \\ &\geq \int \big(\frac{\alpha_{Public}}{0.1}\pi(0.7, 0.6|\varphi) - c(\alpha_{Public}, 0|\varphi)\big)dG(\varphi). \quad (23) \end{aligned}$$

The lower bound above does not account for any additional surplus generated by the option to buy private supplemental insurance to top-up the benefits of the compulsory coverage. Including such an adjustment would provide a tighter lower bound than the bound described in Equation 23, provided that the private supplemental market generates positive surplus. Note that accounting for the opportunity employees have to buy supplemental coverage within the firm would rely on quantifying the value of all the disability plans available to employees, and there is not sufficient variation to directly measure demand for all the plans.[36] For this reason, our primary analysis and discussion centers around the conservative lower bound described above. In Appendix B, we demonstrate how additional assumptions can be leveraged to calculate a tighter lower bound on the value of public disability insurance in the empirical context accounting for the private disability options employees have available at this firm.

3.2 Empirical Strategy and Connection to Welfare Framework

Demand To calculate the welfare measures described above, we need to estimate the demand for supplemental coverage. In the disability insurance setting, an individual's coverage and an individual's premium are defined relative to the individual's annual earnings. In addition, the variation that we use to identify the demand curve is in the relative premium for supplemental coverage. Thus, we estimate the demand curve in relative terms and then discuss how this demand curve is used in the welfare counterfactuals considered below.

Let p_{it}^H and p_{it}^M represent the premium (measured relative to annual earnings) faced by individual i in

[35] As discussed in the following section, we find evidence consistent with no moral hazard induced by the marginal LTD coverage in this setting. In the absence of moral hazard, a private competitive market for supplemental insurance necessarily generates positive surplus. Further analysis reported in Appendix C demonstrates that the incremental coverage moving from a 60% to a 70% replacement rate generates positive surplus even in the presence of moral hazard, when we use moral hazard estimates from prior studies. While there are several mechanisms in place to curb moral hazard in the disability insurance setting (i.e., the disability screening process, the elimination period, etc.), moral hazard may be more prevalent in some other insurance settings with complementary public and private insurance. For example, Cabral and Mahoney (2014) estimate that private Medigap supplemental insurance exerts a large negative externality on the Medicare public insurance system and that this externality is so large that the private supplemental Medigap market is on net welfare-reducing.

[36] Plan L is free for all employees, and thus one cannot directly estimate demand for this lowest generosity plan. However, under some additional assumptions, one can use the demand curve for Plan H (or Plan M) to put bounds on the unobserved willingness-to-pay for Plan L. Additional analysis in Appendix B illustrates how additional assumptions allow us to obtain a tighter lower bound on the value of public coverage for this population accounting for the LTD plans offered within this firm.

year t for Plan H and Plan M, respectively. Let $r_\pi(\theta, \gamma|\varphi) = \frac{\pi(\theta,\gamma|\varphi)}{w_\varphi}$ denote the relative willingness to pay for generosity θ coverage over the outside option of γ coverage. The focus of the empirical analysis is on estimating demand for Plan H relative to Plan M: the distribution of $r_\pi(0.7, 0.6|\varphi)$. To recover this relative demand curve from the empirical setting, we assume that none of these vertically differentiated products are dominated in the ranges of prices observed in the data, so that the relevant price for tracing out the relative demand curve for a product is the price compared to that of the adjacent products. Under this assumption, we can use variation in the relative price for Plan H over Plan M, $(p_{it}^H - p_{it}^M)$, to trace out the distribution of $r_\pi(0.7, 0.6|\varphi)$.

Given rich enough variation in relative prices, one could in principle estimate a nonparametric demand curve for the incremental coverage, giving one a fully nonparametric bound on welfare generated by the compulsory insurance. However, since the price variation in reality is limited in this context, we initially focus attention on estimating linear and linear-log specifications of demand (and later show the results are robust to alternative demand forms). We estimate the demand for Plan H coverage using the following estimation equation:

$$PlanH_{it} = I(r_\pi(0.7, 0.6|\varphi) \geq p_{it}^H - p_{it}^M) = \alpha_0 + \alpha_1(p_{it}^H - p_{it}^M) + \rho X_{it} + \epsilon_{it}, \tag{24}$$

where $PlanH_i$ indicates if individual i selected Plan H and $(p_{it}^H - p_{it}^M)$ is the price of Plan H relative to Plan M. In addition to the linear specification described above, Table 5 also presents the results of a linear-log specification, where the log of the relative price enters rather than the relative price level. The baseline demand estimates include a time trend to control for any unrelated trends in insurance enrollment that may be correlated with the identifying variation. Further analysis presented in Section 4.3 demonstrates that the welfare analysis is robust to several other alternative specifications of demand, including alternative included/excluded covariates (e.g., no controls, individual fixed effects, etc.) and alternative demand forms (e.g., probit demand, logit demand, log-log demand, etc.).

As discussed in Section 2, the main estimation focuses on the demand for Plan H over Plan M because most of the premium variation in this context is in the relative premium for Plan H over Plan M. However, in addition to our main estimates on the demand for Plan H relative to Plan M, we also investigate the demand for Plan M relative to Plan L coverage to provide suggestive evidence on the validity of Assumption 1. To do this, we use the limited variation in the (relative) premium for Plan M compared to Plan L to estimate the demand for Plan M coverage. Thus, we similarly estimate the relative demand for Plan M by estimating the following equation:

$$PlanMorH_{it} = I(r_\pi(0.6, 0.5|\varphi) \geq p_{it}^M) = \theta_0 + \theta_1 p_{it}^M + \phi X_{it} + e_{it}, \tag{25}$$

where $PlanMorH_i$ indicates that the individual chose Plan H or Plan M coverage and p_{it}^M is the price of Plan M relative to Plan L (as $p_{it}^L = 0$).

Costs While cost data is not necessary to measure the willingness-to-pay for disability insurance, to fully evaluate welfare it is necessary to have data on costs. In this regard, the empirical environment has both strengths and weaknesses. A strength of the empirical environment is that there is individual-level administrative cost data, so that at least in principle, one could use the very same price variation used to estimate the demand curves to estimate moral hazard and selection. However, a practical limitation is that it is difficult to obtain precise estimates of selection and moral hazard within the context of disability insurance, as realized costs are noisy proxies for expected costs given the low incidence of disability. In the face of these strengths

and weaknesses, we take the following approach. First, we investigate the possibility of moral hazard in this setting and find evidence consistent with no moral hazard using the price variation used to estimate demand. See Appendix C for this analysis. Second, we proceed with the baseline welfare analysis under the assumption that supplemental disability insurance induces no moral hazard. We then demonstrate that the main findings from the welfare analysis are unchanged when employing alternative assumptions about moral hazard based on estimates from prior studies on disability insurance in other contexts.

As is the case with demand, costs are measured relative to annual earnings. Let c_{it} represent the present discounted value of the realized costs relative to annual earnings associated with providing individual i in year t an incremental 10% replacement rate paid in the event of disability. We use an interest rate of 4% in the baseline estimation. Additional analysis presented in Appendix D illustrates that our results are robust to a range of alternative interest rates. Our focus throughout is on the wage replacement costs for disability insurance (i.e., the costs for which we have data). While we have no data on administrative costs and thus abstract from such costs here, it would be straightforward to calculate alternative welfare measures incorporating administrative costs.[37]

As discussed above, we measure welfare under two hypothetical scenarios with regard to the availability of supplemental insurance: (i) supplemental insurance is not available and (ii) supplemental coverage providing an additional 10% wage replacement is sold in a competitive market. As is clear from the framework discussed above, measuring welfare under the counterfactual scenario in which supplemental insurance is not available simply requires estimating the average cost of all individuals in the population (the mean value of c_{it} under the assumption of no moral hazard). In other words, our primary estimates of the welfare generated by the marginal and inframarginal coverage are independent of the degree of selection in the supplemental insurance market. However, to calculate additional welfare measures under the counterfactual scenario of a competitive supplemental insurance market, we need estimates of the average cost of those who do and those who do not buy insurance in a competitive supplemental insurance market (which will depend on the degree of selection in the supplemental market). Table 5 Panel B displays the mean incremental cost (the mean value of c_{it}) for everyone in the estimation sample and separately for those who do and do not purchase the incremental coverage. Note that the relative price for incremental coverage for each plan is very close to the average cost of those individuals willing to pay for the incremental coverage. In other words, the predicted competitive equilibrium in the market for incremental coverage would be close to the observed pair of average prices and quantities regardless of the shape of the average and marginal cost curves far outside of the variation in the data. Thus, we take the observed mean (relative) costs of the insured and quantity of individuals insured as an approximation of the competitive equilibrium for the incremental coverage, which allows us to use mean cost estimates for those with and without supplemental insurance in the data (reported in Table 5 Panel B) to compute the welfare measures discussed above.[38] Note that a major advantage of relying on this approximation is that the cost estimates that enter the welfare analysis, in this case, are precisely estimated (see Table 5 Panel B). While we use this approach for the baseline welfare analysis, Appendix C presents estimates of selection in this context and illustrates that using alternative

[37] Because the welfare analysis discussed below illustrates that the marginal disability coverage generates substantial surplus, incorporating even quite large administrative costs would not affect the qualitative findings.

[38] Note that several of the counterfactuals we consider are independent of the degree of selection in the supplemental insurance market (Table 6 Panel A row 1 and Panel B rows 1 and 3). However, two of the counterfactuals we consider account for the available private supplemental insurance, and thus depend on the degree of selection within this market (Table 6 Panel A row 2 and Panel B row 2). Table 5 and Appendix Figure A1 demonstrate that the mean observed relative price for Plan H is close to the mean cost of the insured for this incremental coverage. Thus, we investigate counterfactuals related to the private market for supplemental coverage by taking the observed mean (relative) cost of the insured and observed share of individuals enrolled in Plan H as an approximation of a competitive equilibrium in this market, assuming average cost pricing.

methods to predict the counterfactual mean costs in a competitive market yield similar welfare estimates as our baseline strategy.

4 Empirical Application: Estimates and Welfare Analysis

4.1 Estimates

Table 5 displays the demand and cost estimates. Panel A columns (1) through (3) display demand estimates for Plan H coverage (relative to Plan M coverage). Columns (1) and (2) display linear demand estimates for Plan H coverage, without individual fixed effects (column 1) and with individual fixed effects (column 2). For comparison, column (3) displays linear-log demand estimates for Plan H coverage, where the maximum implied willingness-to-pay for the incremental coverage is constrained to be no greater than 1% of annual earnings.[39] Across these specifications, we see that an increase in relative premium for Plan H is associated with a statistically significant decline in the share of individuals with this disability plan. Drawing on the estimates from column (1), the magnitude implies that an increase in the premium for Plan H by 0.1% of annual earnings would lead the share of individuals enrolled in Plan H to decline by 7 percentage points (or 11% of the mean value). Table 5 Panel A also displays the mean implied willingness-to-pay for the incremental coverage based on these regression estimates, along with the associated bootstrapped standard error. Using these estimates from column (1), the implied annual mean willingness-to-pay for the incremental coverage provided by Plan H is 0.31% of annual earnings, or $202 for an individual earning $65K annually (roughly the mean annual earnings in the sample population). This estimate is precisely estimated with a 95% confidence interval that allows us to rule out estimates less than 0.27% of annual earnings or more than 0.34% of annual earnings.[40] Table 5 Panel B displays the mean cost for providing the incremental coverage for everyone and separately for those who do and those who do not purchase supplemental coverage. Using the mean cost estimates to contextualize the mean willingness-to-pay estimates, we see that the mean willingness-to-pay for incremental coverage in the population is 279% of the mean cost of providing this coverage, indicating that the incremental disability coverage generates substantial surplus. The implied valuation of the incremental coverage is similar when employing estimates from a demand specification with individual fixed effects (column 2) or a linear-log specification of demand (column 3).

In addition to presenting the baseline Plan H demand estimates that we use for the welfare analysis below, Table 5 also displays the demand estimates for the incremental coverage provided by Plan M. There are two important things to note about these estimates as compared to analogous estimates for Plan H coverage. First, while the negative premium coefficient estimates in Panel A columns (4) and (5) suggest that the demand for the incremental coverage provided by Plan M is price-sensitive, the coefficient estimates and the corresponding mean implied willingness-to-pay for coverage are much less precisely estimated than the analogous estimates for Plan H coverage. The difference in precision across these estimates is not surprising given the limited variation in Plan M premiums during the sample period relative to the variation in the Plan H premium (see Table 2). Second, though the Plan M demand estimates are statistically imprecise, the general pattern of the mean implied willingness-to-pay estimates is consistent with the assumption we use to derive a lower bound on the total welfare (Assumption 1). As noted in Section 3, Assumption 1 implies that an individual's willingness-to-pay for incremental coverage is decreasing in the replacement rate of the baseline

[39] Because an unconstrained linear-log specification yields somewhat unrealistically high valuations of LTD in the upper tail and our aim is to obtain a lower bound, we constrain the estimated linear-log demand curve such that the maximum willingness-to-pay for Plan H incremental coverage is 1% of annual earnings.

[40] This confidence interval is based on the bootstrapped standard error on the mean willingness-to-pay reported in Table 5, where we utilize 100 bootstrap samples.

coverage, and thus an individual's willingness-to-pay for Plan M is weakly greater than his/her willingness-to-pay for Plan H.[41] Consistent with this assumption, the estimate for the mean willingness-to-pay for Plan M coverage is higher than the estimated mean willingness-to-pay for Plan H coverage, though the Plan M estimates are imprecise enough that these estimates are not statistically distinct. While the baseline welfare analysis below uses the more precise (and more conservative) Plan H demand estimates, Appendix D reports the analogous welfare measures based on the Plan M demand estimates for comparison.

Below, we investigate welfare under several counterfactuals where we use the demand estimates from the baseline linear specification (Table 5 Panel A column 1). In Section 4.3, we show that the main welfare estimates are very similar when using estimates from alternative specifications of the demand for Plan H coverage, including specifications with richer controls (e.g., demographic controls, individual fixed effects, etc.) and specifications with alternative demand forms (e.g., probit demand, logit demand, linear-log demand, log-log demand, etc.).

4.2 Implied Welfare Analysis

Next, we present the welfare analysis based on these estimates. For each counterfactual considered, Table 6 reports the mean willingness-to-pay and welfare per dollar of annual earnings using the demand estimates from Table 5 column 1. In addition, we also present two scaled versions of these measures to ease interpretation. The table reports the mean willingness-to-pay and welfare as a percent of the mean per-capita compulsory insurer's cost associated with the counterfactual. To contextualize the relative estimates in terms of dollars, the table also reports the mean willingness-to-pay and welfare measures scaled by \$65K, approximately the mean annual earnings in the sample. Note that scaling the estimates by the mean annual earnings provides an unbiased estimate of these welfare measures in dollar terms if earnings are uncorrelated with the relative willingness-to-pay for disability coverage.[42]

Table 6 Panel A considers the welfare associated with an extension of the compulsory baseline coverage to include the incremental coverage, where the difference between the counterfactuals in rows (1) and (2) is the benchmark that welfare is measured against: row (1) considers welfare of an extension relative to the benchmark of no available incremental insurance coverage and row (2) considers welfare of an extension relative to the benchmark of a perfectly competitive private market for the incremental insurance. If the relevant benchmark is no supplemental insurance market, then the per-capita mean willingness-to-pay associated with an extension is 0.31% of annual earnings, or 279% of the mean cost of the incremental coverage. The implied per-capita welfare associated with the extension is 0.20% of annual earnings, or \$129 annually when scaled by \$65K. This analysis indicates that individuals in this population highly value the incremental disability coverage starting from a high baseline level of coverage: on average, individuals value the incremental 10 percentage point increase in the replacement rate at more than two and a half times the cost of providing this coverage. Comparing row (1) to row (2), we see that a substantial amount of the value of an insurance extension is crowded-out if individuals would otherwise have access to a competitive market for this incremental insurance coverage. If a competitive private market for the incremental coverage is the relevant benchmark, an extension of compulsory coverage is associated with mean willingness-to-pay of 0.05% of annual earnings, and such an extension would actually reduce welfare according to the baseline estimates.[43] Overall, the marginal welfare analysis reveals that individuals, on average, highly value the in-

[41] Under Assumption 1 used to calculate the welfare bound (from Proposition 1), $\pi(\alpha + \delta, \alpha)$ is decreasing in α.

[42] Note that this scaling will underestimate the value of disability coverage in dollar terms if instead the relative surplus from disability coverage is positively correlated with annual earnings.

[43] While Table 5 Panel B provides evidence consistent with adverse selection in this context, the counterfactuals indicate that the adverse selection is not so dramatic as to call for mandating the purchase of the incremental coverage if individuals would otherwise

cremental coverage moving from a 60% to a 70% replacement rate, and there would be substantial surplus associated with increasing the generosity of compulsory insurance if individuals would not otherwise have access to this incremental insurance coverage.[44]

Next, we turn to the implied inframarginal welfare bound results reported in Table 6 Panel B. As discussed in the prior section, the inframarginal welfare analysis focuses on estimating a lower bound on the value of compulsory disability coverage relative to the benchmark of no available insurance for this risk. Rows (1) and (2) focus on bounding the value of compulsory coverage with a 60% replacement rate, where the difference between these counterfactuals is that row (2) accounts for the additional welfare generated by the optional private supplemental market that exists to top-up the baseline coverage. Without accounting for the opportunities individuals have to buy supplemental insurance, the lower bound on the mean willingness-to-pay for the baseline compulsory coverage is 1.9% of annual earnings, and the welfare associated with the baseline coverage is at least 1.2% annual earnings, or $777 when scaled by $65K. While the lower bound in row (1) is equally applicable regardless of the availability of supplemental insurance, we can obtain a tighter lower bound if we account for the additional surplus generated by the supplemental insurance that exists to top-up the compulsory coverage. Accounting for the possibility to buy supplemental insurance from a competitive market, we obtain that the mean willingness-to-pay for the baseline compulsory coverage is at least 2.1% of annual earnings and the associated net welfare is at least 1.5% of annual earnings, or $947 when scaled by $65K.

While the value generated by compulsory disability insurance with a 60% replacement rate is of interest within the empirical context, more generally we may want to know the value provided by public disability insurance, which has a lower replacement rate than the baseline coverage in the empirical setting. Row (3) presents the welfare bounds associated with disability insurance providing a 34% replacement rate (the mean public disability replacement rate among individuals in the sample) relative to a setting in which individuals have no access to disability insurance. This analysis reveals that the mean willingness-to-pay for public disability coverage is at least 1.1% of annual earnings, and the implied per-capita welfare of this coverage is at least 0.7% of annual earnings, or $440 when scaled by $65K. Overall, this approach produces a meaningful lower bound on the value of the inframarginal disability insurance coverage, which indicates that individuals in this population highly value disability insurance relative to the costs of providing this coverage.

have access to a competitive market for this incremental coverage. Note that the surplus arising from a competitive market for the incremental coverage is substantial, indicating that employees highly value the option to buy the incremental coverage provided by the most generous disability plan offered by this firm. The fact that the estimates indicate that the option to buy a 70% LTD plan generates so much surplus is particularly interesting as this LTD plan is more generous than the vast majority of LTD plans offered by employers nationally. The median replacement rate of employer-provided LTD coverage nationally is 60%, with only 1% of plans providing coverage in excess of 67% and 24% of plans providing coverage less than 60%. (These statistics are reported in Table 30 of the U.S. Department of Labor's National Compensation Survey, Employee Benefits in the U.S., March 2015 publication at: `http://www.bls.gov/ncs/ebs/benefits/2015/ownership/civilian/table30a.pdf`.)

[44] The primary take-away from the marginal welfare analysis is that the incremental disability coverage provides a large amount of surplus in this setting. This general qualitative result, that individuals highly value disability insurance on the margin, is largely in line with the findings from prior studies related to disability insurance. Prior studies by Low and Pistaferri (2015) and Chandra and Samwick (2005) employ calibrated life-cycle models to quantify the welfare associated with disability insurance. While the approaches in these calibration studies is quite different from the approach we employ, these prior studies provide some similar qualitative insights. Calibrations in these studies reveal that extending the generosity of SSDI would improve welfare, indicating that individuals value a marginal extension of coverage at more than the cost of providing this coverage. Given the large stakes associated with disability risk and the incomplete coverage available from other sources, it is perhaps not surprising that individuals would highly value an incremental increase in coverage for this risk. The threat of a career-ending disability is arguably the largest adverse financial shock faced by many individuals. Prior work by Dobkin et al. (2016) has found that while individuals tend to have very good coverage for the medical costs associated with health shocks, individuals have relatively little coverage for the earnings losses associated with severe health shocks. Our baseline results indicate that individuals in this context on average value the incremental 10 percent replacement rate at 279% of the cost of providing this incremental coverage, or approximately $200 annually for a worker with the mean earnings in this population.

4.3 Robustness

Next, we illustrate that the main welfare analysis is robust to a variety of alternative specifications. In the interest of brevity, the discussion focuses on the robustness of the estimated lower bound on the value of public coverage relative to no insurance for this risk. For each alternative specification considered, Table 7 Panel A displays the implied lower bound on mean willingness-to-pay and welfare generated by public insurance, and Table 7 Panel B displays the underlying parameter estimates corresponding to each of the alternative specifications. The first rows of each panel of Table 7 display the baseline welfare measures and linear demand estimates for reference.[45]

Because the welfare analysis requires us to extrapolate outside the empirical variation in premiums, the functional form of the demand curve could meaningfully influence the welfare estimates.[46] Thus, we assess the robustness of our findings by considering alternative demand forms. Rows (2) and (3) present estimates from alternative probit and logit demand specifications, respectively. The welfare bounds based on these alternative estimates are similar in magnitude to the baseline estimates and are statistically indistinguishable. An undesirable property of the linear demand estimates is that the implied willingness-to-pay for those with the lowest willingness-to-pay is below zero. To address this concern, we consider several alternative specifications. Row (4) displays alternative welfare measures based on the baseline linear demand curve truncated at an implied valuation of zero. This alternative piecewise linear demand curve yields a larger lower bound on welfare: 0.9% of annual earnings (compared to the baseline estimate of 0.7% of annual earnings). In addition, we estimate alternative specifications which restrict attention to positive valuations of insurance by including the premium in logs as opposed to levels. Row (5) displays the linear-log specification (based on the estimates from Table 5 discussed above), while Row (6) displays a log-log specification where both quantity and premiums enter in logs. Because unconstrained linear-log and log-log specifications yield somewhat unrealistically high valuations of LTD in the upper tail and our aim is to obtain a lower bound, we constrain the estimated demand curves in Rows (5) and (6) such that the maximum willingness-to-pay for Plan H incremental coverage is 1% of annual earnings. These alternative specifications yield broadly similar welfare bounds as the baseline specification. Overall, we obtain a meaningful lower bound on the value of public coverage across all of these demand specifications, with the lower bound estimates ranging from 0.5% to 0.9% of annual earnings.

Next, we consider the robustness of the analysis with respect to the variation used to estimate demand. Because the identification relies on over-time variation, it is important that individuals in the sample in different years are comparable to one another. In addition to the evidence in Table 4 that individuals across time look similar in terms of observable characteristics, we illustrate that our results are robust to the inclusion/exclusion of a wide range of controls building further confidence in the identifying assumption. The baseline specification estimates Equation 24 where we include a time trend to control for any unrelated trends in coverage that could be correlated with the identifying variation. Row (7) displays the results associated with estimating Equation 24 with no controls, and Row (8) displays the results associated with estimating

[45] For reference, Panel B row (1) displays the baseline estimates (from Table 5 column 1), and Panel A row (1) displays the implied welfare bounds (from Table 6 Panel B row 3).

[46] As discussed in Section 2, the price variation available in this context moves the fraction of people buying the most generous disability plan from 60% to 64%. To evaluate the effect of extending the generosity of compulsory coverage, we need an estimate of the mean willingness-to-pay for this incremental coverage (the integral over the demand curve). Thus, we need to estimate a parametric demand curve to investigate such counterfactuals. In practice, our approach is to estimate many different common parametric demand forms and show that the bottom line welfare estimate is largely unchanged across these specifications. It is worth noting that if one wanted to restrict attention to counterfactuals within the identifying variation, one could still evaluate several such counterfactuals including: (i) what is the deadweight loss associated with asymmetric information in a competitive market for supplemental insurance relative to the first best? or (ii) what is the effect of a moderate-sized subsidy/tax on supplemental insurance?

Equation 24 with additional controls for demographic characteristics.[47] These alternative specifications yield very similar welfare bounds as in the baseline specification. We also estimate a specification of demand that includes individual fixed effects; the identification of this specification comes from individuals switching coverage in response to the change in relative premiums. Row (9) displays the results of the individual fixed effects specification, and these results are very similar to those obtained in the baseline specification. Row (10) illustrates that the welfare bounds look similar when using estimates from a propensity score re-weighted regression, where observations are re-weighted to look demographically alike to individuals in the first year of the sample. Taken together, the results in Table 7 rows (7) through (10) illustrate that the welfare bounds are qualitatively and quantitatively similar with the inclusion of more or fewer controls.

We also consider robustness with respect to the tax treatment of LTD premiums, LTD benefits, and public SSDI benefits. The baseline specification does not adjust for the tax treatment of LTD premiums, LTD benefits, or SSDI benefits. For the LTD plans offered at this employer, premiums are paid by employees with pre-tax dollars, and the benefits individuals receive from their LTD policy in the event of a disability are taxable. In contrast, SSDI benefits are largely exempt from tax.[48] Because it is not clear that employees know the tax treatment of each of these, premiums and the benefits are treated symmetrically in the baseline specification. Row (11) reports the results from an alternative specification that reflects the tax treatment of premiums and benefits, where we assume LTD premiums are paid with pre-tax dollars, LTD benefits individuals receive are taxable, and public SSDI benefits are not taxable (effectively raising the mean SSDI replacement rate from 34% to 52%). For simplicity, in this specification we assume that the same marginal tax rate is applicable to premiums and benefits, and we use a marginal tax rate of 35%, roughly the average marginal tax rate among individuals in this sample accounting for state and federal income taxes. While accounting for the tax treatment of LTD premiums and benefits leads to a lower implied surplus for a fixed extension of LTD coverage, accounting for the tax treatment of public SSDI benefits (and the effectively higher after-tax replacement rate) exactly offsets this giving us the same lower bound on the welfare of public disability insurance as in the baseline specification.[49]

Lastly, we consider robustness with respect to the treatment of moral hazard. The baseline welfare analysis is done under the assumption that supplemental disability insurance induces no moral hazard. While we find evidence consistent with no moral hazard in this context, the empirical setting provides limited statistical power to investigate moral hazard as discussed in Section 3. Thus, we investigate the sensitivity of welfare analysis to the incorporation of moral hazard using estimates from prior studies on disability insurance in other settings: Gruber (2000) and Autor, Duggan and Gruber (2012). See Appendix C for more details on this analysis. Note that the impact of assuming that the observed costs are subject to moral hazard depends

[47] Specifically, the demographic covariates included in this specification are age, sex, and race.

[48] See this SSA website describing the tax treatment of SSDI benefits: `https://www.ssa.gov/planners/taxes.html`.

[49] To see this, suppose that individuals face a constant marginal tax rate τ. In addition, suppose that compulsory public insurance benefits are not taxable and LTD benefits are taxable. Let $WelfareExtn_\tau(10 \times (1-\tau))$ denote the estimated welfare associated with a $(10 \times (1-\tau))$ percentage point increase in the disability replacement rate of compulsory government insurance (relative to no incremental coverage) given that LTD premiums are paid with pre-tax dollars and the marginal tax rate is τ. To account for the tax treatment of premiums, the mean willingness-to-pay should be scaled down by $(1-\tau)$ relative to the baseline estimates if the tax treatment is salient to consumers, and the baseline estimates for the government's costs should be symmetrically scaled down by $(1-\tau)$ to reflect costs associated with the incremental coverage net of tax revenue. Thus, it is easy to see that accounting for taxation leads the welfare associated with a $(10 \times (1-\tau))$ percentage point extension of the replacement rate of government provided disability coverage to be a scaled version of the baseline estimates of an extension not accounting for the tax treatment: $WelfareExtn_\tau(10 \times (1-\tau)) = (1-\tau)WelfareExtn_0(10)$. The lower bound on the value of government provided coverage is unchanged. To see this, let $WelfarePublic(\alpha, \tau)$ represent the welfare associated with the inframarginal public coverage accounting for the tax rate τ and the public insurance is of generosity α. Assuming public benefits are tax-exempt, it is easy to show that the lower bound on welfare is invariant to accounting for the tax treatment: $\{\text{Lower bound on } WelfarePublic(\alpha,\tau)\} = \alpha\frac{WelfareExtn_\tau(10\times(1-\tau))}{10\times(1-\tau)} = \alpha\frac{(1-\tau)WelfareExtn_0(10)}{10\times(1-\tau)} = \alpha\frac{WelfareExtn_0(10)}{10} = \{\text{Lower bound on } WelfarePublic(\alpha,0)\}$.

on the counterfactual considered. Relative to assuming the observed costs are not subject to moral hazard (as in the baseline welfare analysis), incorporating moral hazard reduces welfare in counterfactual scenarios with more insurance than observed in the status quo, and incorporating moral hazard increases welfare in counterfactual scenarios with less insurance than observed in the status quo. Table 7 rows (12) and (13) illustrate that the lower bound on the welfare generated by public coverage is robust to the inclusion of moral hazard.[50] In addition to investigating the robustness of this counterfactual, Appendix Table A3 displays the welfare analysis for each counterfactual we consider under the alternative moral hazard assumptions. Overall, these results illustrate that a marginal extension of disability insurance from a 60% to a 70% replacement rate generates substantial surplus even in the presence of moral hazard, and we obtain a meaningful robust lower bound on the implied welfare generated by the inframarginal coverage regardless of which of these assumptions on moral hazard is employed.

5 Conclusion

This paper outlined an approach to analyze welfare associated with compulsory public insurance relying on data and variation from the private market for complementary insurance. In settings in which private complementary insurance is available, this approach has several nice properties as compared to other methods to estimate welfare associated with compulsory public coverage, as it requires straightforward variation, commonly available data, and minimal assumptions to evaluate a wide range of welfare questions. Specifically, we described how the demand and cost curves from the private market for complementary insurance can be used to investigate counterfactuals associated with a marginal extension of the compulsory public coverage (a counterfactual that is within the set of observed contracts), as well as broader counterfactuals related to the inframarginal coverage provided by compulsory public insurance.

We illustrated this framework by applying it to the setting of disability insurance, employing administrative data on long-term disability policies offered through one large firm. The results indicate that employees at this firm value disability insurance highly. Starting from a baseline replacement rate of 60%, individuals value increasing the replacement rate by 10 percentage points at more than two and a half times the costs of such an extension. Based on our estimates, compulsory public disability insurance is associated with benefits that far exceed the costs of this coverage for this population, and extending the generosity of compulsory public coverage would result in substantial welfare gains for these employees if they would not otherwise have access to employer-provided supplemental disability coverage. More generally, the analysis highlights that individuals' willingness-to-pay for supplemental coverage can provide useful insights into the value of compulsory public insurance.

[50] Note that the welfare bounds in Table 7 are slightly larger when incorporating moral hazard. This is because the counterfactual under consideration (public insurance with no opportunity to buy supplemental coverage) involves less insurance than the status quo, and thus the predicted costs in this counterfactual will be lower if we assume moral hazard has contributed to the costs observed in the status quo. Importantly, an analogous lower bound on the welfare provided by public insurance which accounts for the firm's available LTD policies would be invariant to the degree to which there is moral hazard. (Because we observe costs in the status quo in which individuals have both public coverage and the firm's LTD policies, there is no need to calculate counterfactual costs in this scenario.) See Appendix B for more related discussion.

References

Autor, David, Andreas Ravndal Kostol, and Magne Mogstad. 2015. "Disability Benefits, Consumption Insurance, and Household Labor Supply." *mimeo*.

Autor, David H., and Mark G. Duggan. 2006. "The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding." *Journal of Economic Perspectives*, 20(3): 71–96.

Autor, David, Mark Duggan, and Jonathan Gruber. 2012. "Moral Hazard and Claims Deterrence in Private Disability Insurance." National Bureau of Economic Research Working Paper 18172.

Bundorf, M. Kate, Jonathan D. Levin, and Neale Mahoney. 2012. "Pricing and Welfare in Health Plan Choice." *American Economic Review*, 102(7): 3214–3248.

Cabral, Marika, and Neale Mahoney. 2014. "Externalities and Taxation of Supplemental Insurance: A Study of Medicare and Medigap." National Bureau of Economic Research Working Paper 19787.

Chandra, Amitabh, and Andrew A. Samwick. 2005. "Disability Risk and the Value of Disability Insurance." National Bureau of Economic Research Working Paper 11605.

Chetty, Raj. 2008. "Moral Hazard versus Liquidity and Optimal Unemployment Insurance." *Journal of Political Economy*, 116(2): 173–234.

Chetty, Raj, and Amy Finkelstein. 2013. "Social Insurance: Connecting Theory to Data." *Handbook of Public Economics*, 5: 111.

Chetty, Raj, and Emmanuel Saez. 2010. "Optimal Taxation and Social Insurance with Endogenous Private Insurance." *American Economic Journal: Economic Policy*, 2(2): 85–114.

Dobkin, Carlos Dobkin, Amy Finkelstein, Raymond Kluender, and Matthew J. Notowidigdo. 2016. "The Economic Consequences of Hospital Admissions." *mimeo*.

Eeckhoudt, Louis, and Christian Gollier. 2001. "Which Shape for the Cost Curve of Risk?" *The Journal of Risk and Insurance*, 68(3): 387–401.

Einav, Liran, Amy Finkelstein, and Jonathan Levin. 2010. "Beyond Testing: Empirical Models of Insurance Markets." *Annual Review of Economics*, 2: 311–336.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics*, 125(3): 877–921.

Einav, Liran, and Amy Finkelstein. 2011. "Selection in Insurance Markets: Theory and Empirics in Pictures." *Journal of Economic Perspectives*, 25(1): 115–38.

Finkelstein, Amy, Nathaniel Hendren, and Erzo F.P. Luttmer. 2015. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." National Bureau of Economic Research Working Paper 21308.

Gruber, Jonathan. 1997. "The Consumption Smoothing Benefits of Unemployment Insurance." *The American Economic Review*, 192–205.

Gruber, Jonathan. 2000. "Disability Insurance Benefits and Labor Supply." *Journal of Political Economy*, 108(6): 1162–1183.

Hackmann, Martin, Jonathan T. Kolstad, and Amanda E. Kowalski. 2015. "Adverse Selection and an Individual Mandate: When Theory Meets Practice." *American Economic Review*, 105: 1030–66.

Hendren, Nathaniel. 2013. "Private Information and Insurance Rejections." *Econometrica*, 81(5): 1713–1762.

Kostol, Andreas Ravndal, and Magne Mogstad. 2015. "Earnings, Disposable Income, and Consumption of Allowed and Rejected Disability Insurance Applicants." *The American Economic Review*, 105(5): 137–141.

Liebman, Jeffrey B. 2015. "Understanding the Increase in Disability Insurance Benefit Receipt in the United States." *Journal of Economic Perspectives*, 29(2): 123–50.

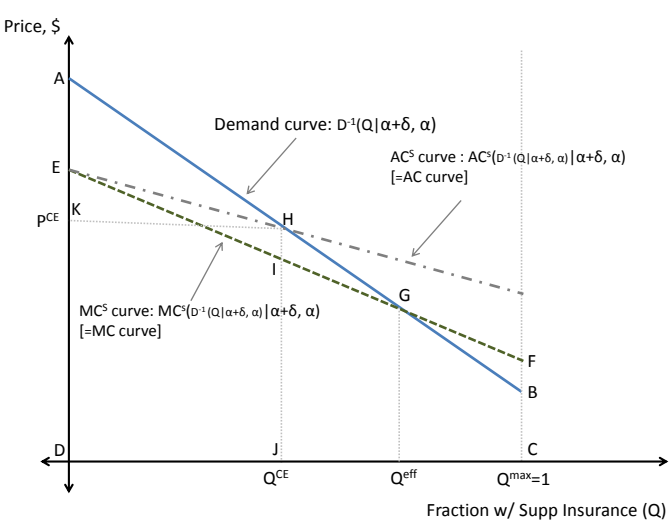
Low, Hamish, and Luigi Pistaferri. 2015. "Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off." *American Economic Review*, 105(10): 2986–3029.

Menezes, C. F., and D. L. Hanson. 1970. "On the Theory of Risk Aversion." *International Economic Review*, 11(3): 481–487.

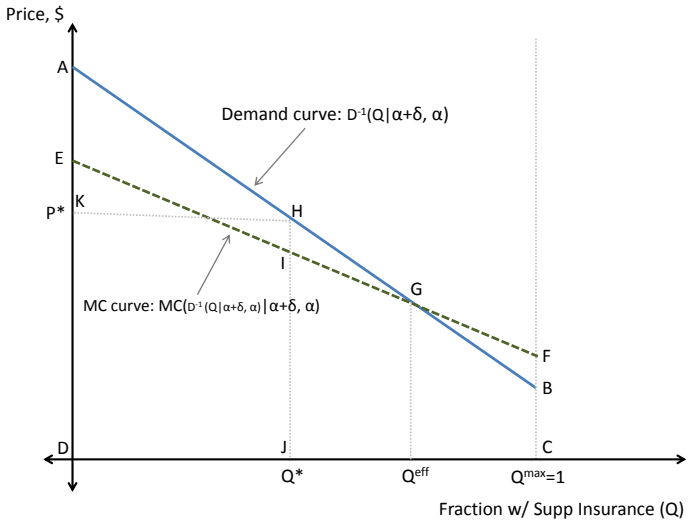
Meyer, Bruce D., and Wallace K.C. Mok. 2013. "Disability, Earnings, Income and Consumption." *National Bureau of Economic Research Working Paper 18869*.

Thomson, Sarah, Robin Osborn, David Squires, and Miraya Jun. 2013. "International Profiles of Health Care Systems, 2013: Australia, Canada, Denmark, England, France, Germany, Italy, Japan, the Netherlands, New Zealand, Norway, Sweden, Switzerland, and the United States." The Commonwealth Fund.

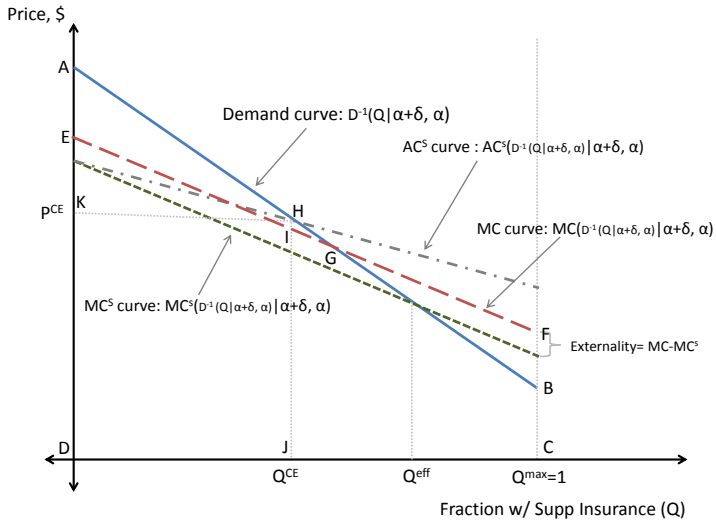
Figure 1: Graphical Illustration: Marginal Counterfactuals



(A) Competitive Eqm, No Externality



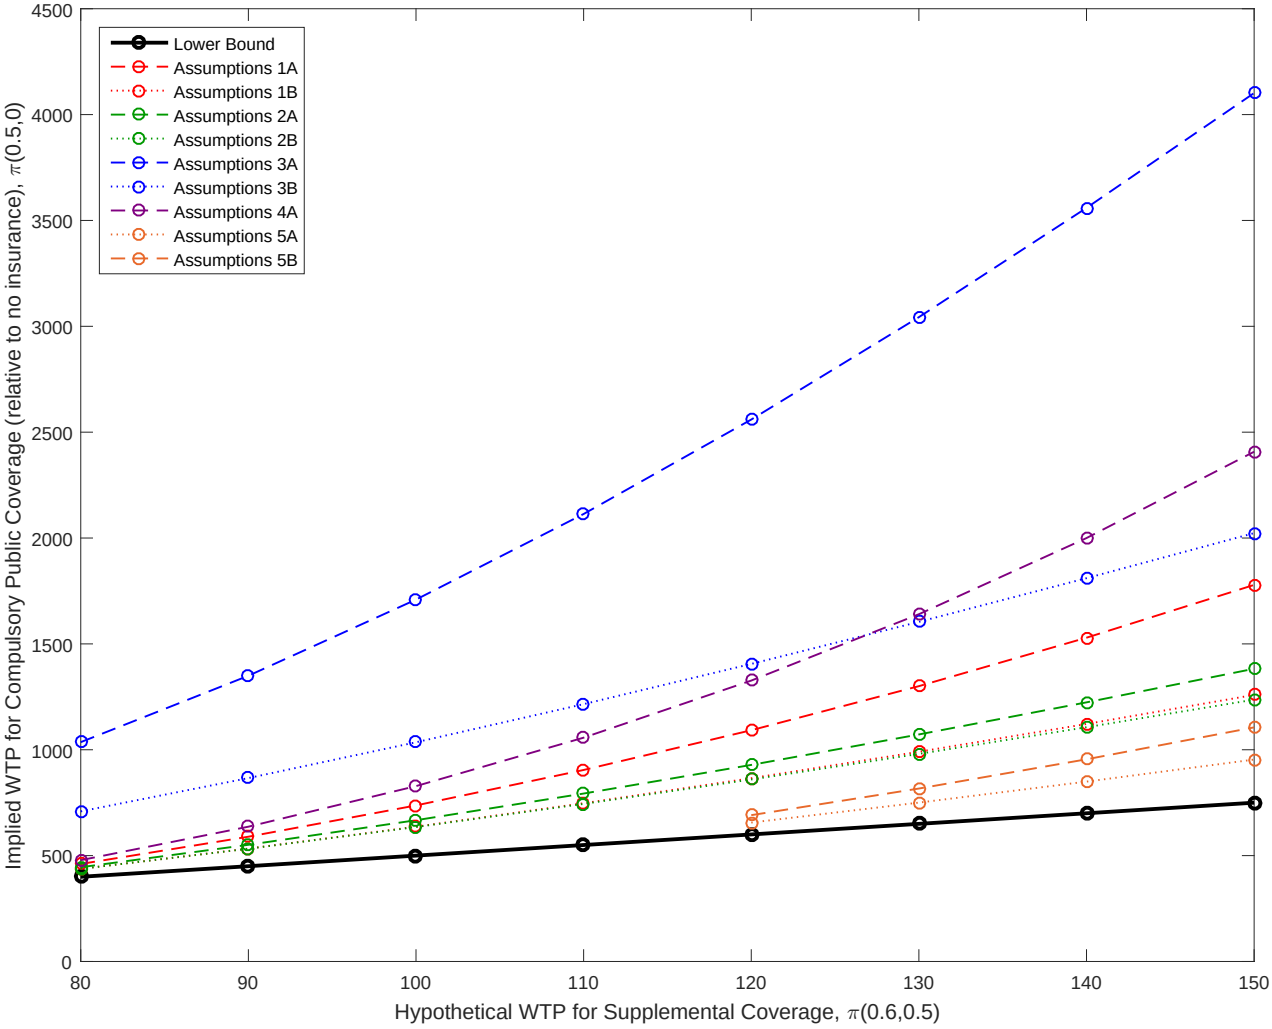
(B) Generic Eqm (P*, Q*)



(C) Competitive Eqm, Negative Externality

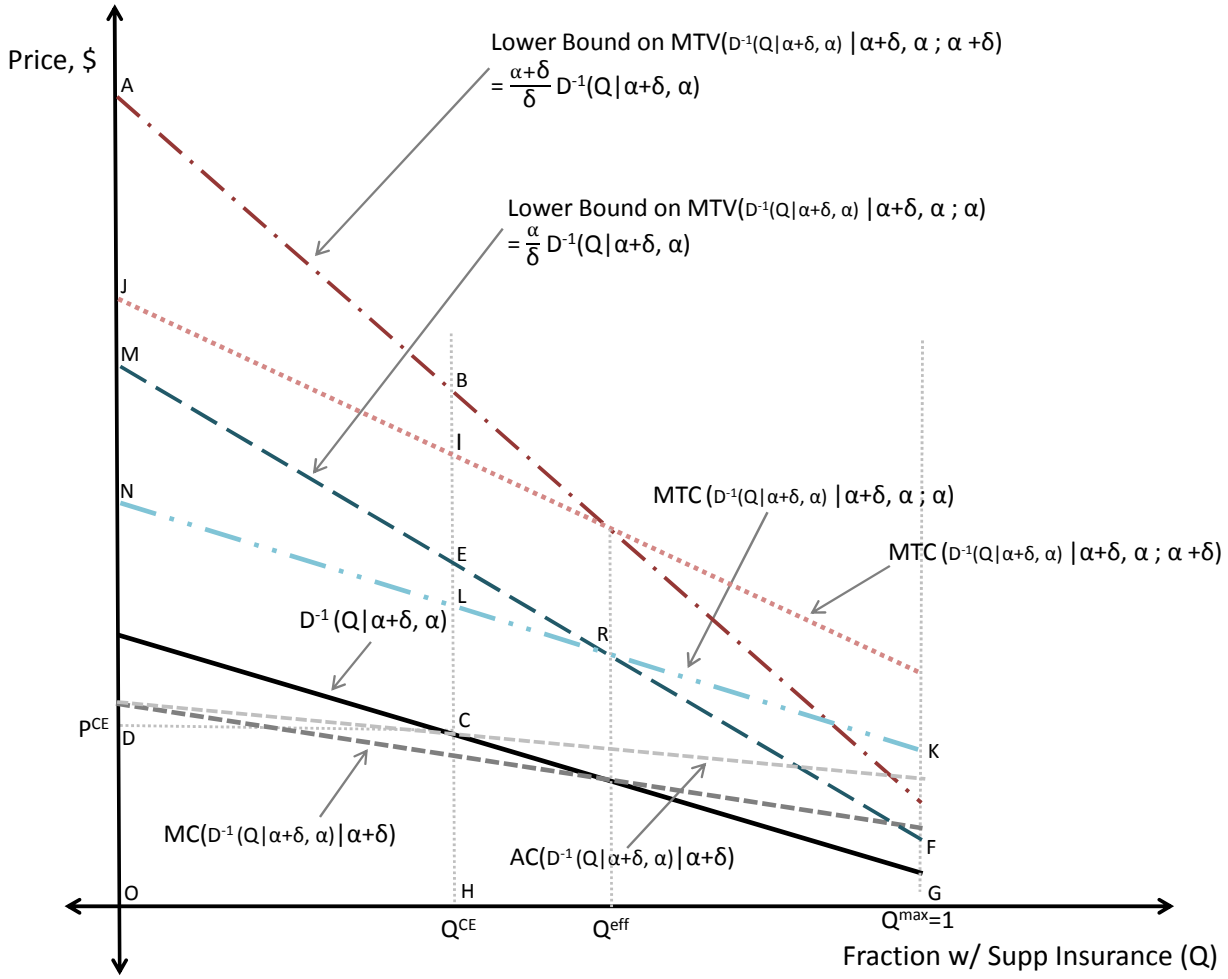
Notes: The above figure illustrates the welfare associated "marginal counterfactuals", or counterfactuals related to the incremental coverage sold in the existing market for supplemental insurance. A marginal counterfactual of particular interest is the welfare associated with extending the generosity of public insurance to include the incremental coverage provided by supplemental insurance. Panel (A) displays the case where there is no moral hazard and considers a competitive equilibrium as the relevant benchmark against which to measure the welfare associated with a public insurance extension. Panels (B) and (C) are similar, except Panel (B) considers a generic equilibrium as the relevant benchmark if public insurance were not extended, and Panel (C) displays the case in which the incremental coverage exerts a negative externality on the primary insurer which is not internalized by the supplemental insurer when setting prices in the benchmark perfectly competitive private supplemental insurance market. The three panels have been labeled such that the welfare measures can be similarly represented across these scenarios. Relative to a benchmark where no supplemental coverage is available, the average willingness-to-pay and the per-capita welfare generated by this extension are: $MeanWTP_Extn(\infty|\alpha+\delta,\alpha)$ = Area ABCD and $Welfare_Extn(\infty|\alpha+\delta,\alpha)$ = Area ABCD − Area EFCD. Relative to a (competitive) private supplemental insurance market, the welfare generated by extending compulsory coverage is: $Welfare_Extn(P^{CE}|\alpha+\delta,\alpha)$ = Area HGI − Area GFB = $Welfare_Extn(\infty|\alpha+\delta,\alpha)$ − Area AHIE, where this second expression makes clear that crowd-out of welfare in this scenario is represented by Area AHIE. The average willingness-to-pay for the extension relative to a (competitive) private supplemental insurance market is $MeanWTP_Extn(P^{CE}|\alpha+\delta,\alpha)$ = Area ABCD − Area AHK.

Figure 2: Numerical Example



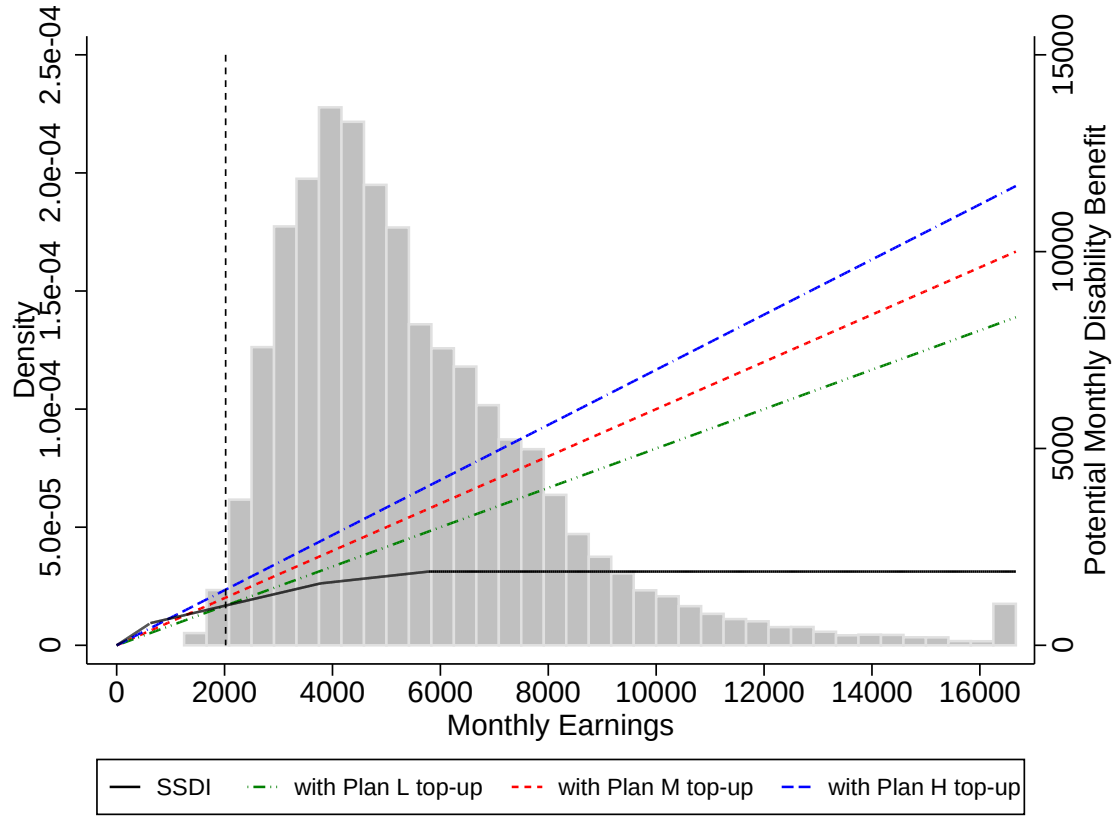
Notes: The above figure displays a numerical example comparing the bounding approach for inframarginal welfare analysis to a more structural analysis that specifies a particular utility function and distribution of risk. Consider an individual who faces some potential loss and has compulsory public insurance that covers 50% of this risk. Further, suppose the individual has the option to purchase private supplemental insurance to cover an additional 10% of the risk on top of the public insurance. Following the definitions in the text, let $\pi(0.5, 0)$ represent the individual's willingness-to-pay for compulsory public insurance relative to no insurance, and let $\pi(0.6, 0.5)$ represent the individual's willingness-to-pay for supplemental insurance to top-up the compulsory public insurance. Suppose the individual's preferences are represented by a univariate utility function $u(c)$, and the individual is risk averse with weakly decreasing absolute risk aversion. Based on Propositions 1 and 2, we get a lower bound on $\pi(0.5, 0)$ for any given value of $\pi(0.6, 0.5)$. The black solid line in the figure plots this lower bound. This bound is applicable regardless of the distribution of uncertainty the individual faces and regardless of what form the utility takes. To contrast this bound with a more structural approach, the figure displays the implied value of $\pi(0.5, 0)$ given a range of potential values of $\pi(0.6, 0.5)$ under various assumptions. Specifically, for each set of assumptions and each value of the willingness-to-pay for supplemental coverage ($\pi(0.6, 0.5)$), Table 1 reports the calibrated utility parameters to match this willingness-to-pay for supplemental insurance and the implied willingness-to-pay for the inframarginal coverage, $\pi(0.5, 0)$, based on these calibrated parameters. Table 1 displays the point estimates and the full set of assumptions underlying this figure.

Figure 3: Graphical Illustration: Inframarginal Counterfactuals



Notes: The above figure illustrates the lower bound on welfare associated with the inframarginal coverage provided by compulsory public insurance relative to the absence of insurance for this risk. Note that with sufficient data and variation in the supplemental insurance market, one can estimate the inverse demand curve for supplemental insurance ($D^{-1}(p|\alpha+\delta,\alpha)$), the marginal cost curve for supplemental insurance ($MC(p|\alpha+\delta,\alpha)$), and the marginal total cost curves as defined in the text ($MTC(p|\alpha+\delta,\alpha;\alpha)$ and $MTC(p|\alpha+\delta,\alpha;\alpha+\delta)$). Based on Proposition 1, under some assumptions the observed demand curve can be used to bound both $MTV(p|\alpha+\delta,\alpha;\alpha)$ and $MTV(p|\alpha+\delta,\alpha;\alpha+\delta)$, as displayed in the figure. Then, the lower bound on the mean willingness-to-pay associated with compulsory coverage given that a competitive supplemental market exists: $MeanWTP_Baseline(P^{CE}|\alpha+\delta,\alpha) \geq$ Area ABCD + Area EFGH, where the first term and second term, respectively, give us the mean willingness-to-pay for compulsory coverage among those who do and who do not purchase supplemental coverage. The analogous lower bound on welfare is: $Welfare_Baseline(P^{CE}|\alpha+\delta,\alpha) \geq$ Area ABIJ + Area ERL − Area RKF. The lower bounds on the mean willingness-to-pay and welfare generated by compulsory public coverage in a world in which supplemental insurance is not available is: $MeanWTP_Baseline(\infty|\alpha+\delta,\alpha) \geq$ Area MFGO and $Welfare_Baseline(\infty|\alpha+\delta,\alpha) \geq$ Area MRN-RKF.

Figure 4: Public and Private Disability Insurance Among Employees



Notes: This figure displays both: (i) the potential monthly disability benefit for public and private insurance for each level of monthly earnings (on the right vertical axis), and (ii) a histogram of the monthly earnings among salaried employees at the firm (on the left vertical axis). Referencing the right vertical axis, the solid line displays the SSDI public disability benefit formula for the year 2003 (the first year of our sample). The remaining three lines display the benefit formula for those who have one of the firm's three available top-up private LTD plans in addition to SSDI public disability insurance. As described in the main text, the estimation sample focuses on individuals with SSDI replacement rate less than the replacement rate of Plan L of 50%. In terms of this figure, these are the individuals whose monthly income lies to the right of the intersection of the "with Plan L top-up" line and the "SSDI" line. This level of monthly earnings based on the 2003 benefit formula is roughly $2,020 and this is indicated in the figure by the vertical reference line. In addition to displaying the public and private potential disability benefits, this figure also displays a histogram of the monthly earnings for all salaried employees (censored at $200,000/12). As can be seen in this figure, the vast majority of salaried employees at the firm have high enough earnings that their implied SSDI replacement rate lies below the 50% threshold used to define the estimation sample.

Table 1: Numerical Example

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Lower bound on π(0.5, 0)	Assumptions 1A		Assumptions 1B		Assumptions 2A		Assumptions 2B	
WTP for supplemental insurance, π(0.6, 0.5)		calibrated parameter	implied π(0.5, 0)	calibrated parameter	implied π(0.5, 0)	calibrated parameter	implied π(0.5, 0)	calibrated parameter	implied π(0.5, 0)
80	$400	0.36	$461	8.66E-06	$438	0.81	$445	1.77E-05	$437
90	$450	0.68	$589	1.63E-05	$533	1.53	$551	3.34E-05	$533
100	$500	0.96	$737	2.32E-05	$637	2.18	$667	4.75E-05	$635
110	$550	1.22	$904	2.94E-05	$747	2.77	$793	6.04E-05	$744
120	$600	1.46	$1,092	3.51E-05	$866	3.31	$928	7.22E-05	$860
130	$650	1.68	$1,300	4.04E-05	$990	3.81	$1,072	8.31E-05	$980
140	$700	1.88	$1,529	4.53E-05	$1,122	4.28	$1,224	9.33E-05	$1,106
150	$750	2.07	$1,779	4.98E-05	$1,259	4.71	$1,383	1.03E-04	$1,237
Assumptions									
utility	u"<0 and u has weakly decreasing absolute risk aversion	CRRA		CARA		CRRA		CARA	
risk, [loss (probability)]	no restrictions	loss=$0 (0.98), -$35k (0.02)		loss=$0 (0.98), -$35k (0.02)		loss=$0 (0.96), -$17.5k (0.04)		loss=$0 (0.96), -$17.5k (0.04)	
wage	no restrictions	wage=$50k		wage=$50k		wage=$50k		wage=$50k	

	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)
	Assumptions 3A		Assumptions 3B		Assumptions 4A		Assumptions 5A		Assumptions 5B	
WTP for supplemental insurance, π(0.6, 0.5)	calibrated parameter	implied π(0.5, 0)	calibrated parameter	implied π(0.5, 0)	calibrated parameter	implied π(0.5, 0)	calibrated parameter	implied π(0.5, 0)	calibrated parameter	implied π(0.5, 0)
80	2.21	$1,036	5.32E-05	$709	0.27	$478				
90	2.53	$1,350	6.09E-05	$866	0.51	$637				
100	2.81	$1,709	6.77E-05	$1,035	0.73	$828				
110	3.07	$2,114	7.39E-05	$1,215	0.92	$1,057	0.126531229	$577	3.05E-06	$567
120	3.30	$2,560	7.96E-05	$1,405	1.10	$1,327	0.363582661	$691	8.76E-06	$656
130	3.52	$3,043	8.49E-05	$1,604	1.26	$1,640	0.582136671	$817	1.40E-05	$750
140	3.72	$3,559	8.97E-05	$1,810	1.42	$2,000	0.784947907	$955	1.89E-05	$850
150	3.91	$4,102	9.43E-05	$2,023	1.56	$2,408	0.974199158	$1,106	2.35E-05	$954
Assumptions										
utility	CRRA		CARA		CRRA		CRRA		CARA	
risk, [loss (probability)]	loss=$0 (0.99), -$35k (0.01)		loss=$0 (0.99), -$35k (0.01)		loss=$0 (0.98), -$35k (0.02)		loss=$0 (0.97), -$35k (0.03)		loss=$0 (0.97), -$35k (0.03)	
wage	wage=$50k		wage=$50k		wage=$40k		wage=$50k		wage=$50k	

Notes: The above table describes a numerical example which compares the bounding approach described in the text to a more structural analysis that defines a particular utility function and distribution of risk. Consider an individual who faces some potential loss and has compulsory public insurance that covers 50% of this risk. Further, suppose the individual has the option to purchase private supplemental insurance to cover an additional 10% of the risk on top of the public insurance. Following the definitions in the text, let $\pi(0.5, 0)$ represent the individual's willingness-to-pay for compulsory public insurance relative to no insurance, and let $\pi(0.6, 0.5)$ represent the individual's willingness-to-pay for supplemental insurance to top-up the compulsory public insurance. Suppose the individual's preferences are represented by a univariate utility function $u(c)$, and the individual is risk averse with weakly decreasing absolute risk aversion. Based on Propositions 1 and 2, we get a lower bound on $\pi(0.5, 0)$ for any given observed $\pi(0.6, 0.5)$. The black solid line in the Figure 2 plots this lower bound. This bound is applicable regardless of the distribution of uncertainty the individual faces and regardless of what form the utility takes. To contrast this bound with a more structural approach, the figure displays the implied value of $\pi(0.5, 0)$ given a value of $\pi(0.6, 0.5)$ under various assumptions. Specifically, for each set of assumptions and each value of the willingness-to-pay for supplemental coverage ($\pi(0.6, 0.5)$), Table 1 reports the calibrated utility parameters to match this willingness-to-pay for supplemental insurance and the implied willingness-to-pay for the inframarginal coverage, $\pi(0.5, 0)$, based on these calibrated parameters. Figure 2 plots the point estimates from this table for easier comparison across these scenarios.

Table 2: Long-Term Disability Plan Details

	Coverage				
	Plan L	Plan M	Plan H		
Replacement Rate, All years	50%	60%	70%		
	Annual Premium (fraction of annual earnings)			Relative Premium (fraction of annual earnings)	
	Plan L	Plan M	Plan H	Plan M-Plan L	Plan H- Plan M
Year					
2003	0	0.00152	0.00315	0.00152	0.00163
2004	0	0.00152	0.00261	0.00152	0.00109
2005	0	0.00151	0.00259	0.00151	0.00108
2006	0	0.00162	0.00273	0.00162	0.00111
	Annual Premium (scaled by reference earnings of $60,000)			Relative Premium (scaled by reference earnings of $60,000)	
	Plan L	Plan M	Plan H	Plan M-Plan L	Plan H- Plan M
Year					
2003	0	91	189	91	98
2004	0	91	157	91	65
2005	0	91	156	91	65
2006	0	97	164	97	67
	Plan Enrollment in Estimation Sample				
	Plan L	Plan M	Plan H		
Year					
2003	18.7%	20.9%	60.3%		
2004	19.5%	16.3%	64.2%		
2005	20.4%	15.0%	64.6%		
2006	22.2%	13.5%	64.3%		

Notes: This table describes the three long-term disability plans available on the menu of plans offered to employees in the estimation sample. In the context of disability, both benefits and costs are stated in relative terms, relative to the annual earnings. The replacement rate for each plan describes the fraction of annual earnings paid out in the case of disability. In this table, premiums are stated as a fraction of annual earnings, and scaled by $60K, roughly the median annual earnings in the estimation sample.

Table 3: Summary Statistics

	Firm Data			March CPS 2004-2007 (corresponding to 2003-2006)			
	All Employees	Salary Employees	Employees in Estimation Sample	All Employees	Manufacturing Employees	Manufacturing White Collar Employees	Manufacturing White Collar Employees with DI RR <0.5 and wage<200k
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Employee-years	203,457	62,539	47,884	307,018	40,685	11,478	10,650
(unique employees)	(69,297)	(20,709)	(15,754)	(228,952)	(38,954)	(11,331)	(10,265)
Fraction Male	0.76	0.72	0.72	0.58	0.71	0.74	0.75
Fraction Salary	0.31	1.00	1.00	0.37	0.29	1.00	1.00
Mean Job Tenure	13.0	13.2	13.5	-	-	-	-
Fraction White	0.78	0.86	0.87	0.69	0.70	0.81	0.81
Age							
Mean	44.1	44.4	44.3	40.8	41.9	43.1	43.2
Median	45	45	45	41	42	43	43
Std Dev	10.3	9.5	9.2	11.1	10.6	9.5	9.4
Wage							
Mean	$44,401	$65,160	$66,668	$44,118	$47,578	$76,145	$73,293
Median	$36,516	$58,794	$59,986	$33,906	$36,813	$63,685	$65,134
Std Dev	$66,696	$28,521	$28,477	$49,209	$45,177	$62,994	$34,366
SSDI Replacement Rate							
Mean	42.2%	34.5%	34.0%	45.2%	43.6%	33.0%	32.0%
Median	44.1%	35.4%	35.4%	44.4%	43.8%	33.6%	33.1%
Std Dev	7.8%	8.8%	8.6%	15.7%	14.5%	12.6%	9.8%

Notes: This table displays summary statistics for employees in the data. All dollar quantities are inflation adjusted to 2005 dollars using the CPI-U. Column (1) displays summary statistics for all employees from the firm, column (2) displays summary statistics for all salary employees at the firm, and column (3) displays summary statistics for employees in the estimation sample. The estimation sample is restricted to salary employees who have the relevant menu of LTD plans offered to them, who have annual earnings less than $200,000, and who have an SSDI replacement rate less than 50%. Columns (4) through (7) display summary statistics for employees from the March Current Population Survey (CPS) for the corresponding years, 2004-2007 (which describe labor market outcomes from 2003-2006). Column (4) displays the summary statistics for all employees in the CPS, while column (5) displays summary statistics for all manufacturing employees in the CPS. Column (6) displays summary statistics for manufacturing employees in white collar occupations, where occupations are classified as white collar if they are managerial or professional occupations. Column (7) displays summary statistics for manufacturing employees with white collar occupations who have implied SSDI replacement rates less than 50% (using their reported monthly earnings as a proxy for their AIME) and annual earnings less than $200,000. As described in the text, the table reports the mean pre-tax SSDI replacement rate for workers in the sample using the SSA formula replacing the average indexed monthly earnings (AIME) with an individual's monthly earnings according to the firm's human resources records (or mean reported monthly earnings in the CPS in columns (4) through (7)). We note this is an imperfect calculation because we don't have data on each worker's full work history which would be needed to precisely calculate the AIME following the SSA formula.

Table 4: Identifying Variation

	Highest Relative Price	Lower Relative Price			Coefficient	p-value
	2003	2004	2005	2006		
	(1)	(2)	(3)	(4)	(5)	(6)
Fraction Male	0.72	0.72	0.73	0.73	6.23E-07	0.39
Mean Age	43.8	44.1	44.4	44.8	4.52E-08	0.38
Mean Job Tenure	13.4	13.5	13.5	13.6	-1.62E-08	0.79
Mean ln(Annual Earnings)	10.97	11.00	11.02	11.05	-1.99E-06	0.12
Fraction White	0.87	0.87	0.87	0.86	-8.14E-08	0.95
Predicted probability of LTD claim	0.0032	0.0032	0.0031	0.0031	1.02E-04	0.38
N	12,687	11,969	11,840	11,388		

Notes: This table displays summary statistics by the variation in premiums used for identification. The row labeled "Predicted probability of LTD claim" reports the mean fitted value from a regression of an indicator of LTD receipt on the demographic characteristics in the table, where we control flexibly for age and wage using indicators for deciles of these distributions. Columns (1) through (4) display summary statistics for employee demographics by year. The remaining columns display how the identifying variation in the relative price of Plan H ($p_{it}^H - p_{it}^M$) is related to the demographics in the table. Column (5) and Column (6) displays the coefficient and p-value associated with a regression of the demographic characteristic on the relative premium for Plan H and a time trend.

Table 5: Demand and Cost Estimates

Panel A: Demand Estimates					
	Plan H	Plan H	Plan H	PlanMorH	PlanMorH
	(1)	(2)	(3)	(4)	(5)
Relative premium Plan H	-71.26	-64.93			
	(11.77)	(12.21)			
Ln (relative premium Plan H)			-0.300		
			(.0073)		
(Relative) premium Plan M				-91.19	-38.15
				(38.30)	(35.34)
Constant	0.719	0.696	-1.347	0.952	0.867
	(.0262)	(.0152)	(.0346)	(.0595)	(.0511)
Controls					
Time Trend	x	x	x	x	x
Individual Fixed Effects		x			x
Dep Var					
Mean	0.63	0.63	0.63	0.80	0.80
Std dev	0.48	0.48	0.48	0.40	0.40
N	47,884	47,884	47,884	47,884	47,884
Implied Mean WTP for Incremental Coverage, $/annual earnings					
Estimate	0.00310	0.00328	0.00289	0.00481	0.00936
Std Error	(.00016)	(.00038)	(.00006)	(.00703)	(.34838)

Panel B: Mean Costs				
	Plan H		PlanMorH	
	Mean	Std Error	Mean	Std Error
Per capita cost, $/annual earnings				
All	0.00111	(.00014)	0.00111	(.00014)
Do Purchase incremental coverage (dep var=1)	0.00155	(.00021)	0.00133	(.00017)
Do not purchase incremental coverage (dep var=0)	0.00036	(.00013)	0.00023	(.00016)
Relative premium, $/annual earnings	0.00124	(1.09E-06)	0.00154	(1.95E-07)

Notes: The above table displays the demand and cost estimates. Panel A columns (1) and (2) display the results of estimating a linear demand specification for Plan H coverage, without individual fixed-effects (column 1) and with individual fixed-effects (column 2). Column (3) reports a linear-log specification for the demand for Plan H coverage, constraining the implied willingness-to-pay to be no more than 1% of annual earnings (i.e., constraining the demand curve to pass through p=0.01, q=0). For comparison, Panel A columns (4) and (5) display estimates of the demand for Plan M coverage (relative to Plan L coverage). There are two important things to note about these estimates as compared to analogous estimates for Plan H coverage. First, while the negative premium coefficient estimates in columns (4) and (5) suggest that the demand for the incremental coverage provided by Plan M is price-sensitive, the coefficient estimates and the corresponding mean implied willingness-to-pay for coverage are much less precisely estimated than the analogous estimates for Plan H coverage. This lack of precision is not surprising given the limited variation in Plan M premiums during the sample period relative to the variation in the Plan H premium (see Table 2). Second, though the Plan M demand estimates are statistically imprecise, the general pattern of the mean implied willingness-to-pay estimates is consistent with the assumptions we use to derive a lower bound on the total welfare (the assumptions within Proposition 1). The bootstrapped standard error on the mean implied willingness-to-pay for the incremental coverage (relative to annual earnings) is calculated using 100 bootstrap samples. Panel B displays the mean cost estimates for everyone and separately for those who do and do not purchase the plan corresponding to each column. As discussed in Section 4, the mean cost reported here is the mean present discounted value of disability claim costs relative to annual earnings. For reference, the relative premium for the plan corresponding to each column is listed below the mean cost estimates.

Table 6: Implied Welfare Analysis

	Mean WTP				Per-Capita Welfare				Mean Compulsory Insurer Cost	
	relative to annual earnings		% of compulsory insurer cost	scaled by 65k	relative to annual earnings		% of compulsory insurer cost	scaled by 65k	relative to annual earnings	
	(1)		(2)	(3)	(4)		(5)	(6)	(7)	(8)
	Est	Std Err			Est	Std Err			Est	Std Err
Panel A: Marginal Welfare Counterfactuals										
Extension of compulsory baseline 60% rep rate coverage to include incremental 10% rep rate coverage….										
(1) relative to no incremental coverage	0.00310	(.00038)	279%	202	0.00199	(.00043)	179%	129	0.00111	(.00014)
(2) relative to competitive private market for incremental coverage	0.00048	(.00037)	43%	31	-0.00063	(.00028)	-57%	-41	0.00111	(.00014)
Panel B: Inframarginal Welfare Counterfactuals										
	Lower Bound on Mean WTP				Lower Bound on Per-Capita Welfare				Mean Compulsory Insurer Cost	
	relative to annual earnings		% of compulsory insurer cost	scaled by 65k	relative to annual earnings		% of compulsory insurer cost	scaled by 65k	relative to annual earnings	
	(1)		(2)	(3)	(4)		(5)	(6)	(7)	(8)
	Est	Std Err			Est	Std Err			Est	Std Err
Compulsory baseline 60% rep rate coverage relative to no disability insurance…										
(1) no supplemental coverage available	0.01860	(.00231)	279%	1,209	0.01195	(.00259)	179%	777	0.00666	(.00083)
(2) competitive private market for supplemental coverage providing additional 10% rep rate	0.02122	(.00294)	319%	1,379	0.01457	(.00326)	219%	947	0.00666	(.00083)
(3) Compulsory disability insurance with mean public insurance rep rate (34%) (and no supplemental insurance) relative to no disability insurance	0.01053	(.0013)	279%	684	0.00677	(.00144)	179%	440	0.00377	(.00047)

Notes: This table displays the implied welfare measures using demand estimates from Table 5 column (1). Panel A displays the implied marginal welfare analysis outlined in Section 3.1.1. Panel B displays the inframarginal welfare analysis, where rows (1) and (2) focus on the baseline coverage offering 60% wage replacement and row (3) focuses on compulsory insurance with 34% wage replacement (the mean SSDI replacement rate in this population). The inframarginal welfare counterfactuals are outlined in Section 3.1.2. As discussed in Section 4, the mean cost reported here is the mean present discounted value of disability claim costs relative to annual earnings. This table reports bootstrapped standard errors using 100 bootstrap samples.

Table 7: Robustness

Panel A: Lower Bound on Value of Public Coverage Based on Alternative Specifications								
	Lower Bound on Mean WTP				Lower Bound on Per-Capita Welfare			
	relative to annual earnings		% of mean cost	scaled by 65k	relative to annual earnings		% of mean cost	scaled by 65k
	(1)		(2)	(3)	(4)		(5)	(6)
	Est	Std Err			Est	Std Err		
1. Baseline	0.01053	(.0013)	279%	684	0.00677	(.00144)	179%	440
Robustness to Alternative Demand								
2. Probit demand	0.01014	(.00125)	269%	659	0.00637	(.0014)	169%	414
3. Logit demand	0.01009	(.00124)	267%	656	0.00633	(.00139)	167%	411
4. Linear demand, wtp=max(linear fitted value,0)	0.01239	(.0021)	328%	805	0.00862	(.00229)	228%	560
5. Linear-log demand, constrained to go through (q=0, p=0.01)	0.00981	(.0003)	260%	638	0.00605	(.00045)	160%	393
6. Log-log demand, constrained to go through (q=0, p=0.01)	0.00827	(.00027)	219%	538	0.00450	(.00046)	119%	293
Robustness to Identifying Variation								
7. No Controls	0.01019	(.00123)	270%	662	0.00643	(.0015)	170%	418
8. Demographics included	0.01055	(.00119)	280%	686	0.00678	(.00135)	180%	441
9. Individual fixed-effects	0.01115	(.00127)	295%	725	0.00738	(.00138)	196%	480
10. Propensity Score Reweighted	0.01094	(.00162)	290%	711	0.00717	(.00176)	190%	466
Robustness to Accounting for Tax Treatment								
11. Baseline specification but incorporating tax treatment of premiums and benefits	0.01053	(.0013)	430%	684	0.00676	(.00144)	276%	439
Robustness to Alternative Moral Hazard Assumptions								
12. Moral Hazard as in Gruber (2000)	0.01053	(.00135)	388%	684	0.00782	(.00138)	288%	508
13. Moral Hazard as in Autor, Duggan, Gruber (2012)	0.01053	(.00135)	363%	684	0.00764	(.00139)	263%	497

Panel B: Parameter Estimates from Alternative Specifications						
	Demand				Mean Cost	
	Relative Premium		Constant			
	Coeff	Std Err	Coeff	Std Err	Mean	Std Err
	(1)	(2)	(3)	(4)	(5)	(6)
1. Baseline	-71.26	(11.77)	0.72	(.0262)	0.00111	(.00014)
Robustness to Alternative Demand						
2. Probit demand	-187.70	(37.67)	0.567	(.0563)	0.00111	(.00014)
3. Logit demand	-303.70	(61.07)	0.913	(.0913)	0.00111	(.00014)
4. Linear demand, wtp=max(fitted linear value,0)					0.00111	(.00014)
5. Linear-log demand, constrained to go through (q=0, p=0.01)	-0.30	(.001)	-1.347	(.0054)	0.00111	(.00014)
6. Log-log demand, constrained to go through (q=0, p=0.01)	-0.21	(.001)	-0.934	(.0037)	0.00111	(.00014)
Robustness to Identifying Variation						
7. No Controls	-75.30	(14.68)	0.726	(.0288)	0.00111	(.00014)
8. Demographics included	-72.63	(14.)	0.342	(.0239)	0.00111	(.00014)
9. Individual fixed-effects	-64.93	(12.21)	0.696	(.0152)	0.00111	(.00014)
10. Propensity Score Reweighted	-65.22	(14.28)	0.709	(.0213)	0.00111	(.00014)
Robustness to Accounting for Tax Treatment						
11. Baseline specification but incorporating tax treatment of premiums and benefits	-109.60	(21.82)	0.719	(.0211)	0.00072	(.0001)
Robustness to Alternative Moral Hazard Assumptions						
12. Moral Hazard as in Gruber (2000)	-71.26	(11.77)	0.72	(.0262)	0.00174	(.00026)
13. Moral Hazard as in Autor, Duggan, Gruber (2012)	-71.26	(11.77)	0.72	(.0262)	0.00163	(.00024)

Notes: This table illustrates that the main results are similar when using alternative functional forms for demand, using alternative controls isolating slightly different variation, accounting for the tax treatment of premiums and benefits, and employing alternative moral hazard assumptions. Panel A displays the implied lower bound on the value of public coverage based on the alternative specifications (analogous to that reported in Table 6 Panel B row (3)); Panel B displays the underlying demand and mean cost estimates associated with the incremental Plan H coverage (relative to Plan M coverage). As discussed in Section 4, the mean cost reported here is the mean present discounted value of disability claim costs relative to annual earnings. Panel A reports bootstrapped standard errors using 100 bootstrap samples. Each row represents the results of a distinct specification. For reference, row (1) displays the baseline estimates (from Table 5 column 1) and implied welfare bounds (from Table 6 Panel B row 3). Rows (2) through (6) display the estimates and implied welfare bounds using alternative specifications of the demand curve. Row (5) reports the results for a linear-log specification where the dependent variable is $PlanH_{it}$ and the main independent variable is $ln(p_{it}^H - p_{it}^M)$; the log-log specification in Row (6) reports results for a specification with the dependent variable $ln(PlanH_{it}+1)$ and the main independent variable $ln(p_{it}^H - p_{it}^M)$. Rows (7) through (10) display the estimates and implied welfare bounds concentrating on slightly different aspects of the identifying variation by including alternative sets of controls. Row (11) displays the estimates and implied welfare bounds accounting for the tax treatment of LTD premiums, LTD benefits, and SSDI benefits. In this specification, we make these adjustments assuming a constant marginal tax rate in this population of 35% (roughly the average marginal tax rate based on the observables we have in the data). Lastly, rows (12) and (13) display the implied welfare bounds when employing alternative moral hazard assumptions based on estimates in prior studies on disability insurance. See Appendix Section C for more details on the welfare analysis incorporating moral hazard.

APPENDIX

A Proposition 2: proof

Proposition 2. Suppose an individual's utility can be represented by the increasing, univariate function $u(c)$, so that $\pi(\theta, \gamma)$ is defined as:

$$Eu(w + (1-\theta)x - \pi(\theta,\gamma)) = Eu(w + (1-\gamma)x), \tag{26}$$

where the expectation is taken over $x \leq 0$, representing the uncertain losses the individual faces. Additionally, suppose the individual is risk averse and his/her utility exhibits (weakly) decreasing absolute risk aversion. Then,

$$\frac{\partial}{\partial\theta}\left(\frac{\pi(\theta,\gamma)}{\theta-\gamma}\right) \leq 0, \text{ and } \frac{\partial}{\partial\gamma}\left(\frac{\pi(\theta,\gamma)}{\theta-\gamma}\right) \leq 0, \tag{27}$$

and thus we obtain the results of Proposition 1.

Proof. Let $z = w + (1-\gamma)x - (\theta-\gamma)x - \pi(\theta,\gamma)$ and $y = \pi(\theta,\gamma) + (\theta-\gamma)x$. Define a function $F(\theta) = Eu(z+\theta y) - Eu(z)$.[51] Obviously, $F(0) = 0$. By definition of $\pi(\theta,\gamma)$, we know that $F(1) = 0$.
We want to show that

$$\frac{\partial}{\partial\gamma}\left(\frac{\pi(\theta,\gamma)}{\theta-\gamma}\right) \leq 0. \tag{28}$$

Note that this is equivalent to,

$$\frac{\pi(\theta,\gamma)}{\theta-\gamma} \leq -\frac{\partial}{\partial\gamma}\pi(\theta,\gamma) = \frac{-Exu'(w+(1-\gamma)x)}{Eu'(w+(1-\theta)x-\pi(\theta,\gamma))}. \tag{29}$$

Let us re-write this as follows:

$$\pi(\theta,\gamma)Eu'(w+(1-\theta)x-\pi(\theta,\gamma)) + (\theta-\gamma)Exu'(w+(1-\gamma)x) \leq 0. \tag{30}$$

Because the individual's utility exhibits decreasing absolute risk aversion, we know by Lemma 1 below that: $Eu'(w+(1-\gamma)x) \geq Eu'(w+(1-\theta)x-\pi(\theta,\gamma))$. So, the LHS of the above is bounded above

$$\pi(\theta,\gamma)Eu'(w+(1-\theta)x-\pi(\theta,\gamma)) + (\theta-\gamma)Exu'(w+(1-\gamma)x) \leq E\Big[\big(\pi(\theta,\gamma)+(\theta-\gamma)x\big)u'(w+(1-\gamma)x)\Big]. \tag{31}$$

Re-writing the RHS term using the above definitions,

$$E\Big[\big(\pi(\theta,\gamma)+(\theta-\gamma)x\big)u'(w+(1-\gamma)x)\Big] = E[yu'(z+y)] = F'(1). \tag{32}$$

Under risk aversion, F is concave. Because $F(0) = F(1) = 0$, we know that $F'(1) \leq 0$. □

Lemma 1. If $u(c)$ exhibits decreasing absolute risk aversion, then $Eu'(w+(1-\gamma)x) \geq Eu'(w+(1-\theta)x-\pi(\theta,\gamma))$

Proof. Decreasing absolute risk aversion means,

$$\frac{\partial}{\partial c}\left(\frac{-u''(c)}{u'(c)}\right) \leq 0. \tag{33}$$

This is equivalent to,

$$\frac{-u'u''' + (u'')^2}{(u')^2} \leq 0. \tag{34}$$

[51] The technique used for this part of the proof draws upon a similar technique used in Eeckhoudt and Gollier (2001).

Multiplying both sides by $\frac{u'}{u''}$ (≤ 0) and rearranging terms we get:

$$\frac{-u'''}{u''} \geq \frac{-u''}{u'}. \tag{35}$$

Notice that $v = -u'$ is a valid utility function as it is increasing in c. Since we know this holds for all c, $-u'$ is more risk averse than u, and thus we can represent $-u'$ as an increasing concave transformation of u: $-u' = \phi(u)$.

The result then follows from showing that the distribution of $u(w + (1-\theta)x - \pi(\theta,\gamma))$ second order stochastically dominates the distribution of $u(w+(1-\gamma)x)$ (and thus every concave utility function, including ϕ, prefers it). Consider two distributions, $F_A(z)$ and $F_B(z)$ with the same mean. Suppose these distributions are defined over a finite domain $[\underline{z}, \overline{z}]$, where $F_A(\underline{z}) = F_B(\underline{z}) = 0$ and $F_A(\overline{z}) = F_B(\overline{z}) = 1$. Showing $F_A(z)$ second order stochastically dominates $F_B(z)$ is equivalent to showing:

$$\int_{\underline{z}}^{y} [F_A(z) - F_B(z)]dz \leq 0, \qquad \forall y \in \{\underline{z}, \overline{z}\}. \tag{36}$$

Applying this to our context, we want to show:

$$\int_{\underline{z}}^{y} [F_x(\frac{u^{-1}(z) - w + \pi(\theta,\gamma)}{1-\theta}) - F_x(\frac{u^{-1}(z) - w}{1-\gamma})]dz \leq 0, \qquad \forall y \in \{\underline{z}, \overline{z}\}. \tag{37}$$

It is easy to show that the cumulative density functions cross exactly once at the point $z^* = u(w + \frac{1-\gamma}{\gamma-\theta}\pi(\theta,\gamma))$:

$$\begin{aligned} F_x(\frac{u^{-1}(z) - w + \pi(\theta,\gamma)}{1-\theta}) - F_x(\frac{u^{-1}(z) - w}{1-\gamma}) \leq 0 \qquad \forall z \in \{\underline{z}, z^*\}, \\ F_x(\frac{u^{-1}(z) - w + \pi(\theta,\gamma)}{1-\theta}) - F_x(\frac{u^{-1}(z) - w}{1-\gamma}) \geq 0 \qquad \forall z \in \{z^*, \overline{z}\}. \end{aligned} \tag{38}$$

Thus, we know that $\int_{\underline{z}}^{y} [F_x(\frac{u^{-1}(z)-w+\pi(\theta,\gamma)}{1-\theta}) - F_x(\frac{u^{-1}(z)-w}{1-\gamma})]dz \leq 0, \qquad \forall y \in \{\underline{z}, z^*\}$. In addition, it is straightforward to show that the fact that the distributions have the same means ($Eu(w + (1-\gamma)x) = Eu(w + (1-\theta)x - \pi(\theta,\gamma))$) implies that $\int_{\underline{z}}^{\overline{z}} [F_x(\frac{u^{-1}(z)-w+\pi(\theta,\gamma)}{1-\theta}) - F_x(\frac{u^{-1}(z)-w}{1-\gamma})]dz = 0$. Thus, it must be the case that $\int_{\underline{z}}^{y} [F_x(\frac{u^{-1}(z)-w+\pi(\theta,\gamma)}{1-\theta}) - F_x(\frac{u^{-1}(z)-w}{1-\gamma})]dz \leq 0, \qquad \forall y \in \{\underline{z}, \overline{z}\}$.

□

B Accounting for Supplemental Insurance: Obtaining a Tighter Lower Bound on the Value of Public Coverage

As discussed in Section 3, if private supplemental insurance only exists because of the existence of public disability coverage, then the lower bound on the value of public coverage described in the text in Equation 23 can be adjusted to account for this. Note that this adjustment will provide a tighter lower bound than the bound described in Equation 23, so long as the private supplemental market generates positive surplus (as is the empirically relevant case in this context). To illustrate this, we next turn to calculating a lower bound on the value of public disability insurance in the empirical context accounting for the private disability plans employees have available at this firm.

To calculate this lower bound, we need to extrapolate from our estimates to measure the value individuals obtain from having the option to purchase any of the disability plans available from the employer, including those plans for which we do not have sufficient variation to trace out the associated demand and cost curves (specifically, Plan M and Plan L). While the variation in the empirical context can inform us of the shape of the demand and cost curves in the market for an incremental 10% replacement rate on top of a baseline level of 60% wage replacement, these estimates do not tell us about the precise shape of the demand curve or the degree of selection within a counterfactual market for alternative linear disability insurance products that complement the baseline disability coverage. Thus, to investigate welfare under counterfactual scenarios

involving alternative insurance contracts, we make the following additional assumption:

Assumption 2. The implied ranking of individuals' willingness-to-pay for the observed disability coverage reflects the ranking of individuals' counterfactual willingness-to-pay for any linear disability insurance coverage.

Assumption 2 above states that the ranking of individuals' revealed willingness-to-pay for the observed incremental disability coverage reflects the ranking of individuals' willingness-to-pay for counterfactual disability insurance products for which there is no variation to identify the associated demand and cost curves. While this assumption allows us to extrapolate from our estimates to define a bound on the welfare provided by public disability insurance in settings involving counterfactual complementary private insurance products, it is important to note that Assumption 2 is potentially quite restrictive. The welfare measures discussed in the text place no across-individual restrictions on utility[52] ; in contrast, Assumption 2 implicitly places across-individual restrictions by restricting heterogeneity in the relative curvature of individuals' underlying utility functions.

Under Assumptions 1 and 2, we can calculate a tighter lower bound on the value of compulsory public coverage among individuals in the sample population in the scenario that the observed supplemental insurance plans would not exist in the absence of public disability insurance:

$$MeanWTP_Public(\varnothing, \text{Observed Supp Policies}) \geq \sum_{j \in \{L,M,H\}} \frac{RR_j}{0.1} E\big[\big(\pi(0.7, 0.6|\varphi)\mathbb{1}\big(\pi(0.7, 0.6|\varphi) \leq \pi_{s_j}\big)\big] - E\big[p_j Plan_j(\varphi, \mathbf{P})\big]$$

$$Welfare_Public(\varnothing, \text{Observed Supp Policies}) \geq \sum_{j \in \{L,M,H\}} \frac{RR_j}{0.1} E\big[\big(\pi(0.7, 0.6|\varphi)\mathbb{1}\big(\pi(0.7, 0.6|\varphi) \leq \pi_{s_j}\big)\big] - E\big[c(RR_j, 0|\varphi)\big) Plan_j(\varphi, \mathbf{P})\big],$$

where RR_j is the replacement rate of Plan j and p_j is the premium for Plan j. In practice, we calculate the bounds above fixing the observed LTD plan enrollment and assuming average cost pricing.[53] Table A2 column (2) displays the results, indicating that the lower bound on the welfare generated by public coverage in this employee population increases by 41% relative to the baseline estimates (from Table 6 Panel B row 3 and displayed in Table A2 column (1) for reference), from 0.7% of annual earnings to 1% of annual earnings, when we account for the additional surplus generated by the opportunity to buy supplemental disability insurance at this firm.

C Moral Hazard and Selection: Estimates and Robustness

C.1 Moral Hazard

Evidence from Empirical Setting As discussed in Section 3, the data and environment has both strengths and weaknesses in terms of capturing the costs associated with disability. A strength of the empirical environment is that there is individual-level administrative cost data, so that at least in principle, one could use the very same price variation used to estimate the demand curves to estimate moral hazard and selection (à la Einav and Finkelstein (2011)). However, a practical limitation is that it is difficult to obtain precise estimates of selection and moral hazard within the context of disability insurance, as realized costs are noisy proxies for expected costs given the low incidence of disability. In the face of these strengths and weaknesses, we take the following approach. First, we leverage the price variation in an instrumental variables framework and demonstrate that there is evidence consistent with no moral hazard in this setting. Second, we proceed

[52]The welfare bound described in Corollary 1 requires no across-individual restrictions. Thus, one could, in principle, estimate this bound non-parametrically making no across-individual restrictions if there is sufficient variation. However, the available variation in premiums (and thus insurance enrollment) in reality is often limited. Because we have a limited range of premium variation in the empirical setting, we proceed by making parametric assumptions on the demand for supplemental coverage to estimate the welfare bound.

[53]Although Plan L coverage is free to employees, it could be that employees are effectively paying for this coverage through wage reductions due to compensating differentials. Implicitly, the calculation of the mean willingness-to-pay bound in this counterfactual assumes that fixing the observed enrollment in LTD plans, beneficiaries are effectively paying the average cost for this coverage either through premiums or wage reductions.

with the welfare analysis under the assumption that the incremental disability coverage induces no moral hazard. We then demonstrate that the welfare analysis is qualitatively unchanged when employing alternative assumptions about moral hazard based on estimates from prior studies on disability insurance in other contexts.

Let c_{it} represent the present discounted value of the realized costs relative to annual earnings associated with providing individual i in year t an incremental 10% replacement rate paid in the event of disability.[54] We investigate the possibility of moral hazard by estimating the following reduced form equation:

$$c_{it} = \beta_0 + \beta_1(p_{it}^H - p_{it}^M) + \beta_2 p_{it}^M + \lambda X_{it} + \epsilon_{it}, \tag{39}$$

where $(p_{it}^H - p_{it}^M)$ and p_{it}^M are the relative prices of Plan H and Plan M, respectively. The demand results reveal that the demand for Plan H coverage and Plan M coverage are both responsive to the respective relative prices. Thus, the test for moral hazard is then a test on whether lower relative prices (and thus more people on more generous coverage) lead to higher claims. In addition to estimating the reduced form specification described above, we also estimate the analogous IV specification where we estimate the relationship between costs and the share of people with Plan H and Plan M coverage, instrumenting these shares with the variation in relative prices.

Table A1 Panel A displays these estimates. Across all the specifications, the estimates are statistically indistinguishable from zero, consistent with the incremental disability coverage inducing no moral hazard in this setting. Note that while the estimates are consistent with no moral hazard associated with the incremental coverage, the estimates are not very precise. This lack of precision is not surprising, given the low incidence of disability claims and the range of variation exploited in this IV strategy. However, based on the estimates in Table A1, we proceed with the baseline welfare analysis maintaining the assumption that the incremental disability coverage induces no moral hazard.[55]

Robustness to Incorporation of Moral Hazard We then investigate the robustness of our welfare analysis to incorporating moral hazard, using estimates obtained in different contexts by prior studies. While many of the prior studies on disability insurance investigate behavioral responses to other margins (such as stringency of the evaluation criteria, wages in the labor force, etc.), there have been few studies investigating the margin of interest in this setting: the response of disability claiming to benefit generosity. There are, however, two prior studies which investigate behavioral responses to disability insurance benefit generosity: Gruber (2000) estimates the impact of an increase in the replacement rate for disability insurance on labor force participation in Canada (leveraging differential variation in disability insurance replacement rates across provinces), and Autor, Duggan and Gruber (2012) use claims data from an LTD insurer to investigate the impact of more generous LTD coverage on LTD claiming rates. Note that neither one of these studies investigates a setting that is directly comparable to the empirical context, and neither of these studies investigates the impact of benefit generosity on the expected relative cost of disability claims (the measure we used as described in the text).[56] Although neither one of these papers estimates completely comparable elasticities, we re-calculate the welfare counterfactuals adjusting for moral hazard based on these prior elasticities in terms of labor force participation and LTD claim rates, assuming that the same semi-elasticity applies to the expected cost of disability claims (the cost measure we use, as described in the text). In this calculation, we assume that this semi-elasticity is constant across individuals and constant across disability benefit levels.

Table A3 displays the results. For reference, the table displays the baseline welfare analysis, under the

[54] Recall that we use an interest rate of 4% in the baseline estimation.

[55] Note that if we had found evidence of moral hazard, it would have been straightforward to incorporate this into the welfare analysis, as discussed in Section 1.

[56] These prior studies focus on different settings than the present context. Gruber (2000) studies the impact of public disability insurance generosity in Canada. Autor, Duggan and Gruber (2012) study the impact of disability insurance using data from LTD plans that look very different than those we study here. As discussed in Section 2, the LTD plans at the firm we study have a six month elimination period before individuals are eligible for benefits. More than two-thirds of the plans in the Autor, Duggan and Gruber (2012) study have a shorter elimination period, with the median plan having a elimination period of 90 days. Given that the Autor, Duggan and Gruber (2012) study finds that the length of the elimination period is an important determinant of claiming behavior, this difference across the settings could meaningfully impact the behavioral response to benefit generosity. In addition to the differences in the setting, these prior studies also estimate different elasticities than the elasticity of interest in the present setting. Gruber (2000) estimates the effect of benefit generosity on labor force participation and Autor, Duggan and Gruber (2012) estimate the effect of benefit generosity on the probability of claiming LTD (which is not directly analogous to the cost measure used here). Despite these differences, we use these prior estimates as a basis for alternative specifications of moral hazard in our setting to evaluate the robustness of the welfare analysis. See Table A3 for further details.

assumption that the incremental disability coverage does not induce moral hazard. In addition, the table reports alternative specifications where costs are adjusted using the elasticities from prior studies as discussed above. As expected, the impact of assuming that the costs we observe are subject to moral hazard depends on the counterfactual considered: Counterfactuals with more insurance than in the status quo are associated with lower welfare than the baseline estimates suggest, and counterfactuals with less insurance than the status quo are associated with greater welfare than the baseline estimates suggest.[57] There are several important things to note about the robustness of the results across these specifications. First, while the welfare associated with an incremental increase in the replacement rate from 60% to 70% is lower when moral hazard is incorporated, we still find that there is substantial welfare gains associated with this incremental coverage. That is, the benefits outweigh the costs for this incremental coverage, even after factoring in the potential moral hazard induced by this incremental coverage. Second, we get a meaningful lower bound on the value of the inframarginal coverage across all of the specifications, and this lower bound is qualitatively and quantitatively similar regardless of which assumption on moral hazard we adopt.

C.2 Selection

As discussed in Section 3, we measure welfare under two hypothetical scenarios with regard to the availability of supplemental insurance: (i) supplemental insurance is not available and (ii) supplemental coverage providing an additional 10% wage replacement is sold in a competitive market. As is clear from the framework discussed in Section 3, measuring welfare under the counterfactual scenario in which supplemental insurance is not available simply requires estimating the average cost of all individuals in the population (the mean value of c_{it} under the assumption of no moral hazard). In other words, our primary estimates of the welfare generated by the marginal and inframarginal coverage are independent of the degree of selection in the supplemental insurance market. However, to calculate additional welfare measures under a counterfactual scenario of a competitive supplemental insurance market, we need estimates of the average cost of those who do and those who do not buy insurance in a competitive supplemental insurance market (and thus will depend on the degree of selection in the supplemental market).

Table 5 Panel B in the text displays the mean incremental cost (the mean value of c_{it}) for everyone in the estimation sample and separately for those who do and do not purchase the incremental coverage. Note that the mean relative price for incremental coverage for each plan is very close to the average cost of those individuals willing to pay for the incremental coverage. In other words, the predicted competitive equilibrium in the market for the incremental coverage would be close to the observed pair of average prices and quantities regardless of the shape of the average and marginal cost curves far outside of the variation in the data. Thus, we take the observed mean (relative) costs of the insured and quantity of individuals insured as an approximation of the competitive equilibrium for the incremental coverage, which allows to use mean cost estimates for those with and without supplemental insurance in the data (reported in Table 5 Panel B) to compute the welfare measures discussed in the text. Note that a major advantage of relying on this approximation is that the cost estimates that enter the welfare analysis, in this case, are precisely estimated (see Table 5 Panel B).

While we use this approach for the baseline welfare analysis, we also explore alternative methods to predict counterfactual mean costs in a competitive market and demonstrate that these yield similar welfare estimates as our baseline strategy. For instance, one way to measure selection in this setting is to use the price variation used to estimate demand to trace out the marginal and average cost curves of the insurer, using the Einav and Finkelstein (2011) approach. Alternatively, another way one could measure selection in this setting would be to assume there is no moral hazard and thus bring in another moment on the cost curve, the average costs of insuring everyone under Plan H (the mean value of c_{it} for the population if there is no moral hazard). Comparing these approaches, the former approach leverages variation in costs among those marginal to the price variation in the data, while the latter approach provides a measure of global selection in this population. Figure A1 plots the average and marginal costs curves obtained from each of these approaches. (Table A1 Panel B column (1) reports the underlying point estimates for the approach leveraging the marginal price variation.) There are at least two things to note in this figure. First, we obtain very similar cost curves representing selection regardless of which method we use to measure selection in this setting. Second, regardless of the shape of the average and marginal cost curves far outside of the variation in the data, the observed average price and quantity across the sample period are close to the predicted

[57] The intuition for this is simple. If there is moral hazard, it will look less attractive to extend the generosity of coverage and more attractive to roll back the generosity of coverage, relative to a setting without moral hazard.

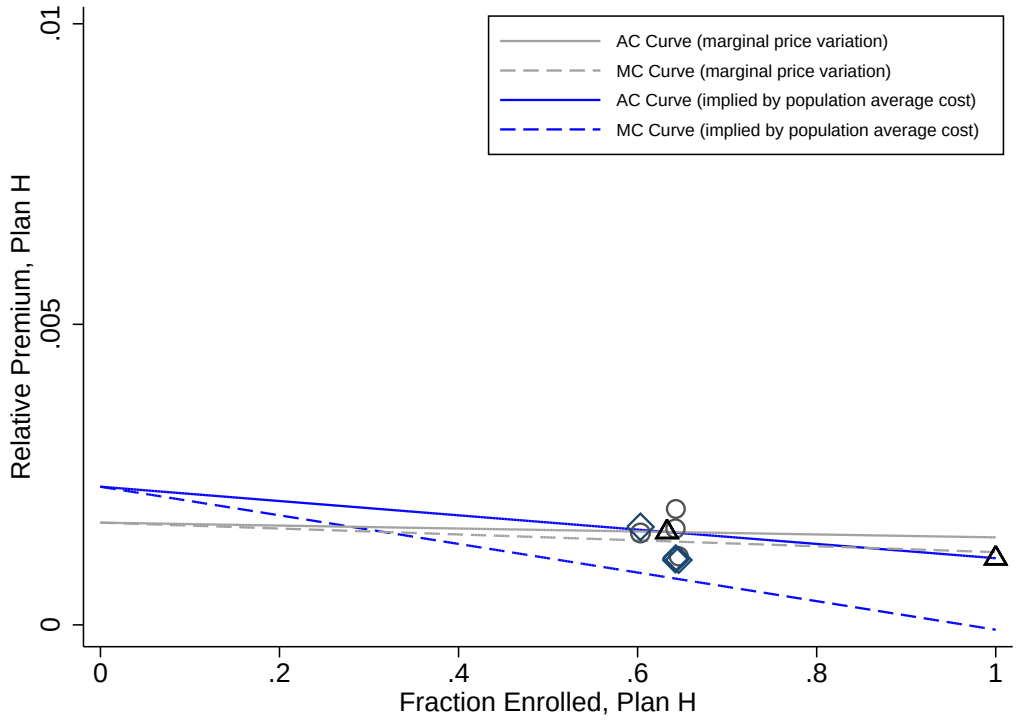
competitive equilibrium using either set of cost curves. Thus, our baseline approach to measuring welfare in counterfactuals that involve a competitive supplemental insurance market is to take the observed average (relative) costs of the insured and quantity as an approximation of the competitive equilibrium price for the incremental coverage, which allows us to use mean cost estimates for those with and without supplemental insurance in the data (reported in Table 5 Panel B in the text) to compute the welfare measures discussed in Section 3.

D Additional Robustness Analysis

In addition to the robustness analysis presented in the text, we examine the robustness of the results to additional alternative specifications. Table A4 displays the results of this additional analysis. Panel A displays the implied lower bound on the value of public coverage based on the alternative specifications, and Panel B displays the underlying regression estimates for the alternative specifications. Each row represents the results of a distinct specification. For reference, row (1) displays the baseline estimates and implied welfare bounds (based on Table 5 Panel A column (1)). Rows (2) through (4) display the estimates and implied welfare bounds using alternative specifications of the cost measures. The baseline analysis uses the present discounted value of costs using an interest rate of 4%, where we use the maximum potential duration for spells censored at five years (as described in the text, Section 2). Row (2) displays an alternative specification where costs are measured only through the first five years following the year of disability onset (the years for which we have complete claims data). This specification yields a less conservative bound than the baseline analysis. Rows (3) and (4) display specifications where costs are defined as the present discounted value of the maximum potential duration for truncated spells, using alternative interest rates to calculate the present discounted value. Comparing these specifications to the baseline estimates, we see that the analysis is not very sensitive to the chosen discount rate within a range of reasonable interest rates.

The baseline welfare analysis utilizes the demand for Plan H, estimated based on the variation described in Table 2. Leveraging the modest amount of variation in Plan M premiums, we can repeat the welfare analysis using the implied Plan M demand curve. Table A4 row (5) displays an alternative specification that relies on the Plan M demand estimates. As discussed in the text, the implied mean valuation of Plan M exceeds that of Plan H, so this alternative specification provides a less conservative lower bound on the value of disability coverage. However, as discussed in the paper, the relatively small amount of variation in Plan M premiums limits the statistical precision of welfare analysis based on the Plan M estimates as compared to the baseline welfare analysis which leverages the Plan H estimates.

Figure A1: Selection: Alternative Estimation Approaches



Notes: This table displays linear cost curves associated with Plan H coverage, where the curves are inferred using two different methods. The first method for estimating the average and marginal cost curves follows Einav and Finkelstein (2011) and uses variation in prices and the induced variation in the fraction of the population with Plan H coverage to trace out how the average expected cost of those on Plan H varies. The second method for inferring the average and marginal cost curves is to leverage the assumption of no moral hazard to interpolate between the average cost if the whole population were insured and the average cost of the actual insured population. These approaches are represented in the figure by the gray and blue lines, respectively. The gray circles represent the mean costs by year for those with Plan H coverage; the black triangles represent the mean costs of those with Plan H and everyone in the population, pooling data across the entire sample period. The blue diamonds represent the observed pairs of Plan H premium and enrollment by year (the variation used to estimate demand).

Table A1: Moral Hazard and Selection

	Panel A: Moral Hazard							
	Realized Costs of All Employees (2SLS)				Realized Costs of All Employees (Reduced Form)			
	Costs	Costs	I(LTD Claim)	I(LTD Claim)	Costs	Costs	I(LTD Claim)	I(LTD Claim)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Plan H	-0.00708	0.00236	-0.0339	-0.04				
	(.01)	(.02)	(.02)	(.04)				
PlanMorH	-0.0155	-0.053	-0.0599	-0.05				
	(.02)	(.08)	(.04)	(.16)				
Relative Premium Plan H					0.157	0.0292	1.102	2.57
					(.58)	(1.22)	(.97)	(1.94)
Premium Plan M					3.41	4.079	13.11	5.44
					(3.57)	(5.15)	(6.78)	(10.15)
Constant	0.018	0.0427	0.0724	0.0675	-0.00435	-0.00516	-0.0184	0.000509
	(.0197)	(.0582)	(.0397)	(.112)	(.00567)	(.00677)	(.0106)	(.00058)
Time Trend Included		x		x		x		x
Dep Var								
Mean	0.0011	0.0011	0.0032	0.0032	0.0011	0.0011	0.0032	0.0032
SD	0.0302	0.0302	0.0561	0.0561	0.0302	0.0302	0.0561	0.0561
N	47,884	47,884	47,884	47,884	47,884	47,884	47,884	47,884
	Panel B: Selection							
	Realized Cost of Employees on Plan H				Realized Cost of Employees on Plan M or Plan H			
	Costs	Costs	I(LTD Claim)	I(LTD Claim)	Costs	Costs	I(LTD Claim)	I(LTD Claim)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Relative Premium Plan H	0.0175	0.609	1.11	3.738				
	(.84)	(1.49)	(1.58)	(2.37)				
Premium Plan M					4.529	4.154	13.49	15.02
					(4.33)	(5.4)	(8.08)	(9.98)
Constant	0.00152	0.000552	0.00269	-0.00163	-0.00566	-0.0051	-0.0171	-0.0194
	(.00105)	(.00223)	(.00197)	(.00344)	(.00662)	(.00811)	(.0123)	(.0149)
Time Trend Included		x		x		x		x
Dep Var								
Mean	0.0015	0.0015	0.0041	0.0041	0.0013	0.0013	0.0037	0.0037
SD	0.0358	0.0358	0.0636	0.0636	0.0329	0.0329	0.0606	0.0606
N	30,306	30,306	30,306	30,306	38,227	38,227	38,227	38,227

Notes: This table displays the moral hazard and selection estimates which leverage the premium variation for LTD plans. Panel A displays the results related to moral hazard, where columns (1) through (4) display the IV results and columns (5) through (8) display the reduced form results. Panel B displays the selection estimates, where columns (1) through (4) focus on selection into Plan H. In addition, columns (5) through (8) focus on selection into at least Plan M coverage (relative to Plan L) leveraging the modest amount of variation Plan M premiums. The outcome variables considered in these regressions are two different cost measures: the present discounted value of costs (relative to annual earnings) incurred using the maximum potential duration for truncated spells as described in the text ("Costs", the baseline cost measure used in the text), and an indicator for having any LTD claim ("I(LTD Claim)").

Table A2: Additional Inframarginal Welfare Counterfactual

Additional Inframarginal Welfare Counterfactuals				
	Compulsory disability insurance with mean public insurance rep rate (34%) (and no supplemental insurance) relative to no disability insurance (1)		Compulsory disability insurance with mean public rep rate (34%) relative to no disability insurance (observed supp policies priced at average cost) (2)	
	Lower Bound Est	Lower Bound Std Err	Lower Bound Est	Lower Bound Std Err
Bound on Mean WTP				
per dollar of annual earnings	0.01053	(0.00130)	0.01336	(0.00193)
scaled by 65k	$684	(84.5)	$868	(125.5)
% of compulsory insurer cost	279%		354%	
Bound on Per-Capita Welfare				
per dollar of annual earnings	0.00677	(0.00144)	0.00959	(0.00214)
scaled by 65k	$440	(93.6)	$623	(139.1)
% of compulsory insurer cost	179%		254%	
Additional Assumptions				
Assumption 1	x		x	
Assumption 2			x	
Mean Costs - per dollar of annual earnings				
Compulsory Insurer Cost	0.00377	(.00047)	0.00377	(.00047)

Notes: This table considers an additional inframarginal welfare counterfactual, beyond those presented in the main text (in Table 6). This analysis draws upon the linear demand estimates in Table 5 column (1). Column (1) repeats the baseline estimates from Table 6 Panel b row (3) representing the lower bound on welfare provided by compulsory public insurance of generosity 34% relative to a setting where no one has access to disability insurance, without accounting for any extra surplus delivered by supplemental insurance. Column (2) presents a lower bound on the welfare provided by compulsory public insurance (of generosity 34%) relative to a setting where no one has access to disability insurance, accounting for the additional surplus generated by the optional observed LTD policies available to employees at the firm. As discussed in Section 4, the mean cost reported here is the mean present discounted value of disability claim costs relative to annual earnings. This table reports bootstrapped standard errors using 100 bootstrap samples.

Table A3: Additional Robustness Analysis: Incorporating Moral Hazard

Panel A: Marginal Welfare Counterfactuals, Baseline and Alternative Specifications

	Extension of compulsory baseline 60% rep rate coverage to include incremental 10% rep rate coverage			
	Relative to no incremental coverage		Relative to competitive private market for incremental coverage	
	(1)		(2)	
	Est	Std Err	Est	Std Err
Mean WTP				
per dollar of annual earnings	0.00310	(.00038)	0.00048	(.00037)
scaled by 65k	$202	(25)	$31	(24)
% of compulsory insurer cost	279%		43%	
Per-Capita Welfare				
Baseline Spec - no moral hazard				
per dollar of annual earnings	0.00199	(.00043)	-0.00063	(.00028)
scaled by 65k	$129	(28)	-$41	(18)
% of compulsory insurer cost	179%		-57%	
Alt Spec - moral hazard as in Gruber (2000)				
per dollar of annual earnings	0.00136	(.00048)	-0.00073	(.00028)
scaled by 65k	$88	(31)	-$47	(18)
% of compulsory insurer cost	78%		-60%	
Alt Spec - moral hazard as in Autor, Duggan, Gruber (2012)				
per dollar of annual earnings	0.00148	(.00047)	-0.00071	(.00028)
scaled by 65k	$96	(31)	-$46	(18)
% of compulsory insurer cost	91%		-60%	
Mean compulsory insurer cost - per dollar of annual earnings				
Baseline Spec - no moral hazard	0.00111	(.00014)	0.00111	(.00014)
Alt Spec - moral hazard as in Gruber (2000)	0.00174	(.00026)	0.00121	(.00018)
Alt Spec - moral hazard as in Autor, Duggan, Gruber (2012)	0.00162	(.00024)	0.00119	(.00018)

Panel B: Inframarginal Welfare Counterfactuals, Baseline and Alternative Specifications

	Compulsory baseline 60% rep rate coverage relative to no disability insurance				Compulsory disability insurance with mean public rep rate (34%) relative to no disability insurance (no supplemental insurance available)	
	No supplemental coverage available to top-up compulsory coverage		Competitive private market for supplemental coverage providing additional 10% rep rate			
	(1)		(2)		(3)	
	Est	Std Err	Est	Std Err	Est	Std Err
Lower Bound on Mean WTP						
per dollar of annual earnings	0.0186	(.00231)	0.02122	(.00294)	0.01053	(.0013)
scaled by 65k	$1,209	(150)	$1,379	(191)	$684	(85)
% of compulsory insurer cost	279%		319%		279%	
Lower Bound on Per-Capita Welfare						
Baseline Spec - no moral hazard						
per dollar of annual earnings	0.01195	(.00259)	0.01457	(.00326)	0.00677	(.00144)
scaled by 65k	$777	(168)	$947	(212)	$440	(94)
% of compulsory insurer cost	179%		219%		179%	
Alt Spec - moral hazard as in Gruber (2000)						
per dollar of annual earnings	0.01246	(.00256)	0.01356	(.00333)	0.00782	(.00138)
scaled by 65k	$810	(166)	$881	(216)	$508	(90)
% of compulsory insurer cost	203%		203%		288%	
Alt Spec- moral hazard as in Autor, Duggan, Gruber (2012)						
per dollar of annual earnings	0.01236	(.00256)	0.01357	(.00333)	0.00764	(.00139)
scaled by 65k	$803	(166)	$882	(216)	$497	(29)
% of compulsory insurer cost	198%		203%		263%	
Mean compulsory insurer cost - per dollar of annual earnings						
Baseline Spec - no moral hazard	0.00666	(.00083)	0.00666	(.00083)	0.00377	(.00047)
Alt Spec - moral hazard as in Gruber (2000)	0.00615	(.00091)	0.00668	(.00100)	0.00272	(.00044)
Alt Spec - moral hazard as in Autor, Duggan, Gruber (2012)	0.00624	(.00093)	0.00667	(.00099)	0.00290	(.00046)
Mean total resource cost (compulsory + supp insurer) - per dollar of annual earnings						
Baseline Spec - no moral hazard	0.00666	(.00096)	0.00761	(.00111)	0.00377	(.00054)
Alt Spec - moral hazard as in Gruber (2000)	0.00615	(.00091)	0.00766	(.00115)	0.00272	(.00044)
Alt Spec - moral hazard as in Autor, Duggan, Gruber (2012)	0.00624	(.00093)	0.00765	(.00115)	0.00289	(.00046)

Notes: This table presents the welfare analysis incorporating moral hazard using estimates from two prior studies. For reference, the table displays the baseline welfare analysis (linear specification from Table 6), under the assumption that the observed costs are not subject to moral hazard. When incorporating moral hazard, we take estimates from two prior studies: Gruber (2000) estimates that a 10 percentage point increase in the replacement rate would lead to a 10% decrease in labor force participation and Autor, Duggan and Gruber (2012) estimate that a 10 percentage point increase in the replacement rate would lead to a 8% increase in the probability of claiming LTD. Although neither one of these papers estimate completely comparable elasticities, we re-calculate the welfare counterfactuals adjusting for moral hazard, assuming that the estimated semi-elasticities in these prior studies apply to the cost measure of interest here as well: the expected present discounted value of disability claim costs relative to annual earnings. In using these prior estimates, we assume that this semi-elasticity is constant across individuals and across all levels of disability insurance generosity. This table reports bootstrapped standard errors using 100 bootstrap samples.

Table A4: Additional Robustness Analysis

Panel A: Lower Bound on Value of Public Coverage Based on Alternative Specifications

	Lower Bound on Mean WTP				Lower Bound on Per-Capita Welfare			
	relative to annual earnings		% of mean cost	scaled by 65k	relative to annual earnings		% of mean cost	scaled by 65k
	(1)		(2)	(3)	(4)		(5)	(6)
	Est	Std Err			Est	Std Err		
1. Baseline	0.01053	(.0013)	279%	684	0.0068	(0.00144)	179%	440
Robustness to Alternative Cost Measures								
2. Alternative Cost Measure, truncated at five years	0.01053	(.0013)	428%	684	0.0082	(0.00136)	332%	531
3. Alternative Cost Measure, discounted at 2%	0.01053	(.0013)	222%	684	0.0060	(0.00150)	127%	390
4. Alternative Cost Measure, discounted at 6%	0.01053	(.0013)	265%	684	0.0067	(0.00144)	170%	437
Robustness to Using Plan M Demand (instead of Plan H)								
5. Plan M Demand (time trend included)	0.01636	(.02363)	416%	1,063	0.0126	(0.02375)	320%	818

Panel B: Parameter Estimates from Alternative Specifications

	Demand				Mean Cost	
	Relative Premium		Constant			
	Coeff	Std Err	Coeff	Std Err	Mean	Std Err
	(1)	(2)	(3)	(4)	(5)	(6)
1. Baseline	-71.26	(11.77)	0.72	(.0262)	0.00111	(.00014)
Robustness to Alternative Cost Measures						
2. Alternative Cost Measure, truncated at five years	-71.26	(11.77)	0.72	(.0262)	0.00069	(.00007)
3. Alternative Cost Measure, discounted at 2%	-71.26	(11.77)	0.72	(.0262)	0.00134	(.00018)
4. Alternative Cost Measure, discounted at 6%	-71.26	(11.77)	0.72	(.0262)	0.00112	(.00014)
Robustness to Using Plan M Demand (instead of Plan H)						
5. Plan M Demand (time trend included)	-91.19	(38.3)	0.952	(.0595)	0.00111	(.00014)

Notes: This table displays additional robustness analysis (beyond that displayed in Table 7). Panel A displays the implied lower bound on the value of public coverage based on the alternative specifications (analogous to that reported in Table 6 Panel B row 3); Panel B displays the underlying demand and mean cost estimates associated with the incremental Plan H coverage (relative to Plan M coverage). As discussed in Section 4, the mean cost reported here is the mean present discounted value of disability claim costs relative to annual earnings. Panel A reports bootstrapped standard errors using 100 bootstrap samples. Each row represents the results of a distinct specification. For reference, row (1) displays the baseline estimates (from Table 5 column 1) and implied welfare bounds (from Table 6 Panel B row 3). Rows (2) through (4) display the estimates and implied welfare bounds using alternative specifications of the cost measure. The baseline analysis uses the present discounted value of costs using a 4% interest rate in this discounting, where we use the maximum potential duration for spells censored at five years (as described in the text, Section 2). Row (2) displays an alternative specification where costs are measured only through the first five years (the years for which we have claims data); rows (3) and (4) display specifications where costs are defined as the present discounted value of the maximum potential duration for truncated spells, using alternative discount rates. Row (5) displays an alternative specification that relies on the less precise Plan M demand estimates (rather than the baseline Plan H demand estimates) to value disability insurance. This table reports bootstrapped standard errors using 100 bootstrap samples.