American Economic Review 2017, 107(7): 1778–1823
https://doi.org/10.1257/aer.20151655

Knowledge of Future Job Loss and Implications for Unemployment Insurance†

By Nathaniel Hendren*

This paper studies the implications of individuals' knowledge of future job loss for the existence of an unemployment insurance (UI) market. Learning about job loss leads to consumption decreases and spousal labor supply increases. This suggests existing willingness to pay estimates for UI understate its value. But it yields new estimation methodologies that account for and exploit responses to learning about future job loss. Although the new willingness to pay estimates exceed previous estimates, I estimate much larger frictions imposed by private information. This suggests privately traded UI policies would be too adversely selected to be profitable, at any price. (*JEL* D82, D83, G22, J22, J64, J65)

The risk of job loss is one of the most salient risks faced by working-age individuals. Job loss leads to drops in consumption and significant welfare losses.[1] Millions of people hold life insurance, health insurance, liability insurance, and many other insurance policies. But there does not exist a thriving private market for insurance against unemployment or job loss. The government provides some unemployment insurance, and individuals may have help from family, savings, or severance if they lose their job. But why don't insurance companies sell policies to provide additional insurance against these risks?[2]

* Department of Economics, Harvard University, Littauer Center, 1805 Cambridge Street, Cambridge, MA 02138 (email: nhendren@fas.harvard.edu). I am very grateful to Daron Acemoglu, Raj Chetty, Amy Finkelstein, Jon Gruber, Luigi Pistaferri, Rob Townsend, and a long list of seminar participants for their comments. I also thank Trevor Bakker, Augustin Bergeron, Lizi Chen, Michael Droste, Robert Fluegge, Frina Lin, Jeremy Majerovitz, Jimmy Narang, Nina Roussille, and Priyanka Shende for helpful research assistance. Early versions of Sections IV and VI of this paper appeared in the second chapter in my MIT PhD thesis and also circulated under the title "Private Information and Unemployment Insurance." Support from NSF Graduate Research Fellowship and the NBER Health and Aging Fellowship, under the National Institute of Aging grant number T32-AG000186 is gratefully acknowledged.

† Go to https://doi.org/10.1257/aer.20151655 to visit the article page for additional materials and author disclosure statement.

[1] See Gruber (1997); Browning and Crossley (2001); Aguiar and Hurst (2005); Chetty (2008); and Blundell, Pistaferri, and Saporta-Ekstein (2012).

[2] Two companies have attempted to sell such policies in the past 20 years. PayCheck Guardian attempted to sell policies from 2008, but stopped selling in 2009 with industry consultants arguing "The potential set of policyholders are selecting against the insurance company, because they know their situation better than an insurance company might." (*New York Times*, http://www.nytimes.com/2009/08/08/your-money/08money.html) More recently, IncomeAssure has partnered with states to offer top-up insurance up to a 50 percent replacement rate for workers in some industries and occupations (https://www.incomeassure.com). Back-of-the-envelope calculations suggest that their markups exceed 500 percent over actuarially fair prices, and it has been criticized for shrouding the true price by not saliently noting that the government provides the baseline 30–40 percent replacement rate (e.g., http://www.mlive.com/jobs/index.ssf/2011/08/get_out_your_calculator_before_you_buy_p.html#).

This paper provides empirical evidence that unemployment or job loss insurance would be too adversely selected to deliver a positive profit, at any price. This market failure provides a potential rationale for government intervention that requires workers to pay into a government UI system.

I begin by documenting several pieces of evidence that individuals have knowledge about their future job loss that could be used to adversely select an insurance contract. First, using data from the Health and Retirement Study (HRS), I show that an individual's elicited probability of losing their job is predictive of subsequent job loss. This remains true even conditional on a wide range of observable characteristics an insurer might use to reduce the information asymmetry. Second, spouses are more likely to enter the labor market when individuals learn they might lose their job. Finally, while it has been shown that consumption expenditure drops 7–10 percent upon the onset of unemployment (e.g., Gruber 1997), I use data from the Panel Study of Income Dynamics (PSID) to show that food expenditure also declines by 2.7 percent in the 1–2 years prior to unemployment. This occurs on a sample who remain employed in these pre-periods with no pretrend in income, and therefore indicates forward-looking savings behavior anticipating potential job loss. Taken together, these patterns suggest individuals have knowledge about their future job loss.[3] Moreover, the labor supply and ex ante consumption responses suggest individuals would prefer to have more financial resources in the event of job loss. This implies they would have a demand for unemployment or job loss insurance and potentially use their knowledge to (adversely) select insurance contracts.

Given these patterns, the primary task of the paper is to ask: "Can this knowledge of future job loss explain why there is not a thriving private market for unemployment insurance?" To assess this, I consider a general model in which individuals face a privately known risk of losing their job and which characterizes when insurance companies can profitably sell insurance. The model extends a similar setup in Hendren (2013) by allowing for both moral hazard and also dynamic consumption and labor supply responses. I show that a market can exist only if the markup that individuals are willing to pay for insurance exceeds the cost imposed by worse risks adversely selecting their contract. The latter cost is measured as the *pooled price ratio* defined in Hendren (2013).[4] Therefore, I analyze the implications of the reduced-form empirical patterns for both the (i) markup that individuals are willing to pay for UI and (ii) the pooled price ratio.

A large literature has attempted to estimate the markup that individuals are willing to pay for UI. The most common approach estimates the impact of unemployment or job loss on a yearly first difference of consumption and then scales this impact by a coefficient of relative risk aversion (Baily 1978; Gruber 1997). However, job loss affects consumption not only at its onset but also in the years

[3] As discussed in the related literature section, these patterns are consistent with previous work documenting knowledge of future job loss (e.g., Stephens 2004). The primary distinction relative to this work is that I isolate the private component of individuals' knowledge after controlling for observables that insurers would use to price insurance, as this is what is relevant for generating adverse selection.

[4] Along the way, the model illustrates that moral hazard does not provide a singular explanation for the absence of an insurance market—a point initially recognized by Shavell (1979). This is because the first $1 of insurance provides first-order welfare gains, but the behavioral response to a small amount of insurance imposes a second-order impact on the cost of insurance.

prior. This suggests the yearly first difference estimate understates the causal effect. To correct for this bias, I develop a two-sample instrumental variables strategy that scales the first difference estimate by the amount of information revealed over the time period encompassing the first difference (i.e., the last year before job loss). Regressing the subjective probability elicitations on the job loss indicator suggests 80 percent of the information about job loss is revealed in the last year prior to unemployment. I show that one can divide the 7–10 percent first difference estimate by this *first stage* of 0.8 to arrive at the average causal effect of 8–13 percent. Scaling by a coefficient of relative risk aversion (e.g., 2) yields the markup individuals are willing to pay for UI (e.g., 16–26 percent).

One potential concern with using the causal effect of the job loss event on consumption to measure willingness to pay is that it requires utility over consumption to be state independent. This could be violated for many reasons. On the one hand, the unemployed may have more time to substitute home production or shop for lower prices (Aguiar and Hurst 2005); on the other hand, unemployment may bring additional job search costs that yield a high value of additional financial resources. To deal with these potential concerns, I provide a new method for valuing UI that exploits the ex ante behavioral responses to learning about future job loss.

Using the changes in consumption prior to unemployment, I provide conditions under which one can scale the 2.7 percent consumption change in the 1–2 years prior to unemployment by the amount of information about unemployment in year t that is revealed between year $t-2$ and $t-1$. Using subjective probability elicitations, I estimate that 10 percent of information about unemployment in year t is revealed between year $t-2$ and $t-1$. This suggests that individuals are willing to pay a markup for UI of 30 percent multiplied by the coefficient of relative risk aversion (i.e., a 60 percent markup for a coefficient of relative risk aversion of 2). Analogously, I show that scaling the ex ante spousal labor supply responses by the elasticity of labor force participation with respect to a change in wages reveals the markup that individuals would be willing to pay for UI. For a semi-elasticity of 0.5, this suggests individuals are willing to pay around a 60 percent markup for UI. Because these approaches exploit behavioral responses while individuals are still employed, they do not require state independence of the utility function and allow for preferences over consumption to depend on leisure or unemployment status.

Are these willingness to pay estimates sufficient to overcome the hurdles imposed by adverse selection? To answer this question, I build on the strategies developed in Hendren (2013) to estimate the pooled price ratio. I use the information contained in the subjective probability elicitations to provide nonparametric lower bounds and semiparametric point estimates on the pooled price ratio. This yields lower bounds that suggest individuals would have to be willing to pay a 70 percent markup and point estimates above 300 percent in order for insurers to profitably sell insurance. The pooled price ratio remains large across varying assumptions about the observables that insurers would use to price the UI policies and also is quite persistent across subgroups. Since these estimates generally exceed the estimates for the willingness to pay for UI, the results suggest a private UI policies would be too heavily adversely selected to deliver a positive profit, at any price.

Related Literature.—This paper is related to a large literature studying the degree to which individuals are insured against unemployment and income shocks[5] and their behavioral response to these adverse events.[6] The methods used in this paper also relate to a previous literature using subjective expectation data (Pistaferri 2001; Manski 2004). In particular, this paper is most closely related to the work of Stephens (2004), who illustrates that subjective probability elicitations in the HRS are predictive about future unemployment status, and to Stephens (2001, 2002), who finds evidence of ex ante consumption drops and spousal labor supply increases in the United States using the PSID.

Relative to previous literature, the primary contribution of this paper is to study the implications of people's knowledge of future job loss for the workings of a private UI market.[7] By using subjective probability elicitations to identify the *supply-side* frictions imposed by private information, the paper utilizes many of the tools developed in Hendren (2013). On the *demand side,* I develop a new methodology to measure willingness to pay for UI when individuals have knowledge of future job loss. Ex ante behavioral responses suggest that methodologies using the response of first differences of consumption tend to understate the value of UI (e.g., Gruber 1997). I provide a correction to this method in Section IVA. In addition, these ex ante behavioral responses open a potentially fruitful new pathway to valuing insurance by exploiting the ex ante responses to measure the value of insurance in Section IVB that avoids requirements of state independence of the utility function employed in previous literature (e.g., Baily 1978; Gruber 1997). In particular, the drop in food expenditure and increased spousal labor supply that occur when people learn they might lose their job suggest job loss is significantly underinsured in the United States, regardless of whether the utility function is state-dependent or individuals have more time to cook or search for lower prices when unemployed.[8]

Finally, this paper also contributes to the growing literature documenting the impact of private information on the workings of insurance markets and the microfoundations for underinsurance. Previous literature often tests for private information by asking whether existing insurance contracts are adversely selected (Chiappori and Salanié 2000; Finkelstein and Poterba 2004). My results suggest this literature has perhaps suffered from a *lamp-post* problem, as forebode in Einav, Finkelstein, and Levin (2010): if private information prevents the existence of entire

[5] In the UI context, see Baily (1978); Acemoglu and Shimer (1999, 2000); Chetty (2006, 2008); Shimer and Werning (2007, 2008); Blundell, Pistaferri, and Preston (2008); and Landais, Michaillat, and Saez (2010). See also Bach (1998) for one of the only papers documenting adverse selection in a UI context in which mortgage insurance companies provide mortgage payments in the event of job loss.

[6] For example, see Guiso, Jappelli, and Terlizzese (1992); Dynan (1993); Hubbard, Skinner, and Zeldes (1994); Carroll and Samwick (1997, 1998); Carroll, Dynan, and Krane (2003); Lusardi (1997, 1998); Engen and Gruber (2001); Guariglia and Kim (2004); Bloemen and Stancanelli (2005); Barceló and Villanueva (2010); Campos and Reggio (2015); and Brown and Matsa (2016). Most closely, Basten, Fagereng, and Telle (2012) and Gallen (2013) find evidence of ex ante savings increases in response to future unemployment in Norway and Denmark.

[7] Along the way, the analysis clarifies the empirical estimands required to answer this question. For example, one needs to know whether the elicitations are predictive of job loss conditional on public information that insurers would use to price insurance. Further, when studying ex ante consumption responses, it is important to restrict to a sample that remains employed in these periods to estimate the demand that would be held by potential insurance customers. To the best of my knowledge, this paper is the first to document that food expenditure drops on the subsample of those who remain employed in the preperiod and experience no drop in consumption.

[8] See Aguiar and Hurst (2005) for a discussion and evidence that expenditure measurements may misstate the impact of retirement and job loss on consumption.

markets, it is difficult to identify its impact by looking for the adverse selection of existing contracts. Combined with the evidence in Hendren (2013) that private information prevents the existence of health-related insurance markets for those with preexisting conditions, the results suggest a broader pattern: the frictions imposed by private information form the boundary to the existence of insurance markets.

The rest of this paper proceeds as follows. Section I discusses the data used in the analysis. Section II presents a series of motivating statistics establishing the presence of private information that individuals would use to (adversely) select insurance. Section III places these patterns in the context of a general model of unemployment risk. A private market will not be profitable if the markup individuals would be willing to pay for UI is less than the pooled price ratio defined in Hendren (2013). Section IV considers the implications of the patterns in Section II for measuring individuals' willingness to pay for UI, and Section V considers the implications for measuring the pooled price ratio. Section VI discusses robustness and alternative theories of market nonexistence. Section VII concludes.

I. Data

I use data from two panel surveys: the Health and Retirement Survey (HRS) and the Panel Study of Income Dynamics (PSID). The HRS provides measures of subjective probability elicitations about future unemployment and measures of spousal labor supply. The PSID does not contain subjective probability elicitations, but includes a panel of information on food expenditures.

A. *HRS*

The HRS sample draws from waves of the Health and Retirement Study spanning the years 1992–2013. The HRS samples individuals over 50 years old and their spouses (included regardless of age). The baseline sample includes everyone under 65 years old in the survey who holds a job in the current survey wave and is not self-employed or in the military.

Subjective Probability Elicitations.—The HRS contains a battery of subjective expectation information about future adverse events. In particular, the survey asks: What is the percent chance (0–100) that you will lose your job in the next 12 months? Figure 1 presents the histogram of these elicitations. As noted in previous literature (Gan, Hurd, and McFadden 2005), the responses concentrate on focal point values, especially 0. Taken literally, these responses of 0 (or 100 percent) imply individuals would be willing to bet an infinite amount of money against the chance of losing (or keeping) ones' job. As a result, at no point will these elicitations be used as true measures of individuals' beliefs. Instead, I follow the approach of Hendren (2013) by treating these elicitations as noisy and potentially biased measures of true beliefs about losing one's job. To maintain this distinction, let Z denote the responses to these survey questions, and let P denote the subjective probability held by the individual that governs their willingness to pay for lotteries and financial contracts, as in Savage (1954). These *true* beliefs, P, are related to an individual's willingness to

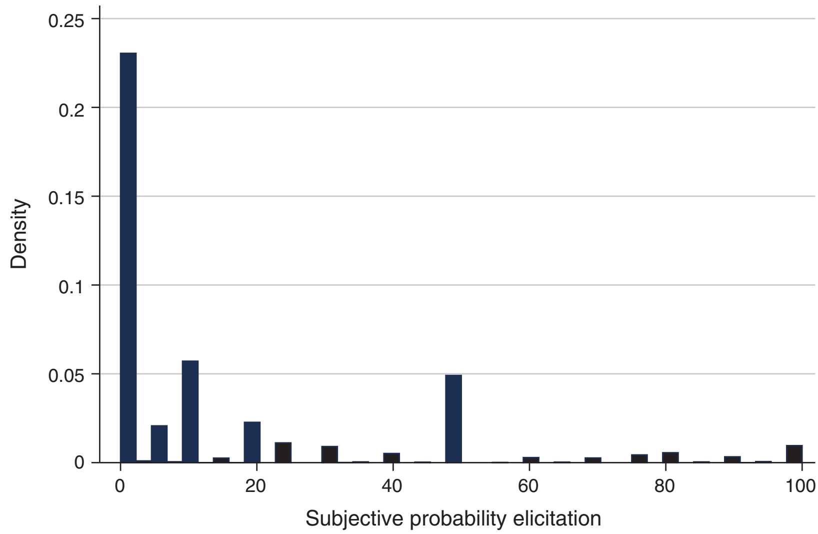

FIGURE 1. HISTOGRAM OF SUBJECTIVE PROBABILITY ELICITATIONS

Notes: This figure presents a histogram of responses to the question "What is the percent chance (0–100) that you will lose your job in the next 12 months?" The figure reports the histogram of responses for the baseline sample (corresponding to column 1 in Table 1). As noted in previous literature, responses tend to concentrate on focal point values, especially $Z = 0$.

pay for an unemployment insurance contract (as will be formalized in Section III), but will be assumed to be unobserved by the econometrician, $Z \neq P$.

Outcomes.—To infer people's knowledge of future job loss, I combine information on subjective beliefs with the subsequent event corresponding to the elicitation. For the baseline analysis, I let U denote an indicator that the individual involuntarily lost their job in the subsequent 12 months from the survey. The subsequent wave asks individuals whether they are working at the same job as the previous wave (roughly two years prior). If not, respondents are asked when and why they left their job (e.g., left involuntarily, voluntarily/quit, or retired). I exclude voluntary quits and retirement in the baseline specifications. Defining U in this manner means that the baseline analysis will estimate the impact of private information on a hypothetical insurance market that pays if the individual loses her job in the subsequent 12 months. I also consider the robustness of the results to alternative definitions of job loss and unemployment below. Changing the definition of U will simulate different hypothetical UI policies. For example, I consider an indicator of job loss in the 6–12 months after the subjective elicitation, which excludes cases where individuals lose their job in the 6 months immediately following the probability elicitation. This will provide an estimate of the impact of private information on an insurance contract that has a six-month waiting period before claims can be exercised.

Public Information.—Estimating private information requires specifying the set of observable information that insurers could use to price insurance policies. Changing the set of observable characteristics simulates how the potential for adverse selection varies with the underwriting strategy of the potential insurer. The data contain a very rich set of observable characteristics that well approximate variables used by insurance companies in disability, long-term care, and life insurance (Finkelstein and McGarry 2006; He 2009; Hendren 2013) and also contain a variety of variables well suited for controlling for the observable risk of job loss. The baseline specification includes a set of these job characteristics, including job industry categories, job occupation categories, log wage, log wage squared, job tenure, and job tenure squared, along with a set of demographic characteristics: census division dummies, gender dummies, age, age squared, and year dummies.[9]

Spousal Labor Supply.—For the subsample of married households in the HRS, I define labor market entry as an indicator for the spouse working for pay in the current wave but not in the previous wave of the survey (two years prior). The primary analysis will focus on labor market entry by previously nonworking spouses, as opposed to total spousal labor supply because of the potential presence of correlated shocks to labor earning opportunities within the household arising from spouses working in the same labor market, industry, or firm.

B. *PSID*

I explore the impact of unemployment on food expenditure in the Panel Study of Income Dynamics (PSID), building upon a large literature (e.g., Gruber 1997; Chetty and Szeidl 2007). I utilize a sample of heads of household between the ages of 25 and 65 who have nonmissing food expenditure. I define food expenditure as the sum of food expenditure in and out of the home, plus food stamps. Following Gruber (1997), I restrict the baseline sample to those with less than a three-fold change in food expenditure relative to the previous year and whose household head is in the labor force (i.e., either employed or unemployed and looking for work). For some specifications, I utilize a measure of household expenditure needs, which the PSID constructs to measure the total expenditure needs given the age and composition of the household. In addition to analyzing food expenditure, I also explore the robustness of the results to the broader consumption expenditure measures available every two years starting after 1997.

I construct an indicator for job loss if the individual was laid off or fired from the job held in the previous wave of the survey.[10] This job loss measure corresponds closely to the definition in the HRS. Ideally, one would measure food expenditure after the onset of a job loss. However, the survey elicits food expenditure only at the time the survey is administered, which may differ from the onset of the job loss. To the extent to which individuals who lose their job find new jobs, food expenditure at the time at which the subsequent survey is administered may understate expenditure

[9] This set is generally larger than the set of information previously used by insurance companies who have tried to sell unemployment insurance. Income Assure, the latest attempt to provide private unemployment benefits, prices policies using a coarse industry classification, geographical location (state of residence), and wages.

[10] I only include job losses coded as "fired or laid off," and do not include cases where the individual quit, the job was seasonal/temporary, or the company "folded/changed hands/moved out of town/went out of business." I do not include this latter case because it includes cases where the individual never loses employment (but had a change in job title because of, for example, a change in management).

immediately after the job loss. To align the timing of food expenditure and employment status, I follow previous literature and define an indicator for whether the individual is currently unemployed at the time of the survey (e.g., Gruber 1997; Chetty and Szeidl 2007). I explore results using both job loss and current unemployment status, but focus on current unemployment status for the baseline analysis.

For the primary sample construction, I select those who are employed in the previous two years of the survey. This aligns with the sample selection in the HRS which requires that individuals are employed at the time of eliciting their job loss probability.

C. *Summary Statistics*

Table 1 presents the summary statistics for the main samples. For the baseline HRS sample, there are 26,640 observations (individual-by-year) in the sample, which correspond to 3,467 unique households. This drops to 2,214 households when using the married subsample. The PSID sample contains 65,450 observations (individuals-by-year) from 9,557 individuals who are heads of household.

Individuals in the samples have relatively similar earnings ($36,000 in the HRS; $40,000 in the PSID), although earnings are slightly higher in the PSID sample of household heads. The household heads in the PSID are also more likely to be male (83 percent versus 40 percent).[11]

The most notable distinction between the HRS and PSID samples is the difference in their age distributions. The HRS sample is older, with a mean age of 56 as opposed to 41. Despite this, the frequency of involuntary job loss and subsequent unemployment is fairly similar, with annualized rates of job loss of 3.1 percent in the HRS and 2.8 percent in the PSID sample. The PSID sample is more likely to be unemployed (2.4 percent versus 1.9 percent in the HRS), perhaps because job losses in the HRS sample are more likely to lead to retirement: the retirement hazard in the HRS sample is 5.3 percent per year versus just 1.7 percent in the PSID sample.[12] As discussed below, the analysis will assume an insurer can distinguish involuntary job loss from retirement; if this is not possible, the knowledge of future retirement plans could present an additional source of adverse selection in a private UI market.

For the married HRS sample, 69.3 percent of spouses are employed and 3.9 percent of spouses make an entry into the labor market (defined as an indicator for being out of the labor force in the previous wave and in the labor force in the current wave of the survey. For the PSID sample, mean household food expenditure is $7,314.).

Finally, the second-to-last set of rows in Table 1 report the summary statistics for the subjective probability elicitations. While 3.1 percent of the HRS sample lose their job in the subsequent 12 months from the survey, the mean subjective probability elicitation is 15.7 percent. Such bias is a common feature of subjective probability elicitations (see, e.g., Hurd 2009). In particular, for low probability events there is a natural tendency for measurement error in elicitations to lead to an upward bias.

[11] Although the HRS is a representative sample of individuals over 50 years old and their spouses, the sample of women exceeds 50 percent because women are more likely than men to be married to someone over age 65 (and those over age 65 are not included in the baseline sample).

[12] I construct the yearly retirement hazard in the HRS by computing the fraction of the sample who is retired in the subsequent wave (two years forward) and dividing by 2.

Table 1—Sample Summary Statistics

	Panel A. HRS full sample		*Panel B. HRS married sample*		*Panel C. PSID sample*	
	Mean	SD	Mean	SD	Mean	SD
Characteristics						
Age	56.1	5.1	56.6	5.0	41.0	9.8
Male	0.40	0.49	0.44	0.50	0.83	0.37
Wage	35,813	150,578	37,395	53,303	40,146	35,503
Unemployment/job loss outcomes						
Involuntary job loss	0.031	0.173	0.029	0.168	0.028	0.165
Unemployed	0.019	0.138	0.016	0.125	0.024	0.152
Retirement hazard rate	0.053	0.153	0.059	0.161	0.017	0.128
Other variables						
Spouse working for pay			0.693	0.461		
Spouse labor market entry			0.039	0.194		
Food expenditure					7,314	3,915
Subjective probability elicitation	15.7	24.8	14.8	24.0		
Sample size						
Observations	26,640		11,049		65,483	
Households	3,467		2,214		9,557	

Notes: This table presents summary statistics for the samples used in the paper. Panel A presents the summary statistics for the baseline HRS sample of individuals; panel B presents the summary statistics for the HRS married subsample used to study spousal labor supply responses; panel C presents the summary statistics for the PSID sample of household heads. The rows present selected summary statistics. Wages and food expenditures are deflated to 2000 US$ using the CPI-U-RS.

This provides further rationale for treating these elicitations as noisy and potentially biased measures of true beliefs, as is maintained throughout the empirical analysis below. An alternative explanation is that individuals hold overly pessimistic beliefs about losing their job. I return to a discussion of the implications of biased beliefs in Section VI.

II. Knowledge of Future Job Loss

This section documents three empirical patterns. First, individuals' subjective probabilities are predictive of future job loss conditional on a wide range of observables that insurers could potentially use to price an insurance policy. Second, when individuals learn they might lose their job, spouses are more likely to enter the labor market. Third, consumption growth differs for those who do versus those who do not lose their job *in the one to two years prior to job loss*. These empirical patterns will then provide a basis for the analysis in Sections III–V that will more formally explore whether private information prevents the existence of a private UI market.

A. *Private Information about Future Job Loss Using Subjective Probability Elicitations*

Do people have private information about their likelihood of losing their job that could be used to adversely select an insurance policy if it were offered? To address this, I ask whether those with higher subjective probability elicitations, Z_{it}, are more

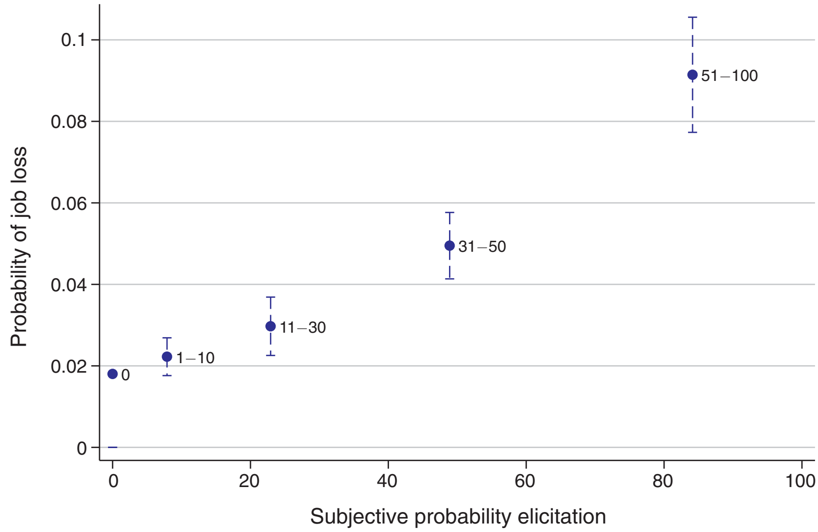

FIGURE 2. PREDICTIVE CONTENT OF SUBJECTIVE PROBABILITY ELICITATIONS: BINNED SCATTER-PLOT OF U VERSUS Z, CONDITIONAL ON X

Notes: This figure reports mean rate of job loss in each elicitation category controlling for demographics, job characteristics, and year controls. To construct this figure, I run the regression in equation (1). The figure plots the coefficients on bins of the elicitations. I omit the lowest bin (corresponding to $Z = 0$) and add back the mean rate of job loss of 1.9 percent to all coefficients. The 5/95 percent confidence intervals are constructed using the standard errors of the regression coefficients, clustering by household.

likely to experience a job loss in the subsequent 12 months, U_{it}, conditional on year dummies, demographic characteristics, and job characteristics, X_{it}. The fact that the elicitations are predictive of subsequent job loss was first shown by Stephens (2004). The analysis in this section expands upon this work by focusing on individuals' knowledge that is not captured by observable characteristics, X_{it}, that insurers might use to price insurance.

To illustrate the predictive content of the elicitations, I partition the range of responses of Z_{it} in the unit interval into five bins, G_j, and construct indicators for $Z_{it} \in G_j$. I regress an indicator for job loss in the subsequent 12 months from the survey, U_{it}, on observable controls at time t, X_{it}, and bin indicators,

$$(1) \qquad U_{it} = \alpha + \sum_{j=1}^{n} \psi_j \mathbf{1}\{Z_{it} \in G_j\} + \Gamma X_{it} + \epsilon_{it}.$$

Figure 2 plots the coefficients, ψ_j, omitting the lowest bin (corresponding to $Z = 0$) and adding back the mean job loss probability for those in the lowest bin of 1.9 percent. The figure reveals an increasing pattern: those with higher subjective probability elicitations are more likely to lose their job, conditional on demographics and job characteristics. This suggests individuals have knowledge about future job loss that is not readily captured by their observable demographic and job characteristics.

Table 2—Presence of Private Information about Future Job Loss

		Alternative controls			Subsamples					
	Baseline (1)	Demo (2)	Health (3)	Ind FE (4)	Age ≤ 55 (5)	Age > 55 (6)	Below median wage (7)	Above median wage (8)	Tenure $\geq$ 5 yrs (9)	Tenure $<$ 5 yrs (10)
Elicitation (Z)	0.0836	0.0956	0.0822	0.0715	0.0716	0.0914	0.0751	0.0957	0.0746	0.0906
Standard error	(0.00675)	(0.00685)	(0.00736)	(0.0107)	(0.0102)	(0.00876)	(0.00812)	(0.0113)	(0.00793)	(0.0120)
Controls (X)										
Year fixed effects	X	X	X	X	X	X	X	X	X	X
Demographics	X	X	X	X	X	X	X	X	X	X
Job characteristics	X		X	X	X	X	X	X	X	X
Health			X							
Individual FE				X						
Mean dep var (U)	0.0307	0.0307	0.0317	0.0307	0.0297	0.0314	0.0369	0.0245	0.0175	0.0575
Observations	26,640	26,640	22,831	26,640	11,134	15,506	13,321	13,319	17,850	8,790
Households	3,467	3,467	3,180	3,467	2,255	3,231	2,824	2,309	2,952	2,437

Notes: This table presents regression coefficients from a linear regression of an indicator for job loss in the next 12 months on the subjective probability elicitation, Z, controlling for observable characteristics, X. Column 1 presents the baseline specification with controls for year dummies, demographics, and job characteristics. Column 2 uses only demographic controls and year dummies. Column 3 uses demographic, job characteristics, and health characteristics. Column 4 adds individual fixed effects to the baseline specification. Columns 5–10 report results for the baseline specification on various subsamples including below and above age 55 (columns 5 and 6), above- and below-median wage earners (columns 7 and 8) and above and below five years of job tenure (columns 9 and 10). Standard errors are clustered by household.

Table 2 presents the results from a linear parameterization of this relationship. I regress the indicator for job loss in the subsequent 12 months, U_{it}, on observable controls at time t, X_{it}, and the subjective probability elicitation, Z_{it}:

$$U_{it} = \alpha + \beta Z_{it} + \Gamma X_{it} + \epsilon_{it}. \tag{2}$$

For the baseline specification with demographics and job characteristics, the estimated slope is 0.0836 (SE 0.00675). For every 1 percentage point (pp) increase in the elicitation, Z, individuals are roughly 0.08pp more likely to lose their job in the subsequent 12 months. Columns 2–4 illustrate the robustness of the estimated coefficient to alternative controls, X_{it}. Dropping job characteristics leads to a slightly higher coefficient of 0.0956 (SE 0.00685); adding additional controls for health characteristics[13] reduces the coefficient to 0.0822 (SE 0.00736).

Column 4 adds individual fixed effects. Of course, an insurer could never actually use an individual fixed effect to price insurance, as it would require using information that is partially realized in the future in order to construct the individual-specific means. Nonetheless, even if an insurer could do so, it would not mitigate the asymmetric information problem: adding individual fixed effects only reduces the coefficient to 0.0738 (SE 0.012). This shows that individuals' risk of job loss is largely time-varying within an individual. This underscores the difficulty that would be

[13] These include indicators for a range of doctor-diagnosed medical conditions (diabetes, a doctor-diagnosed psychological condition, heart attack, stroke, lung disease, cancer, high blood pressure, and arthritis) and linear controls for body mass index (BMI).

faced by an insurance company attempting to use observable variables to reduce the asymmetric information problem.

Columns 5–10 illustrate the presence of private information across a range of subsamples. Although the HRS primarily focuses on older workers, columns 5–6 split the sample into those above and below 55 years old. The coefficients are similar, illustrating the stability of the patterns across the age ranges observed in the data. Columns 7–8 split the sample into those with above- and below-median wage earnings to show the stability of the pattern across the income distribution. Finally, columns 9–10 split the sample into those with more and less than 5 years of job tenure. Again, the pattern is similar with the estimated coefficient ranges around 0.07–0.09. Across demographic subgroups, individuals have knowledge about their potential future job loss.

B. *Spousal Labor Supply Response to Knowledge of Future Job Loss*

Individuals have knowledge about their future job loss. Would they use this information to adversely select an insurance contract if it were offered to them? More generally, how do individuals react to learning they might lose their job? If individuals are underinsured against the risk of job loss, learning today that you might lose your job tomorrow should trigger anticipatory responses to earn more income and cut back on consumption. This subsection explores how spousal labor supply responds to learning about future job loss. This extends a large literature on the added worker effect (e.g., Gruber and Cullen 1996) and builds on Stephens (2002), who shows evidence of an ex ante response by spouses to unemployment shocks in the PSID. Here, I use the subjective probability elicitations in the HRS to provide new evidence that spousal labor supply responds to learning one might lose their job.

Figure 3 shows that spouses of individuals who are likely to lose their job are more likely to enter the labor market. To construct this figure, I replace the job loss variable, U_{it}, in equation (1), with an indicator for spousal labor market entry that is equal to 1 if the spouse not working for pay last wave and working for pay in the current wave of the survey, denoted $Entry_{it}$. Figure 3 plots the resulting coefficients on the elicitation bins, ψ_i. For example, spouses of individuals with $Z > 50$ as opposed to $Z = 0$ are 2 percentage points more likely to enter the labor force. On the one hand, this is a small effect: it suggests roughly 1 in 50 extra spouses are induced into the labor market when the elicitation is above 50 percent. On the other hand, relative to the base entry rate of these spouses of 3.9 percent, it is quite large. For values $Z < 50$, the response is more muted, perhaps consistent with a model in which labor market entry has high fixed cost.

Table 3 linearly parameterizes the relationship in Figure 3 using the regression equation

$$Entry_{it} = \beta Z_{it} + \Gamma X_{it} + \eta_{it}, \tag{3}$$

where β is the increased likelihood of spousal labor force entry for those with a 1pp increase in the subjective probability elicitation, Z. This yields a coefficient of 0.0258 (SE 0.0087) for the baseline specification in column 1. Column 2 restricts

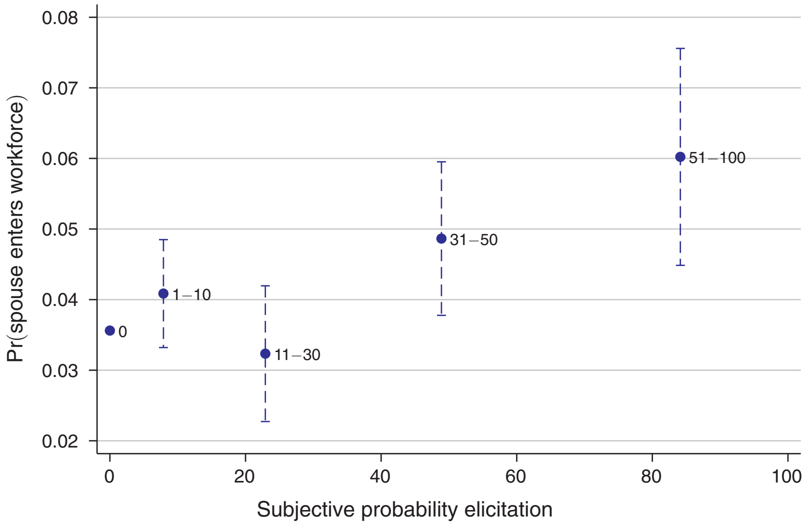

Figure 3. Relationship between Potential Job Loss and Spousal Labor Supply

Notes: This figure presents coefficients from a regression of an indicator for a spouse entering the labor force, defined as an indicator for not working in the previous wave and working in the current wave, on category indicators for the subjective probability elicitations, *Z*, controlling for demographics, job characteristics, and year controls. Figure reports 5/95 percent confidence intervals for each category indicator which are computed by clustering standard errors by household.

Table 3—Spousal Labor Supply Response to Potential Job Loss

	Baseline (1)	Sample without future job loss (2)	Full-time work (3)	2 yr. lagged entry (placebo) (4)	Household fixed effects (5)	Individual fixed effects (6)	Exit (7)	Spouse unemployment (8)
Elicitation (*Z*)	0.0258	0.0256	0.0255	0.00122	0.0243	0.0312	0.0174	0.0213
Standard error	(0.00868)	(0.00898)	(0.00988)	(0.00800)	(0.0114)	(0.0180)	(0.0119)	(0.00966)
Mean dep var	0.0394	0.0389	0.0524	0.0394	0.0394	0.0394	0.0851	0.0296
Observations	11,049	10,726	11,049	11,049	11,049	11,049	11,049	9,079
Households	2,214	2,194	2,214	2,214	2,214	2,214	2,214	1,359

Notes: This table presents the coefficients from a regression of spousal labor entry on the subjective elicitation. I restrict the sample to respondents who are married in both the current and previous wave. I define spousal entry as an indicator for the event that both (i) the spouse was not working for pay in the previous wave (2 years prior) and (ii) the spouse is currently working for pay. I include observations for which the spouse was working for pay in the previous wave (these observations are coded as 0). Column 1 presents a linear regression of an indicator for spousal labor entry on the elicitation, *Z*, and controls for year, demographics, and job characteristics. Column 2 restricts to the subsample that does not lose their job in the subsequent 12 months. Column 3 defines spousal labor force entry using only full-time employment. I define an indicator for the event that both (i) the spouse was not employed full time in the previous wave and (ii) is currently working full time. Column 4 uses the lagged value of *Z* from the previous wave (2 years prior) as a placebo test. Column 5 adds household fixed effects to the specification in column 1. Column 6 adds individual fixed effects to the specification in column 1. Column 7 replaces the dependent variable with an indicator for exit of the spouse from the labor market. I define exit as an indicator for being in the labor force last wave (2 years prior) and out of the labor force this wave. Column 8 replaces the dependent variable with an indicator for spouse unemployment in the subsequent 12 months and restricts the sample to spouses currently in the labor market. All standard errors are clustered by household.

the sample to those who do not end up losing their job in the 12 months after the survey, yielding 0.0256 (SE 0.009). This suggests households are responding to the risk of job loss, even if the realization does not occur. This highlights the value of using the elicitations as opposed to the realization of future job loss as a proxy for knowledge of future job loss. Column 3 uses a specification that defines spousal work as an indicator for full-time employment, as opposed to any working for pay. This definition includes shifts from part-time to full-time work in the definition of labor market entry, and finds a similar slope of 0.0255 (SE 0.0099).

The results suggest that spousal labor supply increases in response to learning one might lose their job. However, the patterns could also reflect a selection effect. For example, perhaps individuals who are more likely to lose their jobs may be more likely to have spouses who have less labor force attachment and are more likely to enter and exit into the labor market. To this aim, column 4 considers a placebo test that uses the lagged value of the elicitation, $Z_{i,t-2}$ instead of $Z_{i,t}$ (where year $t-2$ corresponds to the previous wave of the survey conducted 2 years prior; going forward, commas will separate i and t so that time lags are clearly marked). Here, the coefficient is 0.00122 (SE 0.008) and is not statistically distinct from 0. Columns 5 and 6 add household and individual fixed effects to the specification in column 1, yielding similar coefficients. In short, the results suggest that, in response to learning about future job loss, spouses are more likely to enter the labor market.[14] This suggests not only that individuals and households have private information about future job loss, but that they would act upon this information if presented with opportunities to mitigate this risk, such as private UI markets.

C. *Consumption Response to Knowledge of Future Job Loss*

There is also a large literature studying the impact of unemployment on measures of consumption and food expenditure. For example, Gruber (1997) uses data from the PSID to document that food expenditure drops 6–10 percent upon unemployment. However, if individuals learn about their potential future unemployment prior to its onset, one would expect these cuts to occur before individuals become unemployed. In this section, I find evidence for this ex ante food expenditure drop in the one to two years prior to the job loss or unemployment spell.

Following Gruber (1997), let $g_{i,t} = \log(c_{i,t}) - \log(c_{i,t-1})$ denote yearly food expenditure growth, where $c_{i,t}$ is food expenditure of household i in year t. Let $U_{i,t}$ denote an indicator for being unemployed at the time of the survey in year t. I regress food expenditure growth $g_{i,t}$ on unemployment in year $t-k$,

$$g_{i,t} = a_k + \Delta_k^{FD} U_{i,t-k} + \Gamma_k X_{i,t} + \nu_{i,t}, \tag{4}$$

[14] One may also expect to see fewer spouses leave the labor force in response to learning about future unemployment prospects for the other earner. However, a countervailing force could arise from correlated labor demand shocks (e.g., from spouses working in the same industry). To explore these patterns, column 7 defines labor market exit as an indicator for a spouse working for pay last wave and not working for pay in the current period. The coefficient of 0.0174 (SE 0.0119) is positive, although not statistically significant and suggests that households face correlated unemployment shocks. Further evidence of correlated shocks is presented in column 9, which shows that the elicitation is positively related to spousal unemployment in the subsequent year, with a coefficient of 0.0213 (SE 0.0097). For these reasons, I focus primarily on labor market entry of spouses who are not currently in the labor market and perhaps face greater flexibility in their choice of industry/occupation/firm/etc. when choosing employment.

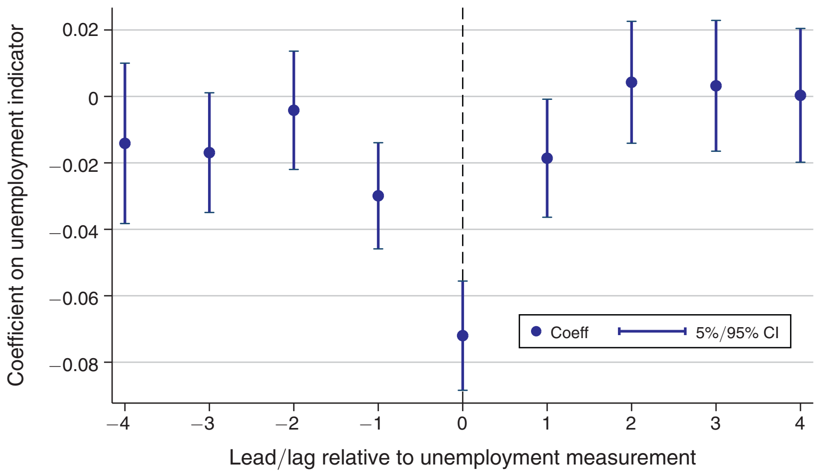

Figure 4. Impact of Unemployment on Consumption Growth

Notes: This figure presents coefficients from separate regressions of leads and lags of the log change in food expenditure on an indicator of unemployment, along with controls for year indicators and a cubic in age. Data are from the PSID with one observation per household per year. Unemployment is defined as an indicator for the household head being unemployed. Following Gruber (1997) and Chetty and Szeidl (2007), food expenditure is the sum of food in the home, food outside the home, and food stamps. The horizontal axis presents the years of the lead/lag for the consumption expenditure growth measurement (i.e., 0 corresponds to consumption growth in the year of the unemployment measurement relative to the year prior to the unemployment measurement). The sample is restricted to household heads who are employed in $t-1$ or $t-2$.

for a range of leads and lags, k. The coefficient Δ_k^{FD} measures the average difference in consumption growth in period t between those who are and are not unemployed in period $t-k$. To control for other life-cycle or aggregate trends in consumption that might affect $g_{i,t}$, I include a cubic in the household head's age and a full set of year dummies in the controls, $X_{i,t}$.

Using the sample of individuals who are employed in the two years prior to the unemployment measurement, Figure 4 plots the coefficients Δ_k^{FD} for $k = -4, -3, \ldots, 0, \ldots, 3, 4$. Consistent with previous literature, food expenditure drops by 7–8 percent at the onset of unemployment. But, there is also a 2–3 percent impact on food expenditure growth in the year prior to unemployment. I also find small but statistically insignificant drops in consumption in the earlier years (e.g., $t-3$ relative to $t-4$).[15] In years after the unemployment measurement, the coefficients Δ_k^{FD} are close to 0. Because all of these regressions are in first differences, the impact of unemployment is consistent with a long-run shock to food expenditure that does not recover, as shown in Stephens (2001).

Table 4 illustrates the robustness of the ex ante drop in food expenditure in the one to two years prior to unemployment, Δ_{-1}^{FD}. Column 1 shows that the baseline specification yields a 2.71 percent (SE 0.975 percent) drop in food expenditure[16] in the

[15] The smaller responses in earlier periods is consistent with evidence discussed below that a large portion (e.g., 10 percent) of knowledge of future job loss in period t is revealed in years $t-2$ relative to $t-1$, but not as much is revealed in the earlier years. See Section V and online Appendix Figure I.

[16] Online Appendix Figure II replicates the baseline regression in Figure 4 using more recent PSID data on total household expenditure on a sample that is surveyed every two years. The broad patterns are similar, although

TABLE 4—EX ANTE DROP IN FOOD EXPENDITURE PRIOR TO UNEMPLOYMENT AND IMPLIED (EX ANTE) WILLINGNESS TO PAY FOR UI

	Baseline (1)	Controls for needs (2)	Under-50 sample (3)	Job loss (4)	Household income controls (5)	Household head income controls (6)
Impact of unemployment on $\log(c_{t-1}) - \log(c_{t-2})$						
Unemployment	−0.0271	−0.0211	−0.0288	−0.0260	−0.0272	−0.028
Standard error	(0.00975)	(0.0105)	(0.0106)	(0.00824)	(0.00969)	(0.00983)
Specification details						
Sample employed in $t-2$ and $t-1$	X	X	X	X	X	X
Controls for change in log needs ($t-2$ versus $t-1$)		X				
Change in log HH income ($t-2$ versus $t-1$) (3rd-order poly)					X	
Change in log HH head income ($t-2$ vs. $t-1$) (3rd-order poly)						X
Mean dependent variable	0.000	−0.001	0.007	0.000	0.000	0.000
Observations	65,483	53,327	52,463	65,556	65,399	64,119
Households	9,557	8,371	8,512	9,560	9,547	9,448

Notes: This table presents estimates of the impact of unemployment in year t on consumption growth in year $t-1$ relative to $t-2$, $\log(c_{t-1}) - \log(c_{t-2})$. Column 1 controls for a cubic in age and year dummies and restricts to the baseline sample of those who are employed in both year $t-2$ and $t-1$. Column 2 adds controls for the change in log expenditure needs (*need_std_p*) between $t-2$ and $t-1$ and the change in total household size between $t-2$ and $t-1$ (this is not available in all years of the survey). Column 3 restricts the sample to those aged 50 and under to the baseline specification. Column 4 replaces the unemployment indicator with an indicator for job loss. Job loss is defined as an indicator for being laid off or fired. Column 5 adds controls to the specification in column 1 for a third-degree polynomial in the household's change in log income between years $t-2$ and $t-1$. Column 6 adds controls to the specification in column 1 for a third-degree polynomial in the household head's change in log income between years $t-2$ and $t-1$. All standard errors are clustered at the household level.

year before unemployment occurs. Column 2 shows that controlling for the change in household size in years $t-2$ versus $t-1$ and the change in expenditure needs delivers a fairly similar coefficient of −2.11 percent (SE 1.05 percent).[17] Column 3 restricts the sample to those under age 50: a group largely not captured by the HRS analysis above. This yields a coefficient of −2.88 percent (SE 1.06 percent), which is similar in magnitude and not statistically distinguishable from the baseline estimate. This suggests that knowledge of future job loss is present across the age distribution in the United States, not only the older sample surveyed by the HRS.

Unemployment versus Job Loss.—The baseline figure studies consumption patterns around unemployment, which may be distinct from a job loss event; column 4 illustrates the ex ante food expenditure response to future job loss (regardless of whether it leads to future unemployment). Individuals who experience a job loss in the future year are likely to have already dropped their consumption by 2.6 percent

the consumption drop upon unemployment is slightly larger (e.g., 12 percent) when using total consumption expenditure as opposed to food expenditure. There is also an ex ante response of 3.6 percent ($p = 0.055$) in $t-2$ relative to $t-4$ in total consumption expenditure, which again suggests that individuals decrease their food expenditure in response to learning about future unemployment.

[17] Restricting to the subsample of 53,327 observations for which the needs variable is available drops the coefficient in column 1 to 2.4 percent, suggesting roughly one-half of the drop in the point estimate is driven by differential sample composition; of course, all of these point estimates are well within 1 standard error of the baseline estimate.

(SE 0.824 percent) relative to those who will not experience a future job loss in the subsequent year.[18] Online Appendix Figure III presents the lead and lag estimates as in Figure 4 using job loss instead of unemployment and finds similar patterns to Figure 4. Overall, the results reveal a similar ex ante drop in food expenditure for both job loss and unemployment.

Forward-Looking Behavior versus Correlated Income Shocks.—While the results are consistent with ex ante responses to learning about future unemployment, a competing hypothesis is that income is dropping prior to unemployment and individuals are simply consuming hand-to-mouth. To explore this, column 5 of Table 4 adds controls for a cubic polynomial of changes in log household income to the baseline specification. This yields a coefficient of 2.72 percent (SE 0.969 percent) nearly identical to the baseline specification in column 1. The results are similar with higher- and lower-order polynomial controls, or restricting to those with small income changes between $t-1$ and $t-2$. Column 6 adds controls for a cubic polynomial of changes in log income of the household head, again yielding a similar coefficient of -2.81 percent (SE 0.983 percent).[19]

To understand why the results are not affected by adding controls for income, online Appendix Figure IV replicates Figure 4 using log household income as the dependent variable as opposed to log food expenditure. For those employed in both $t-2$ and $t-1$, unemployment in period t is not associated with any significant income change in any of the years prior to unemployment.[20] Therefore, the ex ante expenditure drop does not appear to be the result of hand-to-mouth consumption combined with correlated income shocks; it is more consistent with an anticipatory response to learning about future unemployment.[21]

III. Model

Individuals have knowledge about their future job loss and take actions to increase their financial resources available when unemployed. This suggests they would demand an unemployment insurance contract if it were offered. Yet, they would likely use their private information to potentially adversely select the contract. The

[18] Temporary or seasonal work is coded separately and not included in job loss. Therefore, the ex ante consumption responses in the PSID suggest the patterns in the HRS are not driven solely by knowledge about fixed-term temporary contracts.

[19] The sample sizes are slightly lower for these specifications due to nonresponse to income questions. The food expenditure patterns are similar when restricting to a sample with nonmissing income reports.

[20] The absence of an ex ante drop in income differs from the findings of Davis and von Wachter (2011) for plant closings. To be sure, there are subsamples in the PSID for which income does decline; in particular, if one includes those who are unemployed in $t-1$ or $t-2$, then income does decline prior to the unemployment measurement (as shown in Stephens 2001). In this sense, the patterns identified here are similar to those found in Stephens (2001), who shows roughly a 2 percent drop in the year prior to a job loss; the main empirical distinction is that I illustrate that these patterns hold on the sample who remain employed in periods $t-1$ and $t-2$, so that it is not driven by these correlated income shocks in the preperiods.

[21] While the narrative here is that consumption declines because of the future unemployment event, there is a potential for reverse causation that is worth mentioning. An alternative story is that there is a negative shock to the marginal utility of consumption in $t-1$, which leads to an increase in savings, which in turn leads to a wealth effect on job effort and an increase in the likelihood of losing one's job. Although such a story would potentially explain the pattern of consumption, it would yield opposing predictions for spousal labor supply. A decline in the marginal utility of income in period $t-1$ should lead to a decrease in spousal labor supply, which contrasts with the increasing patterns shown in Section IIB.

remainder of this paper explores whether this private information prevents the existence of the private unemployment insurance market.

To do so, I begin by developing a theory of when a private market should exist that can be used to guide the empirical analysis in Sections IV and V. The framework builds upon a model of Hendren (2013) by incorporating dynamics and precautionary responses, and allows for behavioral responses to providing unemployment insurance (i.e., moral hazard). The model will characterize the empirical objects that need to be estimated in order to understand whether private information prevents market existence. Along the way, the model also provides a framework for thinking about alternative explanations for the absence of a private market, including aggregate risk, biased beliefs, and moral hazard, which will be discussed further in Section VI.

A. *Setup*

Individuals face a risk of losing their job in the next year. They (or their households)[22] choose consumption today, c_{pre}, along with a plan for consumption in the event of not losing and losing their job in the future, c_e and c_u. In addition, they also make a set of other choices, denoted by a vector a, which can include spousal labor supply, job search activities, and can be contingent plans for future behavior conditional on other events that may happen in the future. In this sense, the model is both stochastic and dynamic. Let p denote an individual's chance of losing her job in the next 12 months. The model should be thought of as conditioning on a particular observable characteristic, X, so that p reflects the individual's privately known information. This probability p can be also be affected by the choices of the individual—for example, UI may increase the likelihood of job loss (moral hazard).

Individual i makes choices $\{c_{pre}, c_u, c_e, a, p\} \in \Omega_i$ subject to an individual-specific constraint set, Ω_i, to maximize a utility function that satisfies

$$v(c_{pre}) + pu(c_u) + (1-p)\,v(c_e) + \Psi_i(p,a)\,, \tag{5}$$

where $v(c_{pre})$ is the utility over consumption today, $u(c_u)$ and $v(c_e)$ are the utilities over consumption if the individual does and does not lose her job next year, and $\Psi_i(p,a)$ is the (dis)utility from all of the other choices p and a.

The model generalizes the structure in Hendren (2013) in several ways. First, the utility function over consumption is allowed to differ for those who remain employed, v, versus those who lose their job, u. This allows for state-dependent utility. Second, the probability of job loss is allowed to be a choice, so that the problem incorporates moral hazard: more insurance can increase the likelihood of job loss. Third, the model in principle allows for multidimensional heterogeneity (e.g., different individuals, i, may have different utility functions, ψ_i, and face different constraints, Ω_i). In the exposition in the main text, I assume heterogeneity can be fully summarized by the choice of p (i.e., it is unidimensional); but the modeling in online Appendix A illustrates how the analysis readily extends to the case of

[22] Throughout, I use the language of individuals to refer to the maximization decision. But many of the variables will be measured at the household level and can be thought of as the result of joint household decision making. This distinction is without loss of generality as long as the within-household allocations are Pareto efficient.

multidimensional heterogeneity. Finally, the model allows for sources of formal and informal insurance: items such as transfers from friends and family and the current level of government benefits are embodied in the constraints, Ω_i. In particular, the model allows for dynamic responses to underinsurance and thus the ability to match the empirical patterns in Section II.

Consumption Responses.—The model captures the dynamic consumption patterns in Figure 4 by assuming that the constraints, Ω_i, allow the individual to save today to increase consumption in both states of the world tomorrow.[23] Optimization of this savings decision yields the familiar Euler equation,

$$v'\big(c_{pre}(p)\big) \;=\; pu'\big(c_u(p)\big) + (1-p)\,v'\big(c_e(p)\big). \tag{6}$$

The marginal utility of income today equals the expected marginal utility of income in the future. If the marginal utility of income is higher when unemployed, $u' > v'$ (i.e., individuals are underinsured), then learning one might lose their job should cause individuals to cut back on current consumption and save for future consumption. In this sense, the ex ante responses in Figure 4 suggest that individuals are not fully insured against the risk of job loss.

Spousal Labor Supply Responses.—The model also captures the spousal labor supply responses documented in Section IIC. To see this, one can incorporate spousal labor supply into the set of other actions. Let $a = \big(l^{spouse}, a'\big)$ and $\psi_i(p,a) = \kappa_i(p,a') - \eta\big(l^{spouse}(p)\big)$, where a' is the vector of all other actions and η captures the disutility of spousal labor supply. Let $w^{spouse}\,l^{spouse}$ denote the earnings of the spouse with labor supply l^{spouse}. Then, the intratemporal choice of labor supply implies the marginal disutility of labor is equated to the marginal value of consumption multiplied by the wage, $\eta'\big(l^{spouse}(p)\big) = w^{spouse}v'\big(c_{pre}(p)\big)$. So, the Euler equation (6) can be rewritten as

$$\frac{1}{w^{spouse}}\eta'\big(l^{spouse}(p)\big) \;=\; pu'\big(c_u(p)\big) + (1-p)\,v'\big(c_e(p)\big). \tag{7}$$

The marginal disutility of labor divided by the spouse's wage equals the expected marginal utility of income in the future. When individuals are underinsured against job loss (i.e., $u' > v'$), an increase in p should lead to an increase in spousal labor supply because of the increase in the marginal utility of income. Equation (7) illustrates this logic along the intensive margin; the derivation in online Appendix C5 illustrates this logic using extensive margin responses. If individuals have higher marginal utilities of income if they lose their job than if they do not, they will want to increase their income if job loss becomes more likely. In this sense, the model formalizes how the ex ante consumption and labor supply responses suggest

[23] To be specific, I assume that if $\{c_{pre}, c_u, c_e, a, p\} \in \Omega_i$ then $\{c_{pre} - s, c_u + s, c_e + s, a, p\} \in \Omega_i$ for all s. For simplicity, I assume that the interest rate on savings equals 1 to be consistent with the lack of discounting in equation (5). More generally, this Euler equation is obtained as long as the discount rate on utility equals the interest rate on risk-free savings.

$u'(c_u(p)) > v'(c_e(p))$, so that individuals are not fully insured against job loss and would have a demand for a private UI policy.

B. *Existence of Private Markets*

When can a third-party insurance company enter this environment and profitably sell an insurance contract? To begin, I consider a private market for additional insurance on top of what is currently provided by existing formal and informal insurance arrangements. Later, Section VI discusses the alternative (and more difficult) question of whether a private market would arise in a counterfactual world in which the government were to stop providing UI.

Suppose an insurer attempts to sell a policy that pays \$1 in the event the individual loses her job. An individual who has a likelihood p of losing her job is willing to pay $\frac{p}{1-p}\frac{u'(c_u(p))}{v'(c_e(p))}$ from the future state of not losing her job to buy this policy, where $u'(c_u(p))$ and $v'(c_e(p))$ are the marginal utilities of consumption in the event of losing and not losing her job.

The actuarially fair cost of providing this \$1 to a type p is $\frac{p}{1-p}$. If an insurer could sell at this price there would be a profitable insurance market as long as the individual had a higher marginal utility of income when unemployed, $\frac{u'(c_u(p))}{v'(c_e(p))} > 1$. But since p is unobserved to the insurer, the cost to the insurance company depends on who else prefers such an insurance contract. One would expect that those with higher probabilities, p, will tend to prefer this insurance contract. Let P denote the random variable corresponding to the distribution of probabilities chosen by the population. Assuming that individuals with higher p will have a higher demand for insurance, the cost of providing insurance to type p will be determined by the average probability of worse risks, $\frac{E[P|P \geq p]}{1-E[P|P \geq p]}$. Online Appendix A states these assumptions more formally and shows that a private market cannot exist if and only if

$$\frac{u'(c_u(p))}{v'(c_e(p))} \leq T(p) \quad \forall p, \tag{8}$$

where $\frac{u'(c_u(p))}{v'(c_e(p))} - 1$ is the markup over actuarially fair rates that a type p is willing to pay for a small amount of insurance and $T(p) = \frac{E[P|P \geq p]}{1-E[P|P \geq p]}\frac{1-p}{p}$ is the pooled price ratio defined in Hendren (2013). The pooled price ratio is the markup a type p would have to be willing to pay in order to cover the pooled cost of worse risks adversely selecting their insurance contract. In other words, the no-trade condition in equation (8) says that unless someone in the economy is willing to pay the pooled cost of worse risks in order to obtain some insurance, there can be no profitable insurance market.

Moral Hazard.—The responsiveness of behavior to insurance (e.g., an increase in p resulting from insurance) does not affect whether a market exists. This is because the first \$1 of insurance provides first-order welfare gains (given by $\frac{u'(c_u(p))}{v'(c_e(p))}$) but the behavioral response to that first \$1 of insurance imposes only a second-order cost to the insurance company. As a result, moral hazard does not provide a singular explanation for the absence of an insurance market. This insight, initially noted by Shavell (1979), suggests that moral hazard does not affect whether insurers' first \$1 of insurance is profitable.[24] Moral hazard can limit the size of the gains to trade, but does not provide a singular theoretical explanation for the absence of a market. In contrast, the first \$1 of insurance can be adversely selected by strictly worse risks, so that private information can explain the absence of a market.

Multidimensional Heterogeneity and Relation to Akerlof (1970) *and Einav, Finkelstein, and Cullen* (2010).—Equation (8) is similar to the unraveling condition in Akerlof (1970) and Einav, Finkelstein, and Cullen (2010). In those models, the market will fully unravel whenever the average cost curve (average cost of those purchasing the insurance at a given price) lies everywhere below the demand curve (willingness to pay for the marginal purchaser). Loosely, the *demand curve* is given by $\frac{p}{1-p}\frac{u'(c_u(p))}{v'(c_e(p))}$ and the *average cost curve* is given by $\frac{E[P|P \geq p]}{1-E[P|P \geq p]}$.[25] The key distinction of the model in this paper is that it does not exogenously restrict the set of insurance contracts traded. When equation (8) holds, the market for any insurance contract (or menu of insurance contracts, as shown in online Appendix A2) about the job loss will fully unravel in the sense of Akerlof (1970) and Einav, Finkelstein, and Cullen (2010).

While the model allows for endogenous contracts, the assumption that individuals with higher p will always have a higher demand for insurance is more restrictive than the models of Akerlof (1970) and Einav, Finkelstein, and Cullen (2010) because it assumes a single dimension of heterogeneity summarized by p. In online Appendix A.1, I extend the baseline model in Section III to allow two people with the same p to have a different willingness to pay, $\frac{u'}{v'}$. In this setting, ideally one would estimate the joint distribution of each person's willingness to pay for UI and the cost they impose on the insurance company. With this, one could simulate the demand and cost curves of Einav, Finkelstein, and Cullen (2010) for any hypothetical insurance contract and understand whether a market could exist. I extend the model to allow for multidimensional heterogeneity in online Appendix A.1 and derive an equation similar to equation (8) in which one can replace the willingness to pay, $\frac{u'(c_u(p))}{v'(c_e(p))}$, with a more complicated interior quantile of the type space that depends on the joint distribution of the type distribution and willingness to pay. A market can exist only if there exists one of these interior types that is willing to pay the pooled cost of worse risks, $T(p)$, in order to obtain insurance. In this sense, the pooled price ratio remains a key empirical quantity of interest even in the presence

[24] As discussed in Section IVC, one could combine a model of moral hazard and fixed costs to explain the absence of a private market.

[25] Online Appendix Figure V presents a graphical illustration of equation (8).

of multidimensional heterogeneity, and whether individuals are willing to pay the pooled cost of higher risks continues to characterize when a private market can exist.

In short, equation (8) provides a theory of how private information can prevent the existence of a UI market. The next two sections focus on translating the empirical evidence in Section I and II into estimating each side of this equation: the willingness to pay for UI, $\frac{u'(c_u(p))}{v'(c_e(p))}$, and the pooled price ratio, $T(p)$.

IV. Implications for Willingness to Pay

The absence of a private market for UI makes it difficult to know how much people would be willing to pay for it, $\frac{u'(c_u(p))}{v'(c_e(p))}$. This section develops two strategies to estimate this willingness to pay. The first strategy in Section IVA builds on previous literature by inferring the willingness to pay from the size of the causal effect of job loss or unemployment on consumption (i.e., the difference between $c_e(p)$ and $c_u(p)$). The ex ante responses documented in Section II suggest that previous approaches using first difference estimates (e.g., Gruber 1997) understate the true causal effect. I provide a method to correct for this bias.

The second strategy presented in Section IVB provides new methods that exploit these ex ante behavioral responses to infer people's willingness to pay for UI. Under assumptions provided below, the extent to which individuals cut their consumption in response to learning that they might lose their job in the future reveals how much they are willing to pay to have additional resources if they lose their job.

A. *Approach 1: Consumption when Unemployed versus Employed*

Setup.—A common approach in previous literature to measure the willingness to pay for UI is to assume that individuals have state-independent preferences over consumption $(v(c) = u(c))$. Under this assumption, one can follow Chetty (2006) by using a second-order Taylor expansion for u' around $c_e(p)$,[26] to write

$$(9) \qquad \frac{u'(c_u(p))}{v'(c_e(p))} - 1 \approx \sigma \frac{\Delta c}{c}(p)\left[1 + \frac{\gamma}{2}\frac{\Delta c}{c}(p)\right],$$

where $\frac{\Delta c}{c} = \frac{c_e(p) - c_u(p)}{c_e(p)}$ is the causal effect of job loss on type p's percentage difference in consumption, σ is the coefficient of relative risk aversion, $\sigma = -\frac{c_e(p)\, u''(c_e(p))}{u'(c_e(p))}$, and $\gamma = -\frac{c_e(p)\, u'''(c_e(p))}{u''(c_e(p))}$ is the coefficient of relative prudence.[27] For simplicity, I assume a constant coefficient of relative risk aversion $\left(\text{e.g., } u(c) = \frac{c^{1-\sigma}}{1-\sigma}\right)$ so that the coefficient of relative prudence equals the coefficient of relative risk aversion

[26] Formally, $u'(c) \approx u'(c_e(p)) + u''(c_e(p))\,(c - c_e(p)) + \frac{1}{2}u'''(c_e(p))\,(c - c_e(p))^2$.

[27] See online Appendix C.1 for the derivation.

plus 1, $\gamma = \sigma + 1$. Following Gruber (1997), among others, I approximate this percentage change using differences in log consumption:

$$\frac{\Delta c}{c}(p) \approx \log\big(c_e(p)\big) - \log\big(c_u(p)\big). \tag{10}$$

Therefore, estimating the willingness to pay for UI requires estimating the causal impact of job loss on log consumption.

To estimate the causal effect of job loss on consumption using observational data, one must account for the fact that those who lose their job may differ on many other dimensions (e.g., have lower earnings history or less savings) that generate differences in consumption levels. To remove this selection bias, previous literature often uses yearly consumption first differences, as in Δ_0^{FD} in equation (4) (Gruber 1997; Chetty and Szeidl 2007). The first row of Table 5 presents the estimates for Δ_0^{FD} in equation (4) using the baseline sample from the PSID. Consistent with Gruber (1997) and Chetty and Szeidl (2007), the event of unemployment leads to a roughly 7–8 percent lower food expenditure relative to the previous year.[28] Column 1 presents the results for the baseline sample in the PSID, yielding a coefficient of 7.23 percent (SE 0.997 percent).

Correcting Bias in the First Difference Estimator.—Unfortunately, the ex ante responses documented in Section II suggest that Δ_0^{FD} does not capture the full causal effect. Indeed, one can write the first difference estimate as

$$\Delta_0^{FD} = \underbrace{E\big[\log\big(c_e(p)\big) - \log\big(c_u(p)\big)\big]}_{\text{Average causal effect}} - \underbrace{\Big(E\big[\log\big(c_{pre}(p)\big)|U = 0\big] - E\big[\log\big(c_{pre}(p)\big)|U = 1\big]\Big)}_{\text{Bias from ex ante response}}, \tag{11}$$

where the bias term equals the difference in consumption in the year prior between those who subsequently do versus do not lose their job. The divergence in year $t-1$ suggests that Δ_0^{FD} will understate the true average causal effect.

Here, I show that one can recover the average causal effect from the first difference estimate by accounting for the fact that some information about job loss in period t has been revealed at time $t-1$. More precisely, if P reflects the beliefs about the event U in year t measured in year $t-1$, then a fraction $\frac{\mathrm{var}(P)}{\mathrm{var}(U)} = E[P|U = 1] - E[P|U = 0]$ of the information about U has been revealed at the time P is measured.

[28] Ideally, one would measure the impact of job loss on consumption, not necessarily the impact of unemployment on consumption, as this would more closely align with the hypothetical insurance product that pays \$1 in the event of job loss and corresponds to the definition of U and p in Section III. However, as noted in Section IIC, the PSID does not measure food expenditure directly at the precise time a job loss occurs—rather, one must wait for the next survey wave. To be conservative in testing equation (8), I focus on the larger impact of unemployment, as opposed to an indicator for job loss occurring between surveys. This has the added benefit of aligning with much of the previous literature focusing on estimating the impact of unemployment on consumption (e.g., Gruber 1997; Chetty and Szeidl 2007). And I show below that this is likely a conservative path to prevent understating the willingness to pay for insurance.

TABLE 5—IMPACT OF UNEMPLOYMENT ON CONSUMPTION AND IMPLIED WTP FOR UI

							Bound on maximum WTP	
	Baseline (1)	Controls for needs (2)	Under-50 sample (3)	With outliers (4)	With outliers; no food stamps (5)	Job loss (6)	Baseline (7)	With outliers (8)
Reduced-form impact on $E[\log(c_{t-1}) - \log(c_t)]$								
Unemployment	−0.0723	−0.0747	−0.0713	−0.0889	−0.182	−0.0487		
Standard error	(0.00997)	(0.0102)	(0.0108)	(0.0123)	(0.0171)	(0.00860)		
Two-sample IV estimate for $E[\log(c_u) - \log(c_e)]$	−0.090	−0.093	−0.089	−0.111	−0.227	−0.061		
$\Delta^{\min} = \min\{\log(c_e(p)) - \log(c_u(p))\}$							−0.137	−0.146
							(0.02)	(0.011)
Implied $[u'(c_u) - v'(c_e)]/v'(c_e)$ *for various* σ *(percent)*								
$\sigma = 1$	9.8	10.2	9.7	12.3	27.9	6.4	15.6	16.7
$\sigma = 2$	20.4	21.2	20.1	25.8	60.9	13.2	33.2	35.6
$\sigma = 3$	31.9	33.1	31.3	40.6	99.0	20.4	52.6	56.6
Specification details								
Sample employed in $t-1$	X	X	X	X	X	X	X	X
Controls for change in log needs		X						
Mean dependent variable	−0.005	−0.005	0.001	−0.066	−0.007	−0.005	−0.005	−0.006
Observations	65,808	58,086	52,860	66,913	66,913	65,808	65,808	66,913
Households	9,562	9,131	8,524	9,583	9,583	9,562	9,562	9,583

Notes: This table presents two-sample IV estimates of the causal impact of unemployment on consumption and the implied willingness to pay for UI. The first set of rows present the coefficients from a regression of the change in log food expenditure between years $t-1$ and t on an indicator of unemployment in year t. The sample includes all household heads in the PSID who are employed in years $t-1$ and $t-2$. The baseline specification in column 1 controls for a cubic in age and year dummies. Column 2 adds controls for the change in log expenditure needs between $t-1$ and t and the change in total household size between $t-1$ and t. Column 3 restricts the sample to those aged 50 and under for the specification in column 1. Column 4 adds in the outliers in changes in food expenditure that are more than three-fold. Column 5 replaces the dependent variable in the specification in column 4 with the change in log food expenditure excluding expenditure paid with food stamps. Column 6 replaces the unemployment indicator with an indicator for job loss. Job loss is defined as an indicator for being laid off or fired.

The IV estimate for $\log(c_u) - \log(c_e)$ divides the first difference estimate by the estimated amount of information revealed between year $t-1$ and year t. Online Appendix Table I presents this estimate of the first stage. Using the HRS sample, these estimates are constructed using a regression of the subjective probability elicitations, Z, on an indicator for subsequent unemployment in the next 12 months, U. The implied willingness to pay estimates for $[u'(c_u) - v'(c_e)]/v'(c_e)$ are presented for various values of the coefficient of relative risk aversion, σ, and under the assumption that the coefficient of relative prudence is $\sigma + 1$.

Columns 7–8 present estimates for the upper bound on the maximum willingness to pay for UI, Δ^*. Online Appendix Table II presents the estimates of the components comprising these estimates. Column 7 presents estimates for the baseline sample that drops observations with more than a three-fold change in expenditure; column 8 retains these outliers.

All standard errors are clustered by household. Columns 1–6 present analytical standard errors; columns 7–8 present bootstrapped estimates ($N = 500$).

Under the assumptions stated in Proposition 1, the causal effect is recovered by inflating the first difference estimate, Δ_0^{FD}, by the fraction of information about job loss that remains, $1 - \left(E[P|U = 1] - E[P|U = 1]\right)$.[29]

[29] An alternative strategy would be to use longer lags instead of 1-year lagged consumption. However, online Appendix Figure V shows that individuals have predictive (albeit small) information about future unemployment 10 years in advance. This suggests the lags would need to be longer than 10 years to remove the bias in equation (11).

PROPOSITION 1: *Suppose that* (*i*) *the utility function is state-independent* $(u(c) = v(c))$; (*ii*) *the Euler equation* (6) *holds*; *and* (*iii*) *the causal effect of U on consumption is not varying with p*, $\frac{d\left[\log(c_e(p)) - \log(c_u(p))\right]}{dp} = 0$. *Let* $\approx$ *denote an equality up to log-linear consumption approximations and third-order Taylor approximations for u and v. Then, the average causal effect of U on log consumption is given by*

$$E\left[\log(c_e(p)) - \log(c_u(p))\right] \approx \frac{\Delta_0^{FD}}{1 - \kappa\left(E[P|U = 1] - E[P|U = 0]\right)} \equiv \Delta^{IV}, \tag{12}$$

where $\frac{\mathrm{var}(P)}{\mathrm{var}(U)} = E[P|U=1] - E[P|U=0]$ *is the fraction of variance in U that is realized in beliefs P at time* $t-1$ and $\kappa = E\left[\frac{1}{1 + p\frac{u'(c_u(p)) - v'(c_e(p))}{u'(c_e)}}\right] \approx 1$.

If Assumption (*iii*) *is violated,* Δ^{IV} *is greater than* (*less than*) $E\left[\log(c_e(p)) - \log(c_u(p))\right]$ *if higher values of p correspond to larger* (*smaller*) *consumption drops.*

PROOF:

See online Appendix C.2.

If individuals have no knowledge about U, then $E[P|U = 1] = E[P|U = 0]$, so that the denominator equals 1, and the first difference estimate recovers the average causal effect. But, if individuals have ex ante knowledge, one must inflate the first difference estimate to account for the information that has been revealed when measuring c_{t-1}. The correction factor κ adjusts for the fact that the ex ante consumption response is valued using the ex ante marginal utility, whereas the insurance markup is defined relative to the marginal utility in the ex post state of employment. Online Appendix C.2 shows that $\kappa \approx 1$ because p is small on average ($E[p] = 4$ percent). For example, if individuals are willing to pay a 25 percent markup and $E[p] = 4$ percent, then this correction factor is roughly $\kappa = 1.01$. Going forward, I assume $\kappa \approx 1$.

Online Appendix C.2 provides the proof of Proposition 1 and illustrates how Assumptions (i)–(iii) lead to the scaling in equation (12). Assumption (i) is the state independence assumption that allows one to infer differences in marginal utilities from differences in consumption levels multiplied by a coefficient of relative risk aversion. Assumption (ii) is a standard Euler equation that provides a link between the ex ante consumption response of $c_{pre}(p)$ and the causal effect (i.e., difference between $c_e(p)$ and $c_u(p)$). Finally, Assumption (iii) requires that there is no systematic heterogeneity in the causal effect of U on $\log(c)$ that is correlated with p. Assumptions (ii) and (iii) combine to imply that the impact of learning that unemployment is 1 percent more likely (i.e., p increases by 1pp), $\frac{d\log(c_{pre}(p))}{dp}$, equals 1 percent of the causal effect, $\frac{d[\log(c_u) - \log(c_e)]}{dp}$. This equality suggests that the first difference estimate is attenuated by the amount of information about U that has been revealed, $\left(E[P|U = 1] - E[P|U = 0]\right)$, which yields the formula in equation (12).[30]

The results are as follows.

Estimating $E[P|U = 1] - E[P|U = 0]$.—I use the subjective probability elicitations in the HRS to estimate $E[P|U = 1] - E[P|U = 0]$. Regressing an indicator for job loss on the elicitations, Z, yields an estimate of $E[P|U = 1] - E[P|U = 0]$ as long as the measurement error in Z is classical (i.e., $Z - P$ is uncorrelated with U). Because P and U are bounded variables, the classical measurement error assumption is likely violated. Nonetheless, this provides a natural benchmark. The estimates presented in online Appendix Table I suggest $E[P|U = 1] - E[P|U = 0] \approx 0.197$ (SE 0.012), which implies 80.3 percent of the uncertainty in job loss is not known one year in advance. Online Appendix Table I also illustrates that the 80 percent statistic is quite similar across various demographic subgroups.

Baseline Results.—Table 5 shows that the first difference estimate of the impact of unemployment on consumption is $\Delta_0^{FD} = 7.23$ percent. The second set of rows in Table 5 divide the first difference estimate, Δ_0^{FD}, by 0.803. This suggests a 9 percent average causal effect of U on food expenditure. The remaining rows translate the causal effect into a willingness to pay by following equation (9) for common values of risk aversion (e.g., σ ranging from 1 to 3). For $\sigma = 2$, individuals would be willing to pay a 20.4 percent markup for UI. Higher risk aversion (e.g., $\sigma = 3$) raises the willingness to pay to 31.9 percent; if individuals are less risk averse (e.g., $\sigma = 1$), they would be willing to pay a 9.8 percent markup for UI.

Robustness.—The baseline analysis in column 1 makes a couple of specification decisions whose robustness is explored in columns 2–6. Column 2 controls for changes in household size and food needs, yielding a similar willingness to pay estimate. Column 3 explores the pattern for those age 50 and under (analogous to column 3 in Table 4 for Δ_{-1}^{FD}) and again finds a similar coefficient to the baseline specification. Column 4 reintroduces observations with more than a three-fold change in food expenditure, which were dropped to align with the specification of Gruber (1997). Reintroducing these observations increases the first difference estimate to 8.89 percent (SE 1.23 percent).

The baseline definition of food expenditure sums monthly food spending in the house, out of the house, and any spending that occurred through food stamps. While

[30] It is certainly possible that Assumption (iii) is violated so that the counterfactual consumption difference varies with the unobserved likelihood of job loss, p. In this case, one can show that

$$E\left[\log(c_e(p)) - \log(c_u(p))\right] = \frac{\Delta_0^{FD} + \frac{d\left[\log(c_e(p)) - \log(c_u(p))\right]}{dp}\left(E[P|U = 1] - E[P]\right)}{1 - \kappa\left(E[P|U = 1] - E[P|U = 0]\right) - E[P]\,\sigma\frac{d\left[\log(c_e(p)) - \log(c_u(p))\right]}{dp}}.$$

Estimating this more general equation would require an estimate of $\frac{d\left[\log(c_u(p)) - \log(c_e(p))\right]}{dp}$, which in turn requires data on the joint distribution of consumption and beliefs. Because I use belief data in the HRS and consumption data in the PSID, it is difficult to estimate this unobserved heterogeneity, and so I consider the benchmark case where $\frac{d\left[\log(c_u(p)) - \log(c_e(p))\right]}{dp} = 0$.

this follow Zeldes (1989) and Gruber (1997), there are two concerns with including food stamp expenditures. First, individuals may have already included this spending in their report for in- and out-of-house expenditure (although technically this would not be a correct response). Second, the wording of the food stamp question elicits concurrent expenditure for the previous week, whereas the food expenditure measures elicit a *typical* week. Since unemployment is coincident with rises in food stamp use, this differential recall window could lead to an understating of the impact of unemployment on food consumption. To understand the potential impact of this, column 5 expands the specification in column 4 to exclude food stamp expenditure from the food expenditure measure altogether. This yields a larger expenditure drop of 18.2 percent (SE 1.71 percent), and arguably provides an upper bound on the size of the causal effect of 22.7 percent. This would suggest individuals are willing to pay a 60.9 percent markup for UI with a coefficient of relative risk aversion of 2.

Job Loss versus Unemployment.—The analyses in columns 1–6 measure how food expenditure varies with whether the individual is employed at the time of the survey. As discussed in Section IB, one would ideally like to estimate the impact on consumption at the time of the job loss regardless of whether the individual is unemployed at the time of the next survey. To explore this, column 6 shows that food expenditure growth is 4.87 percent (SE 0.860 percent) lower relative to the previous year if a job loss has occurred relative to the previous year. This is slightly lower than the baseline estimate of 7.23 percent in column 1. This is consistent with attenuation from misalignment of the timing of job loss and consumption measurement. This suggests the baseline estimates relying on the impact of unemployment provide a conservatively high estimate of the willingness to pay for insurance against losing one's job.

Heterogeneity.—Formally, equation (8) requires comparing the willingness to pay for all p to the pooled price ratio. Therefore, it is also useful to understand not only the average willingness to pay but also the heterogeneity in the potential willingness to pay across the population. How much might some people be willing to pay for insurance? Estimating a maximum as opposed to an average is generally more difficult. Here, the problem is compounded by the fact that consumption expenditure is generally measured with error (Zeldes 1989; Meghir and Pistaferri 2011). To shed light on the potential heterogeneity in willingness to pay across the population, online Appendix C.3 develops a measurement error model that uses symmetry assumptions to provide an upper bound on the causal effect of unemployment on food expenditure, $\min_p\left\{\log\left(c_u(p)\right) - \log\left(c_e(p)\right)\right\}$, which can be used to construct a maximum willingness to pay for UI. For brevity, the details of this approach are provided in the online Appendix.

Columns 7–8 present the results. For the baseline sample in column 7, the results suggest a maximum causal impact on food expenditure is 13.7 percent, or roughly twice as large as the mean consumption drop. This rises to 14.6 percent on the broader sample that does not drop outliers with greater than a three-fold change in measured food expenditure. The lower rows in Table 5 scale these estimates by various values of risk aversion. For a conservative estimate of 3, it suggests the maximum markup that individuals would be willing to pay is less than 52.6 percent on the baseline sample and 56.6 percent in the broader sample including outliers. In short, the causal impact of job loss on consumption combined with standard risk aversion parameters suggest that the markup which individuals would be willing to pay for UI generally lies below 60 percent.

B. *Approach 2: Exploiting Ex Ante Responses as Evidence of WTP*

Inferring willingness to pay from the causal effect of job loss or unemployment on food expenditure relies heavily on an assumption of state-independence of the utility function over food expenditure.[31] This assumption can lead to an over- or understatement of the true willingness to pay. If unemployed individuals have more time to spend searching for lower-priced consumption goods or have more time to cook at home instead of needing to eat at restaurants (as shown in Aguiar and Hurst 2005), then their marginal utility of additional food expenditure when employed, $v'(c_e)$, may be equal to the marginal utility of additional food expenditure in the event of job loss, $u'(c_u)$, even if $c_e > c_u$. In this case, the causal impact will overstate the willingness to pay for UI. Conversely, individuals may derive greater value from higher expenditure when unemployed. Or the unplanned loss of a job might raise the value of financial resources to help find a new job or cope with other costs that arise. In this case, the size of the causal impact on food expenditure will understate the willingness to pay for UI.

This section presents a new strategy for identifying the willingness to pay for UI that overcomes these potential biases and allows for state dependence of the utility function. Instead of comparing consumption across states of the world, the approach exploits the ex ante behavioral response (while currently employed) to learning one might lose their job in the future. The extent to which individuals take actions in response to learning they might lose their job in the future to generate or save financial resources helps reveal their willingness to pay for UI.

Exploiting Ex Ante Consumption Responses.—To see how ex ante consumption responses can reveal willingness to pay for UI, recall the Euler equation (6):

$$v'\big(c_{pre}(p)\big) = pu'\big(c_u(p)\big) + (1-p)\,v'\big(c_e(p)\big).$$

Individuals who learn today that they will lose their job in the next year will equate their marginal utility of consumption today to the marginal utility of consumption when losing their job, $v'\big(c_{pre}(1)\big) = u'\big(c_u(1)\big)$. Conversely, those who learn today that they won't lose their job in the next year will have a current marginal utility of consumption equal to the marginal utility of consumption when employed next year, $v'\big(c_{pre}(0)\big) = v'\big(c_e(0)\big)$. So, if $c_e(p)$ and $c_u(p)$ do not systematically vary with p, then $v'\big(c_{pre}(0)\big) - v'\big(c_{pre}(1)\big) = u'(c_u) - v'(c_e)$. Taking a Taylor approximation of $v'\big(c_{pre}(p)\big)$ around $\overline{c}_{pre}$ and dividing by $v'(\overline{c}_{pre})$ yields

[31] This is often stated as two distinct assumptions: (i) state independence over consumption ($u(c) = v(c)$ for all c), and (ii) food expenditure as a valid proxy for consumption. But it should be clear that all that matters is how well food expenditure and risk aversion are able to provide a proxy for the marginal utility of consumption.

$\frac{v'(c_{pre}(0)) - v'(c_{pre}(1))}{v'(\overline{c}_{pre})} \approx \sigma \frac{d \log(c_{pre}(p))}{dp}$, where $\sigma = \frac{v''(\overline{c}_{pre})\, \overline{c}_{pre}}{v'(\overline{c}_{pre})}$ is the coefficient of relative risk aversion over consumption (within the employed state of the world) and $\frac{d \log(c_{pre}(p))}{dp}$ is the ex ante response of consumption to learning future job loss is more likely. Therefore, the difference in marginal utilities, $\frac{u'(c_u) - v'(c_e)}{v'(c_e)}$, can be inferred from the size of the response of consumption to an increase in the likelihood of job loss multiplied by the coefficient of relative risk aversion within the state of being employed. Proposition 2 formalizes this result.

PROPOSITION 2: *Suppose* (*i*) *the Euler equation holds*; (*ii*) c_e *and* c_u *do not vary systematically with* p, $\frac{dc_e}{dp} = 0$ *and* $\frac{dc_u}{dp} = 0$; *and* (*iii*) *the coefficient of relative risk aversion in the ex ante employed state is constant*, $\sigma = \frac{-v''(c_{pre}(p))}{v'(c_{pre})} c_{pre}(p)$. *Then*,

$$\frac{u'(c_u) - v'(c_e)}{v'(c_e)} = \frac{\sigma}{\kappa} E\left[\frac{-d \log(c_{pre}(p))}{dp}\right] \tag{13}$$

where $\kappa = E\left[\frac{1}{1 + p\frac{u'(c_u(p)) - u'(c_e(p))}{u'(c_e)}}\right] \approx 1$ *and* $E\left[\frac{-d \log(c_{pre}(p))}{dp}\right]$ *is the average relationship between consumption today*, c_{pre}, *and beliefs about future employment*, p.

PROOF:

See online Appendix C.4.

Proposition 2 shows that one can identify the markup that individuals are willing to pay for UI by scaling the impact of a change in beliefs about future unemployment on consumption today by the coefficient of relative risk aversion over current consumption. Because the consumption response is within the state of being employed, it allows for state dependence (i.e., $u(c) \neq v(c)$). But the ability to allow for state dependence does not come without additional assumptions. In particular, one needs to assume that the levels of future consumption conditional on U not be systematically correlated with beliefs, p: $\frac{dc_e}{dp} = \frac{dc_u}{dp} = 0$. This ensures that the response of c_{pre} to an increase in p is because of the different marginal utilities, u' versus v'. This could be violated if individuals who learn they might lose their job today also tend to have their wages reduced even if they don't end up losing their job. In this case, they would lower their consumption today in response to an increase in p $\left(\text{because } \frac{dc_e}{dp} < 0\right)$. But this would not reflect the value of UI (that moving resources from $v'(c_e)$ to $u'(c_u)$) but rather a desire for more financial resources even if they remain employed. To see how this would lead to an overstatement of the true demand for UI, note that if $\frac{dc_e}{dp} \neq 0$, then differentiating the Euler equation yields

$$\frac{dc_{pre}}{dp} v''(c_{pre}(p)) = u'(c_u(p)) - v'(c_e(p)) + (1-p)\, v''(c_e(p)) \frac{dc_e}{dp}, \tag{14}$$

where the term $(1-p)\,v''(c_e(p))\,\frac{dc_e}{dp}$ reflects an additional reason to save in response to an increase in p. In this case, the ex ante method would overstate the value of UI. In this sense, it generates a conservative estimate for comparing to the pooled price ratio in equation (8). I return to a discussion of this potential bias in Section IVC.

Two-Sample Implementation.—To estimate how $\log(c_{pre}(p))$ varies with p in equation (13), I follow a methodology similar to Section IVA but using ex ante responses. I divide the amount food expenditure changes between $t-2$ and $t-1$, Δ_{-1}^{FD} in equation (4), by the amount of information revealed between $t-2$ and $t-1$. Mathematically,

$$(15)\quad \frac{d\log(c_{pre})}{dp} \approx \frac{\Delta_{-1}^{FD}}{E[P_{t,t-1}-P_{t,t-2}\,|\,U_t = 1] - E[P_{t,t-1}-P_{t,t-2}\,|\,U_t = 0]},$$

where $P_{j,t}$ are the beliefs in period t about job loss in period j. The denominator is the fraction of information which individuals learn about U_t between year $t-2$ and $t-1$.[32] Inflating Δ_{-1}^{FD} by the fraction of information revealed over the time that the first difference is estimated ($t-2$ to $t-1$) yields an estimate of the impact of learning about future job loss on ex ante food expenditure, $c_{pre}(p)$.

Table 4 reported that food expenditure growth is $\Delta_{-1}^{FD} = 2.71$ percent lower in the year prior to the unemployment measurement. To estimate the denominator in equation (15), recall that the average difference in beliefs one year prior to the job loss measure between those who do and do not lose their job is $E[P_{t,t-1}\,|\,U_t = 1] - E[P_{t,t-1}\,|\,U_t = 0] = 19.7$ percent (online Appendix Table I, column 1). In this sense, 20 percent of the information about job loss in year t is already known in year $t-1$. So, one needs to know how much is known in year $t-2$, $E[P_{t-2,t}\,|\,U_t = 1] - E[P_{t-2,t}\,|\,U_t = 0]$. To obtain this, I regress the elicitation, Z, on an indicator for losing one's job in the subsequent 12–24 months after the elicitation. The second row of online Appendix Table I reports this value as $E[P_{t-2,t}\,|\,U_t = 1] - E[P_{t-2,t}\,|\,U_t = 0] = 9.4$ percent.[33] The difference of 10.3 percent is an estimate of the fraction of information about job loss in year t revealed between $t-2$ and $t-1$.

Dividing Δ_{-1}^{FD} by 0.103 yields an estimate of $\frac{d\log(c_{pre})}{dp} = 0.29$ for the baseline specification in column 1 of Table 6; alternative specifications yield similar estimates for $\frac{d\log(c_{pre})}{dp}$ of around 20–30 percent. Scaling these by a coefficient of relative risk aversion of $\sigma = 2$ suggests that individuals are willing to pay around a 50–60 percent markup for unemployment insurance. For higher values of risk aversion ($\sigma = 3$) this increases to 70–90 percent; for lower values ($\sigma = 1$) this decreases to 20–30 percent.

[32] This can also be written as $E[P_{t,t-1}-P_{t,t-2}\,|\,U_t = 1] - E[P_{t,t-1}-P_{t,t-2}\,|\,U_t = 0] = \frac{\mathrm{var}(P_{t,t-1}) - \mathrm{var}(P_{t,t-2})}{\mathrm{var}(U)}$.

[33] Online Appendix Figure I also reports the coefficients for future years of unemployment and obtains estimates of $E[Z_{t-j,t}\,|\,U_t = 1] - E[Z_{t-j,t}\,|\,U_t = 0]$ ranging from 0.1 to 0.05 at $j = 8$, which suggests most of the information in Z is about unemployment in the subsequent year. This is consistent with a relatively flat consumption growth profile for years prior to $t-2$ as shown in Figure 4.

TABLE 6—EX ANTE DROP IN FOOD EXPENDITURE PRIOR TO UNEMPLOYMENT AND IMPLIED (EX ANTE) WILLINGNESS TO PAY FOR UI

	Baseline (1)	Controls for needs (2)	Under-50 sample (3)	Job loss (4)	Household income controls (5)	Household head income controls (6)
Reduced-form impact on $\log(c_{t-2}) - \log(c_{t-1})$						
Unemployment	−0.0271	−0.0211	−0.0288	−0.0260	−0.0272	−0.0281
Standard error	(0.00975)	(0.0105)	(0.0106)	(0.00824)	(0.00969)	(0.00983)
Two-sample IV estimate for $d[\log(c_{pre}(p))]/dp$	−0.29	−0.23	−0.31	−0.30	−0.29	−0.29
Implied $[u'(c_u) - v'(c_e)]/v'(c_e)$ *for various* σ *(percent)*						
$\sigma = 1$	29	23	31	30	29	29
$\sigma = 2$	58	46	62	60	58	59
$\sigma = 3$	87	70	92	90	87	88
Specification details						
Sample employed in $t-2$ and $t-1$	X	X	X	X	X	X
Controls for change in log needs ($t-2$ versus $t-1$)		X				
Change in log HH inc ($t-2$ versus $t-1$) (3rd-order poly)					X	
Change in log HH head inc ($t-2$ versus $t-1$) (3rd-order poly)						X
Mean dependent variable	0.000	−0.001	0.007	0.000	0.000	0.000
Observations	65,483	53,327	52,463	65,556	65,399	64,119
Households	9,557	8,371	8,512	9,560	9,547	9,448

Notes: This table presents the implications of the estimates from Table 4 for individuals' valuation of UI. The first row replicates the first row in Table 4. The second set of rows divide the ex ante first difference estimate by the amount of information revealed between year $t-2$ and $t-1$ for those that do versus do not experience job loss in period t. This is computed as the difference in the coefficient from a regression of the elicitation, Z, on subsequent unemployment in the next year, U, and the coefficient from a regression of Z on an indicator for unemployment in the 12 to 24 months after the survey. Online Appendix Table IV provides the baseline regression results for this first-stage calculation. The third set of rows scales these responses by the coefficient of relative risk aversion to arrive at an estimate of the willingness to pay for UI.

Exploiting Spousal Labor Supply Responses.—Not only does food expenditure respond to learning about future job loss, but spouses are also more likely to enter the labor market. Analogous to scaling food expenditure responses by a coefficient of relative risk aversion, one can also scale the spousal labor supply responses by the semi-elasticity of spousal labor supply to arrive at an alternative measure for the willingness to pay for UI. This alternative measurement has the added benefit that it can be constructed for the HRS sample, which corresponds to the sample used to measure the pooled price ratio, $T(p)$, in Section V. To do so, online Appendix C.5 provides conditions analogous to those in Proposition 1 that enable the willingness to pay for UI in equation (13) to be written as

$$\frac{u'(c_u) - v'(c_e)}{v'(c_e)} \approx \frac{E\left[\frac{dLFP}{dp}\right]}{\epsilon^{semi}}, \tag{16}$$

TABLE 7—SPOUSAL LABOR SUPPLY RESPONSE TO POTENTIAL JOB LOSS AND IMPLIED (EX ANTE) WILLINGNESS TO PAY FOR UI

	Baseline (1)	Sample without future job loss (2)	Full-time work (3)	2-yr lagged entry (placebo) (4)	Household fixed effects (5)	Individual fixed effects (6)
Reduced-form relationship between U and Z						
Elicitation (Z)	0.0258	0.0256	0.0255	0.0012	0.0243	0.0312
Standard error	(0.00868)	(0.00898)	(0.00988)	(0.00800)	(0.0114)	(0.0180)
Measurement error correction						
Total var/signal var ($\mathrm{var}(Z\|X)/\mathrm{cov}(Z,U\|X)$)	12.08	12.08	12.08	23.68	12.08	12.08
bootstrap SE	(1.71)	(1.69)	(1.65)	(6.53)	(1.62)	(1.60)
Implied $dLFP/dp$	0.312	0.309	0.308	0.029	0.293	0.377
Implied $[u'(c_u) - v'(c_e)]/v'(c_e)$ *for various* ϵ^{semi} *(percent)*						
$\epsilon^{semi} = 0.33$	94	93	92	9	88	113
$\epsilon^{semi} = 0.5$	62	62	62	6	59	75
$\epsilon^{semi} = 1$	31	31	31	3	29	38
Mean dependent variable	0.04	0.04	0.05	0.04	0.04	0.04
Observations	11,049	10,726	11,049	11,049	11,049	11,049
Households	2,214	2,194	2,214	2,214	2,214	2,214

Notes: This table scales the regression coefficients in Table 3 to arrive at an estimate of the willingness to pay for UI. The second set of rows presents the total variance of Z relative to the signal variance $(\mathrm{var}(P))$. I estimate the variance of Z given X by regressing Z on the control variables and squaring the RMSE. I estimate the variance of P given X as follows. I regress the future unemployment indicator, U, on the controls and take the residuals. I regress Z on the controls and take those residuals. I then construct the covariance between these two residuals and rescale by $(n-1)/(n-df)$, where df is the number of degrees of freedom in the regression of U on the controls. This provides an estimate of $\mathrm{cov}(Z,L|X)$, which is an approximation to $\mathrm{var}(P|X)$ that is exact under classical measurement error. Standard errors are constructed using 500 bootstrap repetitions, resampling at the household level. The implied impact of p on LFP is constructed by taking the regression coefficient and multiplying by the total variance/signal variance. The next set of rows divides by various assumptions for the semi-elasticity of spousal labor force participation with respect to wages, which yields a willingness to pay for UI.

where $E\left[\frac{dLFP}{dp}\right]$ is the average response of the female labor participation rate to an increase in beliefs, p, and $\epsilon^{semi} = \frac{dLFP}{d\log(w)}$ is the response of the female labor participation rate to a 1 percent increase in wages. Comparing the response of labor supply to the size of the response to a wage increase reveals individuals' implicit valuation of UI. If spousal labor supply is very inelastic, then the finding that many spouses enter the labor market in response to an increase in p suggests they have a higher desire for additional financial resources in the event of losing their job.[34]

[34] In addition to using spousal labor supply for the ex ante method for valuing UI, one could also implement equation (16) using the contemporaneous impact of job loss on spousal labor supply, analogous to the approach in Section IVA. Here, Stephens (2002) finds evidence that spouses increase their labor supply prior to the onset of job loss, but also finds that the impact of job loss in year t on spousal labor supply in year t is considerably more muted relative to the sharp drop in consumption that is observed at the onset of job loss or unemployment. There are two potential interpretations: on the one hand, it could be the case that the willingness to pay is not that high and the large consumption drop upon unemployment reflects state-dependent utility. In this case, the willingness to pay estimates that compare consumption between employed and unemployed will overstate the true willingness to pay. On the other hand, the onset of the job loss could reflect a particularly strong correlated shock to the spouse's job opportunities. In this case, the lack of a spousal labor supply response could indicate the lack of work opportunities as opposed to the lack of desire for additional income.

For the empirical analysis, I use a baseline value of $\epsilon^{semi} = 0.5$, following Kleven, Kreiner, and Saez (2009). I also consider a range of estimates between 0.33 and 1, loosely consistent with the range of estimates found in Blundell et al. (2016).

The first row of Table 7 reports the coefficients, $\beta = \frac{dLFP}{dZ}$, from equation (3) of around 0.025 that were presented in Table 3. Because the elicitations, Z, are noisy measures of true beliefs, P, this slope is likely an attenuated measure of the relationship between true beliefs and spousal labor entry, $\frac{dLFP}{dp}$. To correct for this, I make a couple of assumptions on the distribution of measurement error and beliefs. In particular, I assume that (i) the noise in the elicitations is classical (i.e., $Z - P$ is uncorrelated with P), and (ii) that true beliefs are unbiased ($\Pr(U|P) = P$). The classical measurement error assumption implies that the attenuation will be equal to the ratio of total variance to signal variance, $\frac{\mathrm{var}(Z|X)}{\mathrm{var}(P|X)}$. The unbiasedness of true beliefs implies that $\mathrm{cov}(Z, U|X) = \mathrm{cov}(P, U|X) = \mathrm{var}(P|X)$. Therefore, the true relationship between beliefs and LFP is given by multiplying the coefficient β in equation (3) by the ratio of total to signal variance $\frac{\mathrm{var}(Z|X)}{\mathrm{var}(P|X)} = \frac{\mathrm{var}(Z|X)}{\mathrm{cov}(Z, U|X)}$:

$$\frac{dLFP}{dp} = \frac{\mathrm{var}(Z|X)}{\mathrm{cov}(Z, U|X)}\beta. \tag{17}$$

I estimate $\mathrm{var}(Z|X)$ as the mean square error of a regression of Z on X. I estimate the covariance between Z and U as the covariance between residuals of regressions of Z on X and U on X.[35] This yields an estimate of $\frac{\mathrm{var}(Z|X)}{\mathrm{cov}(Z, U|X)} \approx 12$, as reported in the second set of rows in Table 7.

Multiplying $\beta = 0.025$ by this factor of 12 yields an estimate of $\frac{dLFP}{dp}$ of around 0.3. This suggests that a 10pp increase in the true beliefs, p, will increase LFP by 3pp. Scaling by the semi-elasticity of labor supply, ϵ^{semi}, of 0.5, suggests that individuals would be willing to pay roughly a 60 percent markup for UI. As shown in Table 8, if labor supply is more elastic (e.g., $\epsilon^{semi} = 1$), it suggests a willingness to pay of around 30 percent; if labor supply is less elastic (e.g., $\epsilon^{semi} = 0.33$), it suggests a willingness to pay of around 90 percent.

C. *Discussion*

The valuations exploiting ex ante responses are generally higher than the ex post methods. This could be for several reasons. On the one hand, there could be a violation of state independence of the utility function ($u \neq v$) so that individuals have a higher desire for income after job loss than is suggested by their drop in consumption. In this case, the ex ante methods provide a more accurate measure of the value of UI. On the other hand, individuals who learn they might lose their job might also learn that they have lower earnings prospects even if they do not lose their job. This

[35] I adjust for the degrees of freedom used to estimate the coefficients on X in those regressions.

would suggest consumption in the employed state, c_e, declines in response to an increase in p, which would violate Assumption (iii) in Proposition 2. In this case, one would expect the ex ante methods to overstate the willingness to pay for UI. For the purposes of understanding whether a private market can exist, this potential upward bias in the willingness to pay for UI only makes it more difficult for private information to explain the absence of a private market.

Going forward, the next section will compare both sets of estimates to the pooled price ratio to understand whether this willingness to pay is high enough so that individuals would be willing to pay the pooled cost of worse risks to obtain insurance.

V. Implications for the Pooled Price Ratio

How much of a markup must individuals be willing to pay to cover the pooled cost of worse risks adversely selecting their insurance contract? This section builds on Hendren (2013) by providing two approaches for measuring the pooled price ratio, $T(p)$.

Ideally, one would construct $T(p)$ using the distribution of true beliefs, P, to compute $E[P|P \geq p]$ for each p and form $T(p) = \frac{E[P|P \geq p]}{1 - E[P|P \geq p]} \frac{1-p}{p}$ for each p. Then, one could compare the willingness to pay to the pooled price ratio for each p, as suggested by equation (8). In particular, if insurers know the distribution of P, they would be able to set prices to select an insurance pool that generates the smallest value of $T(p)$, $\inf\ T(p)$. To that aim, Section VB will provide a point estimate of the minimum pooled price ratio using a set of semiparametric assumptions.

Before imposing the semiparametric assumptions, Section VA will provide a lower bound on the average pooled price ratio, $E[T(P)]$, that relies on weaker assumptions. While the minimum pooled price ratio is the relevant statistic if insurers know the distribution of P, online Appendix B.1 shows that the average pooled price ratio provides information about the frictions from adverse selection if insurers do not know the distribution of P. In particular, if insurers enter the market by setting prices in a random fashion (that does not target the type p with the lowest pooled price ratio), then online Appendix B.1 shows that $E[T(P)]$ (as opposed to $\inf T(p)$) will characterize whether firms can profitably sell insurance. In this sense, the nonparametric lower bounds on $E[T(P)]$ and semiparametric point estimates for $\inf T(p)$ will provide complementary evidence on the frictions that would be imposed by private information.

A. *Nonparametric Lower Bounds for the Average Pooled Price Ratio*, $E[T(P)]$

I begin by using the predictive power of the elicitations to form lower bounds on the average pooled price ratio, $E[T(P)]$. Let P_Z denote the predicted values from a regression of U on the elicitations, Z, controlling for X,

$$P_Z = \Pr(U|X, Z).$$

Hendren (2013) shows that the distribution of P_Z is related to the true distribution of beliefs, P, under two assumptions: (i) the elicitations contain no more information about U than does P: $\Pr(U|X,Z,P) = \Pr(U|X,P)$; and (ii) true beliefs are unbiased $\Pr\{U|X,P\} = P$. Assumption (i) is a natural assumption to place on the elicitations, as it is difficult to imagine how someone could report more information than what is contained in their true beliefs. Assumption (ii) is more restrictive, as individuals may have biased beliefs. I return to a discussion of the impact of biased beliefs in Section VI. Under Assumptions (i) and (ii), the true beliefs are a mean-preserving spread of the distribution of predicted values,

$$E[P|X,Z] = P_Z.$$

To construct P_Z, I run a probit of U on X and a third-order polynomial in Z along with indicators for $Z = 1$, $Z = 0.5$, and $Z = 0$. The predicted values yield P_Z. I repeat this probit omitting the Z variables to form $\Pr(U|X)$. Figure 5 then illustrates the predictive content of the elicitations by plotting the cumulative distribution of $P_Z - \Pr(U|X)$. The figure reveals how the information in the elicitations, Z, generate a significant amount of dispersion in the distribution of P_Z. The logic of how a private market may unravel can be seen in Figure 5: would any individual with risk p be willing to pay the pooled cost of those with higher odds of losing their job? The thick upper tail of the distribution suggests insurers would face difficulty from those higher risks aversely selecting an insurance contract.

To quantify these frictions, ideally one would measure the pooled price ratio using the distribution of true beliefs, P, not the predicted values, P_Z. However, it turns out one can use the distribution of P_Z to construct a lower bound on the average pooled price ratio, $E[T(P)]$. For each p, let $m(p) = E[P_Z - p | P_Z \geq p]$ denote the average extent to which the predicted values, P_Z, lie above p. Next, normalize this by the probability of job loss in the population, $\Pr(U)$,

$$T_Z(p) = 1 + \frac{m(p)}{\Pr(U)}.$$

Drawing p from the distribution of predicted values, P_Z, online Appendix B.1 shows that $E[T_Z(P_Z)]$ forms a lower bound on the average pooled price ratio,[36]

$$E[T_Z(P_Z)] \leq E[T(P)]. \tag{18}$$

Equation (18) shows that the predictive content of the elicitations for U form a nonparametric lower bound for the average pooled price ratio. This lower bound, $E[T_Z(P_Z)]$, only requires that the elicitations contain no more information than the true beliefs (Assumption (i)) and that true beliefs are unbiased (Assumption (ii)). To emphasize the weak nature of these assumptions, note that they do not require Z to be a number, nor are they affected by any one-to-one transformation of Z. All that matters for generating the lower bounds is how well Z predicts U, conditional on X.

[36] This lower bound builds upon but is distinct from the results in Hendren (2013). Hendren (2013) shows that $E[m(P_Z)] \leq E[m(P)]$ but does not provide a lower bound on $E[T(P)]$.

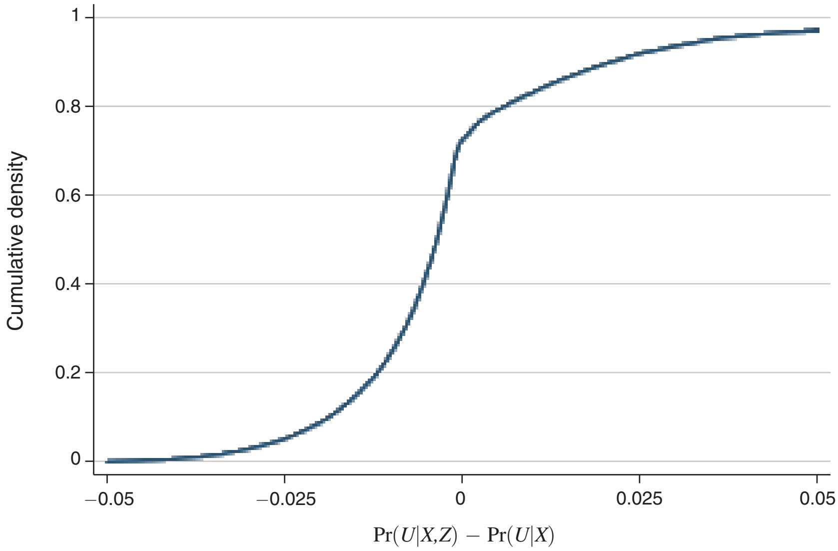

Figure 5. Cumulative Distribution of $\Pr(U|Z,X) - \Pr(U|X)$

Notes: This figure presents the cumulative distribution of $\Pr(U|X,Z) - \Pr(U|X)$. To construct $\Pr(U|X,Z)$, I use the baseline sample in the HRS and regress job loss in the next 12 months, U, on both the observable characteristics (year dummies, demographics, and job characteristics) and the elicitations (a cubic in Z combined with indicators for $Z = 0$, $Z = 0.5$, $Z = 1$ to capture focal point responses). I use a probit specification and construct the predicted values to form an estimate of $\Pr(U|X,Z)$ for each observation in the baseline sample. For $\Pr(U|X)$, I repeat this process but exclude the elicitation variables. These predicted values provide an estimate of $\Pr(U|X)$. I then construct the difference, $\Pr(U|X,Z) - \Pr(U|X)$ for each observation, and plot its cumulative distribution.

Results.—For the baseline demographic and job characteristics controls, the average markup imposed by the presence of worse risks is at least 77 percent (SE 5.2 percent), suggesting $E[T(P)] \geq 1.77$.[37] The top row of Table 8 presents these results and shows that adding health controls or dropping the job characteristic controls do not meaningfully change the estimates (72 percent and 80 percent in columns 2 and 3). To further illustrate how the controls affect the results, panel A of Figure 6 plots estimates of $E[T_Z(P_Z)] - 1$ on the vertical axis against the pseudo R^2 of the model for $\Pr(U|X,Z)$ for specifications with different controls, X. Additional job characteristics help predict job loss, but they do not reduce the average pooled price ratio. This is because the magnitude of $E[P_Z - p | P_Z \geq p]$ depends on the thickness of the upper tail of the distribution of predicted values. This upper tail (shown in Figure 5)

[37] As in Hendren (2013), the construction of $E[T_Z(P_Z)]$ and $E[m_Z(P_Z)]$ is all performed by conditioning on X. To partial out the predictive content in the observable characteristics, I first construct the distribution of residuals, $P_Z - \Pr(U|X)$. I then construct $m_Z(p) = E[P_Z - p | P_Z \geq p]$ for each value of X as the average value of $P_Z - \Pr(U|X)$ above $p + \Pr(U|X)$ for those with observable characteristics X. In principle, one could estimate this separately for each X, but this would require observing a rich set of observations with different values of Z for that given X. In practice, I follow Hendren (2013) and specify a partition of the space of observables, ζ_j, for which I assume the distribution of $P_Z - \Pr(U|X)$ is the same for all $X \in \zeta_j$. This allows the mean of P_Z to vary richly with X, but allows a more precise estimate of the shape by aggregating across values of $X \in \zeta_j$. In principle, one could choose the finest partition, $\zeta_j = \{X_j\}$ for all possible values of $X = X_j$. However, there is insufficient statistical power to identify the entire distribution of P_Z at each specific value of X. For the baseline specification, I use an aggregation partition of five-year age bins by gender. Online Appendix Table III (columns 3–5) documents the robustness of the results to alternative aggregation partitions.

Table 8—Pooled Price Ratio

		Alternative controls			Subsamples					
	Baseline (1)	Demo (2)	Health (3)	Ind FE[1] (4)	Age ≤ 55 (5)	Age > 55 (6)	Below median wage (7)	Above median wage (8)	Tenure ≥ 5 yrs (9)	Tenure < 5 yrs (10)
Lower bound (in percent) for $E[T(P)] - 1$										
$E[T_Z(P_Z) - 1]$	77	80	72	39	70	81	63	100	110	47
Standard error	(5.2)	(4.9)	(5.4)	(4.2)	(9)	(7)	(5.8)	(9.3)	(9.4)	(5.7)
Point estimate (in percent) for $\inf T(p) - 1$										
$\inf T(p) - 1$	336	530	323	N/A[1]	333	344	436	316	474	374
Standard error	(20.3)	(65.5)	(26.8)		(30.6)	(27.9)	(42.2)	(25.6)	(39.2)	(33.6)
$\Pr(U = 1)$	0.0307	0.0307	0.0317	0.0307	0.0297	0.0314	0.0369	0.0245	0.0175	0.0575
Controls (X)										
Year fixed effects	X	X	X	X	X	X	X	X	X	X
Demographics	X	X	X	X	X	X	X	X	X	X
Job characteristics	X		X	X	X	X	X	X	X	X
Health			X							
Individual fixed effects				X						
Observations	26,640	26,640	22,831	26,640	11,134	15,506	13,321	13,319	17,850	8,790
Households	3,467	3,467	3,180	3,467	2,255	3,231	2,824	2,309	2,952	2,437

Notes: This table presents estimates of the nonparametric lower bounds on $E[T(P)]$ and semiparametric point estimates for the minimum pooled price ratio, $\inf[T(p)]$. Each column replicates the specifications in Table 2. All standard errors for $\inf[T_Z(P_Z)]$ and $E[m_Z(P_Z)]$ are constructed using 1,000 bootstrap resamples at the household level for the minimum pooled price ratio and 500 resamples for the lower bound estimate of $E[T(P)]$.

[1] $\inf T(p)$ specification not available for the FE specification because nonlinear fixed effects cannot be partialed out of the likelihood function. The lower bounds on $E[T(P)]$ are constructed using a linear probability model. This model generates some predicted values outside of [0,1], which would induce a nonsensical likelihood in the parametric model used to estimate the point estimate for the distribution of P.

is not removed by changing the set of controls, X.[38] Indeed, even if an insurer could (unrealistically) use individual fixed effects to price insurance, individuals would still on average have to be willing to pay at least a 40 percent markup to cover the pooled cost of worse risks, as shown in column 4 of Table 8.

The remaining columns of Table 8 and panels of Figure 6 explore how the estimated markups vary across subsamples. This yields estimates of $E\left[T_Z(P_Z)\right] - 1$ in excess of 50 percent across occupations (panel B of Figure 6), ages (panel C), and years (panel D). Adverse selection would impose significant barriers to a private market across a wide range of subsamples of the population.

A common strategy which insurers use to mitigate adverse selection is to impose waiting periods on the use of the insurance policy. To explore whether this helps reduce the adverse selection problem, Figure 6 (panel E) presents results for a specification that replaces the baseline definition of U with an indicator for losing one's job in the 6–12 months after the elicitation; this simulates a requirement of a 6-month waiting period.[39] This generates a smaller but still significant lower bound of 57.9 percent ($p < 0.001$). The estimates also remain high for other potential timelines and definitions of U, such as 0–24 and 6–24 month payout windows, and a requirement

[38] Online Appendix Table III explores robustness to various specifications, including linear versus probit error structures, alternative aggregation windows for constructing $E\left[m_Z(P_Z)\right]$, and alternative polynomials for Z. All estimates are quite similar to the baseline and yield lower bounds of $E\left[T_Z(P_Z)\right] - 1$ of around 70 percent.

[39] The lower bound result in equation (18) does not require that Z be an elicitation that perfectly corresponds to an event, U. Therefore the lower bound remains valid even when Z is elicited about future job loss in the next 12 months but U is defined as an indicator of job loss in the 6–12 months after the survey.

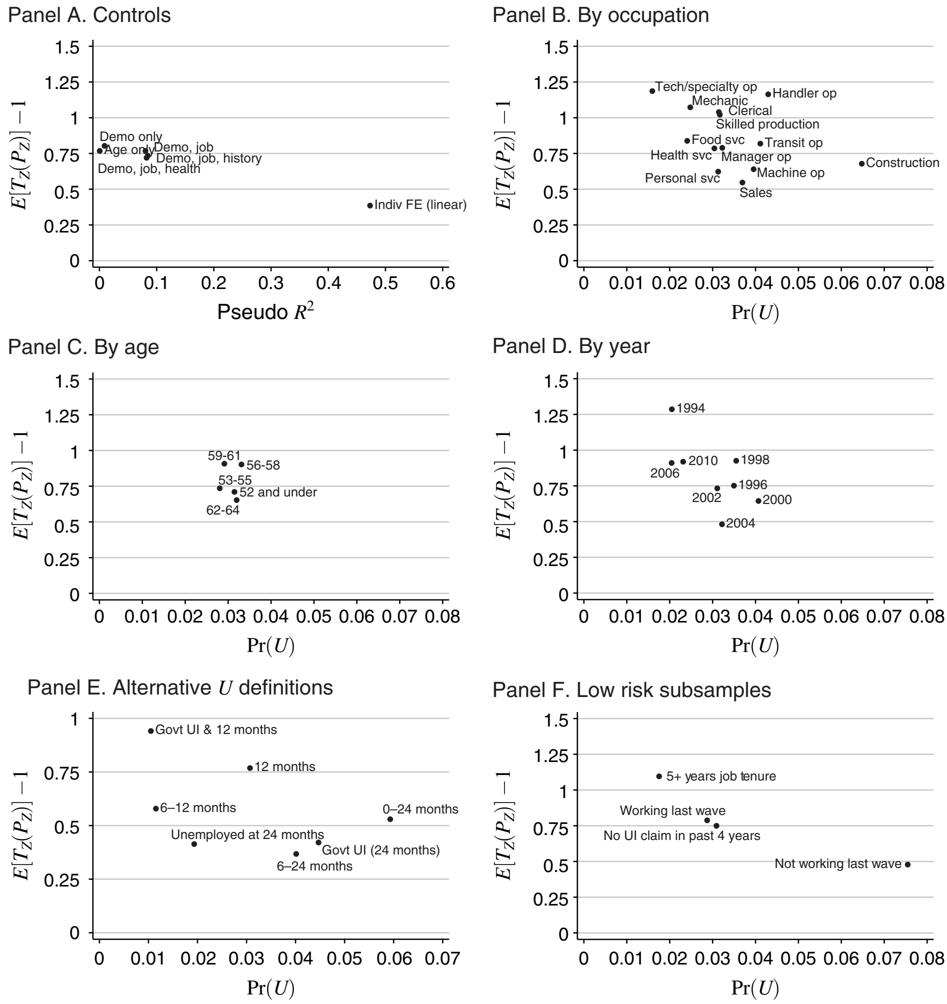

FIGURE 6. LOWER BOUNDS FOR AVERAGE POOLED PRICE RATIO, $E[T(P)] - 1$

Notes: These figures present estimates of the lower bounds on the average pooled price ratio, $E[T_Z(P_Z)]$, using a range of subsamples and controls. Panel A reports estimates of $E[T_Z(P_Z)]$ for a range of control variables, including a specification with individual fixed effects. All specifications use a probit link specification for $\Pr(U \mid X, Z)$ except for the fixed effects specification in panel A, which uses a linear specification because of the presence of fixed effects. The horizontal axis presents the pseudo R^2 of the specification for $\Pr(U \mid X, Z)$. Panel B constructs separate estimates by occupation. Panel C constructs estimates by age group. Panel D constructs separate estimates for each wave of the survey. Panel E reports specifications for alternative definitions of U. These include whether the individual loses her job in the subsequent 6–12 months, 6–24 months, or 0–24 months (instead of 0–12 months). Panel F restricts the sample to varying subsamples, analyzing the relationship between $E[T_Z(P_Z)]$ and restrictions to lower risk subsamples. The horizontal axis in panels B–F report the mean outcome, $\Pr(U)$, for each subsample.

that an individual also files for government UI.[40] Waiting periods and alternative contract lengths would not remove the threat of adverse selection.

[40] This calculation also assumes that individuals are unable to retime their job loss. If individuals can costlessly retime their job loss, this would impose an added cost on the insurer. See Cabral (2013) for an example of this behavior in dental insurance.

Another potential UI policy would pay benefits proportional to the amount of time that the individual spends unemployed. Although duration is not perfectly observed in the HRS, a potential proxy for unemployment duration and severity is whether individuals are unemployed and looking for work in the next survey wave (24 months later). Figure 6 (panel E) shows that this contract would also suffer significant adverse selection, with the lower bound on the average pooled price ratio above 40 percent. Requiring individuals to also file for government UI would also not remove the barriers imposed by private information.

Another underwriting strategy that is common in health-related insurance markets is to sell insurance only to observably low risks. For example, health-related insurance markets generally exclude those with preexisting conditions (Hendren 2013). One potentially analogous strategy in UI would be to sell only to those with long job tenures and steady work histories. Figure 6 (panel F) presents the lower bound estimates on these subsamples. In contrast to the idea that restricting to low risks would help open up an insurance market, if anything the opposite is true: lower risk populations have *higher* values of $E\left[T_Z(P_Z)\right] - 1$. For those with greater than five years of job tenure, I estimate a lower bound of 110 percent. It is true that this population is much less likely to lose their job (less than 2 percent lose their job in the subsequent 12 months). But there is still a presence of some privately known higher risks who impose an especially high cost on the pooled price ratio faced by most of this population.

Overall, the results suggest that common underwriting strategies like imposing waiting periods and restricting insurance to observably low risks will not mitigate the adverse selection problem in a private UI market. Moreover, the size of these lower bounds is generally similar to or larger than the willingness to pay estimates in Section IV.

B. *Semiparametric Point Estimates of* inf $T(p)$

Moving from lower bounds on $E\left[T(P)\right]$ to a point estimate of the minimum pooled price ratio, $\inf T(p)$, requires an estimate of the distribution of beliefs, P. To obtain this, I follow Hendren (2013) by making additional assumptions about the distribution of measurement error in the elicitations. Note that the observed density (probability density function (PDF)/probability mass function (PMF)) of Z and U can be written as[41]

$$f_{Z,U|X}(Z,U|X) = \int_0^1 p^U (1-p)^{1-U} f_{Z|P,X}(Z|P=p,X) f_P(p|X)\, dp, \tag{19}$$

where $f_{Z|P,X}$ is the distribution of elicitations given true beliefs (i.e., elicitation error) and f_P is the distribution of true beliefs in the population. The goal is to use the observed joint distribution of the elicitations and the event conditional on X, $f_{Z,U|X}$, to estimate the true distribution of beliefs, f_P. This can then be used to construct $T(p)$ at each p.

[41] This is obtained by first taking the conditional expectation with respect to p and then using the assumption that $\Pr(U|Z,X,P) = P$.

Without additional assumptions, f_P is not identified from $f_{Z,U|X}$ in equation (19) because the dimensionality of $f_{Z|P,X}$ is too high. Placing parametric structure on the distribution of elicitations given beliefs, $f_{Z|P,X}$, reduces its dimensionality and allows one to identify the distribution of beliefs, $f_P(p|X)$, from $f_{Z,U|X}$. To parameterize $f_{Z|P,X}$, I follow Hendren (2013) and assume elicitations are equal to beliefs plus a noise term, $Z = P + \epsilon$, where ϵ is drawn from a mixture of a censored normal and ordered probit distribution, where the ordered probit captures the excess mass at focal point values of 0, 50, and 100. Because the mechanics of using this approach follow closely to Hendren (2013), I relegate further estimation details to online Appendix B.2.

Results.—The baseline specification yields an estimate of $\inf T(p) - 1$ of 336 percent (SE 20.3 percent), as shown in the second set of rows of Table 8. Including health controls reduces this markup slightly to 323 percent (SE 26.8 percent), and using only demographic controls increases the markup to 530 percent (SE 65.5 percent).[42] The high minimum pooled price ratios are robust across subsamples, as illustrated in columns 4–9 of Table 8. For example, those with longer job tenure have values ofinf $T(p) - 1$ of 474 percent. The minimum pooled price ratio is similar across age groups (333 percent for ages at or below 55 and 344 percent for ages above 55); and it is slightly higher for below-median wage earners (436 percent) than above-median wage earners (316 percent).

Overall, the point estimates far exceed the estimated markups that individuals are willing to pay for UI documented in Section IV. Moreover, even the lower bounds in Section VA generally lie at or above the willingness to pay estimates in Section IV. In short, the results suggest that private information provides an explanation for the absence of a private UI market. If insurers were to try to sell UI, policies would be too heavily adversely selected to deliver a positive profit at any price.

VI. Discussion

The analysis above suggests private information provides an explanation for the absence of a private UI market. Here, I discuss extensions of the model and alternative theories of market nonexistence.

Government UI.—The government is a major provider of UI in the United States. The baseline analysis above considers the hypothetical market for additional UI on top of existing sources of formal and informal insurance, including government UI. The results suggest that this market for additional insurance would unravel because of adverse selection. However, one could also ask an alternative question: if the government were to lower UI benefits, would a private market arise? Answering this requires comparing the willingness to pay for UI to the pooled price ratio, where both are estimated in the counterfactual world with less provision of government UI.

[42] Online Appendix Table IV presents the raw point estimates for α_i and ξ_i. It suggests there is a small (e.g., 10 percent) subsample of the population that has a very high chance of losing their job. The presence of this upper tail drives these high estimated markups.

Using cross-state variation in UI generosity, Gruber (1997) estimates how the consumption impact of unemployment varies with the level of government benefits. Extrapolating his estimates out of sample to a world with no UI, they suggest the willingness to pay could increase by a factor of 3 (Gruber 1997, Table I). This would continue to yield willingness to pay estimates that are smaller than the estimated 300 percent markups that individuals would need to be willing to pay to overcome adverse selection. Assuming the removal of government UI does not significantly affect the distribution of P and estimates of the pooled price ratio, this suggests a private market would not arise even in the absence of government provision of UI.

Moral Hazard and Fixed Costs.—Moral hazard is another common explanation for market nonexistence. As noted in Section III, moral hazard alone cannot make it unprofitable to sell insurance. Although behavior may change when individuals obtain insurance, the behavioral response to a small amount of insurance will be small; and the impact of the small response on the cost of a small insurance policy is second-order—analogous to the logic that the deadweight loss of a tax varies with the square of the tax rate. However, moral hazard can limit the gains to trade. Combined with fixed costs of providing insurance, it could provide a rationale for the absence of a private market. But the results above show that even if firms could costlessly offer insurance policies, private information would render the market unprofitable.[43] In this sense, the presence of private information provides a singular explanation for the absence of a UI market.

Government Regulation.—Another theory of market nonexistence is overly burdensome government regulation. For example, Cochrane (1995) suggests that this is a reason one does not see markets for *reclassification risk* in health insurance markets.[44] Translating this idea to UI, perhaps individuals are willing to pay for UI, but the reason such demand is not satisfied is because the government prevents it from existing. Here, the empirical results in Sections IV and VB suggest that because of adverse selection, firms have no ability to profitably enter the market even absent regulatory hurdles. Regulatory constraints could impose additional costs on insurers, but they are not needed to explain the absence of an insurance market.

Aggregate Risk.—The baseline model in Section III assumes that the insurer is risk neutral or has access to risk-neutral (re)insurance markets. If there is aggregate risk, the cost to the insurer of transferring dollars from employed to unemployed states for a type p may be higher than the risk-neutral cost of $\frac{p}{1-p}$ because of a higher marginal cost of capital in states where people are unemployed. In the limit where unemployment is perfectly correlated across individuals and all individuals have the same willingness to pay for UI, there may be no scope for a profitable insurance market.

[43] It is straightforward to show that in a world with fixed costs of selling insurance contracts, the no-trade condition is sufficient but not necessary for the absence of a private market.

[44] This rationale is given suggestively and the paper does not provide direct evidence of an impact of government regulation.

However, the risk of losing one's job is not purely an aggregate shock. The world does not oscillate between full unemployment and full employment. To be sure, unemployment and job separation rates vary across years; but the R^2 of a regression of job loss on year dummies yields an R^2 of 0.0019. As long as the risk of job loss has an idiosyncratic component, there will be first-order gains to risk pooling and an insurance company should be able to profitably sell UI. If insurers did not want to absorb any aggregate risk, they could in principle condition their insurance contract on the aggregate unemployment rate and fully insure the residual 99.81 percent of job loss risk that is not collinear with the aggregate risk. Doing so would shield the insurance company from aggregate risk. Because my baseline analysis includes time dummies in all regressions, my results suggest that this market for insuring the idiosyncratic component of the risk would be too heavily adversely selected to deliver a positive profit.[45] In this sense, aggregate risk does not readily provide an explanation for the absence of a UI market.

Biased Beliefs.—The baseline model above assumes individuals have accurate beliefs. This is at odds with literature suggesting that individuals may have biased beliefs about their unemployment and job prospects (e.g., Stephens 2004; Spinnewijn 2015).

To extend the model in Section III to include biased beliefs, suppose p is the objective belief for a given individual. Following Kahneman and Tversky (1979), let $q(p)$ denote the reweighted probability function that governs their decisions over financial assets. It is straightforward to show that the no-trade condition becomes

$$\beta(p)\frac{u'(c_u(p))}{v'(c_e(p))} \leq T(p) \quad \forall p, \tag{20}$$

where $\beta(p) = \frac{p}{q}\frac{1-q}{1-p}$ is the distortion in individuals' marginal rate of substitution arising from biases in beliefs (as opposed to differential marginal utilities of income, $\frac{u'}{v'}$). Biased beliefs generate a second reason that individuals would be willing to pay for UI: they may overstate (or understate) the likelihood of job loss, $\beta(p)$.

One can use the estimates above to ask how biased beliefs must be to overcome the frictions imposed by private information. With a baseline minimum pooled price ratio of $T(p) = 1 + 3.3 = 430$ percent and markup individuals are willing to pay less than 75 percent, a market would not be profitable unless individuals believe that job loss is 5.7 times more likely to occur than in reality ($430/75 = 5.7$). Generating a profitable insurance market would require a very large degree of bias to sufficiently inflate demand to overcome the hurdles that would be imposed by adverse selection.

VII. Conclusion

This paper argues that private information prevents the existence of a robust private market for unemployment or job loss insurance. If insurers were to attempt

[45] The fact that the vast majority of unemployment risk is idiosyncratic perhaps also explains why there is not a large private market for households to insure against fluctuations in the aggregate unemployment rate.

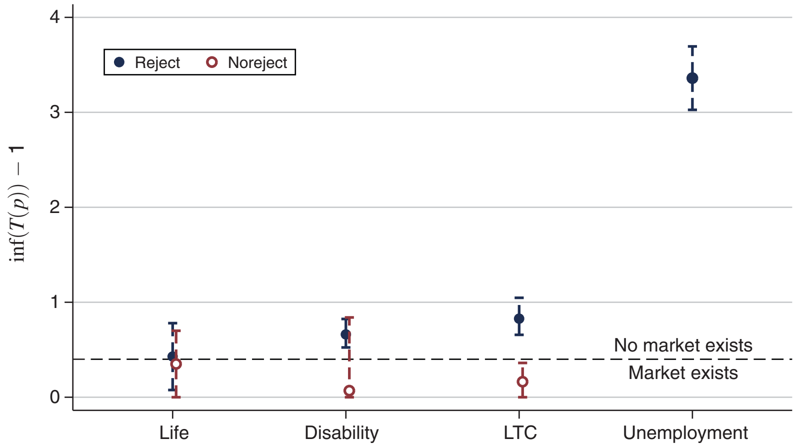

Figure 7. Comparison of $\inf T(p) - 1$ to Other (Non)Existing Insurance Markets

Notes: Hendren (2013) argues private information prevents people with preexisting conditions from purchasing insurance in LTC, Life, and Disability insurance markets. This figure compares the estimates of $\inf T(p) - 1$ for the baseline specification in the unemployment context to the estimates in Hendren (2013) for the sample of individuals who are unable to purchase insurance due to a preexisting condition (circles) and those whose observables would allow them to purchase insurance in each market (hollow circles). The figure reports the confidence interval and the 5/95 percent confidence interval for each estimate in each sample. For the subsamples in LTC, Life, and Disability for which the market exists, one cannot reject the null hypothesis of no private information, $\inf T(p) = 0$. In contrast, subsamples whose observables would prevent them from purchasing insurance tend to involve larger estimates of the minimum pooled price ratio, which suggests the frictions imposed by private information form the boundary of the existence of insurance markets.

to sell such policies, the empirical results suggest that they would be too heavily adversely selected to deliver a positive profit at any price. Similar patterns emerge in other insurance markets. Figure 7 compares the estimates of the minimum pooled price ratio to analogous estimates from health-related insurance markets studied in Hendren (2013). In long-term care insurance, life insurance, and disability insurance, Hendren (2013) finds no statistically significant amounts of private information for those with observable characteristics that allow them to purchase insurance. However, for individuals with preexisting conditions that would cause them to be rejected by an insurance company, the estimated markups are 42 percent for Life, 66 percent for Disability, and 83 percent for Long-Term Care. Combining these patterns with those identified in this paper for UI, the results suggest that the frictions imposed by private information form a boundary to the existence of insurance markets.

While this paper addresses the positive question of why a third-party insurance market for UI does not exist, the paper has not explored the normative implications. In particular, the presence of ex ante knowledge of future job loss suggests that individuals may have demand not only for insurance against losing their job in the future, but also demand for insurance against learning today that they might lose their job in the future. Additionally, it is quite plausible that much of the *private* information documented in the present paper is jointly known to the firm and worker. If this is the case, then one might ask why firms don't provide additional UI or severance[46] and whether additional government UI introduces externalities on the firm's contracting decisions. The normative implications of the patterns documented in this paper for optimal UI design are an interesting direction for future work.

REFERENCES

Acemoglu, Daron, and Robert Shimer. 1999. "Efficient Unemployment Insurance." *Journal of Political Economy* 107 (5): 893–928.

Acemoglu, Daron, and Robert Shimer. 2000. "Productivity Gains from Unemployment Insurance." *European Economic Review* 44 (7): 1195–1224.

Aguiar, Mark, and Erik Hurst. 2005. "Consumption versus Expenditure." *Journal of Political Economy* 113 (5): 919–48.

Akerlof, George A. 1970. "The Market for 'Lemons': Qualitative Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84 (3): 488–500.

Bach, Christine. 1998. *Negativauslese und tarifdifferenzierung im versicherungssektor: Ökonomische modelle und ökonometrische analsysemethoden.* Wiesbaden: Deutscher Universitäts Verlag.

Baily, Martin Neil. 1978. "Some Aspects of Optimal Unemployment Insurance." *Journal of Public Economics* 10: 379–402.

Barceló, Cristina, and Ernesto Villanueva. 2010. "The Response of Household Wealth to the Risk of Losing the Job: Evidence from Differences in Firing Costs." Banco de España Working Paper 1002.

Basten, Christoph, Andreas Fagereng, and Kjetil Telle. 2012. "Saving and Portfolio Allocation before and after Job Loss." KOF Swiss Economic Institute Working Paper 298.

Bloemen, Hans G., and Elena G. F. Stancanelli. 2005. "Financial Wealth, Consumption Smoothing and Income Shocks Arising from Job Loss." *Economica* 72 (287): 431–52.

Blundell, Richard, Monica Costa Dias, Costas Meghir, and Jonathan Shaw. 2016. "Female Labour Supply, Human Capital and Welfare Reform." *Econometrica* 84 (5): 1705–53.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98 (5): 1887–1921.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Ekstein. 2012. "Consumption Inequality and Family Labor Supply." National Bureau of Economic Research Working Paper 18445.

Bollinger, Christopher R. 1998. "Measurement Error in the Current Population Survey: A Nonparametric Look." *Journal of Labor Economics* 16 (3): 576–94.

Brown, Jennifer, and David A. Matsa. 2016. "Boarding a Sinking Ship? An Investigation of Job Applications to Distressed Firms." *Journal of Finance* 71 (2): 507–50.

Browning, Martin, and Thomas F. Crossley. 2001. "Unemployment Insurance Levels and Consumption Changes." *Journal of Public Economics* 80 (1): 1–23.

Cabral, Marika. 2013. "Claim Timing and Ex Post Adverse Selection." Unpublished.

Campos, Rodolfo G., and Iliana Reggio. 2015. "Consumption in the Shadow of Unemployment." *European Economic Review* 78: 39–54.

Carroll, Christopher D., Karen E. Dynan, and Spencer D. Krane. 2003. "Unemployment Risk and Precautionary Wealth: Evidence from Households' Balance Sheets." *Review of Economics and Statistics* 85 (3): 586–604.

Carroll, Christopher D., and Andrew A. Samwick. 1997. "The Nature of Precautionary Wealth." *Journal of Monetary Economics* 40 (1): 41–71.

Carroll, Christopher D., and Andrew A. Samwick. 1998. "How Important Is Precautionary Saving?" *Review of Economics and Statistics* 80 (3): 410–19.

Chetty, Raj. 2006. "A General Formula for the Optimal Level of Social Insurance." *Journal of Public Economics* 90 (10–11): 1879–1901.

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

Chetty, Raj, and Adam Szeidl. 2007. "Consumption Commitments and Risk Preferences." *Quarterly Journal of Economics* 122 (2): 831–77.

Chiappori, Pierre-Andre, and Bernard Salanié. 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy* 108 (1): 56–78.

[46] This could be due to a moral hazard problem (severance reduces effort), or a screening problem (an adverse selection of the type of workers attracted to firms with high severance payments). It could also be that the risk of job loss at the firm level is largely an aggregate risk that firms are unwilling to bear.

Cochrane, John H. 1995. "Time-Consistent Health Insurance." *Journal of Political Economy* 103 (3): 445–73.
Davis, Steven J., and Till M. von Wachter. 2011. "Recessions and the Cost of Job Loss." National Bureau of Economic Research Working Paper 17638.
Dynan, Karen E. 1993. "How Prudent Are Consumers?" *Journal of Political Economy* 101 (6): 1104–13.
Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics* 125 (3): 877–921.
Einav, Liran, Amy Finkelstein, and Jonathan Levin. 2010. "Beyond Testing: Empirical Models of Insurance Markets." *Annual Review of Economics* 2: 311–36.
Engen, Eric M., and Jonathan Gruber. 2001. "Unemployment Insurance and Precautionary Saving." *Journal of Monetary Economics* 47 (3): 545–79.
Finkelstein, Amy, and Kathleen McGarry. 2006. "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market." *American Economic Review* 96 (4): 938–58.
Finkelstein, Amy, and James Poterba. 2004. "Adverse Selection in Insurance Markets: Policyholder Evidence from the UK Annuity Market." *Journal of Political Economy* 112 (1): 183–208.
Gallen, Yana. 2013. "Anticipating Unemployment: Savings Evidence from Denmark." http://yanagallen.com/Gallen_Cost%20of%20Job%20Loss.pdf (accessed May 4, 2017).
Gan, Li, Michael D. Hurd, and Daniel McFadden. 2005. "Individual Subjective Survival Curves." In *Analyses in the Economics of Aging,* edited by David A. Wise, 377–412. Chicago: University of Chicago Press.
Gruber, Jonathan. 1997. "The Consumption Smoothing Benefits of Unemployment Insurance." *American Economic Review* 87 (1): 192–205.
Gruber, Jonathan, and Julie Berry Cullen. 1996. "Spousal Labor Supply as Insurance: Does Unemployment Insurance Crowd Out the Added Worker Effect?" National Bureau of Economic Research Working Paper 5608.
Guariglia, Alessandra, and Byung-Yeon Kim. 2004. "Earnings Uncertainty, Precautionary Saving, and Moonlighting in Russia." *Journal of Population Economics* 17 (2): 289–310.
Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese. 1992. "Earnings Uncertainty and Precautionary Saving." *Journal of Monetary Economics* 30 (2): 307–37.
He, Daifeng. 2009. "The Life Insurance Market: Asymmetric Information Revisited." *Journal of Public Economics* 93 (9–10): 1090–97.
Hendren, Nathaniel. 2013. "Private Information and Insurance Rejections." *Econometrica* 81 (5): 1713–62.
Hendren, Nathaniel. 2017. "Knowledge of Future Job Loss and Implications for Unemployment Insurance: Dataset." *American Economic Review.* https://doi.org/10.1257/aer.20151655.
Hu, Yingyao, and Susanne M. Schennach. 2008. "Instrumental Variable Treatment of Nonclassical Measurement Error Models." *Econometrica* 76 (1): 195–216.
Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1994. "Precautionary Saving and Social Insurance." National Bureau of Economic Research Working Paper 4884.
Hurd, Michael D. 2009. "Subjective Probabilities in Household Surveys." *Annual Review of Economics* 1: 543–62.
Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2): 263–91.
Kleven, Henrik Jacobsen, Claus Thustrup Kreiner, and Emmanuel Saez. 2009. "The Optimal Income Taxation of Couples." *Econometrica* 77 (2): 537–60.
Landais, Camille, Pascal Michaillat, and Emmanuel Saez. 2010. "Optimal Unemployment Insurance over the Business Cycle." National Bureau of Economic Research Working Paper 16526.
Lusardi, Annamaria. 1997. "Precautionary Saving and Subjective Earnings Variance." *Economics Letters* 57 (3): 319–26.
Lusardi, Annamaria. 1998. "On the Importance of the Precautionary Saving Motive." *American Economic Review* 88 (2): 449–53.
Manski, Charles F. 2004. "Measuring Expectations." *Econometrica* 72 (5): 1329–76.
Manski, Charles F., and Francesca Molinari. 2010. "Rounding Probabilistic Expectations in Surveys." *Journal of Business & Economic Statistics* 28 (2): 219–31.
Meghir, Costas, and Luigi Pistaferri. 2011. "Earnings, Consumption and Life Cycle Choices." In *Handbook of Labor Economics*, Vol. 4B, edited by David Card and Orley Ashenfelter, 773–854. New York: Elsevier.
Pistaferri, Luigi. 2001. "Superior Information, Income Shocks, and the Permanent Income Hypothesis." *Review of Economics and Statistics* 83 (3): 465–76.
Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: John Wiley & Sons.

Shavell, Steven. 1979. "On Moral Hazard and Insurance." *Quarterly Journal of Economics* 93 (4): 541–62.

Shimer, Robert, and Iván Werning. 2007. "Reservation Wages and Unemployment Insurance." *Quarterly Journal of Economics* 122 (3): 1145–85.

Shimer, Robert, and Iván Werning. 2008. "Liquidity and Insurance for the Unemployed." *American Economic Review* 98 (5): 1922–42.

Spinnewijn, Johannes. 2015. "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs." *Journal of the European Economic Association* 13 (1): 130–67.

Stephens, Melvin, Jr. 2001. "The Long-Run Consumption Effects of Earnings Shocks." *Review of Economics and Statistics* 83 (1): 28–36.

Stephens, Melvin, Jr. 2002. "Worker Displacement and the Added Worker Effect." *Journal of Labor Economics* 20 (3): 504–37.

Stephens, Melvin, Jr. 2004. "Job Loss Expectations, Realizations, and Household Consumption Behavior." *Review of Economics and Statistics* 86 (1): 253–69.

Zeldes, Stephen P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97 (2): 305–46.

This article has been cited by:

1. Christina Patterson. 2023. The Matching Multiplier and the Amplification of Recessions. *American Economic Review* **113**:4, 982-1012. [Abstract] [View PDF article] [PDF with links]
2. Samir Elsadek Mahmoudi. 2023. Late-career unemployment shocks, pension outcomes and unemployment insurance. *Journal of Public Economics* **218**, 104751. [Crossref]
3. Selçuk Bedük. 2023. Insured Privately? Wealth Stratification of Job Loss in the UK. *Social Inclusion* **11**:1. . [Crossref]
4. Andreas I. Mueller, Johannes Spinnewijn. Expectations data, labor market, and job search 677-713. [Crossref]
5. Jose Garcia-Louzao. 2023. Workers' Job Mobility in Response to Severance Pay Generosity. *SSRN Electronic Journal* **71**. . [Crossref]
6. Brindusa Anghel. 2023. El aumento de los contratos indefinidos y su posible impacto en el gasto. *Boletín Económico* :2023/T1. . [Crossref]
7. Brindusa Anghel, Cristina Barceló, Ernesto Villanueva. 2023. The growth in permanent contracts and its potential impact on spending. *Economic Bulletin* :2023/Q1. . [Crossref]
8. Natalia Emanuel, Emma Harrington. 2023. Working Remotely? Selection, Treatment, and the Market for Remote Work. *SSRN Electronic Journal* **130**. . [Crossref]
9. Bjoern Fischer, Johannes Geyer, Nicolas R. Ziebarth. 2023. Fundamentally Reforming the Di System: Evidence from German Notch Cohorts. *SSRN Electronic Journal* **84**. . [Crossref]
10. Nguyen Van Tran. 2022. Understanding Household Consumption Behaviour: What do we Learn from a Developing Country?. *The B.E. Journal of Economic Analysis & Policy* **22**:4, 801-858. [Crossref]
11. Christos A. Makridis. 2022. The social transmission of economic sentiment on consumption. *European Economic Review* **148**, 104232. [Crossref]
12. Oluyemi T. Adeosun, Kayode E. Owolabi, Idongesit C. Eshiet, Temitope J. Owolabi. 2022. Exploring the transition from informal to formal jobs and its consequent impacts on the livelihood of migrant youths in Lagos metropolis. *Journal of Enterprising Communities: People and Places in the Global Economy* **10**. . [Crossref]
13. Raluca L. Pahontu. 2022. Divisive jobs: three facets of risk, precarity, and redistribution. *Political Science Research and Methods* **10**:3, 507-523. [Crossref]
14. Marika Cabral, Can Cui, Michael Dworsky. 2022. The Demand for Insurance and Rationale for a Mandate: Evidence from Workers' Compensation Insurance. *American Economic Review* **112**:5, 1621-1668. [Abstract] [View PDF article] [PDF with links]
15. Gregory E. Givens. 2022. UNEMPLOYMENT, PARTIAL INSURANCE, AND THE MULTIPLIER EFFECTS OF GOVERNMENT SPENDING. *International Economic Review* **63**:2, 571-599. [Crossref]
16. Jonas R. Jahnert, Hato Schmeiser, Florian Schreiber. 2022. Pricing strategies in the German term life insurance market: An empirical analysis. *Risk Management and Insurance Review* **25**:1, 19-34. [Crossref]
17. Jose Garcia-Louzao. 2022. Workers' job mobility in response to severance pay generosity. *Labour Economics* **75**, 102143. [Crossref]
18. Xuan Zhang. 2022. The effects of physician retirement on patient outcomes: Anticipation and disruption. *Journal of Public Economics* **207**, 104603. [Crossref]
19. Regis Barnichon, Yanos Zylberberg. 2022. A Menu of Insurance Contracts for the Unemployed. *The Review of Economic Studies* **89**:1, 118-141. [Crossref]

20. Jacob S. Hacker, Philipp Rehm. 2022. Reducing Risk as well as Inequality: Assessing the Welfare State's Insurance Effects. *British Journal of Political Science* **52**:1, 456-466. [Crossref]
21. Marcelo Pedroni, Swapnil Singh, Christian A. Stoltenberg. 2022. Advance Information and Consumption Insurance: Evidence from Panel Data. *SSRN Electronic Journal* **102**. . [Crossref]
22. Nicole Immorlica, Andre Sztutman, Robert M. Townsend. 2022. Optimal Credit Scores Under Adverse Selection. *SSRN Electronic Journal* **117**. . [Crossref]
23. Jason M. Lindo, Krishna Regmi, Isaac D. Swensen. 2022. Stable Income, Stable Family. *SSRN Electronic Journal* **100**. . [Crossref]
24. Christian A. Stoltenberg, Arne Uhlendorff. 2022. Consumption Choices and Earnings Expectations: Empirical Evidence and Structural Estimation. *SSRN Electronic Journal* **109**. . [Crossref]
25. Luca Citino, Kilian Russ, Vincenzo Scrutinio. 2022. Manipulation, Selection and the Design of Targeted Social Insurance. *SSRN Electronic Journal* **68**. . [Crossref]
26. David Coyne, Itzik Fadlon, Tommaso Porzio. 2022. Measuring Valuation of Liquidity with Penalized Withdrawals. *SSRN Electronic Journal* **67**. . [Crossref]
27. Liyang Sun, Sarah Abraham. 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics* **225**:2, 175-199. [Crossref]
28. Camille Landais, Johannes Spinnewijn. 2021. The Value of Unemployment Insurance. *The Review of Economic Studies* **88**:6, 3041-3085. [Crossref]
29. Po-Chun Huang, Tzu-Ting Yang. 2021. The welfare effects of extending unemployment benefits: Evidence from re-employment and unemployment transfers. *Journal of Public Economics* **202**, 104500. [Crossref]
30. Pierre Brochu. 2021. A Researcher's Guide to the Labour Force Survey: Its Evolution and the Choice of Public Use versus Master Files. *Canadian Public Policy* **47**:3, 335-357. [Crossref]
31. Pierre Mérel, Ariel Ortiz-Bobea, Emmanuel Paroissien. 2021. How big is the "lemons" problem? Historical evidence from French wines. *European Economic Review* **138**, 103824. [Crossref]
32. Nathaniel Hendren, Camille Landais, Johannes Spinnewijn. 2021. Choice in Insurance Markets: A Pigouvian Approach to Social Insurance Design. *Annual Review of Economics* **13**:1, 457-486. [Crossref]
33. Christoph Wunder, Tugba Zeydanli. 2021. The early costs of plant closures: Evidence on lead effects on workers' subjective and objective outcomes. *Journal of Economic Behavior & Organization* **188**, 489-505. [Crossref]
34. Stefan Pichler, Katherine Wen, Nicolas R. Ziebarth. 2021. Positive Health Externalities of Mandating Paid Sick Leave. *Journal of Policy Analysis and Management* **40**:3, 715-743. [Crossref]
35. Carlo Fezzi, Luisa Menapace, Roberta Raffaelli. 2021. Estimating risk preferences integrating insurance choices with subjective beliefs. *European Economic Review* **135**, 103717. [Crossref]
36. Nathaniel Hendren. 2021. Measuring Ex Ante Welfare in Insurance Markets. *The Review of Economic Studies* **88**:3, 1193-1223. [Crossref]
37. Peter Kuhn, Lizi Yu. 2021. How Costly Is Turnover? Evidence from Retail. *Journal of Labor Economics* **39**:2, 461-496. [Crossref]
38. Camille Landais, Arash Nekoei, Peter Nilsson, David Seim, Johannes Spinnewijn. 2021. Risk-Based Selection in Unemployment Insurance: Evidence and Implications. *American Economic Review* **111**:4, 1315-1355. [Abstract] [View PDF article] [PDF with links]
39. François Gerard, Joana Naritomi. 2021. Job Displacement Insurance and (the Lack of) Consumption-Smoothing. *American Economic Review* **111**:3, 899-942. [Abstract] [View PDF article] [PDF with links]

40. Romain Duval, Prakash Loungani. 2021. Designing Labor Market Institutions in Emerging Market and Developing Economies: A Review of Evidence and IMF Policy Advice. *Comparative Economic Studies* **63**:1, 31-83. [Crossref]
41. Andreas I. Mueller, Johannes Spinnewijn, Giorgio Topa. 2021. Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias. *American Economic Review* **111**:1, 324-363. [Abstract] [View PDF article] [PDF with links]
42. Adam Solomon. 2021. Misperceptions and Contract Distortions in Insurance Markets. *SSRN Electronic Journal* **116**. . [Crossref]
43. Piotr Denderski. 2021. On existence of private unemployment insurance with advance information on future job losses. *SSRN Electronic Journal* **98**. . [Crossref]
44. Catherine Maclean, Stefan Pichler, Nicolas R. Ziebarth. 2021. Mandated Sick Pay: Coverage, Utilization, and Welfare Effects. *SSRN Electronic Journal* **131**. . [Crossref]
45. Liran Einav, Amy Finkelstein, Neale Mahoney. The IO of selection markets 389-426. [Crossref]
46. Nathaniel Pattison. 2020. Consumption smoothing and debtor protections. *Journal of Public Economics* **192**, 104306. [Crossref]
47. Martin Halla, Julia Schmieder, Andrea Weber. 2020. Job Displacement, Family Dynamics, and Spousal Labor Supply. *American Economic Journal: Applied Economics* **12**:4, 253-287. [Abstract] [View PDF article] [PDF with links]
48. Rodolfo G Campos, Iliana Reggio. 2020. Do the Unemployed Pay Lower Prices? A Reassessment of the Value of Unemployment Insurance. *Journal of the European Economic Association* **18**:5, 2135-2181. [Crossref]
49. Roger Douglas, Robert MacCulloch. 2020. A welfare reform for New Zealand: mandatory savings not taxation. *New Zealand Economic Papers* **54**:3, 239-273. [Crossref]
50. Nathaniel Hendren, Ben Sprung-Keyser. 2020. A Unified Welfare Analysis of Government Policies*. *The Quarterly Journal of Economics* **135**:3, 1209-1318. [Crossref]
51. Jacob Gerner Hariri, Amalie Sofie Jensen, David Dreyer Lassen. 2020. Middle Class Without a Net: Savings, Financial Fragility, and Preferences Over Social Insurance. *Comparative Political Studies* **53**:6, 892-922. [Crossref]
52. Piotr Denderski, Christian A. Stoltenberg. 2020. Risk sharing with private and public information. *Journal of Economic Theory* **186**, 104988. [Crossref]
53. Johannes Spinnewijn. 2020. The Trade-Off between Insurance and Incentives in Differentiated Unemployment Policies. *Fiscal Studies* **41**:1, 101-127. [Crossref]
54. Christos Makridis, Tao Wang. 2020. Learning from Friends in a Pandemic: Social Networks and the Macroeconomic Response of Consumption. *SSRN Electronic Journal* **80**. . [Crossref]
55. Breno Braga, Olga Malkova. 2020. Time to Grow Up? Adult Children as Determinants of Parental Labor Supply. *SSRN Electronic Journal* **54**. . [Crossref]
56. Juliana Mesén Vargas, Bruno Van der Linden. 2019. Why Cash Transfer Programs Can Both Stimulate and Slow Down Job Finding. *IZA Journal of Labor Economics* **8**:1. . [Crossref]
57. THERESA KUCHLER, BASIT ZAFAR. 2019. Personal Experiences and Expectations about Aggregate Outcomes. *The Journal of Finance* **74**:5, 2491-2542. [Crossref]
58. Peter Ganong, Pascal Noel. 2019. Consumer Spending during Unemployment: Positive and Normative Implications. *American Economic Review* **109**:7, 2383-2424. [Abstract] [View PDF article] [PDF with links]
59. Bruce D. Meyer, Wallace K.C. Mok. 2019. Disability, earnings, income and consumption. *Journal of Public Economics* **171**, 51-69. [Crossref]

60. Yuri Pettinicchi, Nathanael Vellekoop. 2019. Job Loss Expectations, Durable Consumption and Household Finances: Evidence from Linked Survey Data. *SSRN Electronic Journal* . [Crossref]

61. Pierre Merel, Ariel Ortiz-Bobea, Emmanuel Paroissien. 2019. How Big Is the 'Lemons' Problem? Historical Evidence From French Appellation Wines. *SSRN Electronic Journal* . [Crossref]

62. Ashna Arora. 2018. Juvenile Crime and Anticipated Punishment. *SSRN Electronic Journal* . [Crossref]

63. Raluca L. Pahontu. 2018. Employed in Conflict: Explaining Social Policy Preferences of a Tripartite Labour Market. *SSRN Electronic Journal* **64**. . [Crossref]

64. Christos Andreas Makridis. 2017. Sentimental Business Cycles and the Protracted Great Recession. *SSRN Electronic Journal* . [Crossref]

65. Emily mname Gallagher, Stephen mname Roll, Rourke mname O'Brien, Michal mname Grinstein-Weiss. 2017. Health Insurance and the Earnings Stability of Low-Income Households. *SSRN Electronic Journal* . [Crossref]