

Risk-based Selection in Unemployment Insurance: Evidence and Implications

Camille Landais
LSE

Arash Nekoei
Stockholm University

Peter Nilsson
Stockholm University

David Seim
Stockholm University

Johannes Spinnewijn*
LSE

October 6, 2017

Abstract

This paper studies whether adverse selection can rationalize a universal mandate for unemployment insurance (UI). Building on a unique feature of the unemployment policy in Sweden, where workers can opt for supplemental UI coverage above a minimum mandate, we provide the first direct evidence for adverse selection in UI and derive its implications for UI design. We find that the unemployment risk is more than twice as high for workers who buy supplemental coverage, even when controlling for a rich set of observables. Exploiting variation in risk and prices to control for moral hazard, we show how this correlation is driven by substantial risk-based selection. Despite the severe adverse selection, we find that mandating the supplemental coverage is dominated by a design leaving the choice to workers. In this design, a large subsidy for supplemental coverage is optimal and complementary to the use of a minimum mandate. Our findings raise questions about the desirability of the universal mandate of generous UI in other countries, which has not been tested before.

Keywords: Adverse Selection, Unemployment Insurance, Mandate, Subsidy

*We thank Francesco Decarolis, Liran Einav, Itzik Fadlon, Amy Finkelstein, Francois Gerard, Nathan Hendren, Simon Jaeger, Henrik Kleven, Neale Mahoney, Magne Mogstad, Emmanuel Saez, Frans Spinnewyn and seminar participants at Marseille, Stanford, Ecole Polytechnique, UCL, Cambridge, Bonn, Einaudi, AEA Meetings, IFS, San Diego, NBER PF/Insurance Spring Meeting, Amsterdam, Zurich, CEPR Public Policy Meeting and LSE for helpful comments and suggestions and our discussants Attila Lindner and Florian Scheuer for their valuable input. We also thank Yannick Schindler for excellent research assistance. We acknowledge financial support from the ERC starting grants #679704 and #716485 and FORTE Grant #2015-00490.

1 Introduction

While some features of the unemployment insurance (UI) system vary across countries (for instance, the level and time profile of unemployment benefits), the different UI systems share one striking similarity: the participation to UI is universally mandated and no coverage choice is offered to workers. Workers are forced to pay payroll taxes when employed and receive a set transfer when unemployed, which is not subject to choice. Neither do private markets exist for more comprehensive UI. Why do (almost) all countries mandate UI? Why is no coverage choice available? Are these optimal features of UI design? Despite the large existing literature on UI, these fundamental questions have so far been unanswered.

A universal mandate is seen as the canonical solution to the inefficiencies arising under adverse selection [see Akerlof [1970], Chetty and Finkelstein [2013]]. Indeed, it is well-known that adverse selection hinders efficient market function as low risks leave the market and put upward pressure on equilibrium prices. Adverse selection is arguably the culprit here, but there are two issues with this argument in the context of UI. First, since UI is universally mandated, the role of adverse selection in UI has never been tested before. Second, even when adverse selection is present, the government may do better by using alternative interventions that allow for choice (e.g., subsidies, minimum mandates with choice of supplemental coverage, etc.), which are common practice in all other social insurance programs (health insurance, old-age pensions, etc.). Our paper tries to address both issues. We provide first-time evidence on the presence and severity of risk-based selection into unemployment insurance and we develop a general framework to evaluate the desirability of a universal mandate relative to choice-based interventions using this evidence.

Our empirical analysis exploits the combination of an exceptional setting and rich, administrative data in Sweden. *All* Swedish workers are entitled to a minimum benefit level when becoming unemployed, but can opt to buy more comprehensive UI at a uniform premium set by the government.¹ The comprehensive plan has been heavily subsidized - the premium corresponds to only 25% of the difference between the average cost of providing the comprehensive plan and the average cost of providing the basic plan. This subsidy has encouraged more than 80% of workers to buy the comprehensive plan. We observe the UI choice of the universe of Swedish workers and can link these choices to their unemployment histories registered by the Public Employment Service. We merge this data with a rich collection of household and firm registers, providing extremely detailed information on the determinants of workers' unemployment risk and insurance choices.

We present a set of empirical results, which provide direct and robust evidence that workers have private information about their unemployment risk, and act on this when making their unemployment insurance choice. This would create severe adverse selection in any UI market.

In a first step, following a prominent literature studying insurance markets, we perform so-called positive correlation tests, assessing whether workers who choose to buy comprehensive UI

¹Denmark, Iceland and Finland also run a voluntary UI program, historically administered by trade union-linked funds (the so-called Ghent system). This is the system many countries had in place before switching to compulsory insurance overseen by the government [see Carroll [2005]].

are more likely to be unemployed (see Chiappori and Salanié [2000]). Our estimates indicate that the unemployment risk for workers buying the comprehensive coverage is about 2.3 times the risk for workers who choose to stay on basic coverage. Interestingly, this large difference is robust to various measures of unemployment risk, but also to the introduction of a rich set of controls. Hence, even when absorbing the variation in risk coming from observables, workers' choices remain strongly correlated with unemployment risk, suggesting strong asymmetries in information that cannot be priced. In fact, some controls increase the positive correlation estimate, suggesting that these observables drive advantageous selection into insurance. For example, young workers are more likely to be unemployed, but also less likely to buy comprehensive coverage. In contrast, controlling for unemployment histories does substantially reduce the correlation between current UI choices and future unemployment.

In a second step, we go beyond the positive correlation tests, as the correlations may still be fully driven by moral hazard. We use risk and price variation to provide direct evidence of risk-based selection and estimate empirical moments relevant for welfare analysis:

First, we explain how and under what assumptions risk shifters, that affect individuals' unemployment probability conditional on their own actions, can be used to test for the presence of risk-based selection. This test is similar in spirit to the unused observables test in Finkelstein and Poterba [2014], and can reject that marginal costs curves are flat and unrelated to willingness to pay in both plans. We implement this test by exploiting various features of the Swedish labor market that provide variation in unemployment risk beyond the direct control of individuals. In particular, we focus on firm layoff risk and relative tenure ranking - two key determinants of an individual's unemployment risk due to the strict enforcement of the last-in-first-out principle in Sweden. First, we find that workers who switch firms increase their UI coverage and more so the higher the layoff risk in the new compared to the old firm. Second, we find that workers who are employed by a firm issuing a collective layoff notification, indicative of a shock to layoff risk at the firm level, increase their UI coverage and more so the lower their relative tenure ranking. The large responses confirm the importance of risk-based selection into UI.

Second, we provide additional evidence of risk-based selection following the approach proposed by Einav et al. [2010b], which consists in using price variation to identify *marginal* buyers and compare their unemployment risk to *infra-marginal* buyers of the same insurance plan. Price or policy variation allows estimating how the cost of providing either insurance plan changes, and therefore identifies the risk-based selection, given unpriced heterogeneity, that is relevant for assessing the welfare impact of changing these prices or policies. We exploit a large premium increase (following the first-time election of the right-wing party in Sweden) and provide evidence of significant risk based selection. In particular, we find that the marginal workers who stopped buying comprehensive coverage when the price increased face an unemployment risk that is 30% to 40% higher than inframarginal workers who did not buy comprehensive coverage, neither before nor after the premium increase, when all these workers are observed under the same basic coverage. We show that unpriced observables have a limited role in explaining the magnitude of risk-based

selection revealed by the price variation approach. The 2007 price reform also allows to investigate patterns of selection along other dimensions than unemployment probability. In particular, we use proxies for risk aversion and for the expected value of having unemployment insurance to reveal the presence of significant selection based on risk-preferences.

Our results provide evidence of significant risk-based selection in the Swedish UI system. Yet, a worry may be that observable risks are not priced in the current Swedish UI policy, and that severe risk-based selection could be easily avoided by conditioning the policy on more observables. The comprehensiveness and granularity of our data allows us to go beyond the specificities of the Swedish system and provide compelling evidence that even if a very rich set of observables were priced, severe adverse selection would remain in the UI market.

Despite the severe adverse selection, our estimates indicate that it is not efficient to mandate all Swedish workers to buy comprehensive coverage. Using a revealed preference approach, we can use prices to bound the value of supplemental coverage depending on whether a worker chooses to buy it or not. For workers who choose not to buy the supplemental coverage, its value is exceeded by our estimates of the average cost of providing the supplemental coverage to this group. Hence, we have identified workers for whom the welfare surplus from buying the comprehensive coverage is negative. Mandating them to buy the comprehensive coverage would decrease welfare. This is of course an important conclusion in light of the universal mandates of as comprehensive UI coverage in other countries, the desirability of which has never been tested before.²

We also use our empirical analysis to study the welfare implications of choice-based policy interventions under adverse selection. We build on the seminal work by Einav and Finkelstein [Einav et al. [2010b]], extending their framework to be able to analyze the desirability of both price and coverage interventions. We provide simple sufficient-statistic formulae that highlight the key trade-offs and allow linking our empirical estimates to the theory. A key result that we leverage in deriving these formulae is that the welfare impact of changes in insurance selection is fully captured by the corresponding fiscal externality, which simplifies to the price and cost differential of providing coverage to the marginal workers. The central trade-off when subsidizing comprehensive coverage is between reducing adverse selection into comprehensive insurance - captured by the corresponding fiscal externality - and redistributing to the workers buying comprehensive coverage. A minimum mandate is a complementary policy, as it mitigates the welfare loss from being priced out of comprehensive insurance, but worsens the adverse selection in the supplemental market. The central trade-off when setting the level of the minimum mandate is thus not just between providing insurance and maintaining incentives for those on the basic plan, as it also creates an adverse selection externality by attracting “good” risks away from the comprehensive plan.

Applying our formulae to the Swedish context, we find that the large subsidy for supplemental insurance, covering 75 percent of the difference in average cost of providing the comprehensive vs. basic coverage, is about optimal. The simple reason is that this subsidy is almost equal to the

²Examples of countries mandating UI with similar replacement rates as the voluntary, comprehensive plan in Sweden are Belgium, France, Luxemburg, Netherlands, Portugal, Spain and Switzerland. In other countries like the US and the UK, UI is also compulsory, but at lower replacement rates.

large wedge between average and marginal costs driven by adverse selection. This also implies that for assessing the minimum mandate, the large subsidy basically neutralizes the externality from worsening the adverse selection when increasing the UI benefit level of the minimum mandate. Hence, the minimum mandate can be evaluated using the standard Baily-Chetty formula [Baily [1978], Chetty [2006]], accounting for the fact that workers who opt to stay with the minimum mandate value insurance less, but are under lower coverage as well.

Our work contributes to different strands of literature. First, a large literature has analyzed the role of adverse selection in insurance markets. While the theoretical work dates back to the classical references by Akerlof [1970] and Rothschild and Stiglitz [1976], the surge in empirical work has been recent, pioneered by Chiappori and Salanié [2000] in the context of car insurance and rapidly extended to various insurance markets and settings [see Einav et al. [2010a]]. Our work highlights the advantages of using comprehensive, detailed and population-wide registry data to perform correlation tests, but also proposes new approaches to isolate exogenous risk variation and identify risk-based selection. Second, the lack of private markets and choices related to unemployment insurance, makes that the role of adverse selection in UI has been untested so far. Most related to our paper is the work by Hendren [2017], who analyzes elicited beliefs about job loss and finds that workers' private information on their unemployment risk is sufficient to explain the absence of a private market for supplemental unemployment insurance in the US (in addition to the public UI policy in place). Our paper complements Hendren's evidence with direct evidence based on actual insurance choices and studies the optimality of the public unemployment policy itself. Finally, there is a large literature studying the optimal trade-off between insurance and incentives in determining UI coverage [Baily [1978], Chetty [2006], Schmieder et al. [2012], Kolsrud et al. [2017]], which never considered potential selection effects when allowing for choice. On the other hand, a growing literature starting with the work by Einav and Finkelstein analyzes adverse selection and its welfare consequences, allowing for equilibrium pricing, but taking insurance coverage as given (e.g., Hackmann et al. [2015], Finkelstein et al. [2017]). Our framework tries to bridge these two strands of the literature, allowing not only to evaluate price subsidies, but also the coverage levels themselves. We are also providing applicable insights related to recent work by Veiga and Weyl [2016] and Azevedo and Gottlieb [2017] who characterize equilibria with both endogenous prices and coverages.

Our paper proceeds as follows. In Section 2 we describe the insitutional background and the data we use. In Section 3 we provide the results of our correlation tests relating unemployment risk to unemployment coverage. In Sections 4 and 5 we go beyond the correlation test using risk and price variation to provide direct evidence for risk-based selection. In Section 6 we provide a theoretical framework to analyze the welfare impact of different policy implementations, which we then link to our empirical estimates in the Swedish context. Section 7 concludes.

2 Context and Data

2.1 Institutional Background

Unemployment Insurance Sweden is with Iceland, Denmark and Finland, one of the only four countries in the world to have a voluntary UI scheme derived from the “Ghent system”. In practice, the Swedish UI system consists of two parts.

The first part of the system is mandated and provides basic coverage funded by a payroll tax (that we denote p_0). The benefits that unemployed receive with this basic coverage (b_0) are non-contributory (i.e., do not depend on the unemployed earnings prior to displacement). The benefit level of the basic coverage is low. During our period of analysis (2002-2009) the benefit level remained at 320 SEK per day (≈ 35 USD) which corresponds to a replacement rate of a little less than 20% for the median wage earner.³

The second part of the Swedish UI system is voluntary. By paying an insurance premium \mathbf{p} to UI funds (on top of the payroll tax), workers can opt for more comprehensive coverage.⁴ Upon displacement, workers who have continuously contributed premia for the comprehensive coverage during the past twelve months, get benefits b_1 , that replace 80% of previous earnings up to a cap, *in lieu* of the basic coverage b_0 . Workers are free to opt in or out of the comprehensive UI plan at any time. Apart from the level of benefits, there are no coverage differences between the basic and the comprehensive UI scheme. In particular, the potential duration of benefits b_0 and b_1 is the same, and was unlimited during our period of analysis. Moreover, to be eligible for either benefit upon unemployment, workers must fulfil a labor market attachment criterion, which is that they need to have worked 80 hours per month for six months during the past year.

The administration of the comprehensive UI coverage is done by 27 UI funds (*Kassa's*) but the government, through the Swedish Unemployment Insurance Board (IAF), supervises and coordinates the entire UI system. In particular, both the premia and benefit levels of the basic and comprehensive coverage are fully determined by the government. Importantly, the government does not allow UI funds to charge different prices to different individuals. One exception are union members who get a small rebate of $\approx 10\%$ on the UI premium for the comprehensive coverage.⁵ During our period of study, the government also did not allow premia to differ across UI funds. Premia paid by workers cover only a (small) fraction of benefits paid by the UI funds to eligible unemployed, and the government subsidizes UI funds for the difference out of the general budget.

Until January 1st of 2007, the monthly premium \mathbf{p} for the comprehensive coverage was homogenous across UI funds, at around 100 SEK, and a 40% income tax credit was given for the premia paid. In January 2007, the newly elected right-wing government increased the premium substantially and removed the income tax credit on premia paid to UI funds. It also introduced an

³Benefits are paid per “working day”, which means that there are 5 days of benefits paid per week. Benefits of 320 SEK a day therefore translate into 6960 SEK a month (≈ 765 USD).

⁴We denote the price of the comprehensive coverage as $p_1 = \mathbf{p} + p_0$.

⁵Note that individuals can still continue to contribute to UI funds while unemployed, to build eligibility in case of a future unemployment spell, in which case they are also entitled to paying a reduced premium.

additional fee that partly tied the premium of each UI fund to the average unemployment rate of that fund, starting from July 2008. In our analysis, and partly due to data availability, we focus on the period before July 2008 where insurance premia are homogenous across UI funds.

Historically, with the “Ghent system” in place, labor and trade unions played an important role in providing unemployment insurance in Sweden. Today’s 27 UI funds, which broadly correspond to 27 different industries/occupations, originated from unemployment insurance funds set up by unions. However, since the government overtook the responsibility of supervising the entire UI system in 1948, the links between UI funds and unions have loosened progressively.⁶ In our empirical analysis, we always control for trade union membership to account for the fact that union members face a different UI premium than non-members.

Layoff Notifications and Last-In-First-Out Principle In our analysis, we exploit variation in unemployment risk across individuals within a firm. Under Sweden’s employment-protection law, firms subject to a shock and intending to displace 5 or more workers simultaneously must notify the Public Employment Service in advance. Once a notification is emitted, employers need to come up with the list and dates for the intended layoffs. These layoffs may happen up to 2 years after the original notification has been sent. The list needs to follow the last-in-first-out principle. This means that workers get divided into groups, defined by collective bargaining agreements, and then a tenure ranking within each group is constructed.⁷ The more recent hires are displaced before workers with longer tenure. For firms with multiple establishments, one layoff notification needs to be sent for each establishment intending to layoff workers. And the LIFO principle applies at the level of the establishment.

2.2 Data

We combine data from various administrative registers in Sweden. First, we use UI fund membership information for the universe of workers in Sweden aged 18 and above, from 2002 to 2009, and coming from two distinct sources. The first source is tax data for the period 2002 to 2006, during which workers paying UI premia received a 40% tax credit. This source records the total amount of UI premia paid for each year. From this source, we define a dummy variable V for buying the comprehensive coverage in year t as reporting any positive amount of premia paid in year t . We use this source of information for the positive correlation tests of Section 3, as well as the risk variation analysis in Section 4.2. For the analysis using the price variation of the 2007 reform in Section 5.2, we combine this data with a second source of information, coming from UI fund data that *Kassa*’s sent to the IAF. This data contains a dummy variable indicating whether an individual aged 18 and above in Sweden is contributing premia for the comprehensive coverage as of December of each

⁶The 10% rebate on UI premia for union members is a remnant of the “Ghent system”, but a large ($\approx 20\%$) and growing share of workers are members of an unemployment fund without being members of a union, and a growing share of union members ($\approx 10\%$) do not buy unemployment insurance.

⁷In our data, the collective bargaining agreement that individuals are in is not directly observed. We use detailed occupation codes instead, which are regarded as a good proxy.

year from 2005 until 2009.

We add data on unemployment outcomes coming from the Swedish Public Employment Service, with records for the universe of unemployment spells from 1990 to 2015, and we merge it with the UI benefit registers from the IAF which provides information on all UI benefit payments (for both the basic and comprehensive coverage), information on daily wage for benefit computation, and Kassa membership information for all unemployed individuals. Based on this data, we define unemployment as a spell of non-employment, following an involuntary job loss, and during which an individual has zero earnings, receives unemployment benefits and reports searching for a full-time job. To define the start date of an unemployment spell, we use the registration date at the PES. The end of a spell is defined as finding any employment (part-time or full-time employment, entering a PES program with subsidized work or training, etc.) or leaving the PES (labor force exit, exit to another social insurance program such as disability insurance, etc.).⁸ We define displacement as an involuntary job loss, due to a layoff or a quit following a ‘valid reason’.⁹ In the rest of the paper, we use the terms displacement and layoff as synonyms.

We complement this data with information on earnings, income, taxes and transfers and demographics from the LISA register, and with information on wealth from the wealth tax registers.

Finally, we use two labor market registers. The matched employer-employee register (RAMS), from 1985 to 2015, reports monthly earnings for the universe of individuals employed in establishments of firms operating in Sweden. We use this register to compute tenure and tenure ranking for each employee. We also use the layoff-notification register (VARSEL) which records, for years 2002 to 2012, all layoff notifications emitted by firms.

In Table 1, we provide summary statistics for our main sample of interest over the period 2002 to 2006. To mitigate concerns about younger individuals switching in and out of education, or older individuals close to retirement, we restrict our attention to individuals aged between 25 and 55. The average probability to be displaced in year $t + 1$ conditional on working in year t is 3.35%, (3.56% when including quits) over the period 2002 to 2006. The average probability to be unemployed in year $t + 1$ (unconditional on employment status in year t) is higher, at 4.71%. Note also that the fraction of individuals who are members of a UI fund (i.e., buying the comprehensive UI coverage) is large during the 2002-2006 period, at 86%.

3 Positive Correlation Tests

We first show the presence of a strong positive correlation between an individual’s choice of UI coverage and his or her unemployment risk. This strong correlation is robust to different measures

⁸Note that UI benefits can be received *forever* in Sweden during the period 2002-2006 so the duration spent unemployed is identical to the duration spent receiving unemployment benefits.

⁹Valid reasons for quitting a job are defined as being sick or injured from working, being bullied at work, or not being paid out one’s wage by one’s employer. Quits are reviewed by the Public Employment Service at the moment an individual registers a new spell and if the quit is made because of a valid reason, the individual is eligible for UI and a notification is made in the PES data, allowing us to observe such quits under valid reasons. Involuntary quits are a small fraction of unemployment spells in our sample: 95.0% of unemployment spells observed in our data are due to layoffs. We exclude voluntary quits from our measure of unemployment and displacement.

of the unemployment risk, the addition of a rich set of controls and non-parametric implementations. Correlation tests are a natural first step to investigate adverse selection and common in the insurance literature, but may be confounded by the presence of moral hazard.

3.1 Framework

We start by presenting the conceptual framework for insurance choices that underpins our empirical and theoretical analysis. A Swedish worker faces the choice between two plans: a basic plan (b_0, p_0) and a comprehensive plan (b_1, p_1) with unemployment benefit levels $b_1 > b_0$ and premia $p_1 > p_0$. A worker chooses the plan providing the highest (indirect) expected utility. That is, a worker buys the comprehensive plan ($V = 1$) when her expected utility in the comprehensive plan exceeds her expected utility in the basic plan,

$$\begin{aligned} V &= 1 && \text{if } \mathbf{v} - \mathbf{p} \geq 0, \\ V &= 0 && \text{otherwise,} \end{aligned} \tag{1}$$

where $\mathbf{v} = v_1 - v_0$ and $\mathbf{p} = p_1 - p_0$.¹⁰ In a stylized binary risk setting, the net-value of a plan equals

$$v_k - p_k \equiv \max_{a'} \pi(\theta, a') u(b_k - p_k | \mu, a') + (1 - \pi(\theta, a')) u(w - p_k | \mu, a') \tag{2}$$

where π denotes the probability of unemployment, a denotes effort, and θ and μ are risk and preference parameters. Importantly, not only the value but also the cost of providing the coverage depends on the agent's type and her effort. In the binary-risk setting, the cost of providing plan k equals $c_k = \pi(\theta, a_k) b_k$, depending on the agent's risk type θ and the effort level a_k that she exerts under contract k .

We refer to the group of individuals buying the comprehensive plan by I and those buying the basic plan by U . Throughout the rest of the paper, we will use the notation $E_I(\cdot) = E[\cdot | \mathbf{v} - \mathbf{p} \geq 0]$ and $E_U(\cdot) = E[\cdot | \mathbf{v} - \mathbf{p} < 0]$ for the respective conditional expectations. The individuals at the margin between the two plans are referred to by M .¹¹ Regarding the timing of the model, we stick closely to the structure of the Swedish UI system where individuals become eligible to receive the supplemental benefits when they have been contributing for one year to the comprehensive coverage, and can opt in and out of the comprehensive plan at any time. As a consequence, the valuation \mathbf{v}_t of the comprehensive UI policy in year t depends on unemployment risk π_{t+1} in year $t + 1$. With this in mind, we drop from now on the time subscripts with \mathbf{v} always referring to \mathbf{v}_t and π to π_{t+1} , unless otherwise specified.

¹⁰We focus on valuations that are quasi-linear in the premium p_k as it leads conveniently to a welfare analysis in terms of total surplus. This formulation, although it leaves out income effects, can still easily incorporate distributional concerns through the social welfare function, as we do in Section 6.

¹¹Marginal individuals are defined by the condition $\mathbf{v} = \mathbf{p}$. We further clarify this definition in the context of our price variation experiment in Section 5.2.

3.2 Correlation Tests

The insurance choice model of equation (2) suggests that \mathbf{v} is an increasing function of π and that unpriced heterogeneity in θ leads to risk-based selection into UI. Unless preference heterogeneity undoes this risk-based selection, adverse selection will arise whereby riskier individuals are more likely to buy the comprehensive plan, creating a positive correlation between insurance choice and observed risk. The correlation test consists in comparing the expected risk of individuals conditional on their insurance coverage choice and testing for $E_I(\pi) > E_U(\pi)$.

Linear Probability Model The simplest way to test for $E_I(\pi) > E_U(\pi)$ in practice, is to estimate a simple linear model for various measures Y of realized risk in year $t + 1$, which proxy for π :

$$Y = \gamma \cdot V + \mathbf{X}'\alpha + \epsilon, \quad (3)$$

where V is an indicator for buying the supplemental coverage in year t . The vector \mathbf{X} controls for individual characteristics that affect the unemployment insurance contracts available to each individual. Controlling for these characteristics guarantees that we compare individuals who are facing the same options so that the correlation is driven by demand rather than by supply (different individuals being offered different contracts by the *Kassa*). As explained in Section 2 above, this is strictly regulated by the government. We estimate model (3) over the period 2002-2006, during which UI contracts only differ according to three dimensions: employment history, earnings and union membership.

The first dimension is whether individuals meet the work eligibility requirement or not, for which they need to have worked for at least 6 calendar months within the past 12 months prior to displacement. We therefore include an indicator for having worked at least 6 months in year t in \mathbf{X} .¹² The second dimension of contract differentiation is earnings: the additional daily benefits \mathbf{b} that individuals get when buying the supplemental coverage is a kinked function of daily earnings w . Formally, $\mathbf{b} = b_1 - b_0 = F(w) = (.8 * w - 380) \cdot \mathbf{1}[400 \leq w < 725] + 200 \cdot \mathbf{1}[725 \leq w]$. We therefore include the supplemental benefit function $F(w)$ as a control function in \mathbf{X} to make sure that we compare individuals facing the same benefit level per unit of premium paid. The last dimension of contract differentiation is that union members pay a slightly lower premium than non-union members for the supplemental coverage. We therefore include in \mathbf{X} an indicator variable for union membership. We also include year fixed effects in \mathbf{X} to account for small adjustments to the premium in January every year over the period 2002-2006.

Figure 1 reports the results of specification (3) for four different realized risk outcomes: total UI claims under comprehensive coverage in $t + 1$, total duration spent unemployed in $t + 1$, the

¹²Note that eligibility requires individuals to have worked at least 80 hours per month for 6 calendar months within the past 12 months. While we do not have precise data on monthly hours, to be conservative, we also include a dummy for having earnings above 80 hours \times 6 months \times the negotiated janitor wage. In the absence of an official, legally binding minimum wage in Sweden, the janitor wage is often considered the effective minimum wage in the labor market.

probability of displacement in $t + 1$, and the probability of displacement in $t + 1$ but excluding involuntary quits. The total UI claims are defined as the total amount of UI benefits that individuals would be collecting in $t + 1$ were they to buy the comprehensive coverage.¹³ For each outcome, Figure 1 displays $\hat{\gamma}/\bar{Y}$, that is the semi-elasticity of the realized risk outcome in $t + 1$ with respect to the insurance choice in t .

For all realized risk outcomes, we find a strong and significant positive correlation with UI coverage choice. Individuals who buy the comprehensive coverage in t make UI claims in $t + 1$ that are 161.6% larger than the hypothetical claims under comprehensive coverage by individuals who stick to the basic coverage in t . Their unemployment duration in $t + 1$ is 140.8% longer and they are 131.7% more likely to be displaced in $t + 1$ than individuals who do not buy it.

Alternative risk outcomes All risk outcomes capture ex-post risk realizations rather than ex-ante risks. These realizations reflect in part actions taken by individuals because of their insurance choices. The correlation test amounts to comparing $E_I(\pi(a_1))$ to $E_U(\pi(a_0))$ and estimated correlations could therefore be driven by various sources of moral hazard. Separating risk-based selection from moral hazard is exactly the topic of Sections 4 and 5. Still, simply comparing the magnitude of the correlations across the different realized risk outcomes already sheds light on some margins of moral hazard. A large body of literature has for instance documented that higher unemployment benefits increase the duration of unemployment spells conditional on becoming unemployed (see Schmieder and Von Wachter [2016] for a recent review). Such moral hazard conditional on displacement will increase the correlation between unemployment duration in $t + 1$ and insurance coverage in t . The correlation between displacement probability in $t + 1$ and insurance coverage in t is immune to this particular source of moral hazard. The difference between the two estimates (second vs third bar in Figure 1) captures the presence of moral hazard conditional on displacement, although part of it might also be driven by selection on expected unemployment duration conditional on displacement. The probability of displacement, while immune to moral hazard once displaced, is potentially affected by moral hazard “on the job”. An example of this would be collusion between employers and employees to qualify actual voluntary quits as “quits following a valid reason”, which are eligible for unemployment benefits. The correlation of insurance coverage with displacement probability and with displacement probability excluding “quits following a valid reason” (third and fourth bar in Figure 1) is identical, suggesting that this collusion margin has actually limited impact on the results of the positive correlation test.

Our correlation tests use the risk outcomes in $t + 1$, reflecting the idea that workers need to contribute for a year to be able to get the comprehensive coverage. However, the risk realization in $t + 1$ may fail to fully capture the unemployment risk faced by an individual as she is making her coverage choice at time t , which justifies using risk realizations further into the future. In Figure 2 we report the correlation of the insurance choice in t with displacement outcomes in $t + 1, t + 2, \dots$ up to $t + 8$. For each displacement outcome, the chart displays $\hat{\gamma}_k/\bar{Y}$, that is the semi-elasticity of

¹³For individuals who do not buy the comprehensive coverage in t , we therefore computed the counterfactual benefit claims they would have if they were to receive supplemental benefits.

the realized risk outcomes in $t+k$ with respect to insurance choices in t , from a specification similar to (3) where we also control for all displacement outcomes in previous years ($t+k-1$, $t+k-2$, etc.). The Figure reveals an interesting dynamic pattern. The correlation decreases rapidly as we consider later years, but remains statistically significant up to six years. This pattern could indicate that workers' insurance choices incorporate private information about unemployment risk further into the future (albeit to a decreasing extent), but it may also be affected by moral hazard responses.

Bivariate Probit & Non-parametric Tests While the linear probability model in (3) provides a simple test for the presence of positive correlation between risk and insurance choices, and a straightforward interpretation of its magnitude, it relies on a very limiting functional form and our OLS estimates do not provide correct inference. We now relax these functional form restriction and provide proper inference for the correlation tests. First, we provide results of bivariate probit tests, popularized by Chiappori and Salanié [2000]. We specify both the choice of insurance coverage and the realization of our binary measure of unemployment risk (i.e., the probability of displacement) as probit models:

$$\begin{aligned} V &= \mathbb{1}[\mathbf{X}'\alpha_1 + \epsilon > 0] \\ Y &= \mathbb{1}[\mathbf{X}'\alpha_2 + \eta > 0] \end{aligned} \tag{4}$$

allowing for correlation ρ between the two error terms ϵ and η . The vector of controls \mathbf{X} contains the same variables as in specification (3). We provide in Table 2 estimates of ρ and formal tests of the null that $\rho = 0$. Results confirm the presence of a strong and significant correlation between insurance choices and realized unemployment risk. The functional forms involved in the bivariate probit tests are still relatively restrictive since the latent models are linear and the errors are normal, excluding cross-effects or more complicated non-linear functions of the variables in \mathbf{X} . We therefore also produce results from non-parametric tests as in Chiappori and Salanié [2000]. The procedure of the test consists in partitioning the data into cells where all observations in a given cell have the same value for the variables in \mathbf{X} . The procedure then computes within each cell a Pearson's χ^2 test statistic for independence between V and Y . This test statistic is asymptotically distributed as a $\chi^2(1)$ under the null hypothesis that V and Y are statistically independent (within the cell). We report in the first column of Table 3 results from this non-parametric procedure when cells are defined using the same controls \mathbf{X} as in specification (3) and where our risk measure Y is the probability of displacement. Results again strongly confirm the presence of a positive correlation between insurance choices and unemployment probability.¹⁴

¹⁴In Appendix Figure A.1, we display the empirical distribution of the Pearson's χ^2 test statistics computed from all the cells allows for comparison with a theoretical $\chi^2(1)$ distribution. Taking the largest absolute difference between the theoretical and the empirical distribution gives the Kolmogorov-Smirnov test statistic reported in Table 3.

3.3 The Role of Unpriced Observables

As explained in Section 2, during our period of analysis, the Swedish unemployment system did not allow for price discrimination across individuals based on differences in risks. Many observable characteristics that are known to usually correlate with unemployment risk (age, industry, occupation, gender, etc.) cannot be priced by insurance funds. We briefly explore to what extent the positive correlation between insurance and unemployment risk documented above is directly driven by selection on such unpriced observables and whether the correlation would survive if such characteristics were to be priced. To do so, we start with the baseline positive correlation test from specification (3) where Y is the probability of displacement in $t + 1$, and show how the semi-elasticity $\hat{\gamma}/\bar{Y}$ evolves as we add more characteristics to the vector of controls \mathbf{X} . Results are displayed in Figure 3 panel A. We start by adding (sequentially) demographic controls: age, then gender, and marital status. Interestingly, the estimated correlation increases when adding these covariates, which suggests that these characteristics actually drive the selection to be “advantageous”: they correlate positively with risk but negatively with insurance coverage. Age in particular leads to meaningful advantageous selection as young individuals are more likely to be unemployed, but significantly less likely to buy UI. The correlation does not seem to be affected much by the inclusion of controls for skills and other labor market characteristics. Adding rich sets of controls for education (four categories), industry (1-digit code), occupation (1-digit code) and wealth level (quartiles) decreases the estimated correlation only slightly. But adding controls for past unemployment history (dummies for having been unemployed in $t - 1$, $t - 2$ and up to $t - 8$) has a significant negative effect on the estimate. Past unemployment history is a strong predictor of future unemployment risk and correlates strongly with current insurance choices. Yet, even when controlling very flexibly for past unemployment history and all the other controls, results from Figure 3 show that a large positive correlation remains between insurance choices and probability of displacement.

The results above could be affected by the potentially constraining linear functional form of our specification (3), but are robust to more flexible functional forms and appropriate inference. First, we show in Figure 3 panel B how the correlation from the bivariate probit specification (4) evolves when adding sequentially to the vector \mathbf{X} the same set of characteristics as in panel A. Second, in Table 3, columns (2) to (4), we reproduce the non-parametric Kolmogorov-Smirnov test adding sequentially these same characteristics when partitioning the data into cells. Results confirm that demographics, and age in particular, offer advantageous selection, that past unemployment history creates significant adverse selection, and that a significant positive correlation between insurance and probability of displacement remains even after controlling for all these rich observables.

4 Beyond Correlation Tests: Variation in Risks

The positive correlation tests are a useful starting point, but cannot separate risk-based selection from moral hazard responses. In other words, it is well understood that correlation tests are a joint test of selection and/or moral hazard. Identifying the respective role of selection and moral hazard

is both useful from a descriptive perspective and necessary from a welfare perspective. This section provides further evidence of substantial risk-based selection based on exploiting variation in risks.

4.1 Using Variation in Risk

Moral hazard creates reverse causality from insurance to risk, as individuals choose different levels of optimal actions a_k under each plan $k = 0, 1$. Such reverse causality could in principle fully drive the correlation between observed realized risk Y and insurance. The most direct way to control for moral hazard is to shift individuals' risk, independently of individuals' actions. To do this, one needs to find variables Z that operate as unpriced shifters of individuals' risks, akin to the unused observables test in Finkelstein and Poterba [2014].

We briefly explain how and under what assumptions using such variation in risk can identify the presence of risk-based selection. We start from the framework of equation (2) above, which can be rewritten as:

$$\mathbf{v} - \mathbf{p} \equiv \pi(\theta, a_1)\Delta u_1(\mu, a_1) - \pi(\theta, a_0)\Delta u_0(\mu, a_0) - \Delta u_w(\mu, a_1, a_0) \quad (5)$$

where a_k denotes an individual's optimal action under each plan $k = 0, 1$ and $\Delta u_k(\mu, a_k) = u(b_k - p_k|\mu, a_k) - u(w - p_k|\mu, a_k)$ denotes the utility loss due to unemployment when covered by plan k . We also use the notation $\Delta u_w(\mu, a_1, a_0) = u(w - p_1|\mu, a_1) - u(w - p_0|\mu, a_0)$.

Assume now that we can find risk shifters Z that have the following two properties: $\frac{\partial \pi(\theta, Z, a_k)}{\partial Z} \neq 0$ and $Z \perp \mu$. The first property is equivalent to a first-stage property and guarantees that Z shocks individuals' risk conditional on their actions. The second property can be thought of as an exclusion restriction, and guarantees that Z is uncorrelated with the preference type μ , which would affect willingness-to-pay independent of the change in risk. Under these two assumptions about Z , testing for $\mathbf{v} - \mathbf{p} \perp Z$ is a test for the null of no risk-based selection. In other words, rejecting the null is equivalent to rejecting the "pure moral hazard" model, in which the correlation between realized risk and insurance choice is entirely driven by moral hazard. The "pure moral hazard" model is characterized by the fact that $E(\pi(\theta, a_k)|\mathbf{v}) = \gamma_k$ for both plans k and thus that the average unemployment risk is constant in \mathbf{v} .¹⁵ The positive correlation test will pick up a positive correlation in this model as long as $\gamma_1 \neq \gamma_0$. Note, however, that there will be no correlation between willingness-to-pay and any risk shifter Z satisfying the two properties above in the "pure moral hazard" model. The reason is that if Z affects individuals' risk, but is uncorrelated with individuals' preference type μ , it can only affect the insurance choice through its impact on risk. Hence, a correlation between Z and insurance choices means that unemployment risk can no longer be uncorrelated to the willingness-to-pay \mathbf{v} as is the case in the "pure moral hazard" model.

If the exclusion restriction $Z \perp \mu$ were not to hold, the correlation between Z and preference type μ could in principle exactly offset the direct impact of the risk shifter on the willingness-to-

¹⁵In the "pure moral hazard model", there might still be heterogeneity in individuals' risk types θ or in individuals' actions, but this heterogeneity must be offset by the preference heterogeneity such that the resulting risks are uncorrelated with willingness-to-pay.

pay and leave the resulting risk uncorrelated to the willingness-to-pay. In general, however, any correlation between Z and willingness-to-pay would identify the presence of risk-based selection, either stemming from the direct effect of Z on π , which one could call “direct risk selection”, or from selection on risk-related heterogeneity (including selection on moral hazard).¹⁶ In the rest of this section, we exploit the presence of several unpriced risk shifters Z in the Swedish context, and discuss for each of them whether $\text{Cov}(Z, \mu) = 0$ is a credible assumption.

In practice, we test for correlation between risk shifters Z and willingness to pay by running specifications like:

$$V = \mathbb{1}[\sigma \cdot Z + \mathbf{X}'\alpha_1 + \epsilon > 0] \quad (6)$$

and testing for $\sigma = 0$. For useful comparison with the positive correlation test estimates in the linear case, we also report comparison of the PCT model

$$V = \beta_{OLS} \cdot Y + \mathbf{X}'\alpha + \epsilon \quad (7)$$

with estimates of the two-stage least square model

$$\begin{aligned} V &= \beta_{2SLS} \cdot Y + \mathbf{X}'\alpha_1 + \epsilon \\ Y &= \zeta \cdot Z + \mathbf{X}'\alpha_2 + \eta \end{aligned} \quad (8)$$

The 2SLS model will yield $\hat{\beta}_{2SLS} = 0$ if the OLS PCT estimate $\hat{\beta}_{OLS}$ is fully driven by the pure moral hazard model.¹⁷

4.2 Implementation

The implementation of the “risk-variation approach” relies on finding risk shifters that affect individuals’ risk probability conditional on their own actions, and that are credibly exogenous to preference heterogeneity (μ) governing individuals’ willingness-to-pay for insurance. Our risk shifters exploit two fundamental sources of risk variation, arguably beyond the control of individuals.

The first source is firm level risk, which can vary cross-sectionally, due to permanent differences in turnover across firms, or over time, due to firms experiencing temporary shocks. In Figure 4 panel A, we provide evidence of the role of firm layoff risk as a shifter of individuals’ own displacement probability. For each individual i working in firm j , we define average firm displacement risk $\pi_{-i,j}$ as the average probability of displacement of all other workers within the firm excluding individual i over all years where the firm is observed active in our sample years. We then plot the average firm displacement risk in 20 bins of equal population size, against the individual probability of displacement in $t + 1$. The Figure shows that there is significant heterogeneity in firms’ separation

¹⁶Selection on moral hazard is a form of risk-based selection, but where the choice to buy insurance is related to the difference between e_1 and e_0 . For example, an individual buys more coverage anticipating that he or she will reduce her effort a lot under the extra coverage. This again again creates a correlation between willingness-to-pay and cost of providing the different plans k .

¹⁷While risk-based selection is needed for $\hat{\beta}_{2SLS} \neq 0$, the presence of moral hazard still affects this estimate when individuals exert less effort in response to the extra coverage they buy when their risk is shifted.

rates, and that individuals' unemployment risk is very strongly correlated with firm level risk.

The second source of exogenous risk variation is at the individual level and stems from the strict enforcement of the Last-In-First-Out (LIFO) principle. As explained in Section 2, when a firm wants to downsize, the legal system prescribes that displacement occurs by descending order of tenure within each establishment times occupation group. The tenure ranking of an individual within her establishment and occupation group directly determines her probability to be separated. Figure 4 panel B plots the probability of being displaced in $t + 1$ among individuals working in firms that emit a layoff notification in $t + 1$, as a function of relative tenure ranking within establishment and occupation. The Figure provides clear evidence of a strong negative correlation between relative tenure ranking and individuals' displacement probability. Individuals within the lowest 10 percent of tenure rankings have a probability of being displaced in $t + 1$ larger than .1; this probability declines steadily as tenure ranking increases, and then stays below .02 for individuals in the highest 50 percent tenure rankings.

We combine the sources of variations brought about by these underlying risk shifters (firm level risk and LIFO) into three different identification strategies.

Firm Layoff Risk The first strategy consists in simply using the cross-sectional variation in displacement risk across firms as a risk shifter. In Figure 5 panel A, we group individuals in 50 equal size bins of firm layoff risk, and plot their average firm layoff risk against their average probability of buying supplemental coverage, residualized on the same vector \mathbf{X} of baseline controls affecting UI contracts used in the positive correlation test of Section 3.2. The graph displays a strong positive correlation between firm layoff risk and individuals' probability to buy the comprehensive UI coverage, indicating that there is a clear correlation between Z and willingness to pay ($\sigma \neq 0$). We also report on the graph the coefficient β_{OLS} from an OLS regression of specification (7) and then the estimated coefficient β_{2SLS} from our two-stage least square model (8) where we use $Z = \pi_{-i,j}$ as a risk shifter. In panel B of Figure 5, we replicate the same procedure, but now add to the regression the same rich set of additional controls used in Section 3.3, and find a similar strong positive correlation between insurance choices and firm layoff risk.

The positive and significant coefficient $\beta_{2SLS} = .50$ (.01) rejects that the results of the positive correlation tests of Section 3.2 are solely driven by moral hazard. The relative magnitude of β_{2SLS} and β_{OLS} is also informative. While we anticipate that by controlling for moral hazard β_{2SLS} decreases relative to β_{OLS} , two effects can play in the opposite direction. First, the two-stage least square procedure removes the potential attenuation bias from measurement error in β_{OLS} . Second, risk shifters also introduces some selection, through $\text{Cov}(Z, \mu)$, which has an a priori ambiguous impact on the estimate. $\text{Cov}(Z, \mu)$ will depend on the self-selection of workers into riskier firms: if workers who select to work in riskier firms are more likely to buy UI, selection will be positive. $\text{Cov}(Z, \mu)$ will also depend on the unobserved effect of riskier firm environments on insurance choices: firms with high turnover may have different prevalence of collective bargaining, different firm cultures that can affect individuals' UI choices. Decomposing $\mu = \kappa_i + \rho_j$ into an individual

specific component κ_i and a firm specific component ρ_j , we can think of the the selection introduced by a risk shifter Z as the combination of individual fixed effects and firm fixed effects. We now move to two additional strategies that allow us to control more directly for these two sources of selection introduced by our risk shifters.

Firm Switchers Design In this second strategy, we use the panel dimension of the data to control for the selection introduced by individual specific heterogeneity κ_i . We exploit within individual variation in risk, stemming from both risk shifters: firm risk and tenure ranking. To this end, we focus on individuals who change firms (“switchers”). The employer-employee matched data (*RAMS*) registers all existing labor contracts on a monthly basis. We define a switch as moving from having a labor contract with firm j (the origin firm) to having a contract with firm k (the destination firm), without any recorded non-employment spell between these two contracts. We focus on individuals with more than 1 year of tenure in the origin firm. Switchers experience a change in their layoff risk coming from underlying variation in both risk shifters: their tenure ranking changes, and so does their underlying firm layoff risk.

First, switchers experience a reduction in their relative tenure ranking, as they become the “last-in” when they move to the destination firm. To document the magnitude of the variation in relative tenure ranking and corresponding layoff risk, following a firm switch, we define year $n = 0$ as the year of a firm switch, and run, on the sample of firm switchers, event studies of the form:

$$Y_n = \sum_k \delta_k \cdot \mathbb{1}[n = k] + \mathbf{X}'\alpha + \epsilon \quad (9)$$

where Y_n denotes the risk outcome of interest in event year n , $\mathbb{1}[n = k]$ are a set of event time dummies, and \mathbf{X} is the vector of baseline controls affecting UI contracts. In Appendix Figure A.2, we display the evolution of relative tenure ranking of switchers as a function of event time by plotting the coefficient δ_k , taking event time $n = -1$ as the omitted category. The graph confirms that relative tenure ranking decreases sharply at the moment of the firm switch. Figure 6 panel A explores how this variation in relative tenure ranking affects the probability of displacement over event time n . The graph shows that the displacement rate one year ahead (in year $t + 1$) increases sharply and significantly at the time of the firm switch.

In Figure 7 panel A, we run a similar event study specification with UI coverage as an outcome. The figure shows that the probability of buying the comprehensive coverage increases sharply by about 2.2 percentage points at the time of the firm switch. On the graph, we also display the coefficient from a two-stage least square specification similar to (8) where we use firm switch as risk shifter Z for individual displacement probability and control for individual fixed-effects. β_{2SLS} is positive and strongly significant, which again indicates that the positive correlation tests are not simply picking up moral hazard responses to insurance coverage.

While these event studies control for fixed underlying heterogeneity across individuals that may affect their UI choices (κ_i), one concern with this original implementation of the firm switchers

design is that individuals are somewhat inert, and decide to reoptimize their UI choices only at specific times, like, for instance, when they switch firm.

To mitigate the concern that the surge in UI coverage at the time of the switch is the result of the specific timing of UI choices and not a response to the change in underlying risk, we exploit additional variation in risk in the switchers design coming from changes in underlying firm layoff risk. While all switchers experience an increase in their displacement probability due to the decline in their tenure ranking, the effect of a switch on individual displacement probability exhibits large differences according to whether their destination firm is much riskier (“positive shock”) or a lot less risky (“negative shock”) than their origin firm. We therefore split the population of switchers according to their rank in the distribution of $\Delta_{j,j'}\pi_{-i} = \pi_{-i,j'} - \pi_{-i,j}$, the change in their underlying firm risk when moving from firm j to firm j' . In Figure 6 panel B, we contrast individuals in the bottom decile of $\Delta_{j,j'}\pi_{-i}$ (large negative shock, i.e., individuals who experience a large negative decline in their firm layoff risk, going from a high risk to a low risk firm), and individuals in the top decile of $\Delta_{j,j'}\pi_{-i}$ (large positive shock, i.e., individuals who experience a large increase in their firm layoff risk going from a low risk to a high risk firm). The Figure confirms that individuals experiencing a large positive shock in their firm layoff risk exhibit a significantly larger increase, of about 2 percentage point, in their displacement probability at the time of the switch, relative to individuals experiencing a large negative shock. In panel B of Figure 7, we now compare the evolution of insurance choices around firm switch for individuals experiencing large positive vs large negative shocks. The graph indicates that the increase in the probability to buy UI around firm switch is significantly larger (by about 1.5 percentage point) among individuals moving to significantly more risky firms relative to those moving to less risky firms. We also report on the graph the estimated coefficient $\beta_{2SLS} = .57 (.08)$ of the two-stage model of specification (8) where we now use firm switch interacted with $\Delta\pi_{-i,j}$ as risk shifter Z for individual displacement probability.

Layoff Notification and LIFO The previous strategy may still pick up selection on firm level heterogeneity ρ_j if firm heterogeneity is correlated with $\Delta\pi_{-i,j}$. The final strategy controls jointly for firm level heterogeneity ρ_j and individual level heterogeneity κ_i by exploiting variation in layoff risk both within firm and across individuals over time. To this effect we leverage the fact that under Sweden’s employment-protection law, firms subject to a shock and intending to displace 5 or more workers simultaneously must notify the Public Employment Service in advance. In Figure 8 we report the evolution of the displacement probability of workers around the first layoff notification event in the history of the firm. We define event year $n = 0$ as the year in which a firm emits its first layoff notification, and focus on the panel of workers who are employed in the firm at the date this layoff notification is emitted to the PES. The graph shows that a layoff notification is indeed associated with a sudden and large increase in the displacement probability of workers. Immediately following the layoff notification, the displacement rate of workers jumps by 6 percentage points compared to pre-notification levels, and remains high for about two years,

before decreasing and converging back to pre-notification levels.

Because the panel of workers is selected based on being employed in the firm in year $n = 0$, one may worry that this surge in displacement rates is mechanical, as displacement can only increase after year 0 conditional on all workers being employed in year 0. To mitigate this concern, we follow a matching strategy and create a control panel of workers selected along the same procedure as the original panel. We use nearest-neighbor matching to select a set of firms that are similar, along a set of observable characteristics, to the firms emitting a layoff notification, but never emit a layoff notification.¹⁸ We allocate to the matched firm in the control group a placebo event date equal to the layoff notification date of her nearest-neighbor in the treated group of firms. We then select workers that are in the control firm at the time of the placebo event date to create our matched control panel. Results in Figure 8 show that, pre-event, the displacement risk is very similar in control and treated firms, and that it evolves smoothly in the control firms around the event.

Once a notification is emitted, employers need to come up with the list and dates for the intended layoffs. These layoffs may happen up to 2 years after the original notification has been sent. The list needs to follow the last-in-first-out principle, so that relative tenure ranking within each establishment and occupation group directly determines one’s displacement probability once a notification has been sent. As a consequence, the effect of a layoff notification on displacement probability is strongly heterogenous depending on the relative tenure ranking of workers, as shown in Figure 4 panel B. Workers with relative tenure ranking below .5 have a much higher probability of being laid-off following a layoff notification than workers with relative tenure ranking above .5.

We exploit this additional layer of variation in displacement risk coming from the interaction between a notification event and relative tenure ranking. In Figure 9, we show that insurance choices significantly respond to this shift in displacement risk. Panel A of Figure 9 reports the evolution of UI coverage around the time of the first layoff notification for the panel of workers in the treated group and for workers in our placebo (matching) group, restricting the sample to workers with relative tenure ranking below 50% in year $n = 0$. The Figure shows that UI coverage increases significantly among the treated group, starting one year before the layoff notification is sent, which suggests the existence of some degree of private information among workers regarding the timing of the layoff notification. The Figure displays the estimated coefficient $\beta_{2SLS} = .84 (.21)$ of our two-stage least square model using the layoff event interacted with tenure and a dummy for being in the treatment group as a risk shifter Z . In panel B, we report similar estimates but for the sample of workers with relative tenure ranking above 50% in year $n = 0$. The graph displays no sign of variation in individuals’ insurance coverage around the event.

Taken together, this evidence, which jointly controls for underlying selection on heterogeneity μ coming from firm fixed effects and individual fixed effects, strongly suggests that UI choices do significantly respond to variations in individuals’ underlying unemployment risk. Under the assumption that there is no other heterogeneity varying jointly with the timing of the notification

¹⁸The covariates used for matching are the number of employees, the 4 digit sector codes of the firm, the average earnings and average years of education of workers in the firm.

and individuals' tenure ranking at the time of the notification, $\text{Cov}(Z, \mu) = 0$ in this strategy, and $\beta_{2SLS} \neq 0$ provides evidence for direct risk-based selection.

Predicted Risk Model Previous evidence shows that individuals' displacement probability varies with many different risk shifters that are beyond the control of individuals' own actions. These risk shifters are observable to the econometrician, and therefore, possibly to the insurer as well. We have so far focused on the effect of risk shifters separately, but now combine them in a rich and flexible prediction model of individuals' unemployment risk.

Designing this model serves two purposes. First, by combining flexibly all observable sources of risk shifting together, the model allows for a richer variation in risk than the one obtained from using the risk shifters separately. One can then directly test whether insurance choices react to this richer variation in observable risk. Second, the model is useful from an insurer's perspective, as it allows to test how much positive correlation between insurance and risk would be left when conditioning prices flexibly on the richest set of observable characteristics predicting risk.

The risk outcome that we focus on in our prediction model is the probability of displacement in $t + 1$. The risk shifters included in the model are the average firm layoff risk, the full history of the firm layoff notifications, and the relative tenure ranking of the individual. We also add the full past unemployment history of the individual. Model selection (e.g., deciding on the best set of interactions between risk shifters) follows a cross-validation approach where we minimize misclassification errors in the test sample.

Results are reported in Figure 10. Panel A correlates predicted layoff risk with actual displacement rates in $t + 1$ and provides evidence that the model performs well in predicting individuals' risk. Panel B correlates predicted layoff risk with the probability to buy UI coverage, and suggests again a strong positive correlation between individuals' risk and their probability to buy the supplemental UI coverage. Interestingly, the graph also suggests that the strong positive correlation between risk and insurance coverage is mostly driven by what happens at the bottom of the predicted risk distribution. Between individuals predicted to have a zero risk of displacement and a 2 percent probability of displacement, the fraction buying coverage increases from 60 to 90%. But the fraction buying UI coverage remains flat, and even decreases slightly as the predicted layoff risk increases further.

Based on this prediction model, we now investigate how much positive correlation between insurance and risk there would be if the UI policy used the full extent of this observable information to condition UI coverage prices on individuals' predicted risk. To this effect, we use a non-parametric "Kolmogorov-Smirnov" test similar to the ones used in Section 3.2 above. We partition the data into 50 cells of predicted layoff risk, and within each cell compute a Pearson's χ^2 test statistic for independence between insurance coverage in t and displacement in $t + 1$. Results from this non-parametric procedure are reported in column (5) of Table 3, and indicate that a significant level of correlation between insurance and displacement risk remains once we condition on risk levels predicted from a rich set of observable characteristics.

5 Beyond Correlation Tests: Variation in Prices

Risk shifters that are exogenous to individuals' actions are a direct way to test for the presence of risk-based selection versus a model of "pure moral hazard". Yet, any selection on unpriced risks - whether it is direct or through risk-related preference heterogeneity - is relevant from the insurer's perspective, as selection along these dimensions will affect the cost of providing the insurance coverage. An important drawback of using risk shifters is that they may get rid of the indirect selection that is welfare relevant or introduce selection that is not directly welfare relevant. To overcome this limitation, we follow a second strategy which relies on exploiting exogenous variation in prices \mathbf{p} , following Einav et al. [2010b].

5.1 Using Variation in Prices

With panel data, exogenous price variation allows us to identify marginals who switch coverage in response to the price change and thus to rank individuals in three groups, the insured I , the marginals M and the uninsured U , ordered in terms of their willingness to pay \mathbf{v} : $E_I[\mathbf{v}] > E_M[\mathbf{v}] > E_U[\mathbf{v}]$. This enables the implementation of a direct non-parametric test for selection by correlating willingness to pay with risk of individuals observed *under the same insurance coverage*. Depending on the available variation in prices, \mathbf{p} , one can test for $E_M[Y_0] > E_U[Y_0]$ or $E_I[Y_1] > E_M[Y_1]$, or both, where Y_k denotes individuals' risk observed under contract k .¹⁹ Because individuals are compared under the same coverage, this test is immune to moral hazard and directly identifies selection. This estimate does not separate what directly comes from risk-based selection and what comes from selection along unpriced dimensions of heterogeneity that are correlated with risk, such as for instance selection on MH. Yet, the price variation approach can reveal all the welfare-relevant selection as it helps tracing out the welfare-relevant cost curves given the unpriced heterogeneity in the current UI system. We turn to this in Section 6.

5.2 Implementation

We exploit a sudden and unanticipated increase in the premia paid for the supplemental coverage in 2007. The reform followed the surprise ousting of the Social Democrats from government after the September 2006 general election. With this reform, monthly premia, which had been remarkably stable over the previous years, suddenly increased from 100 SEK to around 320 SEK on January 1st, 2007, as shown in Figure 11. The Figure also shows that the take-up of the supplemental coverage responded significantly to this sharp surge in prices. After staying almost constant around 86%, the fraction of the eligible population buying the comprehensive coverage abruptly dropped to 78% right after the reform. Interestingly, Figure 11 displays little sign of pre-trends or anticipation in the take-up rate of the comprehensive coverage, adding credibility to the assumption that this sudden

¹⁹The outcomes of interest from the insurer's perspective are of course the coverage costs under each contract, $Y_k = c_k$. In principle, we would want to see how the costs under the comprehensive coverage $E_{\mathbf{p}}[c_1]$ and under the basic coverage $E_{\mathbf{p}}[c_0]$ varies with \mathbf{p} . This also allows to account properly for selection on moral hazard. We discuss this further in Section 6.

increase in premia, following the surprise change in political majority, was arguably exogenous to individuals' willingness to pay \mathbf{v} . The unemployment rate was also smoothly decreasing throughout the period, so that the increase in \mathbf{p} cannot be explained by an endogenous pricing response to an increase in the underlying costs of the comprehensive coverage.²⁰

Non-Parametric Tests of Risk-Based Selection The large response to the 2007 exogenous change in prices allows for the identification of marginal individuals, individuals who were buying under the 2006 price regime but are not buying anymore under the 2007 price regime. From equation (1) describing insurance choices, it is immediate to see that these marginal individuals must have a willingness to pay for the supplemental coverage \mathbf{v} lying between the value of the premia \mathbf{p} pre-2007 and the value of the premia \mathbf{p} post-2007.

The identification of marginal individuals therefore enables to rank individuals in three distinct groups by descending order of their willingness-to-pay \mathbf{v} for the supplemental coverage, which we do in Figure 12. First, the group denoted I of individuals insured both in 2006 and 2007: they were buying the supplemental coverage in 2006 under the low premia and continue to buy the supplemental coverage under the high premia, and therefore have the highest level of \mathbf{v} . The marginals M , who were insured in 2006 but switch out in 2007 when premia increase, have a lower willingness-to-pay for insurance than the always insured I . Finally, individuals who were neither insured in 2006 nor in 2007, that we denote by U , have the lowest willingness-to-pay for the supplemental coverage.

Using this ranking, we can now perform direct non-parametric tests for risk-based selection, by correlating willingness-to-pay with various measures of unemployment risk Y . Because the marginals and the uninsured are now observed under the same UI coverage (i.e., the basic plan), the comparison of the average risk of these two groups, $E_M[Y_0] - E_U[Y_0]$ is immune to moral hazard and provides a direct estimate of risk-based selection.²¹ Comparison of the average risk of the marginals and the always insured $E_I[Y_1] - E_M[Y_0]$ will be a combination of selection and moral hazard, as these two groups are now observed under different coverages.

Figure 12 presents the results of such non-parametric tests and provides direct compelling evidence of the presence of risk-based selection into UI. Panel A starts by reporting the average probability of being displaced in 2008 for each group, while panel B reports the average number of days spent unemployed in 2008 for each group.²² Both panels clearly show that the average

²⁰If anything, the 2007 premia reform was combined with a minimal legislated *decrease* in the benefits received in the comprehensive coverage. On January 1st 2007, the cap on the benefits b_1 was slightly decreased for benefits received in the first 20 weeks of an unemployment spell. Given this reform had only a negligible effect on average benefits received, we neglect it in the welfare implementation.

²¹It is worth noting again the timing of the Swedish UI policy: one needs to contribute for at least 12 months in order to become eligible to the comprehensive benefits b_1 . Marginals and uninsured in 2007 did not contribute any premium to the comprehensive plan in 2007. In 2008, if they become unemployed they will therefore get the basic benefits b_0 irrespective of their insurance choice in 2008. In other words, because of their insurance choice in 2007, marginals and uninsured face the exact same coverage in 2008. The difference in their unemployment risk in 2008 cannot be driven by moral hazard due to different coverage choices in 2008.

²²Note that for each risk panel, we report the average outcome of each group conditional on \mathbf{X} , the same vector of controls for contract differentiation that we use in the positive correlation tests. To be precise, we report the average

unemployment risk in 2008 of the marginals is significantly higher (30% to 40% higher) than that of the uninsured, despite both groups being eligible to the same coverage in 2008. This is direct evidence of risk-based selection.

Panels A and B use realized risk outcomes in 2008, i.e., in the year after insurance choices are made and willingness-to-pay is observed. However, the risk realization in 2008 may fail to fully capture the unemployment risk faced by an individual when choosing how much coverage to get in 2007, which justifies using risk realizations further into the future, as done previously in the context of the positive correlation test. In Figure 13 we report the correlation between willingness-to-pay in 2007 and displacement outcomes in $t+1$, $t+2$,... up to $t+5$. We report in panel A for each year $k \in 2008, \dots, 2012$, the semi-elasticity $(E_M[Y_k] - E_U[Y_k])/E_U[Y_k]$ of the displacement rate in year k for the marginals M relative to the uninsured U .²³ The Figure reveals a dynamic pattern comparable to that of Figure 2. The displacement rate of the marginals M is 30% larger than that of the uninsured in the first years, but the semi-elasticity then decreases as we consider later years. Yet, it remains statistically significant up to five years.

Panels A and B of Figure 12 also show that the average unemployment risk in 2008 of the marginals is significantly smaller than that of the always insured. But because marginals and always insured are eligible to different plans in 2008, and because panels A and B focus on realized risk outcomes in 2008, once insurance choices are made, this result might be driven by moral hazard as much as risk-based selection. In panels C and D, we make use of our risk shifters to look at measures of unemployment risks that are independent of individuals' own actions, and therefore immune to individual moral hazard. In panel C, we report for each group the average firm displacement risk of the firm that individuals are in as of 2006, prior to the price variation.²⁴ The graph shows a clear monotonic pattern. Compared to the always insured I , the marginals M are in firms that have a lower layoff risk on average. In other words, individuals are much more likely to switch out of the comprehensive coverage when they are in a lower risk firm. And marginals M are still in firms that are more risky on average than the never insured U .

In panel D, we use the predicted risk model presented in the previous subsection, and compute, using all observable risk shifters, the predicted layoff risk of all individuals based on their observable characteristics as of 2006, prior to the price variation.²⁵ Panel D shows the average individual predicted layoff risk of all three groups and shows again clear monotonicity in the relationship between risk and willingness to pay \mathbf{v} . The always insured I have a significantly larger predicted risk than the marginals M , who also have a significantly larger predicted risk than the uninsured U .

outcome, fixing the average characteristics \mathbf{X} in each group at the same level as that of the uninsured U .

²³In panel B, we report the corresponding semi-elasticity $(E_I[Y_k] - E_U[Y_k])/E_U[Y_k]$ for the insured I .

²⁴For each individual i working in firm j , the average firm displacement risk is defined, following Section 4.2, as the average probability of displacement of all other workers within the firm excluding individual i , $\pi_{-i,j}$ over all years where the firm is observed active in our sample years.

²⁵We fix observable characteristics as of 2006, prior to the price change, as individuals might have changed these characteristics endogenously in 2007 based on their new insurance coverage choice, which would reintroduce potential moral hazard.

Unpriced Observables Figure 12 provides clear and direct evidence of the presence of significant risk-based selection into the comprehensive UI plan. Interestingly, the 2007 price reform also enables to investigate how much of this risk-based selection is driven by selection on unpriced observables correlated with risk, in a similar fashion to what we did in the context of the positive correlation tests in Section 3.3. To this effect, Figure 14 correlates various unpriced observable characteristics with individuals willingness-to-pay, following the same strategy as in Figure 12. In Panel A, we correlate age with willingness-to-pay. The average age of the marginals and the insured is significantly larger than that of the uninsured, which is in line with the evidence from Section 3.3 that younger individuals are less likely to buy the supplemental coverage, despite having a higher unemployment risk on average, so that age is providing advantageous selection on average. Yet, the graph also shows that the magnitude of this selection effect is rather small: the marginals and the insured are less than one year older than the uninsured on average. Panel B of Figure 14 shows relatively little selection along education levels. The differences in terms of total number of years of completed education across groups are quite small (less than .3 years). Panel C correlates the average displacement probability in the industry that individuals are in as of 2006 with willingness-to-pay, and confirms that individuals in riskier industries are more likely to buy the comprehensive UI coverage. Still, the amount of selection based on industry-level risk remains relatively small in magnitude, and confirms the evidence of Section 3.3 that a large and significant amount of risk-based selection remains even when controlling for industry-level variations in unemployment risk. Finally, Panel D investigates the importance of selection based on past unemployment history. The graph correlates the total number of days spent unemployed in the PES registers from 2002 to 2006 with willingness-to-pay. It again suggests the existence of a strong selection along this dimension: individuals who have been unemployed in the past have a significantly higher willingness-to-pay for the comprehensive coverage. Overall, evidence from Figure 14 corroborates previous evidence of Section 3.3 that unpriced observables have a somewhat limited power in explaining the magnitude of risk-based selection in the supplemental UI coverage, although one dimension, past unemployment history, is relatively important in creating adverse selection.

Selection on Preferences The 2007 price reform also allows to investigate patterns of selection along dimensions other than risk. In Figure 15, we examine how characteristics that proxy for preferences and for the expected value of having unemployment insurance correlate with willingness-to-pay for insurance revealed by the 2007 price variation. Panel A correlates the level of individuals' net wealth in 2006 in thousands of SEK with their willingness-to-pay controlling for age. Individuals with larger net wealth have more means to smooth consumption in case of displacement, and as a result, should value extra coverage less. The graph indeed confirms the presence of a clear monotonic relationship between net wealth and \mathbf{v} : the uninsured U have significantly larger net wealth than the marginals M , who have significantly more net wealth than the insured I . In panel B, we probe into the amount of selection based on risk-preferences. To proxy for risk aversion, we use the fraction of total net wealth invested in risky assets (stocks). The graph shows that the

uninsured U have a significantly larger fraction of risky assets in their portfolio than the marginals and the insured, conditional on net wealth. This suggests that individuals with higher risk aversion value the extra coverage more.

Robustness Note that our partition of the population in terms of willingness to pay implicitly assumes that \mathbf{v} is constant over time, or to be more precise that the ranking of individuals' willingness to pay is the same in 2006 and 2007. In practice \mathbf{v} may change over time, due for instance to idiosyncratic shocks to risk, or preferences, thus creating a flow of individuals switching out of the comprehensive plan, even absent price changes. Appendix Figure A.3 provides evidence that the flow of individuals who switch out of the supplemental coverage was in fact very small prior to the 2007 price reform, but experienced a sudden surge in 2007. This alleviates the concern that our ranking of individuals by willingness-to-pay is confounded by underlying changes in individuals' preferences or risk profiles.

Note also that our partition of the population ignores a negligible fourth group of individuals, who were not buying the comprehensive plan in 2006, but switched in the comprehensive plan in 2007. The size of this group is seven times smaller than the group of individuals switching out of the comprehensive plan in 2007. The ranking of this fourth group in terms of willingness-to-pay is ambiguous, as one would need to include idiosyncratic shocks to \mathbf{v} to account for the fact that these individuals switched in the comprehensive coverage in 2007 despite the increase in prices \mathbf{p} . Yet, the behavior and characteristics of this fourth group provide additional evidence of the presence of risk-based selection. We display in Appendix Figure A.3 the evolution of the flow of individuals not buying the comprehensive plan in $t - 1$ but switching in the comprehensive plan in t . The graph shows that this flow of individuals was small prior to 2007, and equivalent in size to the flow of individuals switching out, hence the stability in the fraction of individuals insured displayed in Figure A.3. But the flow of individuals switching in seems to decrease with the 2007 reform. We also note that the average unemployment risk of the workers switching into the comprehensive plan is the highest among the four groups throughout this period.

Summary of empirical evidence We first showed the presence of a large and significant positive correlation between insurance choices and unemployment risk. Second, we have gone beyond positive correlation tests, by exploiting both risk variation and price variation. We have shown, following these two approaches that the positive correlations are not only driven by moral hazard but also by severe adverse selection into the comprehensive UI coverage. Third, we have shown that individuals strongly react to variations in their underlying unemployment risk. Fourth, the current UI policy does not allow prices to be conditioned on many observables correlated with risk. Yet, we provided clear evidence that even if these observables were priced, severe adverse selection would remain. Overall, this combined set of empirical results provides direct and compelling evidence that individuals have private information about their unemployment risk, and act on this asymmetric information when making their unemployment insurance choice, creating large adverse

selection in the UI market. In the next section, we analyze the policy implications of this large adverse selection for the optimal design of UI systems.

6 Policy Implications

This section goes beyond the identification of adverse selection and aims to assess what policy can do to alleviate the potential inefficiencies created by adverse selection. We first provide a direct test of the desirability of mandating insurance, the standard policy response in UI, using the (partial) identification of both value and cost of insurance for the same workers. However, a universal mandate is not the only intervention policy makers can resort to in the face of adverse selection. To study this, we build on the conceptual framework introduced in Section 3.1, but setting this up as a government problem where the government mandates a basic insurance plan and allows individuals to opt for comprehensive insurance at a premium. We use this framework to analyze the two common choice-based interventions in social insurance programs (and in Sweden’s UI policy): minimum mandates and subsidies.²⁶ We provide a simple characterization of their welfare impact, expressed as a function of moments that can be linked to our empirical estimates and allowing for a concrete evaluation of the UI policy in place. Our framework allows for multi-dimensional heterogeneity and moral hazard, in the spirit of Einav et al. [2010b], and our policy instruments involve changes in both insurance premia and coverage, in the spirit of Azevedo and Gottlieb [2017] and Veiga and Weyl [2016]. Our welfare analysis, however, ignores the presence of other frictions or inefficiencies (e.g., behavioural frictions, externalities).²⁷

6.1 Efficiency Consequences of Universal Mandate

From an efficiency perspective, an individual should be insured under the comprehensive plan rather than under the basic plan when the difference in value (i.e., expected utility) for the two plans exceeds the difference in costs, $\mathbf{v} \geq \mathbf{c}$.²⁸ But from a private perspective, when given the choice, an individual chooses to buy the comprehensive plan whenever the difference in value exceeds the difference in price, $\mathbf{v} \geq \mathbf{p}$. When the price differential an individual faces does not reflect the difference in her costs, choice inefficiencies can arise. A choice inefficiency is characterized by either $\mathbf{p} > \mathbf{v} > \mathbf{c}$ or $\mathbf{c} > \mathbf{v} \geq \mathbf{p}$.

Adversely selected markets, like the UI market in Sweden, are prone to such inefficiencies. Under average-cost pricing, the sorting of high-risk individuals into the comprehensive plan drives the price differential up and discourages some individuals from buying the comprehensive coverage, although for some of these individuals the welfare surplus from the extra coverage may be positive

²⁶The Affordable Care Act is a recent policy intervention in the US health insurance market that involved both instruments.

²⁷A number of papers document how insurance choices are subject to important frictions [e.g., Abaluck and Gruber [2011], Handel [2013], Handel and Kolstad [2015]], which can invalidate the use of our revealed preference approach for welfare purposes [see Spinnewijn [2017]].

²⁸Following earlier notation, we continue to use the bolded notation $\mathbf{x} \equiv x_1 - x_0$ to define the difference between the comprehensive and basic UI coverage plan.

($\mathbf{p} > \mathbf{v} > \mathbf{c}$). A standard remedy to this adverse selection problem is to mandate *all* individuals to buy comprehensive coverage. However, absent the mandate, some individuals may be efficiently uninsured (i.e., $\mathbf{c} > \mathbf{v}$ and $\mathbf{p} > \mathbf{v}$). A universal mandate would force even those individuals who efficiently decide not to buy the extra coverage into buying it, creating another source of inefficiency.

By observing individuals' choices (revealing \mathbf{v}) and insurers' costs (determining \mathbf{c}), we can compare value to cost and run simple, non-parametric tests to assess the efficiency consequences of mandating all workers to buy the comprehensive coverage. Such tests have important policy implications. The level of replacement rates in the comprehensive coverage in Sweden is very close to the replacement rates of UI mandates in many countries, especially in Europe. The desirability of mandating such a high level of coverage in UI has never been tested before.

Our non-parametric tests are twofold. First, we focus on individuals from group U , who do not buy the comprehensive coverage at price \mathbf{p} . Using a standard revealed preference argument, we can bound their value of the extra coverage by the price (i.e., $\mathbf{v} < \mathbf{p}$). If the average cost of providing insurance to this group exceeds the price (i.e., $E_U(\mathbf{c}) > \mathbf{p}$), it must be that for some individuals on basic coverage it would be inefficient to mandate comprehensive coverage.

Second, price variation allows to apply a similar test to individuals from group M who are at the margin of buying comprehensive coverage at price \mathbf{p} . It is efficient for marginal buyers to buy the comprehensive coverage if and only if $\mathbf{p} > E_M(\mathbf{c})$.²⁹ If it is inefficient even for the marginal buyer to buy comprehensive coverage, *all* of the uninsured may be inefficiently pushed into buying comprehensive coverage through a universal mandate. However, if the surplus is positive for marginals and negative for some of the uninsured, one needs to trade off the surplus gained for those who were inefficiently priced out of the extra coverage and the surplus lost for those who efficiently decided not to buy the extra coverage, following Einav et al. [2010b].

Implementation Before the 2007 reform in Sweden, workers paid a premium of 100SEK ($\approx 11USD$) per month to get access to comprehensive coverage, but receiving a tax credit of 40 percent of the premium. On average 14 percent of workers do not buy the extra coverage between 2002 and 2006. The low premium provides an upper bound on their valuation \mathbf{v} . We would like to compare this to the cost of providing the supplemental coverage to this group, $E_U(\mathbf{c}) = E_U(c_1 - c_0)$. The challenge is that we only observe these workers under the basic coverage, while the time they spend unemployed (and thus the cost of providing coverage) is expected to increase for more generous UI. We can derive a simple lower bound on the cost of providing the supplemental coverage by ignoring the moral hazard response. That is,

$$E_U(c_{1|0} - c_0) \leq E_U(c_1 - c_0),$$

²⁹A natural reason for the cost of insurance to exceed its value to a risk-averse worker is the presence of moral hazard. People do not internalize the consequences of their decrease in efforts on the insurer. Another reason is that people underestimate the value of insurance, either because they underestimate the coverage or the risk itself.

where $c_{x|y}$ denotes the cost of providing plan x given the behavior under plan y . This lower bound equals 560SEK per year, compared to an annual premium of 720SEK (accounting for the tax paid on UI benefits and the tax credit received on the premium). Since the lower bound on the cost is only just below the premium, the presence of a small moral hazard response of $\varepsilon_{\pi,b} \cong .3$ (in the lower tail of estimates of unemployment elasticities [see Schmieder and Von Wachter [2016]]) would push the costs above so that they would exceed the upper bound on the workers' valuation. We therefore conclude that on average the workers not buying comprehensive coverage are doing this efficiently as they value the coverage less than its costs. Hence, imposing a universal mandate that forces them to buy the comprehensive coverage would be inefficient.

To go beyond this first test and provide finer conclusions on the efficiency consequences of a mandate, we need to estimate how the cost of both insurance plans change with workers' valuation. For this, we need to overcome the challenge that the counterfactual cost of an insurance plan not taken by a worker is unobserved. We do two things. First, the price variation in Section 5.2 allows us to identify an intermediate group of buyers M , closer to the margin, and get a finer assessment of how costs change with valuation than when comparing the group of insured I and uninsured U for one given price. Second, we combine our adverse selection estimates in Section 5.2 and the observed costs of providing the comprehensive plan to I and the basic plan to U respectively to get the counterfactual cost for the other groups of workers.³⁰ In particular, we use our predicted risk model (see Panel D of Figure 12) to extrapolate the observed cost of providing basic coverage $E_U(c_0)$ to the workers with higher valuation and the observed cost of providing comprehensive coverage $E_I(c_1)$ to the workers with lower valuation. This extrapolation requires that the risk-based selection is proportional to the selection on predicted risks and that selection on moral hazard is absent, as we explain in detail in Appendix B.

Our implementation provides us with an estimate of the cost of both the comprehensive plan and the basic plan for each group of workers $X \in \{I, M, U\}$, which we extrapolate linearly to obtain the respective cost curves shown in Panel B of Figure 16. The difference between these two curves gives us the cost of supplemental coverage $E_X(c_1 - c_0)$. Our simple approach accounts for moral hazard responses and implies an unemployment risk elasticity $\varepsilon_{\pi,b} \cong .6$, which falls well in the range of recent estimates in the literature [see Schmieder and Von Wachter [2016]]. The approach, however, assumes that the moral hazard response is constant across workers with different valuations (i.e., there is no selection on moral hazard). We carefully assess in Appendix B the sensitivity of our results to selection on moral hazard and find that they are quite robust.

Panel A of Figure 16 puts the cost curve for supplemental coverage together with the demand curve that it corresponds to. These curves thus show the value and corresponding cost for workers ranked based on their valuation of the supplemental coverage. The graph allows for a direct visual assessment of the welfare impact of mandating comprehensive coverage: value is indeed below cost for a substantial share of workers. Both the cost and demand curves are based on linear extrapola-

³⁰Note that with sufficient price variation (both positive and negative), we could avoid our extrapolation and simply trace out non-parametrically both the cost curve for the comprehensive coverage and the cost curve for the basic coverage.

tions. For the linear demand curve, we use the share of individuals buying the the comprehensive plan before and after the premium increase in 2007. The 8 percent of workers who switch out of the comprehensive plan with the price reform value the supplemental insurance somewhere between the pre-reform price of 720SEK and the average post-reform price of 4,116SEK. The cost curve $E_{\mathbf{p}}(\mathbf{c})$ is a piecewise linear interpolation of our estimates of the costs for the supplemental coverage for the insured, marginals and uninsured. The implied cost for the individuals switching out of the comprehensive plan ranges between 1,442SEK and 1,582SEK. We find that for workers who are marginal at the pre-reform price it is not efficient to buy comprehensive coverage, while it is for workers who are at the margin for the post-reform price. The premium at which value and cost coincide equals 1,477SEK, at which 84 percent of workers would buy insurance.

In sum, our implementation indicates that in Sweden insurance choices are efficient for workers who bought comprehensive coverage at the higher post-reform price and for those who did not at the low pre-reform price. Mandating the latter individuals to buy comprehensive insurance would decrease welfare. This result is interesting. It shows that despite severe adverse selection into comprehensive coverage, a mandate of such generous coverage is inefficient in Sweden. This result has important implications more broadly, as it clearly raises the question whether the universal mandates of generous UI coverages observed in many other countries are socially desirable.

6.2 Choice-based Policy Interventions

When a universal mandate is not desirable, the question remains how to intervene in an adversely selected market while allowing for choice? We analyze two choice-based interventions that are standard in insurance markets and often used in combination: a subsidy, which induces people to buy more coverage and reduces under-insurance due to adverse selection, and a minimum mandate, which increases coverage for people who are potentially priced out of the market and thus mitigates the consequences of adverse selection.

We analyze these two interventions in a setting where the government designs a basic insurance plan (b_0, p_0) and a comprehensive insurance plan (b_1, p_1) , between which workers choose, and aims to maximize social welfare:

$$\mathcal{W} = \int_{\mathbf{v} \geq \mathbf{p}} \omega(v_1 - p_1) dG(\mathbf{v}) + \int_{\mathbf{v} < \mathbf{p}} \omega(v_0 - p_0) dG(\mathbf{v}) + \lambda \{F[p_1 - E_I(c_1)] + (1 - F)[p_0 - E_U(c_0)]\},$$

where λ equals the marginal cost of public funds, pre-multiplying the expected surplus on the unemployment policy, and $F = 1 - G(\mathbf{p})$ equals the share of individuals buying comprehensive insurance (group I). The function $\omega(\cdot)$ maps individuals' consumer surplus into welfare and we denote by $g_j = \omega'(v_j - p_j)$ the marginal social gain from a dollar when buying plan j .

6.2.1 Subsidy

A subsidy lowers the price of coverage and induces more individuals to buy it. A subsidy achieves the same outcome as a universal mandate in the extreme case that it induces everyone to buy insurance. A subsidy, however, can aim to specifically target individuals for whom the surplus is positive.

We consider the determination of a subsidy in our stylized framework where coverage and prices are set by the government. We define the subsidy for supplemental coverage S as the wedge between the difference in average cost of providing comprehensive coverage vs. basic coverage and the difference in prices charged,

$$S \equiv [E_I(c_1) - E_U(c_0)] - [p_1 - p_0].$$

The difference in average cost can be decomposed in three parts: the actual cost of providing the supplemental coverage, a moral hazard response to the supplemental coverage, and risk-based selection into coverage,

$$E_I(c_1) - E_U(c_0) = \underbrace{E_U(c_{1|0} - c_0)}_{\text{Mechanical}} + \underbrace{E_U(c_1 - c_{1|0})}_{\text{MH}} + \underbrace{E_I(c_1) - E_U(c_1)}_{\text{Selection}}. \quad (10)$$

From an efficiency perspective, the premium \mathbf{p} should definitely reflect the cost from providing the supplemental coverage, either due to the mechanical increase in coverage (1st term in (10)) or the behavioral increase due to moral hazard (2nd term in (10)). However, the premium should not reflect the difference in average risks between those who buy the supplemental coverage and those who do not. The subsidy for supplemental coverage allows the government to move a way from average-cost pricing, but also redistributes from those buying basic coverage to those buying comprehensive coverage.

To illustrate the efficiency-equity trade-off, we consider the welfare impact of a small increase in the subsidy dS . This change reduces the price for comprehensive coverage by $dp_1 = -(1 - F) dS$ and increases the price for basic coverage by $dp_0 = F dS$. Absent any behavioural responses, the government's revenues would remain the same. However, the increased subsidy will change the selection into insurance by making comprehensive insurance more attractive.

The small change in the subsidy allows us to invoke the envelope theorem: the impact on agents' welfare only depends on the direct effect of the price changes. The selection response into insurance has a first-order impact on the budget, but not on the agents' welfare. That is, the welfare gain from inducing buyers at the margin (i.e., $\mathbf{v} = \mathbf{p}$) to buy comprehensive insurance simplifies to the corresponding fiscal externality,

$$AS \equiv \mathbf{p} - E_M(\mathbf{c}) = [E_I(c_1) - E_M(c_1)] + [E_M(c_0) - E_U(c_0)] - S.$$

Two insights come out of this simple expression. First, by relating the prices to average costs,

we see that both adverse selection into the comprehensive plan and advantageous selection into the basic plan increase the inefficient wedge between price and marginal costs. That is, with adverse selection, individuals at the margin tend to be cheaper than the average consumer of the comprehensive plan ($E_I(c_1) > E_M(c_1)$), but also more expensive than the average consumer of the basic plan ($E_M(c_0) > E_U(c_0)$). Hence, by inducing the marginals to switch to more comprehensive coverage a budgetary gain is realized at both margins. However, this realized gain in the budgetary costs of both plans comes at the expense of the subsidy already in place. Second, the expression for the fiscal externality AS confirms that moral hazard by itself (i.e., $c_1 > c_{1|0}$) does not contribute to the pricing inefficiency. The impact of moral hazard on the cost differential would be priced efficiently under average-cost pricing if there is no risk-based selection into insurance (i.e., $E_k(c_j)$ does not depend on k for $j = 1, 2$).³¹

The optimal subsidy trades off the fiscal externality and the redistributive consequences, as summarized in the following Proposition:

Proposition 1. *For given coverage levels (b_0, b_1) , the price subsidy S is optimal only if*

$$\left[1 - \frac{E_M(\mathbf{c})}{\mathbf{p}}\right] \times \varepsilon_{\frac{1-F}{F}, \mathbf{p}} = \frac{E_U(g_0) - E_I(g_1)}{\lambda}.$$

The left-hand side of Proposition 1 can be interpreted as the fiscal return to a marginal dollar transferred from the uninsured to the insured due to the reduced adverse selection. This depends on the responsiveness of the demand for supplemental coverage to its price and how much the price exceeds the cost of providing the supplemental coverage to those who respond. The right-hand side can be interpreted as the redistributive cost (or gain if negative) from transferring a marginal dollar from the uninsured to the insured. In the absence of redistributive motives, the optimal subsidy is such that prices reflect the cost of providing insurance to individuals at the margin, i.e., $\mathbf{p} = E_M(\mathbf{c})$. Increasing the subsidy further causes the fiscal externality to be negative (i.e., $\mathbf{p} < E_M(\mathbf{c})$) and can be justified by valuing the redistribution from the uninsured to the insured. However, if value remains below cost for individuals with lower valuation, the redistributive motive from the uninsured to the insured cannot justify inducing everyone to buy comprehensive insurance, as the relative weight of the (negative) fiscal externality goes to infinity when F goes to 1. In other words, a universal mandate can only be the optimal policy if value exceeds cost.

Implementation In Sweden the difference in average costs of providing comprehensive vs. basic coverage was equal to 2,840SEK between 2002 and 2006. We can use our estimates of the observed and counterfactual costs to decompose this difference as in equation (10). We find that 53% of the difference in average costs is driven by adverse selection, while 27% of the difference is due to moral hazard and the remaining 20% is mechanical.³² This decomposition shows the substantial role played by adverse selection and already hints at the desirability of a large subsidy.

³¹As discussed earlier, risk-based selection can be driven by so-called selection on moral hazard (see Einav et al. [2013]), or more generally, any heterogeneity in efforts under either the comprehensive or basic coverage.

³²The details underlying this implementation are again in Appendix B

Before the price reform in 2007 the premium for supplemental coverage paid only for 25% of the difference in average costs. The cost of providing supplemental coverage to individuals at the margin is about twice as high as the premium. The subsidy thus inefficiently induces individuals at the margin to buy insurance causing a negative fiscal externality. In principle, setting the subsidy as high can still be rationalized by valuing the redistribution towards workers buying the comprehensive coverage. The fiscal cost depends on the fiscal externality per marginal worker ($1 - E_M(\mathbf{c})/\mathbf{p} = -1.0$) and the price elasticity of (marginal) workers ($\varepsilon_{\frac{1-F}{F},\mathbf{p}} = .09$), which is estimated to be low due to the modest demand response to the substantial price increase in 2007. The large pre-reform subsidy would thus be optimal if the return from redistributing to the insured equals 9%. The 2007 reform reduced the subsidy dramatically and, in fact, set the premium above the difference in average costs. This indicates that, conditional on the coverage levels, the inelastic demand would still hold this adversely selected market together when pricing at average cost.³³

6.2.2 Minimum mandate

A minimum mandate provides protection against unemployment risk for workers who are priced out of the market for comprehensive coverage. With adverse selection, these workers are the ones with the lowest risk. The challenge is that increasing the basic coverage is particularly attractive to individuals with higher risk and thus makes the selection into comprehensive coverage even more adverse. A minimum mandate thus needs to trade off the protection to individuals who do not buy comprehensive coverage at equilibrium prices against further reducing the share of individuals buying comprehensive coverage. The welfare impact of the latter depends on how generous the subsidy is.

To illustrate this trade-off, we consider a small change in the basic coverage level b_0 . Absent any selection effects, an increase in basic coverage provides more insurance to the group of workers with this coverage, but also reduces their incentives to avoid unemployment. This standard trade-off between insurance and incentives is captured by the well-known Baily-Chetty formula [Baily [1978], Chetty [2006]]. This formula compares the consumption smoothing gains, depending on the relative marginal utility when unemployed $\frac{Eu'(b)-\lambda}{\lambda}$, and the fiscal cost, captured by the unemployment elasticity $\varepsilon_{E\pi,b}$. When considering a change in the minimum benefit level, these moments should be evaluated for workers on basic coverage. On average, these workers value coverage less as they opt out of supplemental coverage, but at the margin, they may value coverage more as they are less covered.

Another key difference with the standard formula comes from the fact that the selection of workers changes in response to the benefit change. An increase in coverage makes the basic insurance plan more attractive and especially so to workers with higher unemployment risk. The corresponding fiscal externality depends on how the cost of supplemental coverage compares to the premium for those who switch to the basic coverage level, i.e., $E_{M(b_0)}(\mathbf{c})/\mathbf{p}$. This effect corresponds

³³Note that the inelastic response may be due to high inertia [see Handel [2013]]. If it were frictional demand that overcomes the pricing inefficiency, this would change the welfare evaluation [see Handel et al. [2015]].

to the sorting effect identified by Veiga and Weyl [2016].³⁴ Note that if workers only differed in their risk, the buyers at the margin responding to a change in UI benefits would be the same as when changing the subsidy. This is no longer true with multi-dimensional heterogeneity.^{35,36} We can state the following result:

Proposition 2. *The minimum mandate b_0 is optimal, for given subsidy S and comprehensive coverage b_1 , only if*

$$\frac{E_U(g_0 \times u'(b_0)) - \lambda}{\lambda} - \varepsilon_{E_U(\pi), b_0} = \left[1 - \frac{E_{M(b_0)}(\mathbf{c})}{\mathbf{p}} \right] \times \frac{\mathbf{p}}{E_U(c_0)} \times \varepsilon_{1-F, b_0}.$$

The proposition delivers two important insights for the design of social insurance more generally:

First, it clearly identifies and provides an implementable characterization of the trade-off between mitigating the consequence of adverse selection at the extensive margin and worsening adverse selection at the intensive margin. This trade-off is also studied by Azevedo and Gottlieb [2017] who provide a general characterization of equilibrium contracts and prices in a competitive market (with multi-dimensional heterogeneity). Their analysis shows how setting a minimum mandate can discourage individuals from buying even more coverage as the mandate effectively pools agents in the lowest range of unemployment risks into the minimum coverage plan and as such reduces its equilibrium price.³⁷ The mechanisms are very related: increasing the minimum mandate provides a more attractive alternative for workers attracted to more comprehensive coverage, either by increasing the available coverage or lowering the price.³⁸ This insight also provides an important perspective on the general absence of private UI, which has been attributed to adverse selection by Hendren [2017]. Indeed, the absence of private UI concerns more specifically the absence of supplemental coverage above and beyond the mandated public UI that is already in place. Adverse

³⁴Veiga and Weyl [2016] characterize the single-traded contract firms offer in an insurance market (with multi-dimensional heterogeneity). Private firms face a similar trade-off in determining the optimal coverage level between creating insurance value and discouraging higher risk from selection into their plan (i.e., cream-skimming). Interestingly, a private firm does not internalize the consequences for competing insurers from attracting different risk types. That is, the negative sorting effect for a private firm increasing b_0 equals $p_0 - E_{M(b_0)}c_0$, which underestimates the total impact, equal to $[p_0 - E_{M(b_0)}c_0 - (p_1 - E_{M(b_0)}c_1)]$. We briefly discuss this further in Appendix C.

³⁵The sorting effect depends on covariance between the marginal utility of coverage and the cost to the insurer (see Veiga and Weyl [2016]). Indeed, we can re-write the marginal cost term as

$$E_{M(b_0)}(\mathbf{c}) = \left[E(\mathbf{c}|\mathbf{p}) + \frac{cov\left(\mathbf{c}, \frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p}\right)}{E\left(\frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p}\right)} \right], \quad (11)$$

illustrating that the marginal cost term will be higher than for the subsidy if costs are positively correlated with marginal value. This would occur naturally with heterogeneity in risks. However, if there is only heterogeneity in risks, the covariance among the marginals equals zero.

³⁶See Ericson et al. [2015] for an analysis of differential responses to changes in coverage and changes in prices depending on the heterogeneity in risks and preferences. In advantageously selected markets, the increase in minimum coverage could as well attract types with lower unemployment risk (but higher risk aversion).

³⁷See also Finkelstein [2004] and Chetty and Saez [2010] for related analyses.

³⁸For completeness, we provide a characterization for the optimal comprehensive contract b_1 in Appendix C. Note that the trade-off faced in setting the comprehensive coverage, either for the social planner or for profit-maximizing firms, will change when the basic coverage will change.

selection may be less constraining if the public mandate entailed a lower coverage level.³⁹

Second, the proposition highlights the complementarity between subsidies and mandates. The welfare impact of the selection effect critically depends on the subsidy in place. The fiscal cost of increasing adverse selection is lower for a higher subsidy, making a higher minimum mandate more desirable. More generally, while the minimum mandate worsens adverse selection into the supplemental market (or even excludes the viability of a private market as in Hendren [2017]), this inefficiency can be countered by a more generous subsidy for supplemental coverage.^{40,41}

Implementation In Sweden the minimum UI benefit equals 320SEK ($\approx 35USD$) a day, which corresponds to 20 percent of median pre-unemployment earnings. The value of increasing this level depends on the consumption smoothing gains provided by the additional coverage. Landais and Spinnewijn [2017] estimate an average consumption drop of around 10% during unemployment for workers on basic coverage, implying a return to a marginal kroner spent on basic coverage of 20 percent for relative risk-aversion $\gamma = 2$.⁴² This insurance value needs to be traded off against the fiscal cost due to moral hazard and adverse selection. A large literature has estimated moral hazard responses finding a large range of elasticities (see Schmieder and Von Wachter [2016]). As mentioned before, our estimates imply an elasticity $\varepsilon_{\pi,b} = .6$, suggesting that the moral hazard cost would well be above the consumption smoothing gain of .2 and thus that the minimum mandate would be set too high. Note, however, that the consumption-based implementation is likely to provide a lower bound on the consumption smoothing gains [see Chetty [2008], Chetty and Finkelstein [2013]].⁴³

More importantly, we still need to correct the cost of increasing the minimum mandate for the fiscal externality due to adverse selection. In an adversely selected market, we would expect this fiscal externality to further increase the cost. However, due the large subsidy pre-reform in Sweden, the fiscal externality is actually negative and thus reduces the gap between insurance value

³⁹As discussed before, we find that a large majority of Swedish workers continues to buy the supplemental coverage when priced just above average cost. Hendren [2013] shows formally that the existence of a (supplemental) market requires that individuals are willing to pay a markup for insurance that exceeds the cost imposed by worse risks adversely selecting their contract. Note that this markup decreases as the basic coverage level increases for risk-averse agents. Moreover, the variation in willingness-to-pay coming from heterogeneity in risk aversion decreases as the basic coverage becomes more generous (see Ericson et al. [2015]). This strenghtens the risk-based selection and thus increases the “pooled price ratio”.

⁴⁰See also the simulations in Azevedo and Gottlieb [2017] showing that subsidizing supplemental coverage increases welfare.

⁴¹A similar complementarity was stated in the seminal paper by Rothschild and Stiglitz [1976] showing that by subsidizing the comprehensive coverage (for the “high-risk” type), the basic coverage can be increased (for the “low-risk” type) while maintaining incentive compatibility.

⁴²Using a Taylor approximation, the difference in marginal utilities when employed vs. unemployed simplifies to the drop in consumption between employment and unemployment, multiplied by the relative risk-aversion [Gruber [1997], Chetty [2006]].

⁴³Using an alternative revealed-preference approach, Landais and Spinnewijn [2017] find that the *average* return of the supplemental coverage \mathbf{b} for workers who choose not to buy it at price \mathbf{p} is bounded from above by

$$\frac{\mathbf{P}}{E_U(c_{1|0} - c_0)} - 1.$$

This upper bound on the average return equals .29 in the Swedish context, but may still be lower than the *marginal* return to supplemental coverage, evaluated at b_0 .

and moral hazard cost (and could reverse it in principle). Taking both the fiscal externality of a subsidy and the demand response to a change in the subsidy as lower bounds, we find a downward correction of the moral hazard cost of at least .09.⁴⁴ This correction, however, would be smaller and eventually change signs for lower subsidies, illustrating the complementarity between the minimum mandate and the subsidy.

7 Conclusion

Seventy five years ago, the Beveridge Report, in its attempt at increasing welfare for all, by recommending a new set of revolutionary social insurance policies, insisted that these insurance policies “must be achieved by co-operation between the State and the individual. (...) The state should not stifle incentive, opportunity, responsibility; in establishing a national minimum, it should leave room and encouragement for voluntary action by each individual to provide more than that minimum for himself and his family”.⁴⁵ In the context of unemployment insurance though, generous mandates have left very little room for choice and voluntary actions by individuals, under the (untested) rationale that offering choice would trigger risk-based selection. Our paper provides the first direct evidence for risk-based selection into unemployment insurance and offers new insights on how to reconcile social insurance design with individual choice in such a context of adverse selection.

Using various empirical strategies and different sources of variation, we robustly find that workers who face higher (ex-ante) unemployment risk select into more comprehensive coverage. We further leverage our rich administrative data to show that the severe adverse selection is expected to survive when workers’ observable risks were to be priced. Despite the severe adverse selection, we find that mandating all workers into comprehensive coverage is not the best policy response. The Swedish workers who choose not buy the comprehensive coverage value it below its cost. This simple observation sheds new light on the universal mandates of comprehensive UI that are in place around the world – the desirability of which has never been tested before. Moreover, the ubiquitous absence of private markets for UI may be precisely because of the mandated public programs in place and the adverse selection into supplemental coverage they cause. As our analysis shows, the impact of adverse selection can be mitigated by subsidizing the premia for more comprehensive coverage, which would increase the desirability of a more generous minimum mandate.

⁴⁴As shown before, the fiscal externality of a subsidy is indeed a lower bound when costs and marginal value of coverage are positively correlated. Regarding the demand elasticity, we note that a risk-neutral worker values the supplemental coverage at $\pi \mathbf{b} - \mathbf{p}$ and thus responds proportionally to a change in coverage or in prices. A risk-averse worker responds relatively more to a change in coverage. Hence,

$$\varepsilon_{1-F,b_0} \geq \varepsilon_{1-F,\mathbf{p}} \frac{\pi b_0}{\mathbf{p}} = F \varepsilon_{\frac{1-F}{F},\mathbf{p}} \frac{E_M(c_0)}{\mathbf{p}}.$$

This implies

$$\frac{\mathbf{p}}{E_U(c_0)} \times \varepsilon_{1-F,b_0} \geq \frac{E_M(c_0)}{E_U(c_0)} F \varepsilon_{\frac{1-F}{F},\mathbf{p}},$$

which corresponds to .088 in our Swedish context.

⁴⁵Beveridge [1942]


The value of providing choice in social insurance programs does rely on individuals using this choice in their best interest. This is an assumption we have maintained throughout our analysis. A rapidly growing literature documents the importance of frictions distorting households' insurance choices. This introduces another important caveat when introducing choice, closely related to the potential for adverse selection. We leave this challenging, but important issue for future research.

References

- Abaluck, Jason and Jonathan Gruber, “Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program,” *American Economic Review*, 2011, 101 (4), 1180–1210.
- Akerlof, G., “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism,” *Quarterly Journal of Economics*, 1970, 84 (3), 488–500.
- Azevedo, Eduardo M. and Daniel Gottlieb, “Perfect Competition in Markets With Adverse Selection,” *Econometrica*, 2017, 85 (1), 67–105.
- Baily, Martin N., “Some Aspects of Optimal Unemployment Insurance,” *Journal of Public Economics*, 1978, 10 (3), 379–402.
- Beveridge, William Henry, *Social Insurance And Allied Services*, HM Stationery Office London, 1942.
- Carroll, Eero, “Voluntary State-Subsidized Social Insurance in the Advanced Industrialized World since the 1890s: The Nordic Experience in Comparative Perspective,” in Olli Kangas and Joakim Palme, eds., *Social Policy and Economic Development in the Nordic Countries*, London: Palgrave Macmillan UK, 2005, pp. 60–87.
- Chetty, Raj, “A General Formula for the Optimal Level of Social Insurance,” *Journal of Public Economics*, 2006, 90 (10-11), 1879–1901.
- , “Moral Hazard versus Liquidity and Optimal Unemployment Insurance,” *Journal of Political Economy*, 2008, 116 (2), 173–234.
- and Amy Finkelstein, *Chapter 3 - Social Insurance: Connecting Theory to Data*, Vol. 5 of *Handbook of Public Economics*, Elsevier,
- and Emmanuel Saez, “Optimal Taxation and Social Insurance with Endogenous Private Insurance,” *American Economic Journal: Economic Policy*, May 2010, 2 (2), 85–114.
- Chiappori, Pierre-André and Bernard Salanié, “Testing for asymmetric information in insurance markets,” *Journal of Political Economy*, 2000, 108 (1), 56–78.
- Einav, L., A. Finkelstein, and J. Levin, “Beyond Testing: Empirical Models of Insurance Markets,” *Annual Review of Economics*, 2010, 2, 311–336.
- , —, and M. Cullen, “Estimating Welfare in Insurance Markets Using Variation in Prices,” *Quarterly Journal of Economics*, 2010, 125 (3), 877–921.
- Einav, Liran, Amy Finkelstein, Stephen Ryan, Paul Schrimpf, and Mark Cullen, “Selection on Moral Hazard in Health Insurance,” *American Economic Review*, 2013, 103 (1), 178–219.
- Ericson, Keith Marzilli, Philipp Kircher, Johannes Spinnewijn, and Amanda Starc, “Inferring Risk Perceptions and Preferences using Choice from Insurance Menus: Theory and Evidence,” Working Paper 21797, National Bureau of Economic Research December 2015.
- Finkelstein, Amy, “Minimum standards, insurance regulation and adverse selection: evidence from


- the Medigap market,” *Journal of Public Economics*, 2004, 88 (12), 2515 – 2547.
- and James Poterba, “Testing for asymmetric information using “unused observables” in insurance markets: Evidence from the UK annuity market,” *Journal of Risk and Insurance*, 2014, 81 (4), 709–734.
- , Nathaniel Hendren, and Mark Shepard, “Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts,” *Mimeo MIT*, 2017.
- Gruber, Jonathan, “The Consumption Smoothing Benefits of Unemployment Insurance,” *American Economic Review*, 1997, 87 (1), 192–205.
- Hackmann, M., J. Kolstad, and A. Kowalski, “Adverse Selection and an Individual Mandate: When Theory Meets Practice,” *American Economic Review*, 2015, 105 (3), 1030–1060.
- Handel, B., “Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts,” *American Economic Review*, 2013, 103 (7), 2643–2682.
- and J. Kolstad, “Health Insurance For Humans: Information Frictions, Plan Choice, and Consumer Welfare,” *American Economic Review*, 2015, 105 (8), 2449–2500.
- , – , and J. Spinnewijn, “Information Frictions and Adverse Selection: Policy Intervention in Health Insurance Markets,” November 2015. NBER Working Paper 21759.
- Hendren, Nathaniel, “Private information and insurance rejections,” *Econometrica*, 2013, 81 (5), 1713–1762.
- , “Knowledge of Future Job Loss and Implications for Unemployment Insurance,” *American Economic Review*, 2017, 107 (7), 1778–1823. Replication code and files.
- Kolsrud, Jonas, Camille Landais, Peter Nilsson, and Johannes Spinnewijn, “The Optimal Timing of UI Benefits: Theory and Evidence from Sweden,” *American Economic Review*, 2017, *forthcoming*.
- Landais, Camille and Johannes Spinnewijn, “The Value of Unemployment Insurance,” 2017.
- Rothschild, M. and J. Stiglitz, “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *Quarterly Journal of Economics*, 1976, 90, 630–649.
- Schmieder, Johannes F and Till Von Wachter, “The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation,” *Annual Review of Economics*, 2016, 8, 547–581.
- Schmieder, Johannes F., Till von Wachter, and Stefan Bender, “The Effects of Extended Unemployment Insurance Over the Business Cycle: Evidence from Regression Discontinuity Estimates Over 20 Years,” *The Quarterly Journal of Economics*, 2012, 127 (2), 701–752.
- Spinnewijn, J., “Heterogeneity, Demand for Insurance and Adverse Selection,” *American Economic Journal: Economic Policy*, 2017, 9 (1), 308–343.
- Veiga, Andre and E. Glen Weyl, “Product Design in Selection Markets,” *The Quarterly Journal of Economics*, 2016, 131 (2), 1007–1056.

Figure 1: POSITIVE CORRELATION TESTS - BASELINE



Notes: The Figure reports estimates of positive correlation tests following specification (3) estimated over the period 2002-2006 for four different realized risk outcomes: total UI claims under comprehensive coverage in $t + 1$, total duration spent unemployed in $t + 1$, the probability of displacement in $t + 1$, and the probability of displacement in $t + 1$ but excluding quits. Total UI benefit claims are defined as the total amount of UI benefits that individuals would be collecting in $t + 1$ were they to buy the comprehensive coverage and receive benefits b_1 . For individuals who do not buy the comprehensive coverage in t , we computed the counterfactual benefit claims they would have if they were to receive the comprehensive benefits. For each outcome, the chart displays $\hat{\gamma}/\bar{Y}$, that is the semi-elasticity of the realized risk outcomes in $t + 1$ with respect to insurance choices in t . Specification (3) controls for year fixed effects and for the limited set of characteristics that affect the unemployment insurance coverage available to individuals: a dummy for whether individuals meet the work eligibility requirement, a dummy for union membership, and earnings level. See text for details. For all realized risk outcomes, we find a strong and significant positive correlation with UI coverage choice.

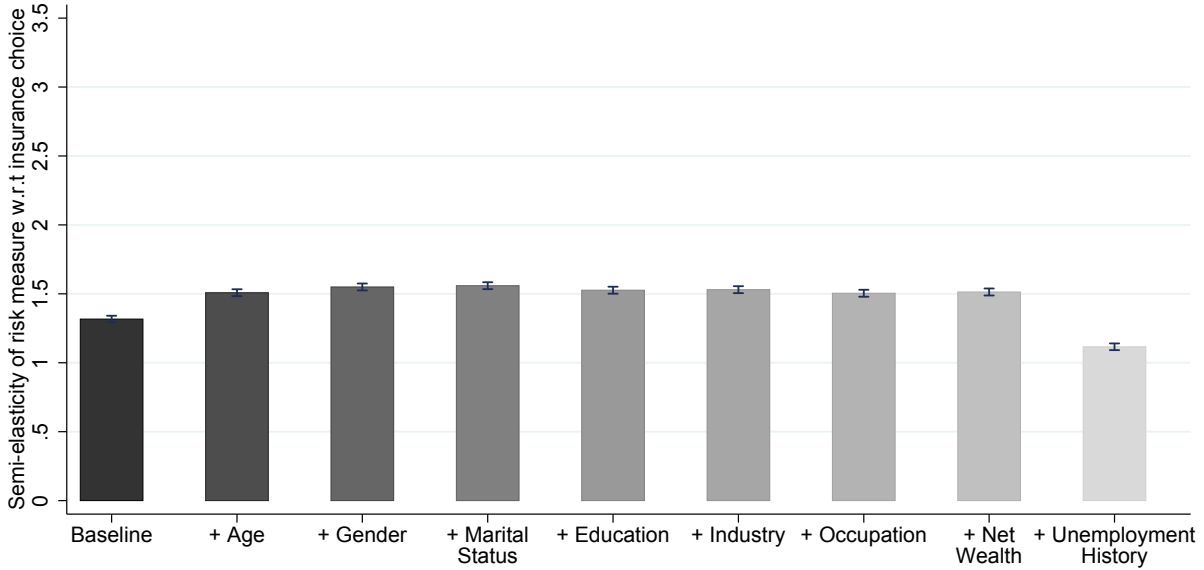
Figure 2: POSITIVE CORRELATION TESTS - DYNAMICS



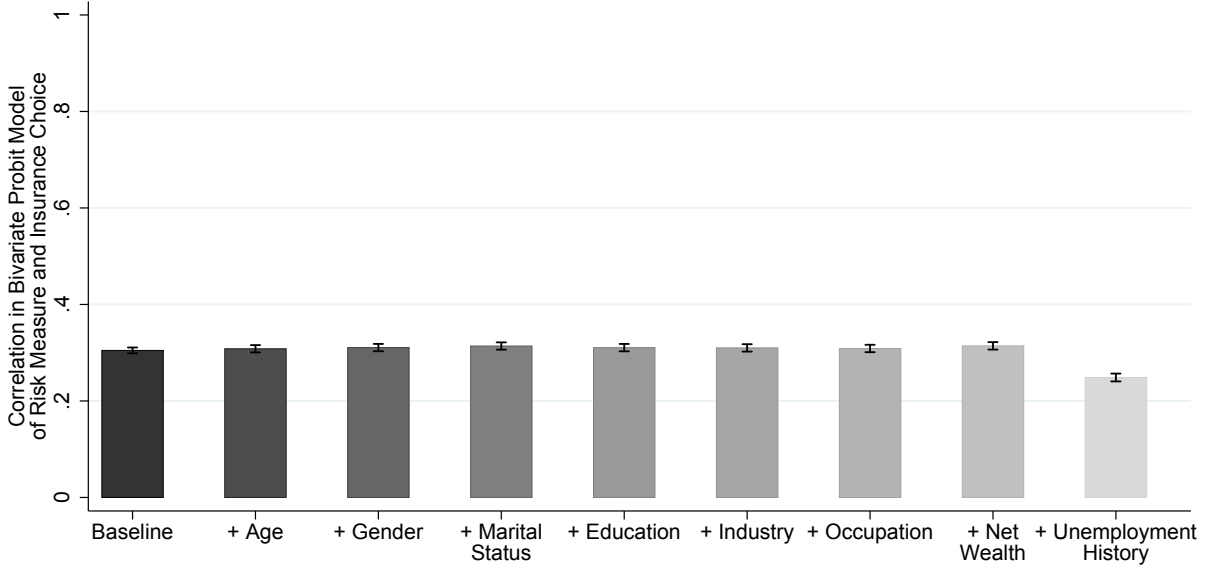
Notes: Risk realization in $t + 1$ may fail to fully capture the unemployment risk faced by an individual as she is making her coverage choice at time t , which justifies using risk realizations for that individual further into the future. This Figure reports the correlation of insurance choice in t with displacement outcomes in $t + 1, t + 2, \dots$ up to $t + 8$. The Figure displays estimates of positive correlation tests following specification (3) estimated over the period 2002-2006. For each outcome, the chart displays $\hat{\gamma}_k/\bar{Y}$, that is the semi-elasticity of the realized displacement rate in $t + k$ with respect to insurance choices in t . For each displacement outcome in year $t + k$, we control for displacement outcomes in previous years ($t + k - 1, t + k - 2$, etc.), for year fixed effects and for the limited set of characteristics \mathbf{X} that affect the unemployment insurance coverage available to individuals. See text for details.

Figure 3: POSITIVE CORRELATION TESTS - ROLE OF UNPRICED HETEROGENEITY

A. Semi-elasticity from linear specification




B. Correlation from bivariate probit



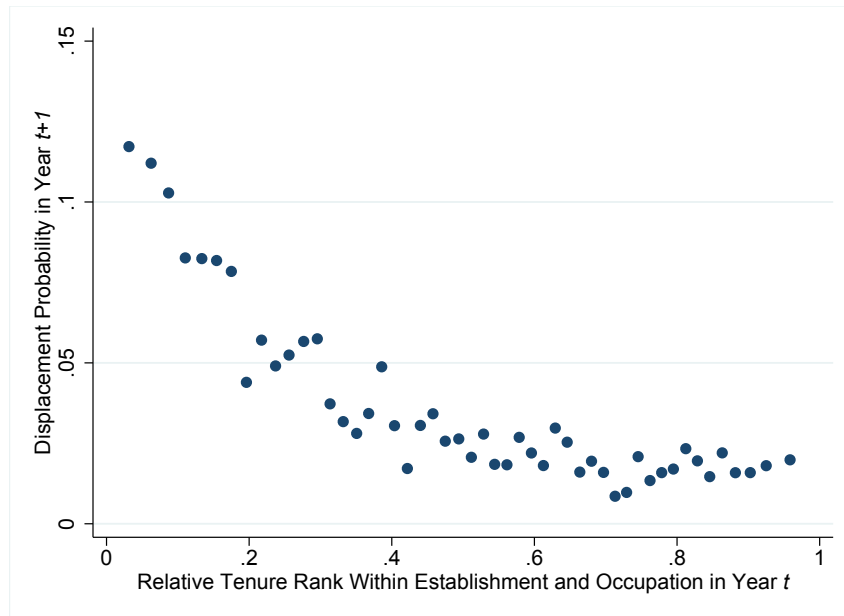
Notes: The Figure explores to what extent the positive correlation between insurance and unemployment risk is driven by selection on observables characteristics that are unpriced in the Swedish UI system. In panel A, we start with the baseline positive correlation test from the linear specification (3) where Y is the probability of displacement in $t + 1$, and show how the semi-elasticity $\hat{\gamma}/\bar{Y}$ evolves as we add sequentially more characteristics to the vector of controls \mathbf{X} . We start by adding demographic controls (age, then gender, and marital status), then controls for skills and other labor market characteristics (controls for education (four categories), industry (1-digit code), occupation (1-digit code) and wealth level (quartiles)). We finally add controls for past unemployment history (dummies for having been unemployed in $t-1$, $t-2$ and up to $t-8$). In panel B, we do a similar positive correlation exercise but using a bivariate probit specification instead. We report how the estimated correlation ρ between the two error terms ϵ and η from the bivariate probit specification (4) evolves when adding sequentially to the vector \mathbf{X} the same set of characteristics as in panel A.

Figure 4: RISK SHIFTERS: FIRM DISPLACEMENT RISK & LAST-IN-FIRST-OUT PRINCIPLE

A. Firm Displacement Risk vs Individual Displacement Probability in $t + 1$



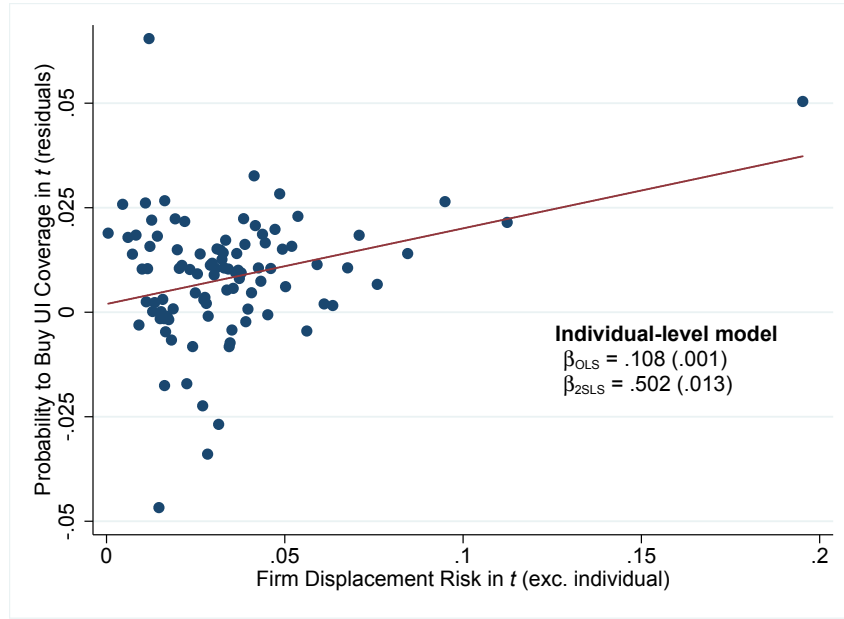
B. Relative Tenure Ranking in year t vs Displacement Probability in year $t + 1$




Notes: The Figure provides evidence of the role of firm level risk and of the Last-In-First-Out (LIFO) principle in creating variation in individuals' unemployment risks. In panel A, we provide evidence of the role of firm layoff risk as a shifter of individuals' own displacement probability. For each individual i working in firm j , we define average firm displacement risk as the average probability of displacement of all other workers within the firm excluding individual i , $\pi_{-i,j}$ over all years where the firm is observed active in our sample years. We then plot the average firm displacement risk in 20 bins of equal population size, against the individual probability of displacement in $t + 1$. The Figure shows that there is significant heterogeneity in firms' separation rates, and that individuals' unemployment risk is very strongly correlated with firm level risk. Panel B plots the probability of being displaced in $t + 1$ among individuals working in firms that emit a layoff notification in $t + 1$, as a function of relative tenure ranking within establishment and occupation in year t . See Section 2 for institutional details. The Figure provides clear evidence of a strong negative correlation between relative tenure ranking and individuals' displacement probability.

Figure 5: FIRM LEVEL RISK AND UI COVERAGE CHOICE

A. Baseline controls for contract space

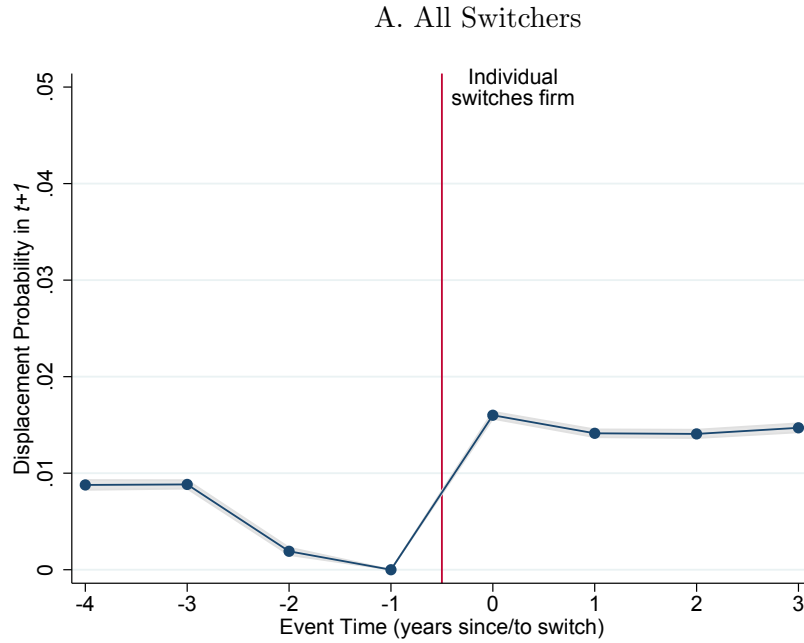


B. With additional demographic controls




Notes: The Figure uses cross-sectional variation in displacement risk across firms as a risk shifter to estimate how UI coverage choices react to variation in risk that is not driven by individual moral hazard. Panel A groups individuals in 50 equal size bins of firm layoff risk, and plot their average firm layoff risk against their average probability of buying supplemental coverage, residualized on the same vector \mathbf{X} of baseline controls affecting UI contracts used in the positive correlation test of Section 3.2. We report on the graph the coefficient β_{OLS} from an OLS regression of specification (7) and then the estimated coefficient β_{2SLS} from our two-stage least square model (8) where we use $Z = \pi_{-i,j}$ as a risk shifter. In panel B, we replicate the same procedure, but now add to the regression the same rich set of additional controls used in Section 3.3, and find a similar strong positive correlation between insurance choices and firm layoff risk.

Figure 6: FIRM SWITCHERS - DISPLACEMENT RATE IN $t + 1$ AS A FUNCTION OF TIME TO/SINCE FIRM SWITCH

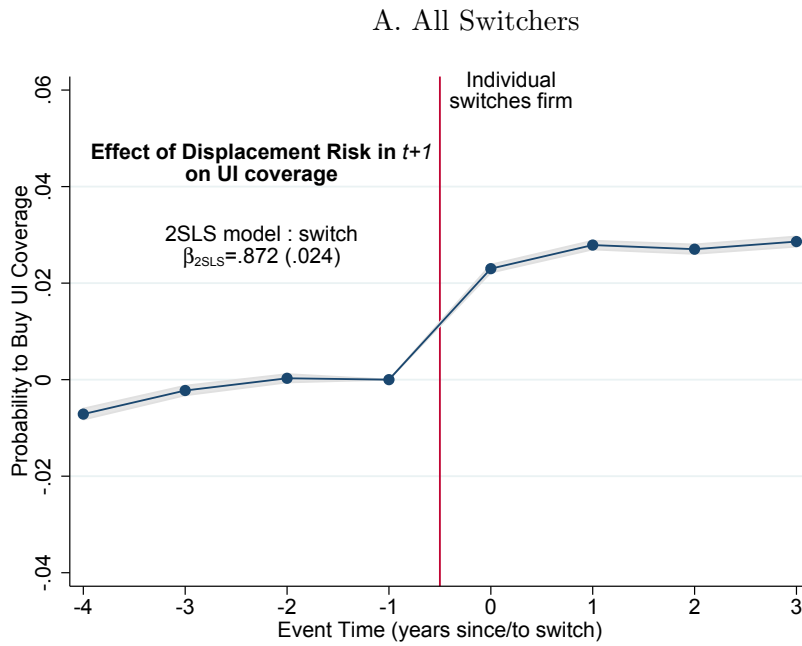


B. Switchers Experiencing Large Positive Firm Layoff Risk Shock vs Large Negative Firm Layoff Risk Shock




Notes: The Figure focuses on “firm switchers”, i.e. individuals moving from having a labor contract with firm j to having a contract with firm k , without any recorded non-employment spell between these two contracts. We focus on individuals with more than 1 year of tenure in the origin firm. Switchers experience a variation in their layoff risk coming from underlying variation in both risk shifters: their tenure ranking changes, and so does their underlying firm layoff risk. In panel A, we display estimates of the event study specification (9) using displacement risk in $t + 1$ as an outcome. The graph shows that the displacement risk increases sharply and significantly at the time of the firm switch. In panel B, we split the population of switchers according to their rank in the distribution of $\Delta_{j,j'}\pi_{-i} = \pi_{-i,j'} - \pi_{-i,j}$, the change in their underlying firm risk when moving from firm j to firm j' . We focus on individuals in the bottom decile of $\Delta_{j,j'}\pi_{-i}$ (large negative shock, i.e., individuals going from a high risk to a low risk firm), and individuals in the top decile of $\Delta_{j,j'}\pi_{-i}$ (large positive shock).

Figure 7: FIRM SWITCHERS - UI COVERAGE CHOICES AS A FUNCTION OF TIME TO/SINCE FIRM SWITCH

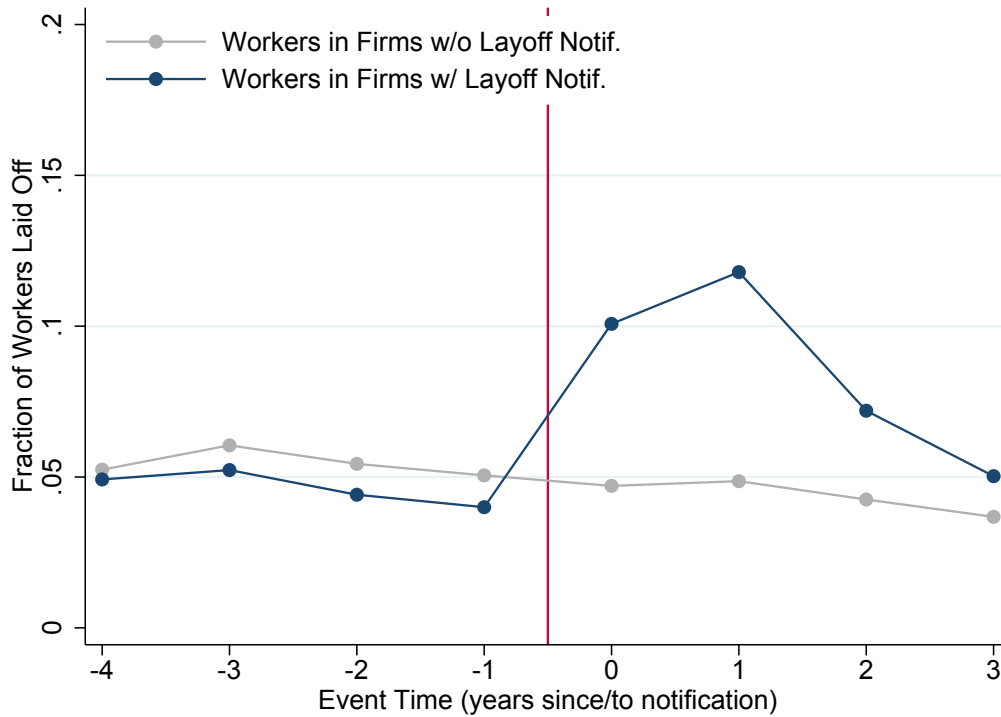


B. Switchers Experiencing Large Positive Firm Layoff Risk Shock vs Large Negative Firm Layoff Risk Shock



Notes: The Figure focuses on “firm switchers”. In panel A, we display estimates of the event study specification (9) using UI coverage V as an outcome. The Figure shows that the probability of buying the comprehensive coverage increases sharply at the time of the firm switch. In panel B, we split the population of switchers according to their rank in the distribution of $\Delta_{j,j'}\pi_{-i} = \pi_{-i,j'} - \pi_{-i,j}$, the change in their underlying firm risk when moving from firm j to firm j' , as in Figure 6 panel B. The graph indicates that the increase in the probability to buy UI around firm switch is significantly larger among individuals moving to significantly more risky firms relative to those moving to less risky firms. On both panels, we display the coefficient from a two-stage least square fixed-effect specification similar to (8) where we use firm switch (and firm switch interacted with shock size) as risk shifter Z for individual displacement probability.


Figure 8: LAYOFF NOTIFICATION AND DISPLACEMENT RISK



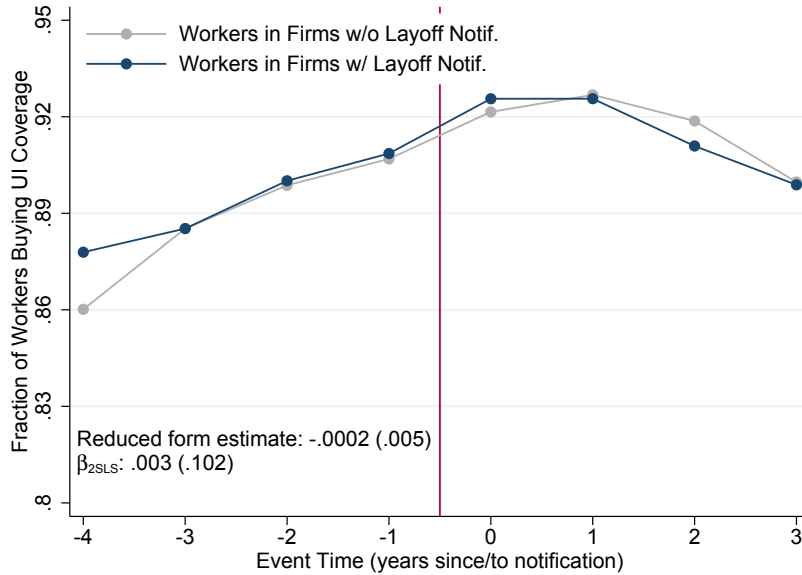
Notes: The Figure reports estimates of the evolution of the displacement probability of workers around the first layoff notification event in the history of the establishment. We define event year $n = 0$ as the year in which an establishment emits its first layoff notification, and focus on the panel of workers who are employed in the establishment at the date this layoff notification is emitted to the PES. The graph shows that a layoff notification is indeed associated with a sudden and large increase in the displacement risk of workers. Because the panel of workers is selected based on being employed in the firm in year $n = 0$, one may worry that this surge in displacement rates is mechanical, as displacement can only increase after year 0 conditional on all workers being employed in year 0. To mitigate this concern, we follow a matching strategy and create a control panel of workers selected along the same procedure as the original panel. We use nearest-neighbor matching to select a set of firms that are similar, along a set of observable characteristics, to the firms emitting a layoff notification, but never emit a layoff notification. We allocate to the matched firm in the control group a placebo event date equal to the layoff notification date of her nearest-neighbor in the treated group of firms. We then select workers that are in the control firm at the time of the placebo event date to create our matched control panel.

Figure 9: LAYOFF NOTIFICATION

A. Workers With Relative Tenure Ranking < .5 at Event Time 0



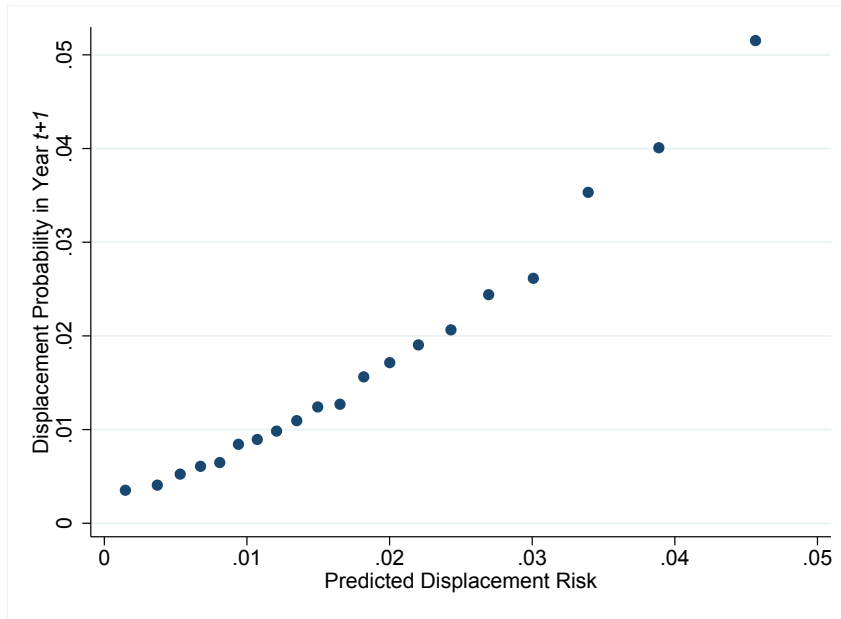
B. Workers With Relative Tenure Ranking $\geq .5$ at Event Time 0




Notes: The Figure uses layoff notification events interacted with relative tenure ranking as a source of variation in displacement risk to investigate how UI coverage choices react to variations in underlying risk. Panel A reports the evolution of UI coverage around the time of the first layoff notification for the panel of workers in the treated group and for workers in our placebo (control) group, restricting the sample to workers with relative tenure ranking below 50% in year $n = 0$. The Figure shows that UI coverage increases significantly among the treated group, starting one year before the layoff notification is sent, which suggests the existence of some degree of private information among workers regarding the timing of the layoff notification. In panel B, we report similar estimates but for the sample of workers with relative tenure ranking above 50% in year $n = 0$. The graph displays no sign of variations in individuals insurance coverage among the event. On both panels, we display the estimated coefficient β_{2SLS} of our two-stage least square model using the layoff event interacted with tenure and a dummy for being in the treatment group as a risk shifter Z .

Figure 10: PREDICTED LAYOFF RISK MODEL

A. Predicted Layoff Risk vs Actual Displacement Rates




B. Predicted Layoff Risk vs UI choices



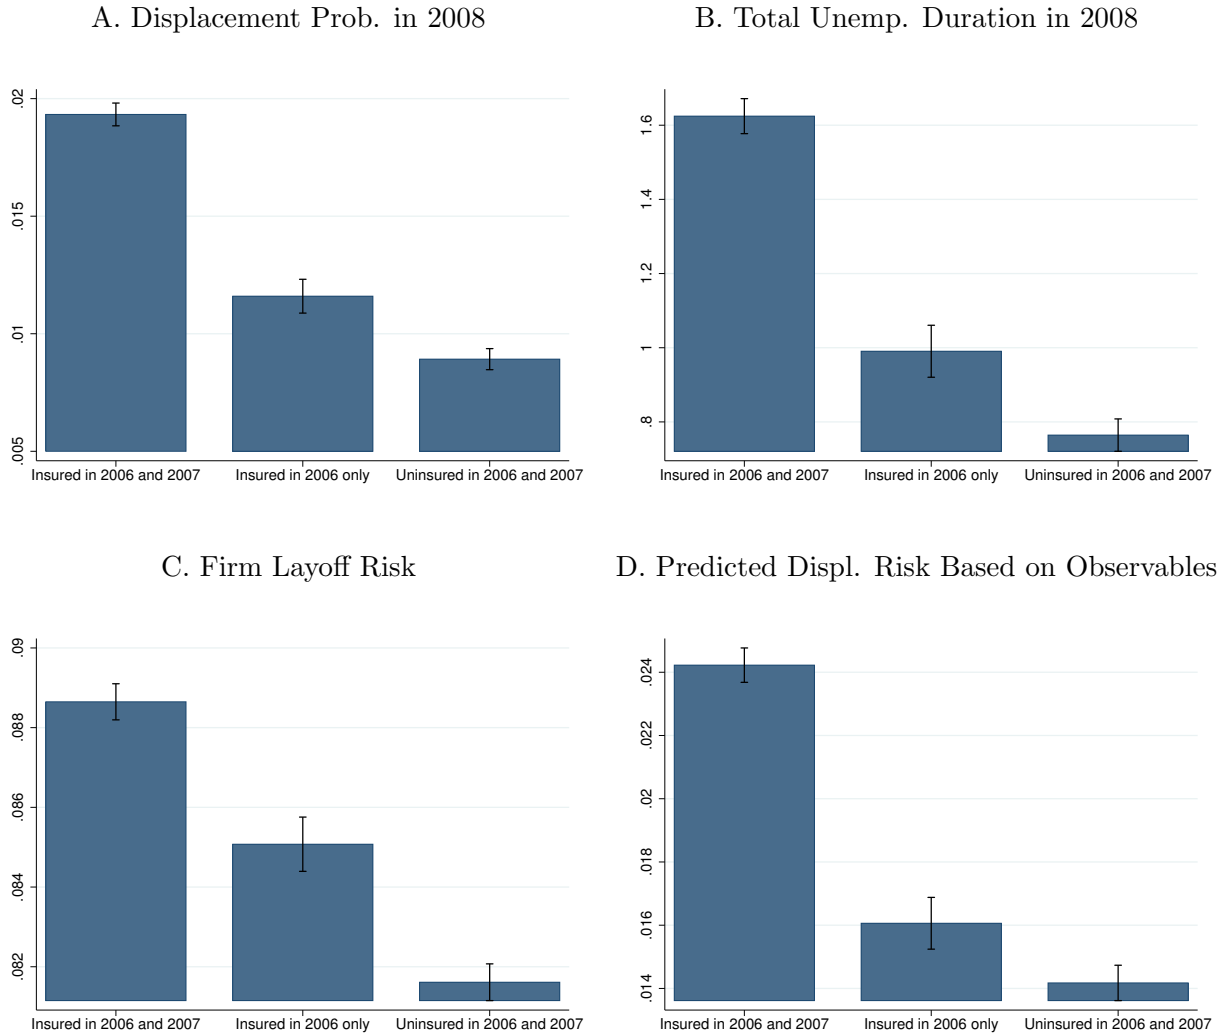
Notes: The Figure reports results from the model of predicted displacement risk in $t + 1$, presented in Section 4.2. The model combines flexibly all observable sources of risk shifting together, and allows for a richer variation in risk than the one obtained from using the risk shifters separately. The risk shifters included in the model are the average firm layoff risk, the full history of the firm layoff notifications, and the relative tenure ranking of the individual. We also add the full past unemployment history of the individual. Model selection follows a cross-validation approach where we minimize misclassification errors in the test sample. Panel A correlates predicted layoff risk with actual displacement rates in $t + 1$ and provides evidence that the model performs well in predicting individuals' risk. Panel B correlates predicted layoff risk with the probability to buy UI coverage in t , and suggests again a strong positive correlation between individuals' risk and their probability to buy the supplemental UI coverage. Interestingly, the graph also suggests that the strong positive correlation between risk and insurance coverage is mostly driven by what happens at the bottom of the predicted risk distribution.

Figure 11: PRICE VARIATION: EVOLUTION OF PREMIA p AND OF THE FRACTION OF WORKERS INSURED AROUND THE 2007 REFORM




Notes: The Figure reports the evolution of monthly premium for the supplemental UI coverage over time. As explained in Section 2, there are no sources of premium differentiation up to 2008, apart from small rebates for union members and for unemployed individuals. Here, we report the value of the premium for employed union members. The Figure shows a large and sudden increase in the premia paid for the supplemental coverage in 2007. This increase followed the surprise ousting of the Social Democrats from government after the September 2006 general election. Note that from July 2008 on, premia started to be differentiated across UI funds. For 2008 and 2009 we therefore report the average monthly premium among unemployed union members across all UI funds. The Figure also shows the evolution of the take-up of the supplemental UI coverage, measured as the sum of all individuals buying the supplemental coverage divided by the total number of individuals aged 18 to 60 meeting the eligibility criteria for receiving UI benefits.

Figure 12: PRICE VARIATION: UNEMPLOYMENT RISK BY WILLINGNESS-TO-PAY \mathbf{v}



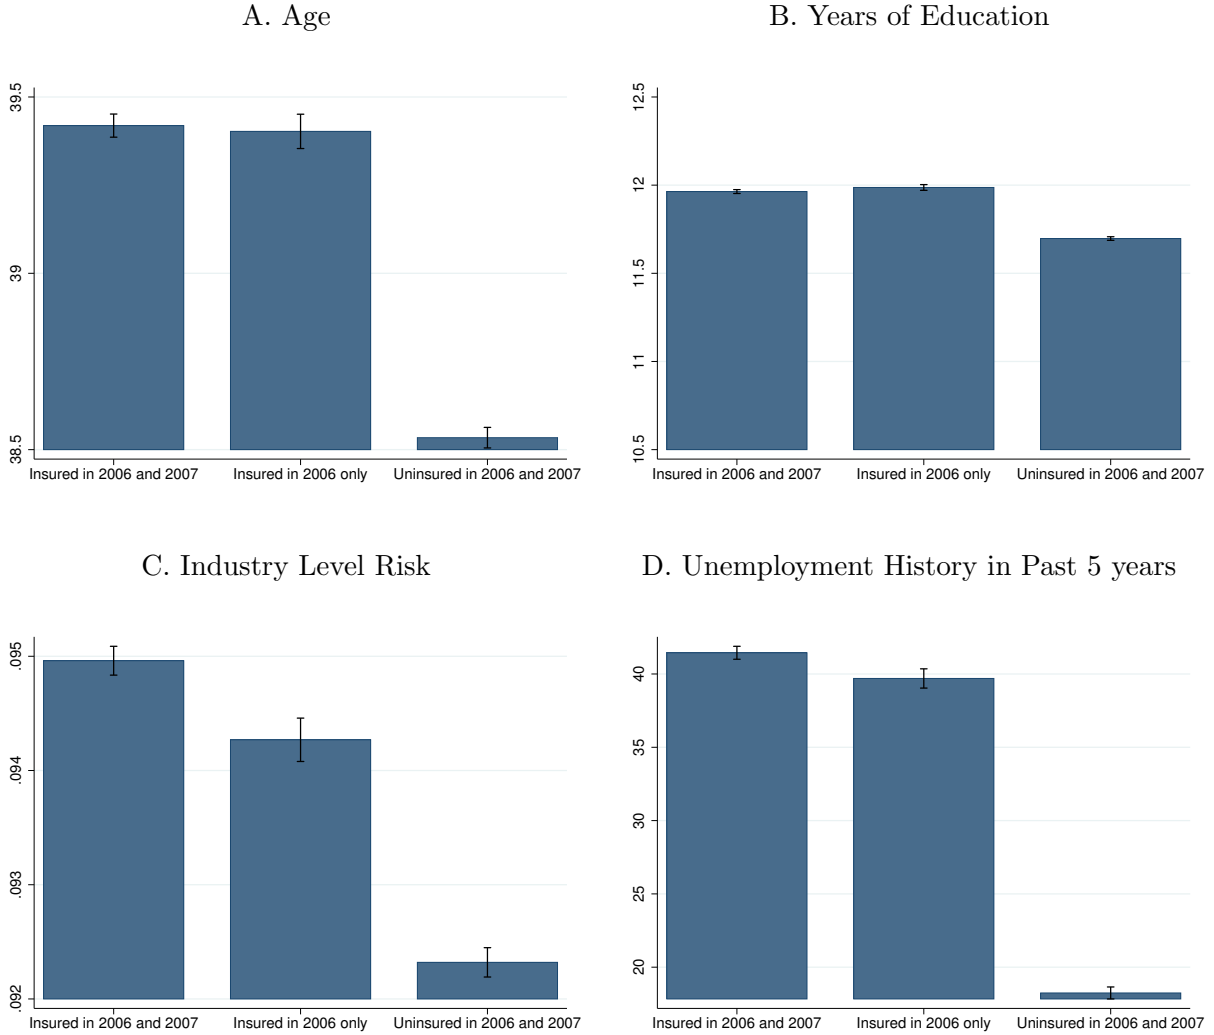
Notes: The Figure uses the 2007 price reform to rank individuals according to their willingness-to-pay for the supplemental coverage \mathbf{v} , and then uses this ranking to correlate willingness-to-pay with various measures of unemployment risk. In each panel, individuals are ranked by decreasing order of their willingness-to-pay. The group on the left (I) are individuals who are insured with the supplemental coverage both in 2006 and 2007 and have the highest level of \mathbf{v} . The middle group corresponds to the marginals (M): individuals who were insured with the supplemental coverage in 2006 but switch out in 2007 when the premium increases. They have a lower level of \mathbf{v} than the always insured (I), but a higher level of \mathbf{v} than the last group on the right (U), of individuals who neither buy the supplemental coverage in 2006, nor in 2007. Using this ranking, we perform direct non-parametric tests for risk-based selection, by correlating willingness-to-pay with various measures of unemployment risk Y . For each risk outcome, we report the average risk outcome of each group controlling for the same vector of characteristics \mathbf{X} affecting contract differentiation, that we use in the positive correlation tests. Panel A reports the average displacement rate in 2008 for each group. Panel B reports the average number of days spent unemployed in 2008 for each group. Panel C and D report for each group the average firm layoff risk and the predicted layoff risk using our risk shifters, as defined in Section 4.2.

Figure 13: PRICE VARIATION: DYNAMICS.




Notes: Because risk realization in 2008 may fail to fully capture the unemployment risk faced by an individual as she is making her coverage choice in 2007, the Figure uses individual risk realizations further into the future, as done previously in the context of the positive correlation test in Figure 2. This Figure reports the correlation between willingness-to-pay in 2007 and realized displacement outcomes in 2008, 2009,.. up to 2012. Panel A reports for each year $k \in 2008, \dots, 2012$, the semi-elasticity $(E_M[Y_k] - E_U[Y_k])/E_U[Y_k]$ of the displacement rate in year k for the marginals M relative to the uninsured U . Panel B, reports the corresponding semi-elasticity $(E_I[Y_k] - E_U[Y_k])/E_U[Y_k]$ for the insured I . Note that the average outcome of each year k and each group G , $E_G[Y_k]$ is conditional on \mathbf{X} , the baseline vector of controls for contract differentiation used in the positive correlation tests. To be precise, we compute the average outcome, fixing the average characteristics \mathbf{X} in each group at the same level as that of the uninsured U .

Figure 14: PRICE VARIATION: UNPRICED OBSERVABLES BY WILLINGNESS-TO-PAY \mathbf{v}



Notes: The Figure uses the 2007 price reform to rank individuals according to their willingness-to-pay for the supplemental coverage \mathbf{v} , and then uses this ranking to correlate willingness-to-pay with various observable characteristics correlated with unemployment risk. In each panel, individuals are ranked by decreasing order of their willingness-to-pay. The group on the left (*I*) are individuals who are insured with the supplemental coverage both in 2006 and 2007 and have the highest level of \mathbf{v} . The middle group corresponds to the marginals (*M*): individuals who were insured with the supplemental coverage in 2006 but switch out in 2007 when the premium increases. They have a lower level of \mathbf{v} than the always insured (*I*), but a higher level of \mathbf{v} than the last group on the right (*U*), of individuals who neither buy the supplemental coverage in 2006, nor in 2007. Using this ranking, we investigate to what extent risk-based selection is driven by selection along unpriced observable characteristics correlated with unemployment risk. For each characteristic, we report the average outcome of each group controlling for the same vector of characteristics \mathbf{X} affecting contract differentiation, that we use in the positive correlation tests. Panel A reports the average age for each group. Panel B reports the years of completed education for each group. Panel C reports the average industry level displacement probability for each group. Panel D reports the total number of days spent unemployed between 2002 and 2006 for each group.


Figure 15: PRICE VARIATION: SELECTION ON PREFERENCES



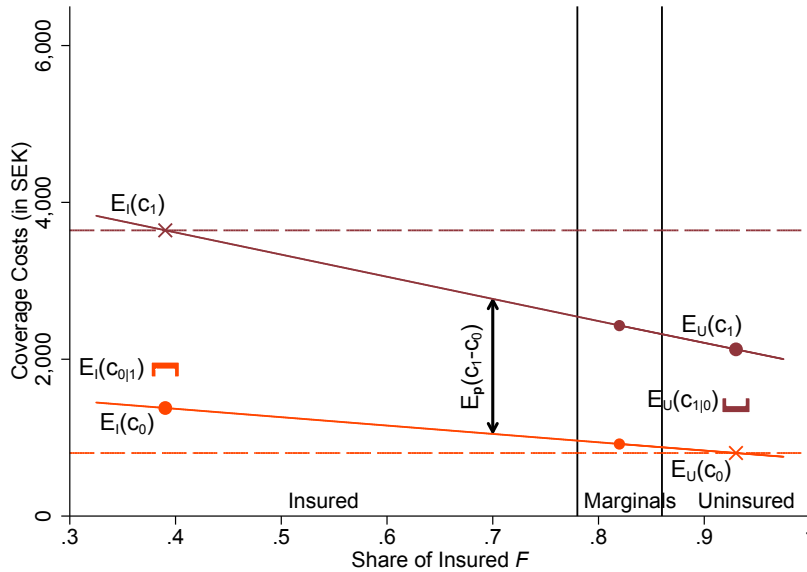
Notes: The Figure uses the 2007 price reform to rank individuals according to their willingness-to-pay for the supplemental coverage v , and uses this ranking to correlate v with proxies for the value of unemployment insurance and risk preferences. In both panels, individuals are ranked by decreasing order of v . The group on the left (I) are individuals who are insured with the supplemental coverage both in 2006 and 2007 and have the highest level of v . The middle group corresponds to the marginals (M): individuals who were insured with the supplemental coverage in 2006 but switch out in 2007 when the premium increases. They have a lower level of v than the always insured (I), but a higher level of v than the last group on the right (U), of individuals who neither buy the supplemental coverage in 2006, nor in 2007. Using this ranking, we correlate in panel A willingness-to-pay with the level of net wealth in 2006. Individuals with higher net wealth have better means to smooth consumption in case of displacement and should have a lower valuation of additional unemployment insurance. We winsorize net wealth and eliminate the bottom and top percentile of the distribution. In panel B, we proxy for risk aversion using the fraction of total net wealth invested in risky assets (stocks). In both panels we report the average outcome of each group controlling for our baseline vector of characteristics X plus a cubic polynomial for age, and a cubic for net wealth in panel B.

Figure 16: WELFARE ANALYSIS

A. Value vs. Cost of Supplemental Coverage



B. Cost Curves of Comprehensive and Basic Plan



Notes: The figure shows the value and cost curves underlying the welfare implementation in Section 6. Panel A compares the value v and cost c of providing supplemental coverage for different workers ranked based on their valuation v . Both coincide for a premium of 1,477 SEK, at which 84 percent of workers would buy insurance. The value is below cost for workers with lower valuation, making it inefficient to mandate them to buy the supplemental coverage. The linear value (or demand) curve is derived based on the share of individuals switching out of the comprehensive plan in response to the premium increase in 2007. The supplemental cost curve is the difference between the cost of providing the comprehensive plan and the basic plan, plotted separately in Panel B. For the comprehensive plan cost curve, we start from the observed $E_I(c_1)$ (marked by 'x') and extrapolate to $E_M(c_1)$ and $E_U(c_1)$ (marked by 'o'). For the basic plan cost curve, we start from the observed $E_U(c_0)$ (again marked by 'x') and extrapolate to $E_M(c_0)$ and $E_I(c_0)$ (again marked by 'o'). The extrapolation uses our estimates from the price variation strategy in Section 5.2 and relies on the equality in (12), as further explained in Appendix B. For the uninsured, we can alternatively calculate $E_U(c_{1|0})$, the counterfactual cost of providing comprehensive coverage, assuming they continue to put in the effort they exert under basic coverage. The difference $E_U(c_{1|0} - c_0)$ provides a lower bound on the cost of providing supplemental coverage $E_U(c)$, indicated by the dashed horizontal line in Panel A.

Table 1: SUMMARY STATISTICS

	Mean	P10	P50	P90
I. Unemployment				
Displacement probability	3.56%	-	-	-
Displacement probability (exc. quits)	3.35%	-	-	-
Unemployment probability	4.71%	-	-	-
Unemployment spell (days)	3.37	0	0	0
Duration of spell (days)	155.99	21	92	328
Fraction receiving layoff notification	.05	-	-	-
Fraction switching firms	.04	-	-	-
II. Union and UI Fund Membership				
Union membership	.75	-	-	-
UI fund membership (V)	.86	-	-	-
III. Demographics				
Age	42.02	28	42	56
Years of education	12.9	11	12	16
Fraction men	.51	-	-	-
Fraction married	.47	-	-	-
IV. Income and Wealth SEK 2003(K)				
Gross earnings	236.7	54.5	228	387.9
Net wealth	490	-178.7	99.7	1148
Bank holdings	50.2	0	0	123

Notes: The Table provides summary statistics for our main sample of interest over the period 2002 to 2006, which comprises the universe of workers in Sweden aged between 25 and 55. Data on unemployment outcomes comes from the Public Employment Service register combined with the IAF register. Unemployment is defined as a spell of non-employment, following an involuntary job loss, and during which an individual has zero earnings, receives unemployment benefits and reports searching for a full time job. We define displacement as an involuntary job loss, due to a layoff or a quit following a ‘valid reason’. Voluntary quits are not included in our measures of displacement and unemployment. See text for details. The probability of displacement is the probability to be displaced in year $t+1$ conditional on working in year t . The unemployment probability is the probability to be unemployed in year $t+1$ unconditional on employment status in year t . The fraction of workers receiving layoff notification comes from the layoff-notification register (VARSEL) and is defined as the fraction of workers that are employed in an establishment emitting a layoff notification in year t . The employer-employee matched data (RAMS) registers all existing labor contracts on a monthly basis. We define a “firm switch” as moving from having a labor contract with firm j (the origin firm) to having a contract with firm k (the destination firm), without any recorded non-employment spell between these two contracts. UI fund membership information comes from tax data for the period 2002 to 2006, during which premia were eligible for a 40% tax credit. The dummy variable V for buying the comprehensive coverage in year t is defined as reporting any positive amount of premia paid in year t . All earnings, income and asset level measures are from wealth and income registers, and are yearly measures in constant k2003SEK. All assets are aggregated at the household level and estimated at their market value. 1SEK2003 \approx 0.11 USD2003

Table 2: POSITIVE CORRELATION TESTS: BIVARIATE PROBITS

	ρ	s.d.	Test $\rho = 0$	
			χ^2	P-Value
Proba. of displacement	.3047	.0030	8842.4	0.00
Proba. of displacement excl. quits	.3056	.0031	8493.9	0.00

Notes: The Table reports positive correlation estimates between insurance and risk using bivariate probit models. We specify both the choice of insurance coverage and the probability of displacement as probit models allowing for correlation ρ between the two error terms ϵ and η . The Table reports estimates of ρ and its standard error. We also report results of formal tests of the null that $\rho = 0$. In the first row, we consider the probability of displacement. In the second row we consider the probability of displacement excluding quits, as some quits may be eligible for UI after a waiting period. See text for details.

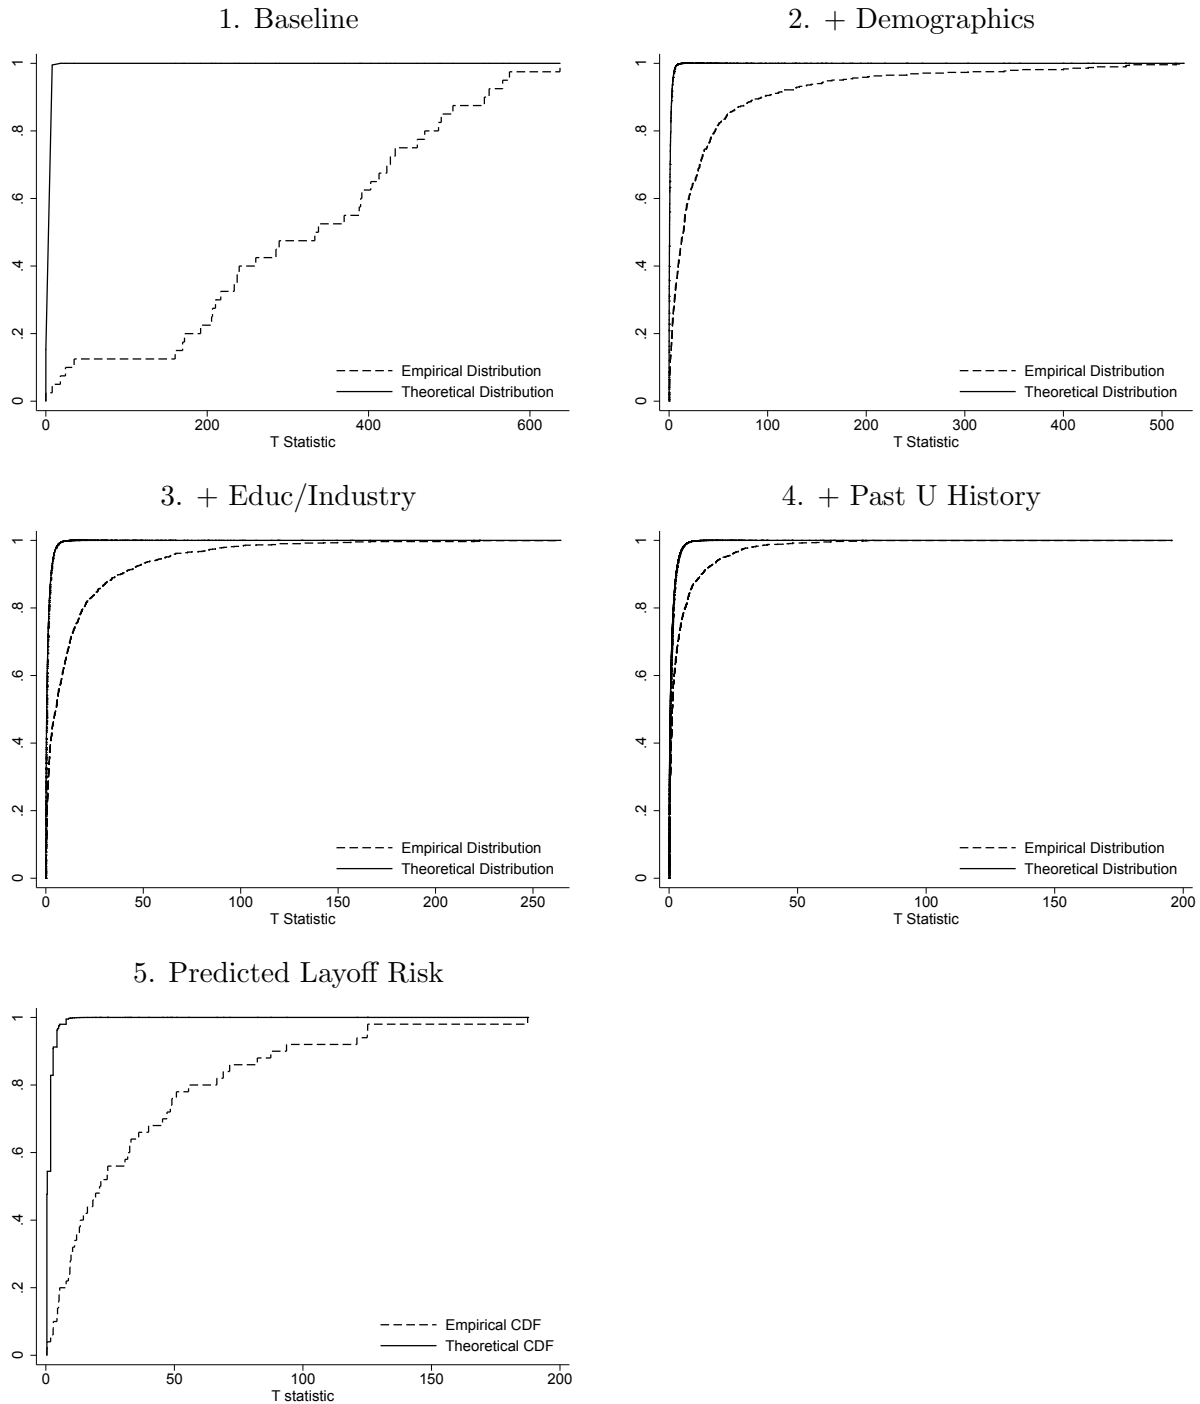
Table 3: POSITIVE CORRELATION TESTS: NON-PARAMETRIC TESTS

	(1)	(2)	(3)	(4)	(5)
Variables included in partitioning the data in cells					
	Baseline	+ Demo- graphics	+ Educ & Industry	+ Past U History	Predicted Layoff Risk
# of cells	40	484	1,124	1,923	50
Average cell size	50,903	3,181	958	415	17,193
Median cell size	35,275	1,270	346	141	17,193
Minimum cell size	14,202	88	6	5	17,193
Fraction of cells too granular	0%	24%	65%	80%	0%
Fraction of rejected cells	98%	74%	53%	28%	88%
Kolmogorov-Smirnov stat.	5.98	15.37	16.20	10.47	6.11
Binomial p-value	0%	0%	0%	0%	0%

Notes: The Table reports results from non-parametric tests of correlation between insurance choices in t and probability of displacement in $t + 1$. The procedure of the test consists in partitioning the data into cells where all observations in a given cell have the same value for the variables in \mathbf{X} . Columns (1) to (4) differ in the control variables included in \mathbf{X} and used to partition the data. The procedure then computes within each cell a Pearson's χ^2 test statistic for independence between UI_t and Y_{t+1} . This test statistic is asymptotically distributed as a $\chi^2(1)$ under the null hypothesis that UI_t and Y_{t+1} are statistically independent (within the cell). The critical values of this statistic for 95% and 99% confidence are 1.36 and 1.63 respectively. The reported Kolmogorov-Smirnov test statistic is scaled by \sqrt{n} where n is the number of cells. When adding a lot of controls to the vector \mathbf{X} , some cells can become too granular to compute the test statistic (division by zero). We therefore also report in the Table the number of cells that are too granular.


Appendix A Additional Graphs And Tables

Figure A.1: POSITIVE CORRELATION TESTS - DISTRIBUTION OF χ^2 TEST STATISTICS FROM ALL CELLS VS THEORETICAL $\chi^2(1)$ DISTRIBUTION - ADDITIONAL CONTROLS



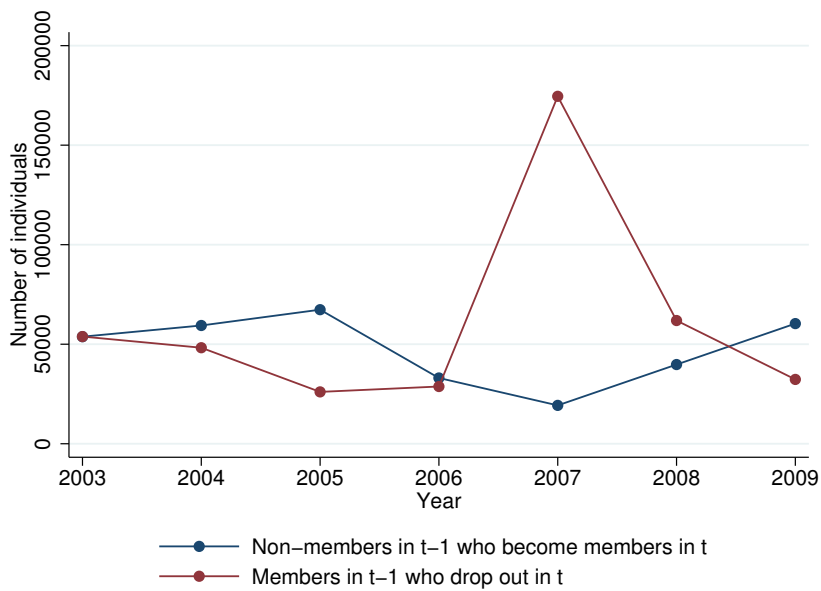
Notes: The Figure displays the empirical distribution of the Pearson's χ^2 test statistics for independence between V (buying the comprehensive coverage) and Y , the probability of layoff in $t + 1$, computed from all the cells where we split individuals in cells corresponding to various unpriced observable characteristics. In panel A, we only use priced characteristics (baseline controls of the positive correlation tests), corresponding to the test implemented in column (1) of Table 3. In panel B, we add controls for demographics (cf. column (2) of Table 3). Panel C and D add education and past unemployment history controls (cf. column (3) and (4) of Table 3), while panel E uses our predicted layoff risk measure based on all observable and risk shifers (cf. column (5) of Table 3). The χ^2 test statistic is asymptotically distributed as a $\chi^2(1)$ under the null hypothesis that V and Y_{t+1} are statistically independent (within the cell). We therefore compare this distribution with a theoretical $\chi^2(1)$ distribution. Taking the largest absolute difference between the theoretical and the empirical distribution gives the Kolmogorov-Smirnov test statistic reported in Table 3.

Figure A.2: SWITCHERS DESIGN: RELATIVE TENURE RANKING AS A FUNCTION OF EVENT TIME



Notes: The Figure focuses on “firm switchers”, i.e. individuals moving from having a labor contract with firm j to having a contract with firm k , without any recorded non-employment spell between these two contracts. We focus on individuals with more than 1 year of tenure in the origin firm. In this Figure we show that switchers experience a variation in their layoff risk coming from underlying variation in their relative tenure ranking. Relative tenure ranking affects displacement probability due to the strict enforcement of the Last-In-First-Out (LIFO) principle in Swedish labor laws. To follow the rules pertaining to the application of LIFO, relative tenure ranking is defined within each establishment times occupation group using the RAMS employer-employee data since 1985. The chart displays estimates of the event study specification (9) using relative tenure ranking as an outcome. The graph shows that relative tenure ranking drops abruptly at the time of the firm switch. Panel A of Figure 6 shows that this drop in tenure ranking translates in a significant increase in displacement risk.

Figure A.3: THE 2007 PRICE REFORM: FLOWS OF INDIVIDUALS SWITCHING IN AND SWITCHING OUT OF THE COMPREHENSIVE COVERAGE OVER TIME



Notes: The Figure reports the evolution of the absolute flows of individuals “switching in” and “switching out” of the comprehensive coverage over time. The sample is restricted to individuals were meeting the work eligibility requirement. Individuals who switch in are individuals who were not buying the comprehensive coverage in year $t - 1$ but are buying in year t (blue curve). Individuals who switch out are individuals who were buying the comprehensive coverage in year $t - 1$ but are no longer buying in year t (red curve). The Figure shows a large and sudden increase in the flow of individuals switching out and a decrease in the flow of individuals switching in, following the large increase in the the premia paid for the supplemental coverage in 2007.

Appendix B Details of Welfare Implementation

This appendix provides further detail on the welfare implementations presented in Section 6.

Cost Curves The cost of providing plan k to a worker of type θ and exerting effort $a_k(\theta)$ equals $c_k = \pi(\theta, a_k(\theta)) \cdot b_k$. Our baseline implementation assumes $c_k = \pi(\theta) \cdot a_k \cdot b_k$ with the exerted effort a_k being independent of the agent’s type θ . We also assume that the agent’s risk $\pi(\theta) = \widehat{\pi(X)} + \epsilon$, with $\widehat{\pi(X)}$ the predicted risk from all observables in our model and ϵ a random term with mean zero, orthogonal to \mathbf{v} . Under these two assumptions, we have that for any two groups $A, B \in \{I, M, U\}$ and any plan k

$$\frac{E_A(c_k)}{E_B(c_k)} = \frac{E_A(\widehat{\pi(X)})}{E_B(\widehat{\pi(X)})}. \quad (12)$$

That is, for both the comprehensive plan and the basic plan, the ratio of the costs for any two groups is equal to the ratio of their predicted risks, estimated using the price variation and shown in panel D of Figure 12.

This allows us to construct the cost curves in panel B of Figure 16. For the comprehensive plan cost curve, we start from the observed $E_I(c_1)$ (marked by ‘x’) and we extrapolate to $E_M(c_1)$ and $E_U(c_1)$ (marked by ‘o’), using the equality in (12). For the basic plan cost curve, we start from the observed $E_U(c_0)$ (again marked by ‘x’) and we extrapolate to $E_M(c_0)$ and $E_I(c_0)$ (again marked by ‘o’), using again equation (12). For the starting points of these extrapolations, we use the average annual cost of providing the respective insurance plans between 2002 and 2006, but scaled down by .20 to account for the taxes paid on the unemployment benefits received.⁴⁶ We locate the observed and imputed average costs for the three groups at the midpoint of the ranges of the corresponding valuation quantiles and then use a piece-wise linear interpolation to construct the cost curves. Note that the resulting cost curves are, in fact, very close to being fully linear. Finally, the difference between the cost curves captures the cost of providing the supplemental coverage, which is compared to the insurance value in the welfare analysis discussed in Section 6. This cost curve is shown in Panel A of Figure 16.

Moral Hazard Our implementation provides an estimate of the moral hazard response to supplemental coverage. For the uninsured, we can calculate $E_U(c_{1|0}) = E_U(\pi(\theta) \cdot a_0 \cdot b_1)$, capturing the cost of providing comprehensive coverage, assuming they continue to put in the effort they exert under basic coverage. This cost estimate is indicated in Panel B of Figure 16 by an upward-facing bracket as it provides a natural bound for the cost of providing comprehensive coverage as it relies on no unemployment response to the change in coverage. The difference with the

⁴⁶Here we follow the assumption in Kolsrud et al. [2017]. Note that the marginal tax rate faced by the workers in our sample in a year that they are observed to be unemployed is on average 22.5%. The average tax rate they pay on their earnings equals 20.7%.

imputed cost of providing comprehensive coverage to the uninsured $E_U(c_1) = E_U(\pi(\theta) \cdot a_1 \cdot b_1)$ provides an estimate of the moral hazard response. (Similarly, for the insured, we can calculate $E_I(c_{0|1}) = E_I(\pi(\theta) \cdot a_1 \cdot b_0)$, indicated by the downward-facing bracket in Figure 16, and compare this to $E_I(c_0) = E_I(\pi(\theta) \cdot a_0 \cdot b_0)$.) Expressed as an elasticity, we find an estimate of $\varepsilon_{\pi,b} = .59$, which falls well in the range of recent estimates in the literature [see Schmieder and Von Wachter [2016]]. These cost estimates also allow for a decomposition of the difference in average cost between providing comprehensive and basic coverage (see equation (10)). Interestingly, our imputation suggests that the cost difference of providing the comprehensive coverage to those who select into comprehensive coverage vs. those who select into basic coverage (i.e., $E_I(c_1) - E_U(c_1)$) is about twice as high as the difference between providing comprehensive coverage to those who behave accordingly vs. those who behave as if they have basic coverage (i.e., $E_U(c_1 - c_{1|0})$).

As stated above, our simple implementation relies on two strong assumptions. First, we require that the risk-based selection is proportional to the selection based on the predicted risk based on observables. We use the predicted risk based on observables rather than the realized risks to avoid attributing the moral hazard response by the workers on comprehensive coverage to risk-based selection. Moreover, our prediction model is not only using our risk shifters, but a comprehensive set of controls, as we do not want to exclude the correlation between risk and value due to other factors affecting both. Second, while we allow for moral hazard to affect the cost differential \mathbf{c} , our implementation assumes that the moral hazard response is orthogonal to the insurance valuation, which excludes so-called selection on moral hazard (see Einav et al. [2013]). We briefly consider how the imputed cost curves would change when relaxing this second assumption.

Selection on moral hazard would imply that workers select into comprehensive coverage anticipating that they will exert less effort to avoid unemployment. Hence, our implementation underestimates the slope of the cost curve for comprehensive coverage and thus overestimates $E_U(c_1)$. To assess the sensitivity of our results to this assumption, we rotate this cost curve downward (keeping $E_I(c_1)$ fixed at its observed value) until the implied $E_U(c_1)$ reaches its lower bound $E_U(c_{1|0})$, which relies on no unemployment response to changes in coverage for the uninsured. As a result, the implied cost curve for the supplemental coverage would be flattened and lower throughout the distribution of workers. Still, the cost continues to exceed the value for the group of uninsured and for the marginal buyers at the pre-reform premium in particular.

Note that it is as plausible that workers select into basic coverage anticipating that they will exert more search efforts. In this case, our implementation would underestimate the slope of the cost curve for basic coverage and thus underestimate $E_I(c_0)$. Again, we can assess the sensitivity of our results by rotating this cost curve upward (keeping $E_U(c_0)$ fixed at its observed value) until the implied $E_I(c_1)$ reaches its upper bound $E_I(c_{0|1})$, implying that the insured's unemployment does not respond to changes in coverage. As a result, the implied cost of providing supplemental coverage is now steeper, but still lower throughout. Again, the cost continues to exceed the value for the group of uninsured. The implied value for the marginal buyers at the pre-reform premium is now just marginally below the costs, which almost fully eliminates the fiscal externality of adverse

selection at the margin.

Appendix C Proofs and Extensions of Welfare Analysis

This appendix provides the proofs of the Propositions in Section 6 and briefly discusses further extensions.

C.1 Proof of Proposition 1

Proof. We consider the welfare impact of the joint price changes $dp_1 = -(1-F)dS$ and $dp_0 = FdS$ and, hence, $d(p_1 - p_0) = -dS$. The share of individuals buying comprehensive insurance is given by $F = \int_{\mathbf{v} \geq \mathbf{p}} dG(\mathbf{v})$. Applying Leibniz rule, we have $\frac{\partial F}{\partial \mathbf{p}} = \frac{\partial F}{\partial p_1} = -\frac{\partial F}{\partial p_0} = -\frac{\partial G(\mathbf{p})}{\partial \mathbf{p}}$. We can re-write the change in cost from providing coverage due to the switchers by

$$\begin{aligned} \frac{\partial}{\partial \mathbf{p}} \left[\int_{\mathbf{v} \geq \mathbf{p}} E(c_1|\mathbf{v}) dG(\mathbf{v}) + \int_{\mathbf{v} < \mathbf{p}} E(c_0|\mathbf{v}) dG(\mathbf{v}) \right] &= E(c_1 - c_0|\mathbf{v} = \mathbf{p}) \frac{\partial G(\mathbf{p})}{\partial \mathbf{p}} \\ &= E_M(\mathbf{c}) \frac{\partial F}{\partial \mathbf{p}}. \end{aligned}$$

Invoking the envelope condition for the individuals at the margin (i.e., $\mathbf{v} - \mathbf{p} = 0$), we can rewrite the welfare impact of the subsidy as

$$\begin{aligned} d\mathcal{W} &= \{F(1-F)E_I(g_1) - F(1-F)E_U(g_0) \\ &\quad - \lambda [(p_1 - p_0) - E_M(c_1 - c_0)] \times \frac{\partial F}{\partial [p_1 - p_0]}\} \times dS, \end{aligned}$$

where $E_I(g_1) = \frac{1}{F} \int_{\mathbf{v} \geq \mathbf{p}} \omega'(v_1 - p_1) dG(\mathbf{v})$ and $E_U(g_0) = \frac{1}{1-F} \int_{\mathbf{v} < \mathbf{p}} \omega'(v_0 - p_0) dG(\mathbf{v})$. Re-arranging terms, we find

$$\frac{d\mathcal{W}}{dS} / F(1-F) = \left\{ E_I(g_1) - E_U(g_0) - \lambda \times \frac{(p_1 - p_0) - E_M(c_1 - c_0)}{p_1 - p_0} \frac{\frac{\partial F}{\partial [p_1 - p_0]} [p_1 - p_0]}{F \times (1-F)} \right\}.$$

Assuming an interior solution, the subsidy can be optimal only if $d\mathcal{W}/dS = 0$. This leads to the optimality condition in the Proposition for $\varepsilon_{\frac{1-F}{F}, \mathbf{p}} = \frac{\partial(1-F)/F}{\partial [p_1 - p_0]} \frac{p_1 - p_0}{(1-F)/F}$. \square

C.2 Proof of Proposition 2

Proof. We consider the welfare impact of an increase in b_0 , for given prices and coverage b_1 . The impact of a change in b_0 on the government's budget depends both on the change in selection into both plans and the direct effect from increasing the coverage,

$$\frac{\partial}{\partial b_0} \left[\int_{\mathbf{v} \geq \mathbf{p}} [p_1 - E(c_1|\mathbf{v})] dG(\mathbf{v}) + \int_{\mathbf{v} < \mathbf{p}} [p_0 - E(c_0|\mathbf{v})] dG(\mathbf{v}) \right]$$

$$\begin{aligned}
&= E \left([\mathbf{p} - \mathbf{c}] \times \frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p} \right) g(\mathbf{p}) - \int_{\mathbf{v} < \mathbf{p}} E \left(\frac{\partial c_0}{\partial b_0} | \mathbf{v} \right) dG(\mathbf{v}) \\
&= \left[\mathbf{p} - E \left(\mathbf{c} \times \frac{\frac{\partial \mathbf{v}}{\partial b_0}}{E \left(\frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p} \right)} | \mathbf{p} \right) \right] \frac{\partial F}{\partial b_0} - (1 - F) \frac{\partial E_U c_0}{\partial b_0} \\
&\equiv [\mathbf{p} - E_{M(b_0)}(\mathbf{c})] \frac{\partial F}{\partial b_0} - (1 - F) \frac{\partial E_U c_0}{\partial b_0}
\end{aligned}$$

By analogy to the subsidy change, we decompose the change in cost from providing coverage due to the change in selection as the demand effect $\frac{\partial F}{\partial b_0}$ multiplied by the fiscal externality $\mathbf{p} - E_{M(b_0)}(\mathbf{c})$, caused by the switching of individuals who respond to the coverage change. This fiscal externality differs from the fiscal externality of the subsidy as different individuals respond to a change in coverage depending on their marginal value of basic coverage $\frac{\partial \mathbf{v}}{\partial b_0}$, explaining the weights put on the costs of the different marginals with valuation equal to \mathbf{p} .⁴⁷

Invoking now the envelope condition for the individuals at the margin (i.e., $\mathbf{v} - \mathbf{p} = 0$), we find

$$d\mathcal{W} = (1 - F) \frac{\partial E_U \omega_0}{\partial b_0} - \lambda (1 - F) \frac{\partial E_U c_0}{\partial b_0} + \lambda [\mathbf{p} - E_{M(b_0)}(\mathbf{c})] \frac{\partial F}{\partial b_0},$$

where

$$\begin{aligned}
\frac{\partial E_U \omega_0}{\partial b_0} &= \frac{1}{1 - F} \int_{\mathbf{v} < \mathbf{p}} E \left(\omega'(v_0 - p_0) \frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{v} \right) dG(\mathbf{v}) \\
&= \frac{1}{1 - F} \int_{\mathbf{v} < \mathbf{p}} E \left(\omega'(v_0 - p_0) \pi u'(b_0) | \mathbf{v} \right) dG(\mathbf{v}) \\
&= E_U(\pi) E_U \left(g_0 \times \frac{\pi}{E_U(\pi)} \times u'(b_0) \right).
\end{aligned}$$

Using $\frac{\partial \mathbf{v}}{\partial b_0} = \pi \partial u'(b_0)$ and $E_U(c_0) = E_U(\pi) b_0$ in the static unemployment model with unemployment probability π , we can re-write

$$\begin{aligned}
\frac{\partial E_U \omega_0}{\partial b_0} &= E_U(\pi) E_U \left(\frac{\pi}{E_U(\pi)} g_0 \times u'(b_0) \right), \\
\frac{\partial E_U c_0}{\partial b_0} &= E_U(\pi) + \frac{\partial E_U(\pi)}{\partial b_0} b_0 = E_U(\pi) [1 + \varepsilon_{E_U(\pi), b_0}].
\end{aligned}$$

⁴⁷Following the arguments in Veiga and Weyl [2016] and Handel et al. [2015], we provide a formal derivation in the technical appendix showing that

$$\frac{\partial}{\partial b_0} \int_{\mathbf{v}(b_0) \geq \mathbf{p}} E(z | \mathbf{v}(b_0)) dG(\mathbf{v}(b_0)) = E \left(z \frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p} \right) g(\mathbf{p}).$$

Hence, the welfare impact becomes

$$\frac{d\mathcal{W}}{\lambda(1-F)E_U(\pi)} = \frac{E_U\left(\frac{\pi}{E_U(\pi)}g_0 \times u'(b_0)\right) - \lambda}{\lambda} - \varepsilon_{E_U(\pi),b_0} - \left[1 - \frac{E_{M(b_0)}(\mathbf{c})}{\mathbf{p}}\right] \frac{\mathbf{p}}{E_U(c_0)} \varepsilon_{1-F,b_0}.$$

At an (interior) optimum, we need $d\mathcal{W} = 0$ and thus the Proposition follows. \square

C.3 Baily-Chetty Representation

Our optimal policy characterization considers one-dimensional changes in the policy and compares the impact on agents' welfare relative to the impact on the government's budget. Alternatively, we can consider a joint change in policy instruments that keeps the budget fixed. In particular, consider the budgetary impact of a joint change in b_0 and p_0 ,

$$dB = \left[(1-F) + [\mathbf{p} - E_M(\mathbf{c})] \frac{\partial F}{\partial p_0}\right] dp_0 + \left[-(1-F) \frac{\partial E_U c_0}{\partial b_0} + [\mathbf{p} - E_{M(b_0)}(\mathbf{c})] \frac{\partial F}{\partial b_0}\right] db_0.$$

A budget-balanced change requires

$$\begin{aligned} \frac{dp_0}{db_0} &= E_U(\pi) [1 + \varepsilon_{E_U(\pi),b_0}] - [\mathbf{p} - E_{M(b_0)}(\mathbf{c})] \frac{\partial F}{\partial b_0} \frac{1}{1-F} - [\mathbf{p} - E(\mathbf{c})] \frac{\partial F}{\partial p_0} \frac{1}{1-F} \frac{dp_0}{db_0} \\ &\equiv E_U(\pi) \left[[1 + \varepsilon_{E_U(\pi),b_0}] + [\mathbf{p} - E_M(\mathbf{c})] \frac{\varepsilon_{1-F,[b_0,p_0]}}{E_U(c_0)} - \frac{\text{cov}\left(\mathbf{c}, \frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p}\right)}{E\left(\frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p}\right)} \frac{\varepsilon_{1-F,b_0}}{E_U(c_0)} \right], \end{aligned}$$

where the $\varepsilon_{1-F,[b_0,p_0]}$ corresponds to the demand elasticity with respect to an actuarially fair increase in basic coverage (accounting for the changed selection). Intuitively, in an adversely selected market, the increase in basic coverage is relatively cheap, given the low risk types buying basic coverage, and thus likely to induce individuals who buy comprehensive coverage to switch to basic coverage.

The welfare impact equals

$$\frac{d\mathcal{W}/db_0}{(1-F)} = \frac{\partial E_U(\omega)}{\partial b_0} - \frac{\partial E_U(\omega)}{\partial p_0} \frac{dp_0}{db_0}.$$

Now to move closer to the original Baily-Chetty formula, we use $\omega = \pi u(b-p) + (1-\pi)u(w-p)$ for an agent's contribution to (utilitarian) social welfare, and assume homogeneity in either risks

or in marginal utilities, so we find that at the optimum,

$$\frac{E_U(u'(b_0)) - E_U(u'(w - p_0))}{E_U(\pi u'(b_0) + (1 - \pi)u'(w - p_0)) / (1 - E_U(\pi))} = \text{varepsilon}_{E_U(\pi), b_0} + [\mathbf{p} - E_M(\mathbf{c})] \frac{\varepsilon_{1-F, [b_0, p_0]}}{E_U(c_0)} - \frac{\text{cov}\left(\mathbf{c}, \frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p}\right)}{E\left(\frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p}\right)} \frac{\varepsilon_{1-F, b_0}}{E_U(c_0)},$$

where the left-hand side is determined by the relative difference in marginal utilities of consumption when unemployed vs. employed, just like in the original Baily-Chetty formula.

C.4 Comprehensive Coverage and Private Markets

In the main text, we only provided the characterization of the basic coverage level, because of its immediate relation to a minimum mandate. We can derive a characterization of the socially optimal level of comprehensive coverage in an analogous manner:

Proposition 3. *The comprehensive coverage level b_1 is optimal, for given subsidy S and basic coverage b_0 , only if*

$$\frac{E_I(g_1 \times u'(b_1)) - \lambda}{\lambda} - \varepsilon_{E_I(\pi), b_1} = - \left[1 - \frac{E_{M(b_1)}(\mathbf{c})}{\mathbf{p}} \right] \times \frac{\mathbf{p}}{E_I(c_1)} \times \varepsilon_{F, b_1}.$$

Note that here the fiscal externality enters the optimality condition with opposite sign, since an increase in comprehensive coverage attracts more individuals to this plan. To provide intuition, it is, however, useful to consider again an actuarially fair increase in b_1 (funded by those buying comprehensive coverage). By analogy, we find

$$\frac{E_I(u'(b_1)) - E_I(u'(w - p_1))}{E_I(\pi u'(b_1) + (1 - \pi)u'(w - p_1)) / (1 - E_I(\pi))} = \varepsilon_{E_I(\pi), b_1} - [\mathbf{p} - E_M(\mathbf{c})] \frac{\varepsilon_{F, [b_1, p_1]}}{E_I(c_1)} + \frac{\text{cov}\left(\mathbf{c}, \frac{\partial \mathbf{v}}{\partial b_1} | \mathbf{p}\right)}{E\left(\frac{\partial \mathbf{v}}{\partial b_1} | \mathbf{p}\right)} \frac{\varepsilon_{F, b_1}}{E_I(c_1)}.$$

Here, again, the fiscal externality enters with the opposite sign. However, in an adversely selected market, an increase in comprehensive coverage will be costly given the high risk of buying it and will likely discourage individuals at the margin from buying it. This suggests that both for setting the basic coverage and the comprehensive coverage, an increase in the coverage level in an actuarially fair way is likely to make the supplemental market more adversely selected.

While the equilibrium characterization for a market with profit-maximizing insurance is beyond the scope of this paper, we can see how profit-maximizing firms will face a similar trade-off when deciding how to set coverage and price. Two important differences arise compared to the social planner's trade-off. First, private insurers may care only about the utility increase for individuals at the margin of buying their plan. This would change the left-hand side of the above characterization.

Second, private insurers are likely to only care about how the changed selection affects their own profits, not how it affects the selection out of other plans. This would change the right-hand side of the above characterization. In particular, private insurers deciding how to set b_1 when the government provides minimum coverage b_0 , will account for $[p_1 - E_{M(b_1)}c_1]$ rather than $[\mathbf{p} - E_{M(b_1)}\mathbf{c}]$. In an adversely selected market with average-cost pricing, this would reduce the impact of the selection effect. As selection effects tend to put downward pressure on the optimal coverage levels, private insurers may end up over-insuring as a consequence.

C.5 Sorting Effect

The fiscal externality in both Propositions 1 and 2 depends on how many individuals change in response to the policy (as captured by the demand elasticity) and the cost characteristics of those who switch. Here we develop formally the argument why the cost characteristics of the switchers in response to a change in coverage is different than for a change in price under multi-dimensional heterogeneity.⁴⁸ In particular, we show that

$$\frac{\partial}{\partial b_0} \left[\int_{\mathbf{v} \geq \mathbf{p}} E(c_1 | \mathbf{v}) dG(\mathbf{v}) \right] = E \left(c_1 \frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p} \right) g(\mathbf{p}).$$

An analogue derivation applies for the cost of providing basic coverage.

We use notation $\mathbf{v}' \equiv \frac{\partial \mathbf{v}}{\partial b_0}$. Using iterated expectations, we can re-write

$$\begin{aligned} \frac{\partial}{\partial b_0} \left[\int_{\mathbf{v} \geq \mathbf{p}} E(c_1 | \mathbf{v}) dG(\mathbf{v}) \right] &= \frac{\partial}{\partial b_0} \left[\int_{\mathbf{v} \geq \mathbf{p}} \int E(c_1 | \mathbf{v}, \mathbf{v}') f(\mathbf{v}' | \mathbf{v}) g(\mathbf{v}) d\mathbf{v}' d\mathbf{v} \right] \\ &= \int \frac{\partial}{\partial b_0} \left[\int_{\mathbf{v}_\varepsilon \geq \mathbf{p} - \mathbf{v}'[b_0 - b_\varepsilon]} E(c_1 | \mathbf{v}_\varepsilon, \mathbf{v}') g^\varepsilon(\mathbf{v}_\varepsilon | \mathbf{v}') d\mathbf{v}_\varepsilon \right] f(\mathbf{v}') d\mathbf{v}' \end{aligned}$$

The second equality follows from using $f(\mathbf{v}' | \mathbf{v}) g(\mathbf{v}) = g(\mathbf{v} | \mathbf{v}') f(\mathbf{v}')$, approximating $\mathbf{v}(b_0) \cong \mathbf{v}(b_\varepsilon) + \mathbf{v}' \times [b_0 - b_\varepsilon]$, and substituting the variable in the integral $\mathbf{v}(b_0) (\equiv \mathbf{v})$ by $\mathbf{v}(b_\varepsilon) (\equiv \mathbf{v}_\varepsilon)$, where $d\mathbf{v} = d\mathbf{v}_\varepsilon$, conditional on \mathbf{v}' . We can now apply Leibniz rule and find after re-substituting,

$$\begin{aligned} \frac{\partial}{\partial b_0} \left[\int_{\mathbf{v} \geq \mathbf{p}} E(c_1 | \mathbf{v}) dG(\mathbf{v}) \right] &= \int [E(c_1 \mathbf{v}' | \mathbf{p}, \mathbf{v}') f(\mathbf{v}' | \mathbf{p}) d\mathbf{v}'] g(\mathbf{p}) \\ &= E \left(c_1 \frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p} \right) g(\mathbf{p}). \\ &= E(c_1 | \mathbf{p}) E \left(\frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p} \right) g(\mathbf{p}) + cov \left(c_1, \frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p} \right) \end{aligned}$$

Note also that the effect on the share of individuals buying comprehensive coverage equals

$$\frac{\partial}{\partial b_0} \left[\int_{\mathbf{v} \geq \mathbf{p}} dG(\mathbf{v}) \right] = E \left(\frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p} \right) g(\mathbf{p}) \equiv \frac{\partial F}{\partial b_0}.$$

⁴⁸The derivation follows Handel et al. [2015], which is a slight variation of the approach in Veiga and Weyl [2016].

Hence,

$$\frac{\partial}{\partial b_0} \left[\int_{\mathbf{v} \geq \mathbf{p}} E(c_1 | \mathbf{v}) dG(\mathbf{v}) \right] = E \left(c_1 \frac{\frac{\partial \mathbf{v}}{\partial b_0}}{E \left(\frac{\partial \mathbf{v}}{\partial b_0} | \mathbf{p} \right)} | \mathbf{p} \right) \frac{\partial F}{\partial b_0}$$

as used in the proof of Proposition 2.