

Knowledge of Future Job Loss and Implications for Unemployment Insurance

Nathaniel Hendren*

December, 2016

Abstract

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.

1 Introduction

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

*Harvard University, 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 4 and 6 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.

¹See Gruber (1997), Browning and Crossley (2001), Aguiar and Hurst (2005), Chetty (2008), and Blundell et al. (2012).

²Two 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" (<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% replacement rate for work-

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 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% upon the onset of unemployment (e.g. Gruber (1997)), I use data from the Panel Study of Income Dynamics to show that food expenditure also declines by 2.7% in the 1-2 years prior to unemployment. This occurs on a sample who remain employed in these pre-periods with no pre-trend 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 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 individuals are willing to pay for insurance exceeds the cost imposed by worse risks adversely selecting their contract – where this 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 (a) markup individuals are

ers in some industries and occupations (<https://www.incomeassure.com>). Back-of-the envelope calculations suggest their markups exceed 500% 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% replacement rate (e.g. http://www.mlive.com/jobs/index.ssf/2011/08/get_out_your_calculator_before_you_buy_p.html#).

³As discussed below in the related literature, 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 individual's knowledge after controlling for observables insurers would use to price insurance, as this is what is relevant for generating adverse selection.

⁴Along 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 dollar of insurance provide first-order welfare gains, but the behavioral response to a small amount of insurance imposes a second order impact on the cost of insurance.

willing to pay for UI and (b) the pooled price ratio.

A large literature has attempted to estimate the markup 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 (1976); Gruber (1997)). However, job loss affects consumption not only at its onset but also in the years 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 elicitation on the job loss indicator suggests 80% of the information about job loss is revealed in the last year prior to unemployment. I show that one can divide the 7-10% first difference estimate by this “first stage” of 0.8 to arrive at the average causal effect of 8-13%. 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%).

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 shopping 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% 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 elicitation, I estimate that 10% 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% multiplied by the coefficient of relative risk aversion (i.e. a 60% 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 individuals would be willing to pay for UI. For a semi-elasticity of 0.5, this suggests individuals are willing to pay of around a 60% 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 elicitation to provide non-parametric lower bounds and semi-parametric point estimates on the pooled price ratio. This yields lower bounds that are generally in excess of 70% and point estimates in excess of 300%. The pooled price ratio remains large across varying assumptions about the observables 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 shocks⁵ 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 elicitation in the HRS are predictive about future unemployment status, and to Stephens (2001) and Stephens (2002) who find evidence of ex-ante consumption drops and spousal labor supply increases in the U.S. 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 elicitation 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 methodologies using the response of first-differences of consumption tend to under-state the value of UI (e.g. Gruber (1997)). I provide a correction to this method in Section 5.1. In addition, these ex-ante behavioral responses open a po-

⁵In the UI context, see Baily (1976); Acemoglu and Shimer (1999, 2000); Chetty (2006); Shimer and Werning (2007); Blundell et al. (2008); Chetty (2008); Shimer and Werning (2008); Landais et al. (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.

⁶For example, see Guiso et al. (1992); Dynan (1993); Hubbard et al. (1994); Carroll and Samwick (1997, 1998); Carroll et al. (2003); Lusardi (1997, 1998); Engen and Gruber (2001); Guariglia and Kim (2004); Bloemen and Stanca (2005); Barceló and Villanueva (2010); Campos and Reggio (2015); Brown and Matsa (2015). Most closely, Basten et al. (2012) and Gallen (2013) find evidence of ex-ante savings increases in response to future unemployment in Norway and Denmark.

⁷Along the way, the analysis clarifies the empirical estimands required to answer this question. For example, one needs to know whether the elicitation are predictive of job loss conditional on public information 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 growth drops on the sub-sample of those who remain employed in the pre-period and experience no drop in consumption.

tentially fruitful new pathway to valuing insurance by exploiting the ex-ante responses to measure the value of insurance in Section 5.2 that avoids requirements of state independence of the utility function employed in previous literature (e.g. Baily (1976); Gruber (1997)). In particular, the drop in food expenditure and increased spousal labor supply that occurs when people learn they might lose their job suggests job loss is significantly under-insured in the U.S., 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 micro-foundations for under-insurance. 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 et al. (2010): If private information prevents the existence of entire markets, it is difficult to identify its impact by looking for the adverse selection of existing contracts. Combining with the evidence in Hendren (2013) that private information prevents the existence of health-related insurance markets for those with pre-existing 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 2 discusses the data used in the analysis. Section 3 presents a series of motivating statistics establishing the presence of private information that individuals would use to (adversely) select insurance. Section 4 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 5 considers the implications of the patterns in Section 3 for measuring individual’s willingness to pay for UI, and Section 6 considers the implications for measuring the pooled price ratio. Section 7 discusses robustness and alternative theories of market non-existence. Section 8 concludes.

2 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 elicitation about future unemployment and measures of spousal labor supply. The PSID does not contain subjective probability elicitation, but includes a panel of information on food expenditures.

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

2.1 HRS

The HRS sample draws from waves of the Health and Retirement Study (HRS) spanning years 1992-2013. The HRS samples individuals over 50 and their spouses (included regardless of age). The baseline sample includes everyone under 65 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 I presents the histogram of these elicitations. As noted in previous literature (Gan et al. (2005)), the responses concentrate on focal point values, especially zero. Taken literally, these responses of zero (or 100%) 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, I will let Z denote the responses to these survey questions, and I will 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 individual's willingness to pay for an unemployment insurance contract (as will be formalized in Section 4), but will be assumed to be unobserved to the econometrician, $Z \neq P$.

Outcomes To infer people's knowledge of future job loss, we 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, denoted U . The subsequent wave asks individuals whether they are working at the same job as the previous wave (roughly 2 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 6-month waiting period before claims can be exercised.

Public Information Estimating private information requires specifying the set of observable information 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 sub-sample 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 (2 years prior). The primary analysis will focus on labor market entry by previously non-working 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.

2.2 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 to heads of household between the ages of 25 and 65 who have non-missing 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 threefold 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.

⁹This 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.

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 the subsequent survey is administered may understate expenditure 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 that requires that individuals are employed at the time of eliciting their job loss probability.

2.3 Summary Statistics

Table I 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 (\$36K in the HRS; \$40K 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% versus 40%).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% in the HRS and 2.8% in the PSID sample. The PSID sample is more likely to be unemployed (2.4% versus 1.9% 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% per year

¹⁰I 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, for example, of a change in management).

¹¹Although the HRS is a representative sample of individuals over 50 and their spouses, the sample of women exceeds 50% 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).

versus just 1.7% 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% of spouses are employed and 3.9% 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 elicitation. While 3.1% of the HRS sample loses their job in the subsequent 12 months from the survey, the mean subjective probability elicitation is 15.7%. Such bias is a common feature of subjective probability elicitation (see, e.g., Hurd (2009)). In particular, for low probability events there is a natural tendency for measurement error in elicitation to lead to an upward bias. This provides further rationale for treating these elicitation as noisy and potentially biased measures of true beliefs, as is maintained throughout the empirical analyses 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 7.

3 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 1-2 years prior to job loss*. These empirical patterns will provide a basis for considering the implications for the workings of a private UI market, discussed in Sections 4-6.

3.1 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 elicitation, Z_{it} , are more 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 elicitation are predictive of subsequent job loss was first

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

shown by Stephens (2004). The analysis in this section expands upon this work by focusing on individual’s knowledge that is not captured by observable characteristics, X_{it} , that insurers might use to price insurance.

To illustrate the predictive content of the elicitation, 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,

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

Figure II 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%. The figure reveals an increasing pattern: those with higher subjective probability elicitation 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 II 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} :

$$U_{it} = \alpha + \beta Z_{it} + \Gamma X_{it} + \epsilon_{it} \quad (2)$$

For the baseline specification with demographics and job characteristics, the estimated slope is 0.0836 (s.e. 0.00675). For every 1pp 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 (s.e. 0.00685); adding additional controls for health characteristics¹³ reduces the coefficient to 0.0822 (s.e. 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 (s.e. 0.012). This highlights that individuals’ risk of job loss is largely time-varying within an individual. This underscores the difficulty that would

¹³These 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 bmi

be 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.

3.2 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 under-insured 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 III 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 III 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%. On the other hand, relative to the base entry rate of these spouses of 3.9%, 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 III linearly parameterizes the relationship in Figure III 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 (s.e. 0.0087) for the baseline specification in Column (1). Column (2) restricts the sample to those who do not end up losing their job in the 12 months after the survey, yielding 0.0256 (s.e. 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 elicitation 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 (s.e. 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 that have less labor force attachment and are more likely to come and go into the labor market. To this aim, Column (4) considers a placebo test that uses the lagged value of the elicitation, Z_{it-2} instead of $Z_{i,t}$ (where year $t-2$ corresponds to the previous wave of the survey conducted 2 years prior). Here, the coefficient is 0.00122 (s.e. 0.008) and is not statistically distinct from zero. 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.

3.3 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% 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

¹⁴One 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 (s.e. 0.0119) is positive, although not statistically significant and suggests 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 (s.e. 0.0097). For these reasons, I focus primarily on labor market entry of spouses not currently in the labor market and perhaps face greater flexibility in their choice of industry/occupation/firm/etc when choosing employment.

1-2 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_{it} denote an indicator for being unemployed at the time of the survey in year t , U_{it} . I regress food expenditure growth $g_{i,t}$ on unemployment in year $t - k$,

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

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 2 years prior to the unemployment measurement, Figure IV plots the coefficients Δ_k^{FD} for $k = -4, -3, \dots, 0, \dots, 3, 4$. Consistent with previous literature, food expenditure drops by 7-8% at the onset of unemployment. But, there is also a 2-3% impact on food expenditure growth in the the year prior to unemployment. We 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 zero. 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 IV illustrates the robustness of the ex-ante drop in food expenditure in the 1-2 years prior to unemployment, Δ_{-1}^{FD} . Column (1) shows that the baseline specification yields a -2.71% (s.e. 0.975%) drop in food expenditure¹⁶ in the 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% (s.e. 1.05%).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% (s.e. 1.06%), which is similar in magnitude but not statistically distinguishable from the baseline estimate. This suggests knowledge of future job loss is present

¹⁵The smaller responses in earlier periods is consistent with evidence discussed below that a large portion (e.g. 10%) 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 6 and Online Appendix Figure I.

¹⁶Online Appendix Figure II replicates the baseline regression in Figure IV using more recent PSID data on total household expenditure on a sample that is surveyed every two years. The broad patterns are similar, although the consumption drop upon unemployment is slightly larger (e.g. 12%) when using total consumption expenditure as opposed to food expenditure. There is also an ex-ante response of 3.6% ($p = 0.055$) in $t - 2$ relative to $t - 4$ in total consumption expenditure, which again suggests individuals decrease their food expenditure in response learning about future unemployment.

¹⁷Restricting to the subsample of 53,327 observations for which the needs variable is available drops the coefficient in column (1) to 2.4%, suggesting roughly 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.

across the age distribution in the U.S, 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% (s.e. 0.824%) 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 IV using job loss instead of unemployment and finds similar patterns to Figure IV. Overall, the results reveal a similar ex-ante drops 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 IV adds controls for a cubic polynomial of changes in log household income to the baseline specification. This yields a coefficient of 2.72% (s.e. 0.969%) nearly identical to the baseline specification in Column (1). The results are similar with higher and lower order polynomial controls, or restricting 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% (s.e. 0.983%).¹⁹

To understand why the results are not affected by adding controls for income, Online Appendix Figure IV replicates Figure IV 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

¹⁸Temporary 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.

¹⁹The sample sizes are slightly lower for these specifications due to non-response to income questions. The food expenditure patterns are similar when restricting to a sample with non-missing income reports.

²⁰The 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 sub-samples 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% 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 pre-periods.

²¹While the narrative here is that consumption declines because of the future unemployment event, there is a

4 Model

Individuals have knowledge about their future job loss and take actions to increase their financial resources available when unemployed. This suggests they would have a demand for an unemployment insurance contract if it were offered. Yet they would likely use their private information to potentially adversely select the contract. This 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 5 and 6. 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 7.

4.1 Setup

Individuals face a risk of losing their job in the next year. They (or their households²²) choose consumption today, c_{pre} , along with a plan for consumption in the event of not losing and losing her 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 characteristics, 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

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 3.2.

²²Throughout, 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.

set, Ω_i , to maximize a utility function that satisfies

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

where $v(c_{pre})$ is the utility over consumption today, $u(c_u)$ and $v(c_e)$ are the utility 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 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 cost to the insurer of providing that insurance. Third, the model in principle allows for multiple dimensional 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 uni-dimensional); but the modeling in Appendix A illustrates how the analysis readily extends to the case of multi-dimensional 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 under-insurance and thus the ability to match the empirical patterns in Section 3.

Consumption Responses The model captures the dynamic consumption patterns in Figure IV 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:

$$v'(c_{pre}(p)) = pu'(c_u(p)) + (1-p)v'(c_e(p)) \quad (6)$$

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 under-insured), 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 IV suggest individuals are not fully insured against the risk of job loss.

²³To 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.

Spousal Labor Supply Responses The model also captures the spousal labor supply responses documented in Section 3.3. To see this, one can incorporate spousal labor supply into the set of other actions. Let $a = (l^{spouse}, a')$ and $\psi_i(p, a) = \kappa_i(p, a') - \eta(l^{spouse}(p))$, 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, $\eta'(l^{spouse}(p)) = w^{spouse}v'(c_{pre}(p))$. So, the Euler equation (6) can be re-written as

$$\frac{1}{w^{spouse}}\eta'(l^{spouse}(p)) = pu'(c_u(p)) + (1-p)v'(c_e(p)) \quad (7)$$

The marginal disutility of labor divided by the spouse's wage equals the expected marginal utility of income in the future. When individuals are under-insured 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; Appendix C.5 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 income if job loss becomes more likely. In this sense, the model formalizes how the ex-ante consumption and labor supply responses suggest $u'(c_u(p)) > v'(c_e(p))$, so that individuals are not fully insured against job loss and would have a demand for a private UI policy.

4.2 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 7 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, $\frac{E[P|P \geq p]}{1 - E[P|P \geq p]}$. Appendix A states these assumptions more formally and shows that a private market cannot exist if and only if

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

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 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 dollar 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 dollar 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 moral hazard does not affect whether insurers' first dollar 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 dollar of insurance can be adversely selected by strictly worse risks, so that private information can explain the absence of a market.

Multi-Dimensional 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 et al. (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.

²⁴As discussed in Section 5.3, one could combine a model of moral hazard and fixed costs to explain the absence of a private market.

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

When Equation (8) holds, the market for any insurance contract (or menu of insurance contracts, as shown in Appendix A.2) about the job loss will fully unravel in the sense of Akerlof (1970) and Einav et al. (2010).

While the model allows for endogenous contracts, the assumption that individuals with higher p will always have a higher demand for insurances is more restrictive than the models of Akerlof (1970) and Einav et al. (2010) because it assumes a single dimension of heterogeneity summarized by p . In Appendix A.1, I extend the baseline model in Section 4 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 et al. (2010) for any hypothetical insurance contract and understand whether a market could exist. I extend the model to allow for multi-dimensional heterogeneity in 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 of multi-dimensional 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 2 and 3 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, $T(p)$.

5 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 5.1 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 3 suggests that previous approaches using first-difference estimates (e.g. Gruber and Cullen (1996)) understate the true causal effect. I provide a method to correct for this bias.

The second strategy presented in Section 5.2 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 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.

5.1 Approach #1: Consumption when unemployed versus employed

5.1.1 Setup

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

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

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 constant coefficient of relative risk aversion (e.g. $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$) so that the coefficient of relative prudence equals the coefficient of relative risk aversion plus 1, $\gamma = \sigma + 1$. Following Gruber (1997), it is common to approximate this percentage change using log consumption,

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

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 V 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% 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% (s.e. 0.997%).

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

²⁷See Appendix C.1 for the derivation.

²⁸Ideally, 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 4. However, as noted in Section (3.3), 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

5.1.2 Correcting Bias in First Difference Estimator

Unfortunately, the ex-ante responses documented in Section 3 suggest Δ_0^{FD} does not capture the full causal effect. Indeed, one can write the first difference estimate as:

$$\Delta_0^{FD} = \underbrace{E[\log(c_e(p)) - \log(c_u(p))]}_{\text{Average Causal Effect}} - \underbrace{(E[\log(c_{pre}(p)) | U = 0] - E[\log(c_{pre}(p)) | U = 1])}_{\text{Bias from ex-ante response}} \quad (11)$$

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 under-state 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 $\frac{\text{var}(P)}{\text{var}(U)} = E[P|U = 1] - E[P|U = 0]$ of the information about U has been revealed at the time P is measured. 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, $1 - (E[P|U = 1] - E[P|U = 0])$.²⁹

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 , $\frac{d[\log(c_e(p)) - \log(c_u(p))]}{dp} = 0$. 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*

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

where $\frac{\text{var}(P)}{\text{var}(U)} = E[P|U = 1] - E[P|U = 0]$ 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) $E[\log(c_e(p)) - \log(c_u(p))]$ if higher values of p correspond to larger (smaller) consumption drops.

Proof. See 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

indicator for job loss occurring between surveys. This has the added benefit of aligning with much previous literature focusing on estimating the impact of unemployment on consumption (e.g. Gruber (1997); Chetty and Szeidl (2007)). And, I show below this is likely a conservative path to prevent under-stating the willingness to pay for insurance.

²⁹An alternative strategy would be to use longer lags instead of 1-year lagged consumption. However, Online Appendix Figure V shows that individuals have (albeit small) predictive 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).

individuals have ex-ante knowledge, one must estimate 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. Appendix C.2 shows that $\kappa \approx 1$ because p is small on average ($E[p] = 4\%$). For example, if individuals are willing to pay a 25% markup and $E[p] = 4\%$, then this correction factor is roughly $\kappa = 1.01$. Going forward I assume $\kappa \approx 1$.

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% more likely (i.e. p increases by 1pp), $\frac{d\log(c_{pre}(p))}{dp}$, equals 1% 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, ($E[P|U = 1] - E[P|U = 0]$), which yields the formula in equation (12).³⁰

5.1.3 Results

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 Appendix Table I suggest $E[P|U = 1] - E[P|U = 0] \approx 0.197$ (s.e. 0.012), which implies 80.3% of the uncertainty in job loss is not known 1 year in advance. Appendix Table I also illustrates that the 80% statistic is quite similar across various demographic subgroups.

³⁰It 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

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

Estimating this more general equation would require an estimate of $\frac{d[\log(c_u(p)) - \log(c_e(p))]}{dp}$, 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 $\frac{d[\log(c_u(p)) - \log(c_e(p))]}{dp} = 0$.

Baseline Results Table V shows that the first difference estimate of the impact of unemployment on consumption is $\Delta_0^{FD} = 7.23\%$. The second set of rows in Table V divide the first difference estimate, Δ_0^{FD} , by 0.803. This suggests a 9% 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% markup for UI. Higher risk aversion (e.g. $\sigma = 3$) raises the willingness to pay to 31.9%; if individuals are less risk averse (e.g. $\sigma = 1$), they would be willing to pay a 9.8% markup for UI.

Robustness The baseline analysis in column (1) makes a couple of specification decisions whose robustness are 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 IV for Δ_{-1}^{FD}) and again finds a similar coefficient to the baseline specification. Column (4) re-introduces observations with more than a threefold change in food expenditure, which were dropped to align with the specification of Gruber (1997). Re-introducing these observations increases the first difference estimate to 8.89% (s.e. 1.23%).

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 this follows Zeldes (1989); 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 measure elicit a “typical” week. Since unemployment is co-incident with rises in food stamp use, this differential recall window could lead to an under-stating 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% (s.e. 1.71%), and provides arguably an upper bound on the size of the causal effect of 22.7%. This would suggest individuals are willing to pay a 60.9% markup for UI with a coefficient of relative risk aversion of 2.

Job Loss versus Unemployment The analysis in Columns (1)-(6) measure how food expenditure varies with whether or not the individual is employed at the time of the survey. As discussed in Section 2.2, 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% (s.e. 0.860%) 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% in Column (1). This is consistent with attenuation from mis-alignment 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, 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, $\min_p \{ \log(c_u(p)) - c_e(p) \}$, which can be used to construct a maximum willingness to pay for UI. For brevity, the details of this approach are provided in the 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%, or roughly twice as large as the mean consumption drop. This rises to 14.6% on the broader sample that does not drop outliers with greater than a threefold change in measured food expenditure. The lower rows in Table V scale these estimates by various values of risk aversion. For a conservative estimate of 3, it suggests the maximum markup individuals would be willing to pay is less than 52.6% on the baseline sample and 56.6% 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 individuals would be willing to pay for UI generally lies below 60%.

5.2 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 under-statement of the true willingness to pay. If

³¹This is often stated as two distinct assumptions: (1) state independence over consumption ($u(c) = v(c)$ for all c), and (2) 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.

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):

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

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, $v'(c_{pre}(1)) = u'(c_u(1))$. Conversely, those that 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, $v'(c_{pre}(0)) = v'(c_e(0))$. So, if $c_e(p)$ and $c_u(p)$ do not systematically vary with p , then $v'(c_{pre}(0)) - v'(c_{pre}(1)) = u'(c_u) - v'(c_e)$. Taking a Taylor approximation to $v'(c_{pre}(p))$ around \bar{c}_{pre} and dividing by $v'(\bar{c}_{pre})$ yields $\frac{v'(c_{pre}(0)) - v'(c_{pre}(1))}{v'(\bar{c}_{pre})} \approx \sigma \frac{d \log(c_{pre}(p))}{dp}$, where $\sigma = \frac{v''(\bar{c}_{pre})\bar{c}_{pre}}{v'(\bar{c}_{pre})}$ 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 (a) the Euler equation holds, (b) c_e and c_u do not vary systematically with p , $\frac{dc_e}{dp} = 0$ and $\frac{dc_u}{dp} = 0$, and (c) 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,

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

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 Appendix C.4. □

Proposition 2 shows that one can identify the markup 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 to 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 (because $\frac{dc_e}{dp} < 0$). 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 over-statement of the true demand for UI, note that if $\frac{dc_e}{dp} \neq 0$, then differentiating the Euler equation yields

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

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 over-state 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 5.3.

2-Sample Implementation To estimate how $\log(c_{pre}(p))$ varies with p in equation 13, I follow a methodology similar to Section 5.1 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,

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

where $P_{j,t}$ is the beliefs in period t about job loss in period j . The denominator is the fraction of information that 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 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 IV reported that food expenditure growth is $\Delta_{-1}^{FD} = 2.71\%$ 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 $E[P_{t,t-1}|U_t = 1] - E[P_{t,t-1}|U_t = 0] = 19.7\%$ (Appendix Table I, Column (1)). In this sense, 20% 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 Appendix Table I reports this value as $E[P_{t-2,t}|U_t = 1] - E[P_{t-2,t}|U_t = 0] = 9.4\%$.³³ The difference of 10.3% is an estimate of the fraction of information about job loss in year t is 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 VI; alternative specifications yield similar estimates for $\frac{d\log(c_{pre})}{dp}$ of around 20-30%. Scaling these by a coefficient of relative risk aversion of $\sigma = 2$ suggests individuals are willing to pay around a 50-60% markup for unemployment insurance. For higher values of risk aversion ($\sigma = 3$) this increases to 70-90%; for lower values ($\sigma = 1$) this decreases to 20-30%.

Exploiting Spousal Labor Supply Responses Not only does food expenditure respond to learning about future job loss, but also spouses are 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 6. To do so, 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

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

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

³³Online 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 IV.

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% 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

For the empirical analysis, I use a baseline value of $\epsilon^{semi} = 0.5$, following Kleven et al. (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 VII reports the coefficients, $\beta = \frac{dLFP}{dZ}$, from equation (3) of around 0.025 that were presented in Table III. Because the elicitation, 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 elicitation 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, $\frac{var(Z|X)}{var(P|X)}$. The unbiasedness of true beliefs implies that $cov(Z, U|X) = cov(P, U|X) = var(P|X)$. 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 $\frac{var(Z|X)}{var(P|X)} = \frac{var(Z|X)}{cov(Z, U|X)}$,

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

I estimate $var(Z|X)$ 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 $\frac{var(Z|X)}{cov(Z, U|X)} \approx 12$, as reported in the second set of rows in Table VII.

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-

³⁴In 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 5.1. 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 over-state 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.

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

elasticity of labor supply, ϵ^{semi} , of 0.5, suggests that individuals would be willing to pay roughly a 60% markup for UI. As shown in Table VII, if labor supply is more elastic (e.g. $\epsilon^{semi} = 1$), it suggests a willingness to pay of around 30%; if labor supply is less elastic (e.g. $\epsilon^{semi} = 0.33$), it suggests a willingness to pay of around 90%.

5.3 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 that learn they might lose their job might also learn that they have lower earnings prospects even if they do not lose their job. This 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.

6 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 6.2 will provide a point estimate of the minimum pooled price ratio using a set of semi-parametric assumptions.

Before imposing the semi-parametric assumptions, Section 6.1 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 , 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 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 non-parametric lower bounds on $E[T(P)]$ and semi-parametric point estimates for $\inf T(p)$ will provide complementary evidence on the frictions that would be imposed by private information.

6.1 Non-parametric Lower Bounds for the Average Pooled Price Ratio, $E[T(P)]$

I begin by using the predictive power of the elicitation 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 elicitation, 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: (a) the elicitation contains no more information about U than does P : $\Pr\{U|X, Z, P\} = \Pr\{U|X, P\}$ and (b) true beliefs are unbiased $\Pr\{U|X, P\} = P$. Assumption (a) is a natural assumption to place on the elicitation, as it is difficult to imagine how someone could report more information than what is contained in their true beliefs. Assumption (b) is more restrictive, as individuals may have biased beliefs. I return to a discussion of the impact of biased beliefs in Section 7. Under Assumptions (a) and (b), 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 3rd 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 V then illustrates the predictive content of the elicitation by plotting the cumulative distribution of $P_Z - \Pr\{U|X\}$. The figure reveals how the information in the elicitation, 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 V: 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 adversely 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 , Appendix B.1 shows that $E[T_Z(P_Z)]$ forms a lower bound on the average pooled price ratio³⁶:

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

Equation (18) shows that the predictive content of the elicitation for U form a non-parametric lower bound for the average pooled price ratio. This lower bound, $E[T_Z(P_Z)]$, only requires that the elicitation contain no more information than the true beliefs (Assumption (a)) and that true beliefs are unbiased (Assumption (b)). 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 .

Results For the baseline demographic and job characteristics controls, the average markup imposed by the presence of worse risks is at least 77% (s.e. 5.2%), suggesting $E[T(P)] \geq 1.77$.³⁷ The top row of Table VIII presents these results and shows that adding health controls or dropping the job characteristic controls do not meaningfully change the estimates (72% and 80% in Columns (2) and (3)). To further illustrate how the controls affect the results, Figure VI, Panel A, plots estimates of $E[T_Z(P_Z)] - 1$ on the vertical axis against the pseudo-R squared 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

³⁶This 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 $E[T(P)]$.

³⁷As 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 5 year age bins by gender. Appendix Table III (Columns (3)-(5)) documents the robustness of the results to alternative aggregation partitions.

values. This upper tail (shown in Figure V) 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% markup to cover the pooled cost of worse risks, as shown in Column (4) of Table VIII.

The remaining columns of Table VIII and panels of Figure VI explore how the estimated markups vary across subsamples. This yields estimates of $E[T_Z(P_Z)] - 1$ in excess of 50% across occupations (Figure VI, Panel B), 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 insurers use to mitigate adverse selection is to impose waiting periods on the use of the insurance policy. To explore whether this help reduce the adverse selection problem, Figure VI (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% ($p < 0.001$). 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 that an individual also files for government UI.⁴⁰ Waiting periods and alternative contract lengths would not remove the threat of adverse selection.

Another potential UI policy would pay benefits proportional to the amount of time 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 VI (Panel E) shows that this contract would also suffer significant adverse selection, with the lower bound on the average pooled price ratio above 40%. 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 only sell insurance only to observably low risks. For example, health-related insurance markets generally exclude those with pre-existing 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 VI

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

³⁹The 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.

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

(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 $E[T_Z(P_Z)] - 1$. For those with greater than 5 years of job tenure, I estimate a lower bound of 110%. It is true that this population is much less likely to lose their job (less than 2% 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 5.

6.2 Semi-parametric Point Estimates of $\inf T(p)$

Moving from lower bounds on $E[T(P)]$ 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 elicitation. Note that the observed density (p.d.f./p.m.f.) of Z and U can be written as⁴¹

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

where $f_{Z|P,X}$ is the distribution of elicitation 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 elicitation 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 .

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 elicitation 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 elicitation 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 Appendix B.2.

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

Results The baseline specification yields an estimate of $\inf T(p) - 1$ of 3.36 (s.e. 0.203), as shown in the second set of rows of Table VIII. Including health controls reduces this markup slightly to 323% (s.e. 26.8%), and using only demographic controls increases the markup to 530% (s.e. 65.5%).⁴² The high minimum pooled price ratios are robust across subsamples, as illustrated in Columns (4)-(9) of Table VIII. For example, those with longer job tenure have values of $\inf T(p) - 1$ of 474%. The minimum pooled price ratio is similar across age groups (333% for ages at or below 55 and 344% for ages above 55); and it is slightly higher for below-median wage earners (436%) than above-median wage earners (316%).

Overall, the point estimates far exceed the estimated markups individuals are willing to pay for UI documented in Section 5. Moreover, even the lower bounds in Section 6.1 generally lie at or above the willingness to pay estimates in Section 5. 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.

7 Discussion

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

Government UI The government is a major provider of UI in the U.S. The baseline analysis above considers the hypothetical market for a 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 question 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.

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, p196). This would continue to yield willingness to pay estimates that are smaller than the estimated 300% markups individuals would need to be willing to pay to overcome adverse selection. Assuming the removal of government UI does not significantly affect

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

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 non-existence. As noted in Section 4, 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 non-existence 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 Section 5 and 6.2 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 4 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.

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

⁴³It 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.

⁴⁴This theory is made suggestively, and the paper does not provide direct evidence of an impact of government regulation.

separation rates vary across years; but the R-squared of a regression of job loss on year dummies yields an R-squared 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% 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 individuals may have biased beliefs about their unemployment and job prospects (e.g. Stephens (2004); Spinnewijn (2015)).

To extend the model in Section 4 to include biased beliefs, suppose p is the objective belief for a given individual. But, following Kahneman and Tversky (1979), let $q(p)$ denote the re-weighted probability function that governs their decisions over financial assets. It is straightforward to show that the no trade condition becomes:

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

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 individuals would be willing to pay for UI: they may over (or under) state 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 $T(p) = 1 + 3.3 = 430\%$ and markup individuals are willing to pay less than 75%, 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.

⁴⁵The 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.

8 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 to sell such policies, the empirical results suggest that they would be too heavily adversely selected to deliver a positive profit at any price. Figure VII compares the estimates of the minimum pooled price ratio to 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 pre-existing conditions that would cause them to be rejected by an insurance company, the estimated markups are 42% for life, 66% for disability, and 83% 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 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 severance⁴⁶ 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, D. and R. Shimer (1999). Efficient unemployment insurance. *Journal of Political*. 5
- Acemoglu, D. and R. Shimer (2000). Productivity gains from unemployment insurance. *European Economic Review* 44(7), 1195–1224. 5
- Aguiar, M. and E. Hurst (2005). Consumption versus expenditure. *Journal of Political Economy* 113(5), 919–948. 1, 1, 8, 5.2

⁴⁶This 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.

- Akerlof, G. (1970). The market for lemons: Qualitative uncertainty and the market mechanism. *Quarterly Journal of Economics* 84(3), 488–500. 4.2
- Bach, K. (1998). Negativauslese und tarifiedifferenzierung im versicherungs-sektor. *DUV, Schesslitz*. 5
- Baily, M. (1976). Some aspects of optimal unemployment insurance. *Journal of Public Economics* 10, 379–402. 1, 5, 1
- Barceló, C. and E. Villanueva (2010). The response of household wealth to the risk of losing the job: Evidence from differences in firing costs. 6
- Basten, C. C., A. Fagereng, and K. Telle (2012). Saving and portfolio allocation before and after job loss. *KOF Swiss Economic Institute Working Paper* (298). 6
- Bloemen, H. G. and E. G. Stancanelli (2005). Financial wealth, consumption smoothing and income shocks arising from job loss. *Economica* 72(287), 431–452. 6
- Blundell, R., M. Costa Dias, C. Meghir, and J. Shaw (2016, Feb). Female labour supply, human capital and welfare reform. 5.2
- Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *The American Economic Review*, 1887–1921. 5
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2012). Consumption inequality and family labor supply. *National Bureau of Economic Research Working Paper No. 18445*. 1
- Bollinger, C. R. (1998). Measurement error in the current population survey: A nonparametric look. *Journal of Labor Economics* 16(3), 576–594. 56
- Brown, J. and D. A. Matsa (2015). Boarding a sinking ship? an investigation of job applications to distressed firms. *Journal of Finance (Forthcoming)*. 6
- Browning, M. and T. Crossley (2001). Unemployment insurance levels and consumption changes. *Journal of Public Economics* 80, 1–23. 1
- Cabral, M. (2013). Claim timing and ex post adverse selection. *Working Paper*. 40
- Campos, R. G. and I. Reggio (2015). Consumption in the shadow of unemployment. *European Economic Review*. 6
- Carroll, C. D., K. E. Dynan, and S. D. Krane (2003). Unemployment risk and precautionary wealth: Evidence from households' balance sheets. *Review of Economics and Statistics* 85(3), 586–604. 6
- Carroll, C. D. and A. A. Samwick (1997). The nature of precautionary wealth. *Journal of Monetary Economics* 40(1), 41–71. 6
- Carroll, C. D. and A. A. Samwick (1998). How important is precautionary saving? *Review of Economics and Statistics* 80(3), 410–419. 6
- Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics* 90(10), 1879–1901. 5, 5.1.1

- Chetty, R. (2008). Moral hazard vs. liquidity and optimal unemployment insurance. *Journal of Political Economy* 116(2), 173–234. 1, 5
- Chetty, R. and A. Szeidl (2007). Consumption commitments and risk preferences. *Quarterly Journal of Economics* 122(2). 2.2, 5.1.1, 28
- Chiappori, P. and B. Salanié (2000). Testing for asymmetric information in insurance markets. *Journal of Political Economy*, 56–78. 1
- Cochrane, J. H. (1995). Time-consistent health insurance. *Journal of Political Economy*, 445–473. 7
- Davis, S. J. and T. M. von Wachter (2011). Recessions and the cost of job loss. *NBER Working Paper*. 20
- Dynan, K. E. (1993). How prudent are consumers? *Journal of Political Economy*, 1104–1113. 6
- Einav, L., A. Finkelstein, and M. Cullen (2010). Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics* 125(3), 877–921. 4.2
- Einav, L., A. Finkelstein, and J. Levin (2010, Sep). Beyond testing: Empirical models of insurance markets. *Annual Review of Economics* 2(1), 311–336. 1
- Engen, E. M. and J. Gruber (2001). Unemployment insurance and precautionary saving. *Journal of Monetary Economics* 47(3), 545–579. 6
- Finkelstein, A. and K. McGarry (2006). Multiple dimensions of private information: Evidence from the long-term care insurance market. *American Economic Review* 96(4), 938–958. 2.1
- Finkelstein, A. and J. Poterba (2004, Jan). Adverse selection in insurance markets: Policyholder evidence from the UK annuity market. *Journal of Political Economy*. 1
- Gallen, Y. (2013). Anticipating unemployment: Savings evidence from denmark. *Northwestern University*. 6
- Gan, L., M. Hurd, and D. McFadden (2005, Aug). Individual subjective survival curves. *Analyses in the Economics of Aging*, ed. D. Wise. 2.1, B.2
- Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review* 87(1), 192–205. 1, 1, 1, 2.2, 3.3, 5.1.1, 5.1.1, 28, 5.1.3, 7
- Gruber, J. and J. B. Cullen (1996). Spousal labor supply as insurance: Does unemployment insurance crowd out the added worker effect? *NBER Working Paper*. 3.2, 5
- Guariglia, A. and B.-Y. Kim (2004). Earnings uncertainty, precautionary saving, and moonlighting in russia. *Journal of Population Economics* 17(2), 289–310. 6
- Guiso, L., T. Jappelli, and D. Terlizzese (1992). Earnings uncertainty and precautionary saving. *Journal of Monetary Economics* 30(2), 307–337. 6
- He, D. (2009, Jan). The life insurance market: Asymmetric information revisited. *Journal of Public Economics*. 2.1
- Hendren, N. (2013). Private information and insurance rejections. *Econometrica* 81(5), 1713–1762. 1, 1, 2.1, 2.1, 4, 4.1, 4.2, 6, 6.1, 36, 37, 6.1, 6.2, 6.2, 8, A, A.2, B.1.2, B.2, B.2

- Hu, Y. and S. Schennach (2008). Instrumental variable treatment of nonclassical measurement error models. *Econometrica* 76(1), 195–216. 56
- Hubbard, R. G., J. Skinner, and S. P. Zeldes (1994). Precautionary saving and social insurance. *NBER Working Paper*. 6
- Hurd, M. (2009). Subjective probabilities in household surveys. *Annual Review of Economics*, 543–562. 2.3
- Kahneman, D. and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, 263–291. 7
- Kleven, H. J., C. T. Kreiner, and E. Saez (2009). The optimal income taxation of couples. *Econometrica* 77(2), 537–560. 5.2
- Landais, C., P. Michaillat, and E. Saez (2010). Optimal unemployment insurance over the business cycle. 5
- Lusardi, A. (1997). Precautionary saving and subjective earnings variance. *Economics Letters* 57(3), 319–326. 6
- Lusardi, A. (1998). On the importance of the precautionary saving motive. *American Economic Review*, 449–453. 6
- Manski, C. F. (2004). Measuring expectations. *Econometrica*, 1329–1376. 1
- Manski, C. F. and F. Molinari (2010). Rounding probabilistic expectations in surveys. *Journal of Business and Economic Statistics* 28(2), 219–231. B.2
- Meghir, C. and L. Pistaferri (2011). Earnings, consumption and life cycle choices. *Handbook of Labor Economics* 4, 773–854. 5.1.3, C.3
- Pistaferri, L. (2001). Superior information, income shocks, and the permanent income hypothesis. *Review of Economics and Statistics* 83(3), 465–476. 1
- Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons Inc., New York. 2.1
- Shavell, S. (1979). On moral hazard and insurance. *The Quarterly Journal of Economics*, 541–562. 4, 4.2, A
- Shimer, R. and I. Werning (2007). Reservation wages and unemployment insurance. *The Quarterly Journal of Economics*, 1145–1185. 5
- Shimer, R. and I. Werning (2008). Liquidity and insurance for the unemployed. *The American Economic Review* 98(5), 1922–1942. 5
- Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association* 13(1), 130–167. 7
- Stephens, M. (2001). The long-run consumption effects of earnings shocks. *Review of Economics and Statistics* 83(1), 28–36. 1, 3.3, 20
- Stephens, M. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics* 20(3), 504–537. 1, 3.2, 34

- Stephens, M. (2004). Job loss expectations, realizations, and household consumption behavior. *Review of Economics and statistics* 86(1), 253–269. 3, 1, 3.1, 7
- Zeldes, S. P. (1989). Consumption and liquidity constraints: an empirical investigation. *The Journal of Political Economy*, 305–346. 5.1.3, 5.1.3, C.3

TABLE I
Sample Summary Statistics

Variable	Panel 1: HRS Full Sample		Panel 2: HRS Married Sample		Panel 3: PSID Sample	
	mean	std dev	mean	std dev	mean	std dev
<i>Characteristics</i>						
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
<i>Unemployment/Job Loss Outcomes</i>						
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
<i>Other Variables</i>						
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		
<i>Sample Size</i>						
Number of Observations		26,640		11,049		65,483
Number of Households		3,467		2,214		9,557

Notes: This table presents summary statistics for the samples used in the paper. Panel 1 presents the summary statistics for the baseline HRS sample of individuals; Panel 2 presents the summary statistics for the HRS married subsample used to study spousal labor supply responses; Panel 3 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 dollars using the CPI-U-RS.

TABLE II
Presence of Private Information about Future Job Loss

	Alternative Controls				Sub-Samples					
	Baseline (1)	Demo (2)	Health (3)	Ind FE (4)	Age ≤ 55 (5)	Age > 55 (6)	Below Median Wage (7)	Above Median Wage (8)	Tenure ≥ 5 yrs (9)	Tenure < 5 yrs (10)
Elicitation (Z)	0.0836*** (0.00675)	0.0956*** (0.00685)	0.0822*** (0.00736)	0.0715*** (0.0107)	0.0716*** (0.0102)	0.0914*** (0.00876)	0.0751*** (0.00812)	0.0957*** (0.0113)	0.0746*** (0.00793)	0.0906*** (0.0120)
s.e.										
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
Num of Obs.	26640	26640	22831	26640	11134	15506	13321	13319	17850	8790
Num of HHs	3467	3467	3180	3467	2255	3231	2824	2309	2952	2437

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. Columns (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 5 years of job tenure (Columns 9 and 10). Standard errors are clustered by household.

TABLE III
Spousal Labor Supply Response to Potential Job Loss

<i>Specification:</i>	Baseline (1)	Sample without Future Job Loss (2)	Full Time Work (3)	2yr Lagged Entry ("Placebo") (4)	Household Fixed Effects (5)	Individual Fixed Effects (6)	Exit (7)	Spouse Unemployment (8)
Elicitation (Z)	0.0258*** (0.00868)	0.0256*** (0.00898)	0.0255*** (0.00988)	0.00122 (0.00800)	0.0243*** (0.0114)	0.0312* (0.0180)	0.0174 (0.0119)	0.0213** (0.00966)
s.e.								
Mean Dep Var	0.0394	0.0389	0.0524	0.0394	0.0394	0.0394	0.0851	0.0296
Num of Obs.	11049	10726	11049	11049	11049	11049	11049	9079
Num of HHs	2214	2194	2214	2214	2214	2214	2214	1359

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 (a) the spouse was not working for pay in the previous wave (2 years prior) and (b) 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 zero). 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 (a) the spouse was not employed full time in the previous wave and (b) 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.

TABLE IV
Ex-Ante Drop in Food Expenditure Prior to Unemployment and Implied (Ex-Ante) Willingness to Pay for UI

Specification:	Baseline	Controls for Needs	Under 50 Sample	Job Loss	Household Income Controls	Household Head Income Controls
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Impact of Unemployment on $\log(c_{i,t})-\log(c_{i,t-2})$</i>						
Unemp	-0.0271*** (0.00975)	-0.0211** (0.0105)	-0.0288*** (0.0106)	-0.0260*** (0.00824)	-0.0272*** (0.00969)	-0.0281*** (0.00983)
s.e.						
<i>Specification Details</i>						
Sample Employed in t-2 and t-1	X	X	X	X	X	X
Controls for change in log needs (t-2 vs t-1)		X				
Change in log HH inc (t-2 vs t-1) (3rd order poly)		X	X		X	
Change in log HH head inc (t-2 vs t-1) (3rd order poly)						X
Mean Dep Var	0.000	(0.001)	0.007	0.000	0.000	0.000
Num of Obs.	65483	53327	52463	65556	65399	64119
Num of HHs	9557	8371	8512	9560	9547	9448

Notes: This Table presents estimates of the impact of unemployment in year t on consumption growth in year t-1 relative to t-2, $\log(c_{i,t})-\log(c_{i,t-2})$. 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 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.

TABLE V
Impact of Unemployment on Consumption and Implied WTP for UI

Specification:	Controls for Needs			With Outliers		Job Loss	Bound on Maximum WTP	
	Baseline (1)	Under 50 Sample (2)	Under 50 Sample (3)	With Outliers (4)	With Outliers; no food stamps (5)		Baseline (7)	With Outliers (8)
<i>Reduced Form Impact on $E[\log(c_{t+1})-\log(c_t)]$</i>								
Unemp	-0.0723*** (0.00997)	-0.0747*** (0.0102)	-0.0713*** (0.0108)	-0.0889*** (0.0123)	-0.182*** (0.0171)	-0.0487*** (0.00860)		
s.e.								
2-Sample IV Estimate for $E[\log(c_{t+1})-\log(c_t)]$	-0.090	-0.093	-0.089	-0.111	-0.227	-0.061		
$\Delta^{min} = \min\{\log(c_e(p))-\log(c_e(p))\}$							-0.137*** (0.02)	-0.146*** (0.011)
<i>Implied $[u'(c_{t+1}) - v'(c_t)]/v'(c_t)$ for various σ</i>								
$\sigma = 1$	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%
<i>Specification Details</i>								
Sample Employed in t-1	X	X	X	X	X	X	X	X
Controls for change in log needs		X						
Mean Dep Var	(0.005)	(0.005)	0.001	(0.006)	(0.007)	(0.005)	(0.005)	(0.006)
Num of Obs.	65808	58086	52860	66913	66913	65808	65808	66913
Num of HHs	9562	9131	8524	9583	9583	9562	9562	9583

Notes: This Table presents 2-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 50 and under for the specification in Column (1). Column (4) adds in the outliers in changes in food expenditure that are more than threefold. 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 $\log(c_{t+1})-\log(c_t)$ divides the first difference estimate by the estimated amount of information revealed between year t-1 and year t. Appendix Table I presents this estimate of the first stage. Using the HRS sample, this estimates are constructed using a regression of the subjective probability elicitation, Z , on an indicator for subsequent unemployment in the next 12 months, U . The implied willingness to pay estimates for $[u'(c_{t+1})-v'(c_t)]/v'(c_t)$ 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^* . 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 threefold 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).

TABLE VI
Ex-Ante Drop in Food Expenditure Prior to Unemployment and Implied (Ex-Ante) Willingness to Pay for UI

<i>Specification:</i>	Controls for			Job Loss	Household Income Controls	Household Head Income Controls
	Baseline	Needs	Under 50 Sample			
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Reduced Form Impact on $\log(c_{t-2}) - \log(c_{t-1})$</i>						
Unemp	-0.0271*** (0.00975)	-0.0211** (0.0105)	-0.0288*** (0.0106)	-0.0260*** (0.00824)	-0.0272*** (0.00969)	-0.0281*** (0.00983)
s.e.						
<i>2-Sample IV Estimate for $d[\log(c_{pne}(p))]/dp$</i>						
	-0.29	-0.23	-0.31	-0.30	-0.29	-0.29
<i>Implied $[u'(c_{it}) - v'(c_e)]/v'(c_e)$ for various σ</i>						
$\sigma = 1$	29%	23%	31%	30%	29%	29%
$\sigma = 2$	58%	46%	62%	60%	58%	59%
$\sigma = 3$	87%	70%	92%	90%	87%	88%
<i>Specification Details</i>						
Sample Employed in t-2 and t-1		X	X	X	X	X
Controls for change in log needs (t-2 vs t-1)			X			
Change in log HH inc (t-2 vs t-1) (3rd order poly)				X		
Change in log HH head inc (t-2 vs t-1) (3rd order poly)					X	
Mean Dep Var	0.000	(0.001)	0.007	0.000	0.000	0.000
Num of Obs.	65483	53327	52463	65556	65399	64119
Num of HHs	9557	8371	8512	9560	9547	9448

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 12-24 months after the survey. 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.

TABLE VII
Spousal Labor Supply Response to Potential Job Loss and Implied (Ex-Ante) Willingness to Pay for UI

Specification:	Baseline	Sample without Future Job Loss	Full Time Work	2yr Lagged Entry ("Placebo")	Household Fixed Effects	Individual Fixed Effects
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Reduced Form Relationship between U and Z</i>						
Elicitation (Z)	0.0258*** (0.00868)	0.0256*** (0.00898)	0.0255** (0.00988)	0.00122 (0.00800)	0.0243** (0.0114)	0.0312* (0.0180)
s.e.						
<i>Measurement Error Correction</i>						
Total Var / Signal Var (var(Z X)/cov(Z,U X))	12.08 (1.71)	12.08 (1.69)	12.08 (1.65)	23.68 (6.53)	12.08 (1.62)	12.08 (1.60)
bootstrap s.e.						
Implied dLFP/dp	0.312	0.309	0.308	0.029	0.293	0.377
Implied $[u'(c_0) - v'(c_e)]/v'(c_e)$ for various ϵ^{semi}						
$\epsilon^{\text{semi}} = 0.33$	94%	93%	92%	9%	88%	113%
$\epsilon^{\text{semi}} = 0.5$	62%	62%	62%	6%	59%	75%
$\epsilon^{\text{semi}} = 1$	31%	31%	31%	3%	29%	38%
Mean Dep Var	0.04	0.04	0.05	0.04	0.04	0.04
Num of Obs.	11049	10726	11049	11049	11049	11049
Num of HHs	2214	2194	2214	2214	2214	2214

Notes: This table scale 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 (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 $\text{Cov}(Z,U|X)$, which is an approximation to $\text{var}(P|X)$ 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.

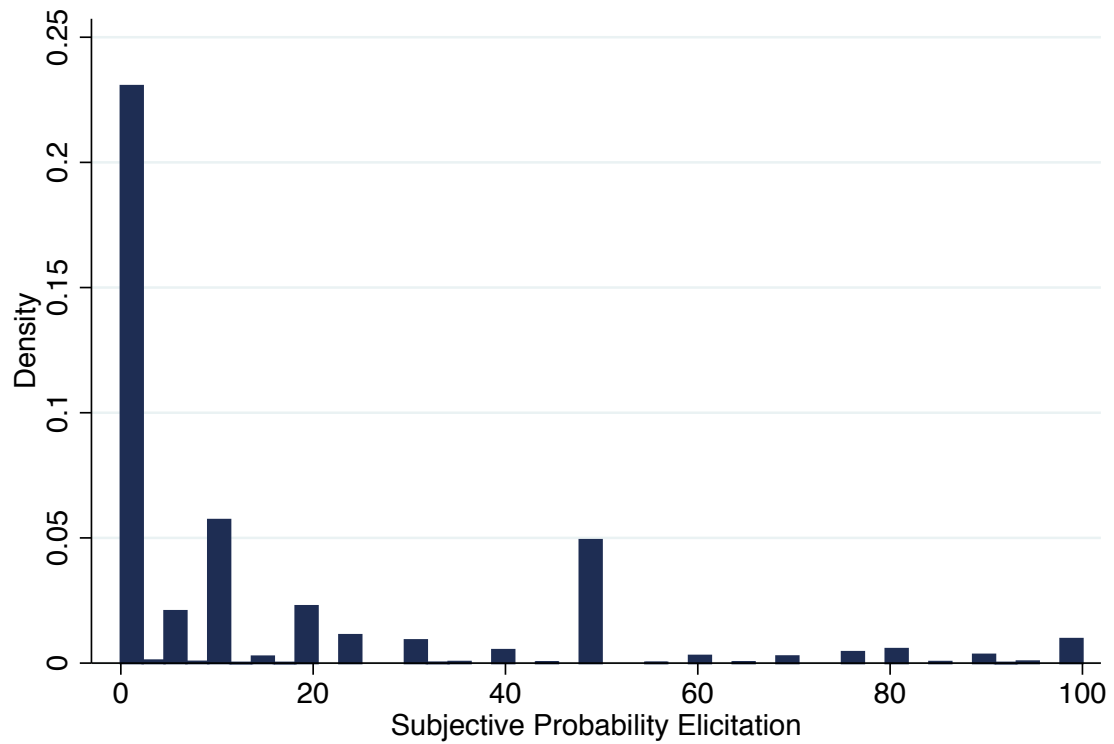
TABLE VIII
Pooled Price Ratio

Specification:	Alternative Controls				Sub-Samples					
	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)
<i>Lower Bound for E[T(P)]-I</i>										
E[T ₂ (P ₂)-I]	77%	80%	72%	39%	70%	81%	63%	100%	110%	47%
s.e.	(5.2%)	(4.9%)	(5.4%)	(4.2%)	(9%)	(7%)	(5.8%)	(9.3%)	(9.4%)	(5.7%)
<i>Point Estimate for inf T(p)-I</i>										
Inf T(p) - I	336%	530%	323%	N/A ¹	333%	344%	436%	316%	474%	374%
s.e.	(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 FE				X						
Num of Obs.	26640	26640	22831	26640	11134	15506	13321	13319	17850	8790
Num of HHS	3467	3467	3180	3467	2255	3231	2824	2309	2952	2437

Notes: Table presents estimates of the nonparametric lower bounds on E[T(P)] and semi-parametric 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(PZ) and E[mZ(PZ)] are constructed using 1000 bootstrap resamples at the household level for the minimum pooled price ratio and 500 resamples for the lower bound estimate of E[T(P)].


¹inf T(p) specification not available for the FE specification because nonlinear fixed effects cannot be partialled 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.

FIGURE I: 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 $Z = 0$.

FIGURE II: Predictive Content of Subjective Probability Elicitations: Binned Scatterplot of U versus Z , conditional on X



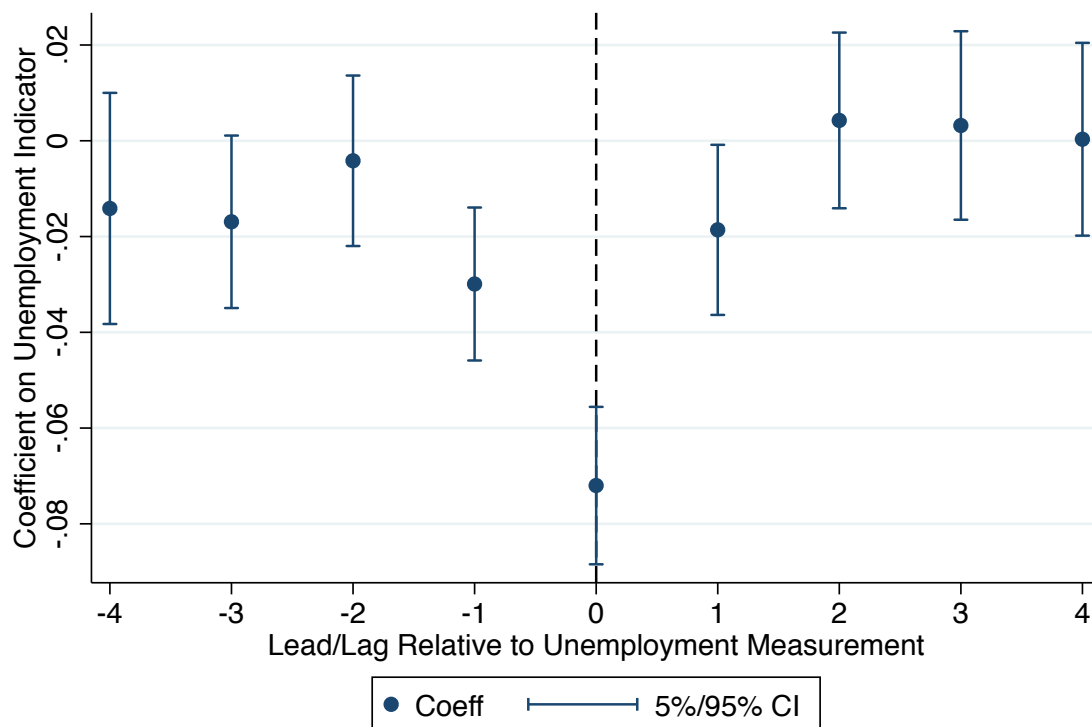
Notes: This figure reports mean unemployment rate in each elicitation category controlling for demographics, job characteristics, and year controls. To construct this figure, I run the regression in Equation (1). The figure plots the coefficients on bins of the elicitations. I omit the lowest bin (corresponding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

FIGURE III: Relationship between Potential Job Loss and Spousal Labor Supply



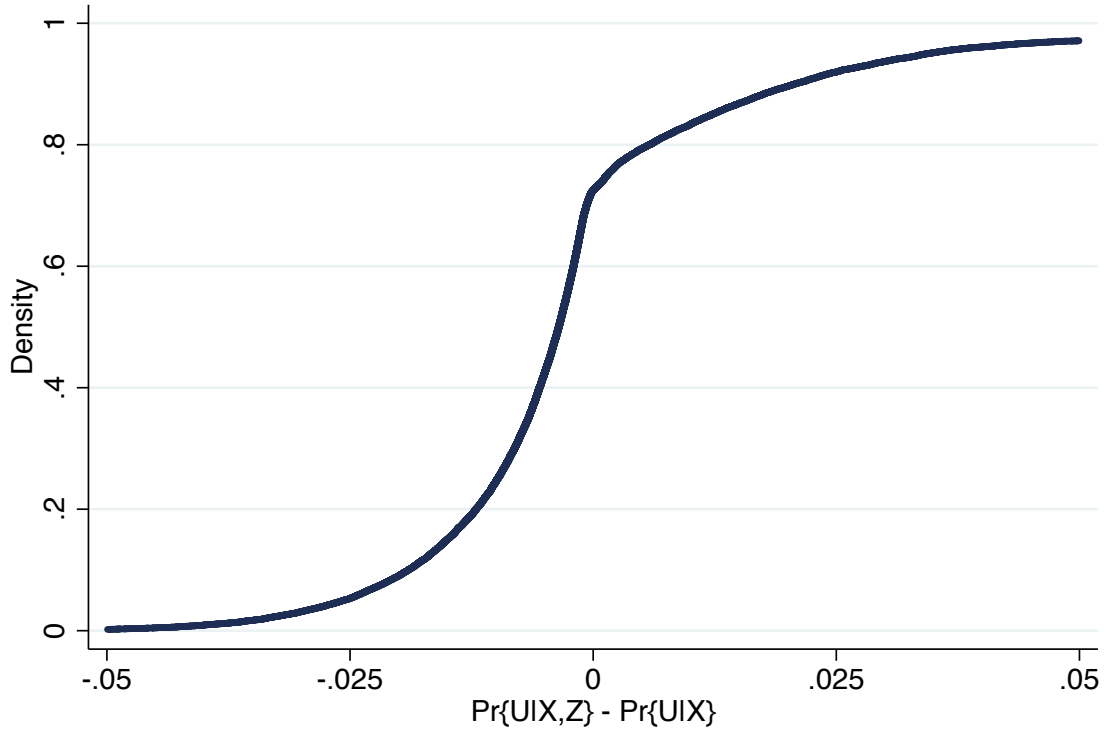
Notes: The figure present coefficients from a regression of an indicator for a spouse entering the labor force, defined as an indicator for not working in the previous wave and working in the current wave, on category indicators for the subjective probability elicitation, Z , controlling for demographics,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

FIGURE IV: Impact of Unemployment on Consumption Growth




Notes: These figures present 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 is 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 et al. (2005), 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 $t - 2$.

FIGURE V: 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 elicitation (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.

FIGURE VI: 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 sub-samples 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 $\Pr\{U|X, Z\}$ 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 $\Pr\{U|X, Z\}$. 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 sub-samples, 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 unemployment probability, $\Pr\{U\}$, for each sub-sample.

FIGURE VII: Comparison of $\inf T(p) - 1$ to Other (Non-)Existing Insurance Markets



Notes: Hendren (2013) argues private information prevents people with pre-existing 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 pre-existing condition (blue circles) and those whose observables would allow them to purchase insurance in each market (red hollow circles). Figure reports the confidence interval and the 5 / 95% confidence interval for each estimate in each sample. For the sub-samples 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, sub-samples 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.

ONLINE APPENDIX: Not For Publication

A No Trade Condition

This section provides a more formal exposition of the no trade condition in Section 4. To provide a general treatment, I begin by relaxing the condition of uni-dimensional heterogeneity in the population. Individuals are indexed by a heterogeneity parameter, θ , and make choices $\{c_{pre}(\theta), c_u(\theta), c_e(\theta), a(\theta), p(\theta)\} \in \Omega_\theta$, where the constraint set varies arbitrarily across types.

Consider a policy that provides a small payment, db , in the event of losing one's job that is financed with a small payment in the event of remaining employed, $d\tau$, offered to those with observable characteristics X . By the envelope theorem, the utility impact to type θ of buying such a policy will be given by

$$dU = -(1 - p(\theta)) v'(c_e(\theta)) d\tau + p(\theta) u'(c_u(\theta)) db$$

which will be positive if and only if

$$\frac{p(\theta) u'(c_u(\theta))}{(1 - p(\theta)) v'(c_e(\theta))} \geq \frac{d\tau}{db} \quad (21)$$

The LHS of equation (21) is a type θ 's willingness to pay (i.e. marginal rate of substitution) to move resources from the event of remaining employed to the event of job loss.⁴⁷ The RHS of equation (21), $\frac{d\tau}{db}$, is the cost per dollar of benefits of the insurance policy.

Let $\bar{\Theta}(\frac{d\tau}{db})$ denote the set of all individuals, θ , who prefer to purchase the additional insurance at price $\frac{d\tau}{db}$ (i.e. those satisfying equation (21)). An insurer's profit from a type θ is given by $(1 - p(\theta))\tau - p(\theta)b$. Hence, the insurer's marginal profit from trying to sell a policy with price $\frac{d\tau}{db}$ is given by

$$d\Pi = \underbrace{E \left[1 - p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right]}_{\text{Premiums Collected}} d\tau - \underbrace{E \left[p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right]}_{\text{Benefits Paid}} db - \underbrace{\left(dE \left[p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right] \right)}_{\text{Moral Hazard}} (\tau + b)$$

The first term is the amount of premiums collected, the second term is the benefits paid out, and the third term is the impact of additional insurance on its cost. If more insurance increases the probability of job loss, $dE[p(\theta)] > 0$, then it reduces premiums collected, τ , and increases benefits paid, b .⁴⁸

However, for the first dollar of insurance when $\tau = b = 0$, the moral hazard cost to the insurer is zero. This insight, initially noted by Shavell (1979), suggests moral hazard does not affect whether insurers' first dollar of insurance is profitable – a result akin to the logic that deadweight loss varies with the square of the tax rate.

The first dollar of insurance will be profitable if and only if

$$\frac{d\tau}{db} \geq \frac{E \left[p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right]}{E \left[1 - p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right]} \quad (22)$$

If inequality (22) does not hold for any possible price, $\frac{d\tau}{db}$, then providing private insurance will not be profitable at any price. Under the natural assumption⁴⁹ that profits are concave in b and τ , the inability to profitably sell a small amount of insurance also rules out the inability to sell larger insurance contracts.

Equation (22) characterizes no trade under an arbitrary dimensionality of unobserved heterogeneity, θ . To provide a clearer expression of how demand relates to underlying fundamentals, such as marginal rates of substitution and beliefs, it is helpful to impose a dimensionality reduction on the unobserved heterogeneity.

Assumption A1. (*Uni-dimensional Heterogeneity*) Assume the mapping $\theta \rightarrow p(\theta)$ is 1-1 and continuously differentiable in b and τ in an open ball around $b = \tau = 0$. Moreover, the marginal rate of substitution, $\frac{p}{1-p} \frac{u'(c_u(p))}{v'(c_e(p))}$, is increasing in p .

Assumption A1 states that the underlying heterogeneity can be summarized by one's belief, $p(\theta)$. In this case, the adverse selection will take a particular threshold form: the set of people who would be attracted to a contract for which type $p(\theta)$ is indifferent will be the set of higher risks whose probabilities exceed $p(\theta)$. Let P denote the

⁴⁷Note that, because of the envelope theorem, the individual's valuation of this small insurance policy is independent of any behavioral response. While these behavioral responses may impose externalities on the insurer or government, they do not affect the individuals' willingness to pay.

⁴⁸To incorporate observable characteristics, one should think of the expectations as drawing from the distribution of θ conditional on a particular observable characteristic, X .

⁴⁹See Appendix A.3 for a micro-foundation of this assumption.

random variable corresponding to the distribution of probabilities chosen in the population in the status quo world without a private unemployment insurance market, $b = \tau = 0$.⁵⁰ And, let $c_u(p)$ and $c_e(p)$ denote the consumption of types $p(\theta)$ in the unemployed and employed states of the world. Under Assumption A1, equation (22) can be re-written as:

$$\frac{u'(c_u(p))}{v'(c_e(p))} \leq T(p) \quad \forall p \quad (23)$$

where $T(p)$ is the pooled price ratio defined in Hendren (2013):

$$T(p) = \frac{E[P|P \geq p]}{E[1 - P|P \geq p]} \frac{1 - p}{p}$$

The market can exist only if there exists someone who is willing to pay the markup imposed by the presence of higher risk types adversely selecting her contract. Here, $\frac{u'(c_u(p))}{v'(c_e(p))} - 1$ is the markup individual p would be willing to pay and $T(p) - 1$ is the markup that would be imposed by the presence of risks $P \geq p$ adversely selecting the contract. This suggests the pooled price ratio, $T(p)$, is the fundamental empirical magnitude desired for understanding the frictions imposed by private information.

The remainder of this Appendix further discusses the generality of the no trade condition. A.1 discusses multi-dimensional heterogeneity. Appendix A.2 also discusses the ability of the firm to potentially offer menus of insurance contracts instead of a single contract to screen workers. Appendix A.3 illustrates that while in principle the no trade condition does not rule out non-marginal insurance contracts (i.e. b and $\tau > 0$), in general a monopolist firm's profits will be concave in the size of the contract; hence the no trade condition also rules out larger contracts.

A.1 Multi-Dimensional Heterogeneity and Robustness to Outlier Willingness to Pay

In reality, there are many reasons beyond one's chance of job loss that drive differences in willingness to pay. To understand the impact of multidimensional heterogeneity, this section solves for the no-trade condition in the case where there is an (unbounded) distribution of $\frac{u'(c_u(\theta))}{v'(c_e(\theta))}$ among the set of those with the same risk type, $p(\theta)$. In this case, there is heterogeneous willingness to pay for additional UI for different types θ with the same $p(\theta)$.

I show that there exists a mapping, $f(p) : A \rightarrow \Theta$, that maps $A \subset [0, 1]$ into the interior of the type space, Θ , such that the no trade condition reduces to testing

$$\frac{u'(c_u(f(p)))}{v'(c_e(f(p)))} \leq T(p) \quad \forall p \quad (24)$$

In this sense, even though some types are willing to pay an unboundedly high amount for UI, their extreme willingness to pay does not directly affect the no trade condition. Rather, one needs to search through an interior subset of the type space. Hence, if there are sufficiently many people of risk type p with very high willingness to pay, then one would expect the type $f(p)$ to be willing to pay the pooled cost of worse risks, so that equation (24) will not hold. But, the results illustrate that a simple addition of individuals with outlier willingness to pay for UI will not open up a market unless there are sufficiently many other types with the similar risk type that are also willing to pay the pooled cost of worse risks.

I prove this result as follows. First, I assume for simplicity that the distribution of $p(\theta)$ has full support on $[0, 1]$ and the distribution of $\frac{u'(c_u(\theta))}{v'(c_e(\theta))}$ has full support on $[0, \infty)$ (this is not essential, but significantly shortens the proof – note this allows for some individuals with unboundedly high willingness to pay). Now, fix a particular policy, $\frac{d\tau}{db}$, and consider the set of θ that are willing to pay for this policy:

$$E \left[p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right]$$

Without loss of generality, there exists a function $\tilde{p} \left(\frac{d\tau}{db} \right)$ such that

$$E \left[p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right] = E \left[p(\theta) \mid p(\theta) \geq \tilde{p} \left(\frac{d\tau}{db} \right) \right]$$

⁵⁰In other words, the random variable P is simply the random variable generated by the choices of probabilities, $p(\theta)$, in the population.

so that the average probability of the types selecting $\frac{d\tau}{db}$ is equal to the average cost of all types above $\tilde{p}(\frac{d\tau}{db})$. Without loss of generality, one can assume that \tilde{p} is strictly increasing in $\frac{d\tau}{db}$ so that \tilde{p}^{-1} exists.⁵¹

I construct $f(p) : A \rightarrow \Theta$ as follows. Define A to be the range of \tilde{p} when taking values of $\frac{d\tau}{db}$ ranging from 0 to ∞ . For each p , define $f(p)$ to be a value(s) of θ such that the willingness to pay equals $\tilde{p}^{-1}(p)$:

$$\frac{p}{1-p} \frac{u'(c_e(f(p)))}{v'(c_u(f(p)))} = \tilde{p}^{-1}(p)$$

Now, suppose $\tilde{p}^{-1}(p) \leq T(p)$ for all p . One needs to establish that inequality (22) does not hold for any $\frac{d\tau}{db}$:

$$\frac{d\tau}{db} \leq \frac{E[p(\theta) | \theta \in \bar{\Theta}(\frac{d\tau}{db})]}{E[1-p(\theta) | \theta \in \bar{\Theta}(\frac{d\tau}{db})]}$$

To see this, note that

$$\frac{E[p(\theta) | \theta \in \bar{\Theta}(\frac{d\tau}{db})]}{E[1-p(\theta) | \theta \in \bar{\Theta}(\frac{d\tau}{db})]} = \frac{E[p(\theta) | p(\theta) \geq \tilde{p}(\frac{d\tau}{db})]}{1 - E[p(\theta) | p(\theta) \geq \tilde{p}(\frac{d\tau}{db})]}$$

so that we wish to show that

$$\frac{E[p(\theta) | p(\theta) \geq \tilde{p}(\frac{d\tau}{db})]}{1 - E[p(\theta) | p(\theta) \geq \tilde{p}(\frac{d\tau}{db})]} \geq \frac{d\tau}{db} \quad (25)$$

for all $\frac{d\tau}{db}$. Note that the set A is generated by the variation in $\frac{d\tau}{db}$, so that testing equation (25) is equivalent to testing this equation for all p in the range of A :

$$\frac{E[p(\theta) | p(\theta) \geq p]}{1 - E[p(\theta) | p(\theta) \geq p]} \geq \tilde{p}^{-1}(p) \quad \forall p \in A$$

which is equivalent to

$$\frac{E[p(\theta) | p(\theta) \geq p]}{1 - E[p(\theta) | p(\theta) \geq p]} \geq \frac{p}{1-p} \frac{u'(c_e(f(p)))}{v'(c_u(f(p)))} \quad \forall p \in A$$

which proves the desired result.

Intuitively, it is sufficient to check the no trade condition for the set of equivalent classes of types with the same willingness to pay for $\frac{d\tau}{db}$. Within this class, there exists a type that one can use to check the simple uni-dimensional no trade condition.

A.2 Robustness to Menus

Here, I illustrate how to nest the model into the setting of Hendren (2013), then apply the no trade condition in Hendren (2013) to rule out menus in this more complex setting with moral hazard. I assume here that there are no additional choices, a , other than the choice p , although the presence of such additional choices should not alter the proof as long as they are not observable to the insurer. With this simplification, the only distinction relative to Hendren (2013) is the introduction of the moral hazard problem in choosing p . This section shows that allowing p to be a choice doesn't make trade any easier than in a world where $p(\theta)$ is exogenous and not affected by the insurer's contracts; hence the no trade condition results from Hendren (2013) can be applied to rule out menus.

I consider the maximization program of a monopolist insurer offering a menu of insurance contracts. Whether there exists any implementable allocations other than the endowment corresponds to whether there exists any allocations other than the endowment which maximize the profit, π , subject to the incentive and participation constraints.

Without loss of generality, the insurer can offer a menu of contracts to screen types, $\{\nu(\theta), \Delta(\theta)\}_{\theta \in \Gamma}$ where $\nu(\theta)$ specifies a total utility provided to type θ , $v(\theta) = p(\theta)u(c_u(\theta)) + (1-p(\theta))v(c_e(\theta)) - \Psi(p; \theta)$, and $\Delta(\theta)$ denotes the difference in utilities if the agent becomes unemployed, $\Delta(\theta) = u(c_u(\theta)) - v(c_e(\theta))$. Note that $\nu(\theta)$ implicitly contains the disutility of effort.

Given the menu of contracts offered by the insurer, individuals choose their likelihood of unemployment. Let $\hat{q}(\Delta, \theta)$ denote the choice of probability of employment for a type θ given the utility difference between employment and unemployment, Δ , so that the agent's effort cost is $\Psi(\hat{q}(\Delta; \theta); \theta)$. Note that a type θ that accepts a contract containing Δ will choose a probability of employment $\hat{q}(\Delta; \theta)$ that maximizes their utility. I assume that \hat{q} is weakly increasing in Δ for all θ .

⁵¹If \tilde{p} is not strictly increasing (e.g. because of "advantageous selection"), it will be strictly more profitable to an insurance company to sell the insurance at a higher price. Hence, one need not test the no trade condition for such intermediate values of $\frac{d\tau}{db}$ where \tilde{p} is decreasing in p .

Let $C_u(x) = u^{-1}(x)$ and $C_e(x) = v^{-1}(x)$ denote the consumption levels required in the employed and unemployed state to provide utility level x . Let $\pi(\Delta, \nu; \theta)$ denote the profits obtained from providing type θ with contract terms ν and Δ , given by

$$\pi(\Delta, \nu; \theta) = \hat{q}(\Delta; \theta) (c_e^e - C_e(v - \Psi(\Delta; \theta))) + (1 - \hat{q}(\Delta; \theta)) (c_u^e - C_u(\nu - \Delta - \Psi(\Delta; \theta)))$$

Note that the profit function takes into account how the agents' choice of p varies with Δ .

Throughout, I maintain the assumption that profits of the monopolist are concave in (ν, Δ) . Such concavity can be established in the general case when u is concave and individuals do not choose p (see Hendren (2013)). But, allowing individuals to make choices, p , introduces potential non-convexities into the analysis. However, it is natural to assume that if a large insurance contract would be profitable, then so would a small insurance contract. In Section A.3 below, I show that global concavity of the firm's profit function follows from reasonable assumptions on the individuals' utility function. Intuitively, what ensures global concavity is to rule out a case where small amounts of insurance generate large increases in marginal utilities (and hence increase the demand for insurance).

I prove the sufficiency of the no trade condition for ruling out trade by mapping it into the setting of Hendren (2013). To do so, define $\tilde{\pi}(\nu, \Delta; \theta)$ to be the profits incurred by the firm in the alternative world in which individuals choose p as if they faced their endowment (i.e. face no moral hazard problem). Now, in this alternative world, individuals still obtain total utility ν by construction, but must be compensated for their lost utility from effort because they can't re-optimize. But, note this compensation is second-order by the envelope theorem. Therefore, the marginal profitability for sufficiently small insurance contracts is given by

$$\pi(\nu, \Delta; \theta) \leq \tilde{\pi}(\nu, \Delta; \theta)$$

Now, define the aggregate profits to an insurer that offers menu $\{\nu(\theta), \Delta(\theta)\}_\theta$ by

$$\Pi(\nu(\theta), \Delta(\theta)) = \int \pi(\nu(\theta), \Delta(\theta); \theta) d\mu(\theta)$$

and in the world in which p is not affected by Π ,

$$\tilde{\Pi}(\nu(\theta), \Delta(\theta)) = \int \tilde{\pi}(\nu(\theta), \Delta(\theta); \theta) d\mu(\theta)$$

So, for small variations in ν and Δ , we have that

$$\Pi(\nu(\theta), \Delta(\theta)) \leq \tilde{\Pi}(\nu(\theta), \Delta(\theta))$$

because insurance causes an increase in p . Now, Hendren (2013) shows that the no trade condition implies that $\tilde{\Pi} \leq 0$ for all menus, $\{\nu(\theta), \Delta(\theta)\}$. Therefore, the no trade condition also implies $\Pi \leq 0$ for local variations in the menu $\{\nu(\theta), \Delta(\theta)\}$ around the endowment. Combining with the concavity assumption, this also rules out larger deviations.

Conversely, if the no trade condition does not hold, note that the behavioral response is continuous in Δ , so that sufficiently small values of insurance allow for a profitable insurance contract to be traded.

A.3 Concavity Assumption and Sufficient Conditions for Concavity

The presence of moral hazard in this multi-dimensional screening problem induces the potential for non-convexities in the constraint set. Such non-convexities could potentially limit the ability of local variational analysis to characterize the set of implementable allocations. To be specific, let $\pi(\Delta, \mu; \theta)$ denote the profit obtained from type θ if she is provided with total utility μ and difference in utilities Δ ,

$$\pi(\Delta, \mu; \theta) = (1 - \hat{p}(\Delta; \theta)) (c_e^e - C_v(\mu - \Psi(1 - \hat{p}(\Delta; \theta)))) + \hat{p}(\Delta; \theta) (c_u^e - C_u(\mu - \Delta - \Psi(1 - \hat{p}(\Delta; \theta))))$$

To guarantee the validity of our variational analysis for characterizing when the endowment is the only implementable allocation, it will be sufficient to require that $\pi(\Delta, \mu; \theta)$ is concave in (Δ, μ) .

Assumption. $\pi(\Delta, \mu; \theta)$ is concave in (Δ, μ) for each θ

This assumption requires the marginal profitability of insurance to decline in the amount of insurance provided. If the agents choice of p is given exogenously (i.e. does not vary with Δ), then concavity of the utility functions, u and v , imply concavity of $\pi(\Delta, \mu; \theta)$. Assumption A.3 ensures that this extends to the case when p is a choice and can respond to θ .

Claim. If $\Psi'''(q; \theta) > 0$ for all θ and $\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$ then π is globally concave in (μ, Δ) .

For simplicity, we consider a fixed θ and drop reference to it. Profits are given by

$$\pi(\Delta, \mu) = \hat{q}(\Delta) (c_e^e - C_e(\mu - \Psi(\hat{q}(\Delta)))) + (1 - \hat{q}(\Delta)) (c_u^e - C_u(\mu - \Delta - \Psi(\hat{q}(\Delta))))$$

The goal is to show the Hessian of π is negative semi-definite. I proceed in three steps. First, I derive conditions which guarantee $\frac{\partial^2 \pi}{\partial \Delta^2} < 0$. Second, I show that, in general, we have $\frac{\partial^2 \pi}{\partial \mu^2} < 0$. Finally, I show the conditions provided to guarantee $\frac{\partial^2 \pi}{\partial \Delta^2} < 0$ also imply the determinant of the Hessian is positive, so that both eigenvalues of the Hessian must be negative and thus the matrix is negative semi-definite.

A.3.1 Conditions that imply $\frac{\partial^2 \pi}{\partial \Delta^2} < 0$

Taking the first derivative with respect to Δ , we have

$$\begin{aligned} \frac{\partial \pi}{\partial \Delta} &= \frac{\partial \hat{q}}{\partial \Delta} (c_e^e - c_u^e + C_u(\mu - \Delta - \Psi(\hat{q}(\Delta)))) \\ &\quad - (1 - \hat{q}(\Delta)) C'_u(\mu - \Delta - \Psi(\hat{q}(\Delta))) - \hat{q}(\Delta) C'_e(\mu - \Psi(\hat{q}(\Delta))) \end{aligned}$$

Taking another derivative with respect to Δ , applying the identity $\Delta = \Psi'(\hat{p}(\Delta))$, and collecting terms yields

$$\begin{aligned} \frac{\partial^2 \pi}{\partial \Delta^2} &= - \left[(1 - \hat{q}(\Delta)) (1 + \Delta)^2 C''_u(\mu - \Delta - \Psi(\hat{q}(\Delta))) + \hat{q}(\Delta) (\Delta \hat{q}'(\Delta))^2 C''(\mu - \Psi(\hat{q}(\Delta))) \right] \\ &\quad + \frac{\partial \hat{q}}{\partial \Delta} \left[(1 - \hat{q}(\Delta)) C'(\mu - \Delta - \Psi(\hat{q}(\Delta))) + \hat{q}(\Delta) C'(u - \Psi(\hat{q}(\Delta))) - (2 + 2\Delta \hat{q}'(\Delta)) C'(\mu - \Delta - \Psi(\hat{q}(\Delta))) \right] \\ &\quad + \frac{\partial^2 \hat{q}}{\partial \Delta^2} \left[c_e^e - c_u^e + C(\mu - \Delta - \Psi(\hat{q}(\Delta))) - C(\mu - \Psi(\hat{q}(\Delta))) + (1 - \hat{q}(\Delta)) \Delta C'(\mu - \Delta - \Psi(\hat{q}(\Delta))) + \hat{q}(\Delta) C'(\mu - \Psi(\hat{q}(\Delta))) \right] \end{aligned}$$

We consider these three terms in turn. The first term is always negative because $C'' > 0$. The second term, multiplying $\frac{\partial \hat{q}}{\partial \Delta}$, can be shown to be positive if

$$(1 + \hat{q}(\Delta)) C'(\mu - \Delta - \Psi(\hat{q}(\Delta))) \geq \hat{q}(\Delta) C'(\mu - \Delta)$$

which is necessarily true whenever

$$\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$$

This inequality holds as long as people are willing to pay less than a 100% markup for a small amount of insurance, evaluated at their endowment.

Finally, the third term is positive as long as $\Psi''' > 0$. To see this, one can easily verify that the term multiplying $\frac{\partial^2 \hat{q}}{\partial \Delta^2}$ is necessarily positive. Also, note that $\frac{\partial^2 \hat{q}}{\partial \Delta^2} = \frac{-\Psi'''}{(\Psi'')^2}$. Therefore, if we assume that $\Psi''' > 0$, the entire last term will necessarily be negative. In sum, it is sufficient to assume $\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$ and $\Psi''' > 0$ to guarantee that $\frac{\partial^2 \pi}{\partial \Delta^2} < 0$.

A.3.2 Conditions that imply $\frac{\partial^2 \pi}{\partial \mu^2} < 0$

Fortunately, profits are easily seen to be concave in μ . We have

$$\frac{\partial \pi}{\partial \mu} = - (1 - \hat{q}(\Delta)) C'(\mu - \Delta - \Psi(\hat{q}(\Delta))) - \hat{q}(\Delta) C'(\mu - \Psi(\hat{q}(\Delta)))$$

so that

$$\frac{\partial^2 \pi}{\partial \mu^2} = - (1 - \hat{q}(\Delta)) C''(\mu - \Delta - \Psi(\hat{q}(\Delta))) - \hat{q}(\Delta) C''(\mu - \Psi(\hat{q}(\Delta)))$$

which is negative because $C'' > 0$.

A.3.3 Conditions to imply $\frac{\partial^2 \pi}{\partial \Delta^2} \frac{\partial^2 \pi}{\partial \mu^2} - \left(\frac{\partial^2 \pi}{\partial \Delta \partial \mu} \right) > 0$

Finally, we need to ensure that the determinant of the Hessian is positive. To do so, first note that

$$\frac{\partial^2 \pi}{\partial \mu \partial \Delta} = (1 - \hat{q}(\Delta)) C''(\mu - \Delta - \Psi(\hat{q}(\Delta))) (1 + \Delta \hat{q}'(\Delta)) + \hat{q}(\Delta) C''(\mu - \Psi(\hat{q}(\Delta))) \Delta \hat{q}'(\Delta)$$

Also, we note that under the assumptions $\Psi''' > 0$ and $\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$, we have the inequality

$$\frac{\partial^2 \pi}{\partial \Delta^2} < - \left[(1 - \hat{q}(\Delta)) (1 + \Delta)^2 C_u'' (\mu - \Delta - \Psi(\hat{q}(\Delta))) + \hat{q}(\Delta) (\Delta \hat{q}'(\Delta))^2 C'' (\mu - \Psi(\hat{q}(\Delta))) \right]$$

Therefore, we can ignore the longer terms in the expression for $\frac{\partial^2 \pi}{\partial \Delta^2}$ above. We multiply the RHS of the above equation with the value of $\frac{\partial^2 \pi}{\partial \mu^2}$ and subtract $\left(\frac{\partial^2 \pi}{\partial \Delta \partial \mu} \right)^2$. Fortunately, many of the terms cancel out, leaving the inequality

$$\begin{aligned} \frac{\partial^2 \pi}{\partial \Delta^2} \frac{\partial^2 \pi}{\partial \mu^2} - \left(\frac{\partial^2 \pi}{\partial \Delta \partial \mu} \right)^2 &\geq (1 - \hat{q}(\Delta)) \hat{q}(\Delta) (1 + \Delta \hat{q}'(\Delta))^2 C'' (\mu - \Delta - \Psi(\hat{q}(\Delta))) C'' (\mu - \Psi(\hat{q}(\Delta))) \\ &\quad + \hat{q}(\Delta) (1 - \hat{q}(\Delta)) (\Delta \hat{q}'(\Delta))^2 C'' (\mu - \Psi(\hat{q}(\Delta))) C'' (\mu - \Delta - \Psi(\hat{q}(\Delta))) \\ &\quad - 2(1 - \hat{q}(\Delta)) \hat{q}(\Delta) (1 + \Delta \hat{q}'(\Delta)) \Delta \hat{q}'(\Delta) C'' (\mu - \Delta - \Psi(\hat{q}(\Delta))) C'' (\mu - \Psi(\hat{q}(\Delta))) \end{aligned}$$

which reduces to the inequality

$$\frac{\partial^2 \pi}{\partial \Delta^2} \frac{\partial^2 \pi}{\partial \mu^2} - \left(\frac{\partial^2 \pi}{\partial \Delta \partial \mu} \right)^2 \geq \hat{q}(\Delta) (1 - \hat{q}(\Delta)) C'' (\mu - \Delta - \Psi(\hat{q}(\Delta))) C'' (\mu - \Psi(\hat{q}(\Delta))) K(\mu, \Delta)$$

where

$$\begin{aligned} K(\mu, \Delta) &= (1 + \Delta \hat{q}'(\Delta))^2 + (\Delta \hat{q}'(\Delta))^2 - 2\Delta \hat{q}'(\Delta) - 2(\Delta \hat{q}'(\Delta))^2 \\ &= 1 \end{aligned}$$

So, since $C'' > 0$, we have that the determinant must be positive. In particular, we have

$$\frac{\partial^2 \pi}{\partial \Delta^2} \frac{\partial^2 \pi}{\partial \mu^2} - \left(\frac{\partial^2 \pi}{\partial \Delta \partial \mu} \right)^2 \geq \hat{q}(\Delta) (1 - \hat{q}(\Delta)) C'' (\mu - \Delta - \Psi(\hat{q}(\Delta))) C'' (\mu - \Psi(\hat{q}(\Delta)))$$

A.3.4 Summary

As long as $\Psi''' > 0$ and $\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$, the profit function is globally concave. Empirically, I find that $\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$. Therefore, the unsubstantiated assumption for the model is that the convexity of the effort function increases in p , $\Psi''' > 0$. An alternative statement of this assumption is that $\frac{\partial^2 \hat{q}}{\partial \Delta^2} < 0$, so that the marginal impact of Δ on the employment probability is declining in the size of Δ . Put differently, it is an assumption that providing utility incentives to work has diminishing returns.

Future work can derive the necessary conditions when individuals can make additional actions, $a(\theta)$, in response to unemployment. I suspect the proofs can be extended to such cases, but identifying the necessary conditions for global concavity would be an interesting direction for future work.

B Details of Empirical Approach

B.1 Details on Lower Bounds on Average Pooled Price Ratio

This section provides details on the estimation of the lower bounds on the average pooled price ratio. I begin by providing theoretical motivation for the average pooled price ratio by showing it characterizes when an insurer can earn positive profits if it enters with a particular random pricing strategy. Then, I provide conditions under which the average pooled price defined by the predicted values provides a lower bound on the average pooled price ratio, $T_Z \leq E[T(P)]$.

B.1.1 Motivating the Average Pooled Price Ratio when Insurers don't know P

To see the theoretical relevance of $E[T(P)]$, suppose an insurer seeks to start an insurance market by randomly drawing an individual from the population and, perhaps through some market research, learns exactly how much this individual is willing to pay. The insurer offers a contract that collects \$1 in the event of being employed and pays an amount in the unemployed state that makes the individual perfectly indifferent to the policy. If p is the probability

this individual will become unemployed, then all risks $P \geq p$ will choose to purchase the policy as well. The profit per dollar of revenue will be

$$r(p) = \frac{u'(c_u(p))}{v'(c_e(p))} - T(p)$$

So, if the original individual was selected at random from the population, the expected profit per dollar would be positive if and only if

$$E \left[\frac{u'(c_u(p))}{v'(c_e(p))} \right] \geq E[T(P)]$$

If the insurer is randomly choosing contracts to try to sell, the average pooled price ratio, $E[T(P)]$, provides information on whether or not a UI market would be profitable.

B.1.2 Conditions under which $E[T_Z(P_Z)] \leq E[T(P)]$

Here, I provide conditions under which $E[T_Z(P_Z)]$ provides a lower bound on the average pooled price ratio, $E[T(P)]$. To begin, assume that (a) the elicitation, Z , have no more information about U than do true beliefs, P . Stated formally: $\Pr\{U|X, Z, P\} = \Pr\{U|X, P\}$. Second, assume that beliefs are unbiased, so that $\Pr\{U|X, P\} = P$. Hendren (2013) shows that these two assumptions imply that $E[m(P_Z)] \leq E[m(P)]$. This suggests that $E[T_Z(P_Z)] \leq 1 + \frac{E[m(P)]}{E[P]}$. So, what remains to show is that $1 + \frac{E[m(P)]}{E[P]} \leq E[T(P)]$. For this, we make one additional assumption that $\text{cov}\left(\frac{m(P)}{P}, P\right) \leq 0$. This is a natural assumption because $m(p)$ is, on average, a decreasing function in p (because $m(1) = 0$), so dividing by P renders it an even more strongly decreasing function in P . Indeed, I have been unable to find a random variable P for which $\text{cov}\left(\frac{m(P)}{P}, P\right) > 0$.

Given these assumptions note that

$$\begin{aligned} E[T(P)] &= E_p \left[\frac{E[P|P \geq p]}{p} \frac{1-p}{1-E[P|P \geq p]} \right] \\ &\geq E_p \left[1 + \frac{m(p)}{p} \right] \\ &\geq 1 + E \left[\frac{m(P)}{P} \right] \end{aligned}$$

where E_p represents the expectation with respect to drawing p from the distribution of P . Note the second line follows from the fact that $E[P|P \geq p] \geq p$.

So, it suffices to show that $E \left[\frac{m(P)}{P} \right] \geq \frac{E[m(P)]}{E[P]}$. Clearly

$$E[m(P)] = E \left[\frac{m(P)}{P} \right] E[P] + \text{cov} \left(P, \frac{m(P)}{P} \right)$$

so that

$$E \left[\frac{m(P)}{P} \right] = \frac{E[m(P)] - \text{cov} \left(P, \frac{m(P)}{P} \right)}{E[P]}$$

Imposing $\text{cov} \left(\frac{m(P)}{P}, P \right) \leq 0$ yields $E[T(P)] \geq 1 + \frac{E[m(P)]}{E[P]}$, which in turn implies $E[T(P)] \geq E[T_Z(P_Z)]$.

B.2 Specification for Point Estimation

I follow Hendren (2013) by assuming that $Z = P + \epsilon$, where ϵ has the following structure. With probability λ , individuals report a noisy measure of their true belief P that is drawn from a $[0, 1]$ -censored normal distribution with mean $P + \alpha(X)$ and variance σ^2 . With this specification, $\alpha(X)$ reflects potential bias in elicitation and σ represents the noise. While this allows for general measurement error in the elicitation, it does not produce the strong focal point concentrations shown in Figure 1 and documented in existing work (Gan et al. (2005); Manski and Molinari (2010)). To capture these, I assume that with probability $1 - \lambda$ individuals take their noisy report with the same bias $\alpha(X)$ and variance σ^2 , but censor it into a focal point at 0, 50, or 100. If their elicitation would have been below κ , they report zero. If it would have been between κ and $1 - \kappa$, they report 50; and if it would have been above $1 - \kappa$, they report 1. Hence, I estimate four elicitation error parameters: $(\sigma, \lambda, \kappa, \alpha(X))$ that capture the patterns of noise and bias in the relationship between true beliefs, P , and the elicitation reported on the surveys, Z .

Specifically, the p.d.f./p.m.f. of Z given P is given by

$$f(Z|P, X) = \begin{cases} (1 - \lambda) \Phi\left(\frac{-P - \alpha(X)}{\sigma}\right) + \lambda \Phi\left(\frac{\kappa - P - \alpha(X)}{\sigma}\right) & \text{if } Z = 0 \\ \lambda \left(\Phi\left(\frac{1 - \kappa - P - \alpha(X)}{\sigma}\right) - \Phi\left(\frac{\kappa - P - \alpha(X)}{\sigma}\right) \right) & \text{if } Z = 0.5 \\ (1 - \lambda) \Phi\left(\frac{1 - P - \alpha(X)}{\sigma}\right) + \lambda \left(1 - \Phi\left(\frac{1 - \kappa - P - \alpha(X)}{\sigma}\right) \right) & \text{if } Z = 1 \\ \frac{1}{\sigma} \phi\left(\frac{Z - P - \alpha(X)}{\sigma}\right) & \text{if } o.w. \end{cases}$$

where ϕ denotes the standard normal p.d.f. and Φ the standard normal c.d.f. I estimate four elicitation error parameters: $(\sigma, \lambda, \kappa, \alpha(X))$. σ captures the dispersion in the elicitation error, λ is the fraction of focal point respondents, κ is the focal point window. I allow the elicitation bias term, $\alpha(X)$, to vary with the observable variables, X . This allows elicitation to be biased, but maintains the assumption that true beliefs are unbiased. This approach builds upon Manski and Molinari (2010) by thinking of the focal point responses as “interval data” (i.e. 50/50 corresponds to some region around 50%, but not exactly 50%). However, the present approach differs from Manski and Molinari (2010) by allowing the response to be a noisy and potentially biased measure of this response (as 50/50 corresponds to a region around 50% for the noisy Z measure, not the true P measure).

Ideally, one would flexibly estimate the distribution of P given X at each possible value of X . This would enable separate estimates of the minimum pooled price ratio for each value of X . However, the dimensionality of X prevents this in practice. Instead, I again follow Hendren (2013) and adopt an index assumption on the cumulative distribution of beliefs, $F(p|X) = \int_0^p f_P(\tilde{p}|X) d\tilde{p}$,

$$F(p|X) = \tilde{F}(p|\Pr\{U|X\}) \quad (26)$$

where I assume $\tilde{F}(p|q)$ is continuous in q (where $q \in \{0, 1\}$ corresponds to the level of $\Pr\{U|X\}$). This assumes that the distribution of private information is the same for two observable values, X and X' , that have the same observable unemployment probability, $\Pr\{U|X\} = \Pr\{U|X'\}$. Although one could perform different dimension reduction techniques, controlling for $\Pr\{U|X\}$ is particularly appealing because it nests the null hypothesis of no private information ($F(p|X) = 1\{p \leq \Pr\{U|X\}\}$).⁵²

Beta versus Point-Mass Distribution Hendren (2013) flexibly approximates $F(p|q)$ using mixtures of Beta distributions. In the current context, a key difficulty with using functions to approximate the distribution of P is that much of the mass of the distribution is near zero. Continuous probability distribution functions, such as Beta distributions, require very high degrees for the shape parameters to acquire a good fit. Therefore, I approximate P as a sum of discrete point-mass distributions.⁵³ Formally, I assume

$$\tilde{F}(p|q) = w1\{p \leq q - a\} + (1 - w) \sum_i \xi_i 1\{p \leq \alpha_i\}$$

where α_i are a set of point masses in $[0, 1]$ and ξ_i is the mass on each point mass. I estimate these point mass parameters using maximum likelihood estimation. For the baseline results, I use 3 mass points, which generally provides a decent fit for the data. Appendix Table IV presents the raw estimates for these point mass distributions.

Given the estimate of $\tilde{F}(p|q)$, I then compute the pooled price ratio at each mass point and report the minimum across all values aside from the largest mass point. Mechanically, this has a value of $T(p) = 1$. As noted in Hendren (2013), estimation of the minimum $T(p)$ across the full support of the type distribution is not feasible because of an extremal quantile estimation problem. To keep the estimates “in-sample”, I report values for the mean value of $q = \Pr\{U\} = 0.031$; but estimates at other values of q are similarly large.

⁵²Moreover, it allows the statistical model to easily impose unbiased beliefs, so that $\Pr\{U|X\} = E[P|X]$ for all X .

⁵³This has the advantage that it does not require integrating over high degree of curvature in the likelihood function. In practice, it will potentially under-state the true variance in P in finite sample estimation. As a result, it will tend to produce lower values for $T(p)$ than would be implied by continuous probability distributions for P since the discrete approximation allows all individuals at a particular point mass to be able to perfectly pool together when attempting to cover the pooled cost of worse risks.

C Willingness to Pay Metrics

C.1 Taylor Expansion

Note that $u(c) = v(c)$ so that

$$\begin{aligned}
\frac{u'(c_u(p))}{u'(c_e(p))} &\approx \frac{u'(c_e(p)) + u''(c_e(p))(c_u(p) - c_e(p)) + \frac{1}{2}u'''(c_e(p))(c_u(p) - c_e(p))^2}{u'(c_e(p))} \\
&= 1 + \frac{u''(c_e(p))}{u'(c_e(p))} \frac{c_u(p) - c_e(p)}{c_e(p)} \\
&= 1 + \frac{-c_e(p) u''(c_e(p))}{u'(c_e(p))} \frac{c_e(p) - c_u(p)}{c_e(p)} + \frac{c_e(p) u'''(c_e(p))}{u''(c_e(p))} \frac{c_e(p) u''(c_e(p))}{u'(c_e(p))} \frac{1}{2} \left(\frac{c_e(p) - c_u(p)}{c_e(p)} \right)^2 \\
&= 1 + \sigma \frac{\Delta c}{c}(p) \left[1 + \frac{\gamma}{2} \frac{\Delta c}{c}(p) \right]
\end{aligned}$$

And, under an assumption of constant relative risk aversion, we have $\gamma = \sigma + 1$

$$\begin{aligned}
\frac{u'(c_u(p))}{u'(c_e(p))} - 1 &\approx \sigma \frac{\Delta c}{c}(p) \left[1 + \frac{\sigma + 1}{2} \frac{\Delta c}{c}(p) \right] \\
&= \sigma \frac{\Delta c}{c}(p) + \sigma \frac{\sigma + 1}{2} \left(\frac{\Delta c}{c}(p) \right)^2
\end{aligned}$$

C.2 Proof of Proposition 1

Note under state independence, the Euler equation implies

$$u'(c_{pre}(p)) = pu'(c_u(p)) + (1-p)u'(c_e(p))$$

so that

$$u''(c_{pre}(p)) \frac{dc_{pre}}{dp} = u'(c_u(p)) - u'(c_e(p)) + pu''(c_u(p)) \frac{dc_u}{dp} + (1-p)u''(c_e(p)) \frac{dc_e}{dp}$$

Dividing,

$$u'(c_{pre}(p)) \frac{u''(c_{pre}(p))}{u'(c_{pre}(p))} \frac{dc_{pre}}{dp} = u'(c_e) \frac{u'(c_u(p)) - u'(c_e(p))}{u'(c_e(p))} + pu'(c_u(p)) \frac{u''(c_u(p))}{u'(c_u(p))} \frac{dc_u}{dp} + (1-p)u'(c_e(p)) \frac{u''(c_e(p))}{u'(c_e(p))} \frac{dc_e}{dp}$$

or

$$u'(c_{pre}(p)) \sigma \frac{-d \log(c_{pre})}{dp} = u'(c_e) \sigma [\log(c_e) - \log(c_u)] + pu'(c_u(p)) \sigma \frac{-d \log(c_u(p))}{dp} + (1-p)u'(c_e(p)) \sigma \frac{-d \log(c_e(p))}{dp}$$

So, dividing by $u'(c_e(p))$ yields:

$$\frac{u'(c_{pre}(p))}{u'(c_e)} \sigma \frac{-d \log(c_{pre})}{dp} = \sigma [\log(c_e) - \log(c_u)] + p \frac{u'(c_u(p))}{u'(c_e(p))} \sigma \frac{-d \log(c_u(p))}{dp} + (1-p) \sigma \frac{-d \log(c_e(p))}{dp}$$

And, using the Euler equation, $pu'(c_u(p)) + (1-p)u'(c_e(p)) = u'(c_{pre}(p))$,

$$\begin{aligned}
\frac{pu'(c_u(p)) + (1-p)u'(c_e(p))}{u'(c_e(p))} \sigma \frac{-d \log(c_{pre}(p))}{dp} &= \frac{u'(c_u(p)) - u'(c_e(p))}{u'(c_e(p))} + p \frac{u'(c_u(p))}{u'(c_e(p))} \sigma \frac{-d \log(c_u(p))}{dp} + (1-p) \sigma \frac{-d \log(c_e(p))}{dp} \\
\left(p \frac{u'(c_u(p))}{u'(c_e(p))} + 1 - p \right) \sigma \frac{-d \log(c_{pre}(p))}{dp} &= \frac{u'(c_u(p)) - u'(c_e(p))}{u'(c_e(p))} + \left(p \frac{u'(c_u(p))}{u'(c_e(p))} + 1 - p \right) \sigma \frac{-d \log(c_e(p))}{dp} + p \frac{u'(c_u(p))}{u'(c_e(p))} \sigma \left(\frac{d[\log(c_e) - \log(c_u)]}{dp} \right)
\end{aligned}$$

so that

$$\sigma \frac{-d \log(c_{pre})}{dp} = \frac{\frac{u'(c_u(p))}{u'(c_e(p))} - 1}{1 + p \left(\frac{u'(c_u(p))}{u'(c_e(p))} - 1 \right)} + \sigma \frac{-d \log(c_e(p))}{dp} + \frac{p \frac{u'(c_u(p))}{u'(c_e(p))}}{p \frac{u'(c_u(p))}{u'(c_e(p))} + 1 - p} \sigma \left(\frac{d[\log(c_e) - \log(c_u)]}{dp} \right)$$

or

$$\frac{-d\log(c_{pre})}{dp} = \frac{\frac{1}{\sigma} \left(\frac{u'(c_u(p))}{u'(c_e(p))} - 1 \right)}{1 + p \left(\frac{u'(c_u(p))}{u'(c_e(p))} - 1 \right)} + \frac{-d\log(c_e(p))}{dp} + \frac{p \frac{u'(c_u(p))}{u'(c_e(p))}}{p \frac{u'(c_u(p))}{u'(c_e(p))} + 1 - p} \left(\frac{d[\log(c_e) - \log(c_u)]}{dp} \right)$$

Note that the assumption is maintained that $\log(c_{pre})$ is linear in p , in addition to $\log(c_e)$ and $\log(c_u)$ being linear in p . This is of course an approximation in practice, as the equation above illustrates this cannot simultaneously be true for all p . Therefore, I assume it is true only in expectation, so that

$$\frac{-d\log(c_{pre})}{dp} = \frac{1}{\sigma} \left(\frac{u'(c_u(p))}{u'(c_e(p))} - 1 \right) E \left[\frac{1}{1 + p \left(\frac{u'(c_u(p))}{u'(c_e(p))} - 1 \right)} \right] + \frac{-d\log(c_e(p))}{dp} + E \left[\frac{p \frac{u'(c_u(p))}{u'(c_e(p))}}{p \frac{u'(c_u(p))}{u'(c_e(p))} + 1 - p} \left(\frac{d[\log(c_e) - \log(c_u)]}{dp} \right) \right]$$

which if it holds for all p must also hold for the expectation taken with respect to p . Let $\kappa = E \left[\frac{1}{1 + p \left(\frac{u'(c_u(p))}{u'(c_e(p))} - 1 \right)} \right]$.

Note also that

$$\frac{u'(c_u(p))}{u'(c_e(p))} - 1 \approx \sigma E[\log(c_e(p)) - \log(c_u(p))]$$

which implies

$$\frac{-d\log(c_{pre})}{dp} = E[\log(c_e(p)) - \log(c_u(p))] \kappa + \frac{-d\log(c_e(p))}{dp} + E \left[\frac{p \frac{u'(c_u(p))}{u'(c_e(p))}}{p \frac{u'(c_u(p))}{u'(c_e(p))} + 1 - p} \left(\frac{d[\log(c_e) - \log(c_u)]}{dp} \right) \right]$$

Now, consider the impact of unemployment on the first difference of consumption. Define Δ^{FD} as the estimated impact on the first difference in consumption:

$$\Delta^{FD} = E[\log(c) - \log(c_{-1}) | U = 1] - E[\log(c) - \log(c_{-1}) | U = 0]$$

Adding and subtracting $E[\log(c_e) | U = 1]$ yields

$$\Delta^{FD} = E[\log(c) | U = 1] - E[\log(c_e) | U = 1] + E[\log(c_e) | U = 1] - E[\log(c) | U = 0] - (E[\log(c_{-1}) | U = 1] - E[\log(c_{-1}) | U = 0])$$

Note that $c = c_u$ for those with $U_t = 1$ and $c = c_e$ for those with $U = 0$. The following three equations help expand Δ^{FD} :

$$E[\log(c_{-1}) | U = 1] - E[\log(c_{-1}) | U = 0] = \frac{d\log(c_{pre})}{dp} \frac{\text{var}(P)}{\text{var}(U)}$$

and

$$\begin{aligned} E[\log(c) | U = 1] - E[\log(c_e) | U = 0] &= E[\log(c_u) | U = 1] - E[\log(c_e) | U = 1] \\ &= E[\log(c_u) - \log(c_e)] + \frac{d[\log(c_u) - \log(c_e)]}{dp} (E[P | U = 1] - E[P]) \end{aligned}$$

and

$$\begin{aligned} E[\log(c_e) | U = 1] - E[\log(c) | U = 0] &= E[\log(c_e) | U = 1] - E[\log(c_e) | U = 0] \\ &= \frac{d\log(c_e)}{dp} \frac{\text{var}(P)}{\text{var}(U)} \end{aligned}$$

So, substituting these into Δ^{FD} yields:

$$\Delta^{FD} = E[\log(c_u) - \log(c_e)] - \frac{\text{var}(P)}{\text{var}(U)} \left[\kappa (E[\log(c_e) - \log(c_u)]) + E[P] \frac{u'(c_u)}{u'(c_e)} \frac{d[\log(c_e) - \log(c_u)]}{dp} \right] + \left[\frac{d[\log(c_e) - \log(c_u)]}{dp} \right] (E[P | U = 1] - E[P])$$

Let $\frac{d\Delta}{dp} = \frac{d[\log(c_e) - \log(c_u)]}{dp}$ denote how the consumption drop varies with p . Solving for $E[\log(c_e) - \log(c_u)]$ yields

$$E[\log(c_u) - \log(c_e)] = \frac{\Delta^{FD} + \frac{d\Delta}{dp} (E[P | U = 1] - E[P])}{1 - \frac{\text{var}(P)}{\text{var}(U)} \kappa - \bar{p} \sigma \frac{d\Delta}{dp}}$$

where κ = which yields the desired result. Note that if the consumption drop does not vary with p , then this reduces to

$$E[\log(c_u) - \log(c_e)] = \frac{\Delta^{FD}}{1 - \frac{\text{var}(P)}{\text{var}(U)} \kappa} \equiv \Delta^{IV}$$

More generally, if the size of the consumption drop is increasing with p , then $E[\log(c_u) - \log(c_e)] > \Delta^{IV}$.

C.3 Maximum Willingness to Pay

While the analysis to this point estimates the average causal effect of unemployment, Equation (8) requires comparing the willingness to pay for all p to the pooled price ratio. Therefore, it is also useful to understand the heterogeneity in the potential willingness to pay across the population. How much might some people be willing to pay for insurance?

Estimating minima and maxima is always more difficult than estimating means; but this section attempts to make a bit of progress to help shed light on this important question. Let Δ^{min} denote the largest causal effect of unemployment in the population,

$$\Delta^{min} = \min_p [\log(c_u(p)) - \log(c_e(p))] \quad (27)$$

Following equation (9), note that the willingness to pay satisfies

$$\frac{u'(c_u) - u'(c_e)}{u'(c_e)} \geq \sigma \left(-\Delta^{min} \right) \left(1 + \frac{\gamma}{2} \left(-\Delta^{min} \right) \right) \quad (28)$$

Therefore, Δ^{min} generates an upper bound on the willingness to pay (note that $\Delta^{min} < 0$).

This motivates the question: How big can the causal impact of unemployment be? To address this, note that one can write the causal effect as the sum of two first differences:

$$\log(c_u(p)) - \log(c_e(p)) = \log(c_u(p)) - \log(c_{t-1}(p)) - (\log(c_e(p)) - \log(c_{t-1}(p)))$$

where the first term captures consumption change if unemployed and the second term is the consumption change if employed. Let $\Delta_e^{max} = \max_p \{\log(c_e(p)) - \log(c_{t-1}(p))\}$ denote the maximum consumption change experienced by those who did not lose their job. And, let $\Delta_u^{min} = \min_p \{\log(c_u(p)) - \log(c_{t-1}(p))\}$ denote the minimum consumption change experienced by those who lose their job. Note that we expect $\Delta_e^{max} > 0$ and $\Delta_u^{min} < 0$. The Euler equation ((6)) combined with the assumption of CRRA preferences implies that c_{t-1} lies between $c_u(p)$ and $c_e(p)$, $c_{t-1}(p) \in [c_u(p), c_e(p)]$ for all p . Under this natural assumption, the causal impact of unemployment is bounded below by the difference between these drops:

$$\Delta^{min} \geq \Delta_u^{min} - \Delta_e^{max}$$

Therefore, one can bound the causal effect of unemployment on consumption by the largest consumption drop minus the smallest consumption increase. The question now becomes: how large can the consumption drop upon unemployment be, Δ_u^{min} ? And, how large can the consumption increase upon learning that you didn't lose your job be, Δ_e^{max} ?

If one observed consumption directly, one could estimate the full distribution of first differences in consumption for those who become unemployed, $\log(c_u(p)) - \log(c_{t-1}(p))$, and remain employed, $\log(c_e(p)) - \log(c_{t-1}(p))$. Then, one could in principle find Δ_u^{min} and Δ_e^{max} directly from the data.

However, consumption data in the PSID and other datasets is quite noisy in practice (see for example Zeldes (1989); Meghir and Pistaferri (2011)).⁵⁴ Therefore, I proceed as follows. Note that the Euler equation implies that $c_{t-1}(p) \in [c_u(p), c_e(p)]$ for all p . In particular, this implies that the log consumption change should always drop for those who lose their job, $c_{t-1}(p) \geq c_u(p)$. Therefore, I use the extent to which one observes consumption increases for those who become unemployed to provide information about how the consumption change distribution and its minimum, Δ_u^{min} , is affected by measurement error.

I begin by removing systematic variation in consumption changes due to life cycle and year effects. In particular, I regress the consumption change on the observables, X , in Equation (4) (an age cubic and year dummies) and let Δ_{it}^* denote the residuals.⁵⁵ Online Appendix Figure VI plots the distributions of Δ_{it}^* for those with $U_{it} = 1$ and $U_{it} = 0$. As one can see, the wide dispersion is suggestive of considerable measurement error, as noted in previous literature.

Let $Q(\alpha, U)$ denote the α -quantile of the distribution of Δ_{it}^* as a function of unemployment status, U_t . Appendix Table II reports that 41.7% of the sample who become unemployed have $\Delta_{it}^* > 0$ (i.e. $Q(58.3, 1) = 0$), even after controlling for age and year effects. Because the Euler equation suggests consumption changes should not be positive,

⁵⁴This “left-hand-side” measurement error was necessarily not a problem for estimating the mean consumption impact of unemployment (assuming the error is classical). But, for estimating properties of the distribution of consumption changes such as minima and maxima, this measurement error becomes a significant limitation.

⁵⁵This residualization can be formalized by assuming there are known time and year preference shocks affecting the marginal utility of consumption that are common across individuals. Note the residuals now satisfy the ex-ante Euler equation, $E[\Delta_{it}^*] = 0$. But, the means of the residuals will differ for those who do and do not experience unemployment, $E[\Delta_{it}^*|U_{it} = 1] \leq 0$ and $E[\Delta_{it}^*|U_{it} = 0] \geq 0$.

it suggests the excess dispersion is the result of measurement error. The key assumption I impose is that the impact of measurement error is symmetric across the distribution of consumption changes. In particular, I assume that the probability that the observed consumption change lies above the maximum plausible consumption change of 0 is less than or equal to the probability that the observed consumption change is below the minimum actual consumption change, $\Pr \{\Delta_{it}^* \leq \Delta_u^{min} | U_{it} = 1\} \geq \Pr \{\Delta_{it}^* \geq 0 | U_{it} = 1\}$, where $\Pr \{\Delta_{it}^* \geq 0 | U_{it} = 1\} = 41.7\%$. Appendix Figure V shows that the observed distribution of consumption changes is fairly symmetric, which would be consistent with the underlying assumptions needed for Equation (29) to hold.⁵⁶ With this assumption,

$$\Delta_u^{min} \geq Q(\Pr \{\Delta_{it}^* \geq 0 | U_{it} = 1\}, 1) \quad (29)$$

Because $\Pr \{\Delta_{it}^* \geq 0 | U_{it} = 1\} = 41.7\%$, one can bound the consumption drop by the 41.7th quantile of the observed consumption drop distribution. This equals -13.7% , as shown in Appendix Table II.

Similarly, one can impose an analogous assumption on the distribution of consumption changes for the employed that the observed fraction of the population that experiences a consumption decline when remaining employed is less than or equal to the fraction who experience a consumption increase that is above Δ_e^{max} , $\Pr \{\Delta_{it}^* \leq 0 | U_{it} = 0\} = \Pr \{\Delta_{it}^* \geq \Delta_e^{max} | U_{it} = 0\}$. Under this assumption,

$$\Delta_e^{max} \leq Q(\Pr \{\Delta_{it}^* \leq 0 | U_{it} = 0\}, 0) \quad (30)$$

Appendix Table II shows this maximal consumption increase equals 0.5% . Combining equations (29) and (30) yields the lower bound on the causal impact of unemployment on consumption:

$$\Delta^{min} \geq Q(\Pr \{\Delta_{it}^* \geq 0 | U_{it} = 1\}, 1) - Q(\Pr \{\Delta_{it}^* \leq 0 | U_{it} = 0\}, 0)$$

where the right hand side equals $-13.7\% - 0.5\% = 14.2\%$ (s.e. 1.1%), as reported in Table V, Column (7). Therefore, the maximum causal impact on food expenditure is 14.2% , or roughly twice as large as the mean consumption drop. The lower rows in Table V scale this estimate by various values of risk aversion. With a conservative estimate of 3, it suggests the maximum markup individuals would be willing to pay is 54.7% .

C.4 Proof of Proposition 2

Differentiating the Euler equation under assumption (b) yields

$$u'(c_u) - v'(c_e) = v''(c_{pre}(p)) \frac{dc_{pre}}{dp}$$

Now, dividing by $v'(c_e)$ yields

$$\frac{u'(c_u) - v'(c_e)}{v'(c_e)} = \frac{v''(c_{pre}(p))}{v'(c_e)} \frac{dc_{pre}}{dp}$$

and expanding the RHS yields

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

And, imposing the Euler equation to replace $v'(c_{pre}(p))$ in the numerator on the RHS ($v'(c_{pre}(p)) = pu'(c_u) + (1-p)v'(c_e)$) yields,

$$\frac{u'(c_u) - v'(c_e)}{v'(c_e)} = \left[p \frac{u'(c_u)}{v'(c_e)} + (1-p) \right] \frac{c_{pre}(p)}{v'(c_{pre}(p))} \frac{v''(c_{pre}(p))}{v'(c_{pre}(p))} \frac{1}{c_{pre}(p)} \frac{dc_{pre}}{dp}$$

Dividing by $p \frac{u'(c_u)}{v'(c_e)} + (1-p)$ and taking expectations over p yields

$$\kappa \frac{u'(c_u) - v'(c_e)}{v'(c_e)} = E \left[\frac{c_{pre}(p)}{v'(c_{pre}(p))} \frac{v''(c_{pre}(p))}{v'(c_{pre}(p))} \frac{1}{c_{pre}(p)} \frac{dc_{pre}}{dp} \right]$$

Now, imposing $\sigma = \frac{-c_{pre}(p)v''(c_{pre}(p))}{v'(c_{pre}(p))}$ for all p , noting that $\frac{1}{c_{pre}(p)} \frac{dc_{pre}}{dp} = \frac{d \log(c_{pre})}{dp}$, and dividing by κ yields:

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

⁵⁶While the symmetry assumption is not directly testable, it can be micro-founded from many common assumptions on measurement error distributions. For example, if the true distribution of consumption changes is symmetric and the distribution of measurement error is symmetric and unbiased, then it is straightforward to show that $\Pr \{\Delta_{it}^* \leq \Delta_u^{min} | U_{it} = 1\} \geq \Pr \{\Delta_{it}^* \geq x | U_{it} = 1\}$, where $x \leq 0$ is the maximum consumption change for those who become unemployed. Symmetric and median-unbiasedness is a common assumption measurement error models (see, e.g., Bollinger (1998); Hu and Schennach (2008)).

C.5 Ex-ante labor supply derivation

This section illustrates how to use the spousal labor supply response, combined with known estimates of the spousal labor response to labor earnings, to estimate the ex-ante willingness to pay for UI.

To begin, one needs a model of extensive margin labor supply response. I assume spousal labor force participation generates income, y , but has an additively separable effort cost, $\eta(\theta)$. I assume a spouse labor supply decision, $l \in \{0, 1\}$, is a binary decision and is contained in the set of other actions, a . Formally, let utility be given by

$$v(c_{pre}) + pu(c_u) + (1-p)v(c_e) + 1\{l=1\}\eta(\theta) + \tilde{\Psi}_i(1-p, \tilde{a}, \theta)$$

where $\eta(\theta)$ is the disutility of labor for type θ , distributed F_η in the population.

Let $k(y, l, p)$ denote the utility value to a type p of choosing l to obtain income y when they face an unemployment probability of p . The labor supply decision is

$$k(y, 1, p) - k(0, 0, p) \geq \eta(\theta)$$

so that types will choose to work if and only if it increases their utility. This defines a threshold rule whereby individuals choose to work if and only if $\eta(\theta) \leq \bar{\eta}(y, p)$ and the labor force participation rate is given by $\Phi(y, p) = F(\bar{\eta}(y, p))$.

Now, note that

$$\frac{d\Phi}{dp} = f(\bar{\eta}) \frac{\partial \bar{\eta}}{\partial p} = f(\bar{\eta}) \left[\frac{\partial k(y, 1, p)}{\partial p} - \frac{\partial k(0, 0, p)}{\partial p} \right]$$

and making an approximation that the impact of the income y does not discretely change the instantaneous marginal utilities (i.e. because it will be smoothed out over the lifetime or because the income is small), we have

$$\frac{d\Phi}{dp} \approx f(\bar{\eta}) \frac{\partial^2 k}{\partial p^2} y$$

Finally, note that $\frac{\partial k}{\partial y} = v'(c_{pre}(p))$ is the marginal utility of income. So,

$$\frac{d\Phi}{dp} \approx f(\bar{\eta}) \frac{d}{dp} [v'(c_{pre}(p))] y$$

and integrating across all the types p yields

$$E_p \left[\frac{d\Phi}{dp} \right] \approx E_p \left[f(\bar{\eta}) \frac{d}{dp} v'(c_{pre}(p)) y \right]$$

To compare this response to a wage elasticity, consider the response to a \$1 increase in wages

$$\frac{d\Phi}{dy} = f(\bar{\eta}) \frac{\partial k}{\partial y}$$

so,

$$E_p \left[\frac{d\Phi}{dp} \right] \approx E_p \left[\frac{d\Phi}{dy} y \frac{\frac{d}{dp} v'(c_{pre}(p))}{v'(c_{pre}(p))} \right]$$

Now, let $\epsilon^{semi} = \frac{d\Phi}{d \log(y)}$ denote the semi-elasticity of spousal labor force participation. This yields

$$\frac{E_p \left[\frac{d\Phi}{dp} \right]}{\epsilon^{semi}} \approx E_p \left[\frac{\frac{d}{dp} v'(c_{pre}(p))}{v'(c_{pre}(p))} \right]$$

so that the ratio of the labor supply response to p divided by the semi-elasticity of labor supply with respect to wages reveals the average elasticity of the marginal utility function. Assuming this elasticity is roughly constant and noting that a Taylor expansion suggests that for any function $f(x)$, we have $\frac{f(1)-f(0)}{f(0)} \approx \frac{d}{dx} \log(f)$,

$$\frac{E_p \left[\frac{d\Phi}{dp} \right]}{\epsilon^{semi}} \approx \frac{v'(1) - v'(0)}{v'(0)}$$

Now, how does one estimate $\frac{d\Phi}{dp}$? Regressing labor force participation, l , on Z will generate an attenuated coefficient because of measurement error in Z . If the measurement error is classical, one can inflate this by the ratio of the variance of Z to the variance of P , or

$$\frac{v'(1) - v'(0)}{v'(0)} \approx \beta \frac{1}{\epsilon^{semi}} \frac{\text{var}(Z)}{\text{var}(P)}$$

APPENDIX TABLE I
Information Realization Between t-2 and t-1 ("First Stage")

Full Sample	Male (2)	Female (3)	Age > 55 (4)	Age ≤ 55 (5)	Year ≤ 1997 (6)	Year > 1997 (7)	Male, Age ≤ 55, Year ≤ 1997 (7)
<i>Panel 1: Regression of Job Loss on Elicitations</i>							
Job Loss (Next 12 months)	0.1968	0.1978	0.2079	0.1806	0.2316	0.1829	0.2089
s.e.	0.0120	0.0156	0.0159	0.0195	0.0246	0.0140	0.0624

<i>Panel 2: Regression of Job Loss in Subsequent 12-24 Months on Z</i>							
Job Loss (12-24 months)	0.0937	0.1199	0.0893	0.0994	0.1080	0.0847	0.0454
s.e.	0.0120	0.0156	0.0159	0.0195	0.0246	0.0140	0.0624
Difference	0.1031	0.0779	0.1186	0.0812	0.1236	0.0982	0.1635
bootstrap s.e.	0.0120	0.0156	0.0159	0.0195	0.0246	0.0140	0.0624
Num of Obs.	26,640	15,900	15,506	11,134	8,571	18,069	1,210

Note: This table presents estimates from regressions of the elicitation, Z , on unemployment measured in both (a) the subsequent 12 months and (b) the subsequent 12-24 months. The first row corresponds to the first stage for the ex-post welfare measure; the final row ("Difference") corresponds to the first stage for the ex-ante welfare measure. Column (1) uses the baseline HRS sample. Columns (2)-(7) explore the heterogeneity in the estimates by subgroup. Columns (2)-(3) restrict the sample to males and females. Columns (4)-(5) restrict the sample to those above and below age 55. Columns (6)-(7) restrict the sample to before and after 1997. Standard errors are computed using 500 bootstrap repetitions resampling at the household level.

APPENDIX TABLE II
Maximum Causal Effect of Unemployment on Food Expenditure

	Baseline Sample
	(1)
Estimate for $\max\{\log(c_u(p))-\log(c_e(p))\}$, Δ^{\min}	-0.137
s.e.	(0.02)
Lower bound for drop when unemployed, Δ_u^{\min}	-0.138
s.e.	(0.02)
Upper bound for increase when unemployed, Δ_c^{\max}	-0.001
s.e.	(0.002)
Fraction unemployed with positive consumption change	0.415
Fraction employed with negative consumption change	0.499
Num of Obs.	65,808
Num of HHs	9,562

Note: This table presents the calculation for the maximum causal effect of unemployment on food expenditure. The first row presents the resulting estimate, Δ^{\min} . The second rows present the estimates for the lower bound on the consumption change when unemployed, Δ_u^{\min} . The third rows present the estimates for the upper bound for the increase in consumption when employed, Δ_c^{\max} . The fourth row presents the fraction of people who are unemployed, $U_{it}=1$, who experience a positive consumption change, $\Delta_{it}^* > 0$. The fifth row presents the fraction of people who are employed, $U_{it}=0$, who experience a negative consumption change, $\Delta_{it}^* < 0$. All standard errors are constructed using 1000 replications.

APPENDIX TABLE III
Alternative Lower Bound Specifications

Specification:	Baseline		Alternative Aggregation Windows			Alternative Subj. Prob. Spec			Alternative Outcomes		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	
$E[T_Z(P_Z)-1]$ s.e.	0.7687 (0.058)	0.6802 (0.051)	0.7716 (0.05)	0.7058 (0.048)	0.7150 (0.048)	0.7462 (0.051)	0.7681 (0.054)	0.5296 (0.033)	0.3675 (0.04)	0.5790 (0.086)	
$E[m_Z(P_Z)]$ s.e.	0.0239 (0.002)	0.0209 (0.002)	0.0237 (0.002)	0.0217 (0.002)	0.0220 (0.002)	0.0229 (0.002)	0.0236 (0.002)	0.0314 (0.002)	0.0147 (0.002)	0.0067 (0.001)	
p-value	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	
$Pr\{U=1\}$	0.0310	0.0307	0.0307	0.0307	0.0307	0.0307	0.0307	0.0593	0.0401	0.0115	
Controls											
Demographics	X	X	X	X	X	X	X	X	X	X	
Job Characteristics	X	X	X	X	X	X	X	X	X	X	
Elicitation Specification											
Polynomial Degree	3	3	3	3	3	1	4	3	3	3	
Focal pt dummies (0, 50, 100)	X	X	X	X	X	X	X	X	X	X	
Aggregation Window	Age x Gender	Age x Gender	Constant	Age x Gender x Industry	Age x Gender x Occupation	Age x Gender	Age x Gender	Age x Gender	Age x Gender	Age x Gender	
Unemployment Outcome Window	12 months	12 months	12 months	12 months	12 months	12 months	12 months	24 months	6-24 months	6-12 months	
Error Specification	Probit	Linear	Probit	Probit	Probit	Probit	Probit	Probit	Probit	Probit	
Num of Obs.	25516	26640	26640	26640	26640	26640	26640	26640	26640	26640	
Num of HHs	3467	3467	3467	3467	3467	3467	3467	3467	3467	3467	

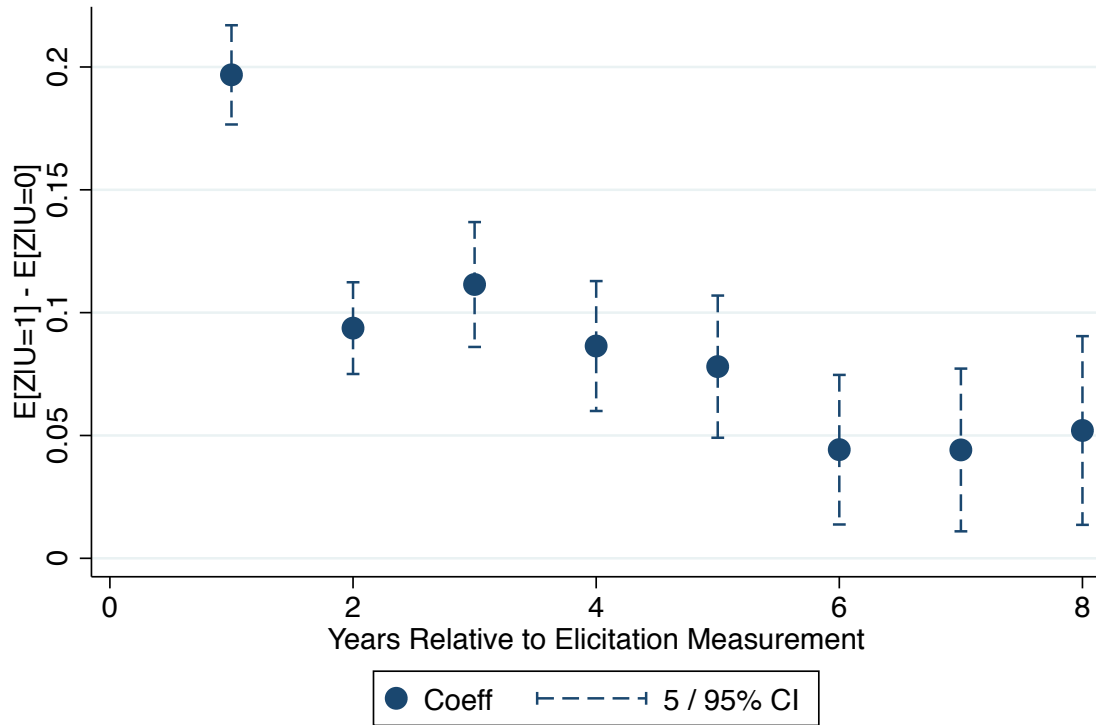
Notes: Table reports robustness of lower bound estimates in Table II to alternative specifications. Column (1) replicates the baseline specification in Table II (Column (1)). Column (2) constructs the predicted values, $Pr\{U|X,Z\}$, using a linear model instead of a probit specification. Columns (3)-(5) consider alternative aggregation windows for translating the distribution of predicted values into estimates of $E[m_Z(P_Z)]$. While Column (1) constructs $m_Z(P_Z)$ using the predicted values within age-by-gender groups, Column (3) aggregates the predicted values across the entire sample. Column (4) uses a finer partition, aggregating within age-by-gender-by-industry groups. Column (5) aggregates within age-by-gender-by-occupation groups. Columns (6)-(7) consider alternative specifications for the subjective probability elicitations. Column (6) uses only a linear specification in Z combined with focal point indicators at $Z=0, Z=50$, and $Z=100$, as opposed to the baseline specification that also includes a polynomial in Z . Column (7) adds a third and fourth order polynomial in Z to the baseline specification. Columns (8)-(10) consider alternative outcome definitions for U . Column (8) defines unemployment, U , as an indicator for involuntary job loss at any point in between survey waves (24 months). Column (9) defines unemployment as an indicator for job loss in between survey waves excluding the first six months after the survey (i.e. 6-24 months). Finally, Column (10) defines unemployment as an indicator for job loss in the 6-12 months after the survey wave.

APPENDIX TABLE IV
Estimation of $F(p|X)$

Specification	Alternative Controls			Sub-Samples					
	Baseline	Demo	Health	Age <= 55	Age > 55	Below Median Wage	Above Median Wage	Tenure > 5 yrs	Tenure <= 5 yrs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1st mass									
Location	0.001	0.012	0.002	0.001	0.002	0.007	0.000	0.000	0.022
s.e.	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)	(0.000)	(0.003)
Weight	0.446	0.713	0.449	0.437	0.461	0.530	0.452	0.422	0.612
s.e.	(0.024)	(0.071)	(0.054)	(0.035)	(0.030)	(0.032)	(0.034)	(0.036)	(0.034)
T(p)	63.839	6.301	39.032	101.038	36.986	12.413	262.088	6.9E+08	5.052
s.e.	6.1E+06	1.7E+00	1.8E+06	1.0E+07	1.1E+06	3.2E+00	7.6E+07	2.5E+08	6.0E-01
2nd mass									
Location	0.031	0.031	0.032	0.030	0.031	0.037	0.024	0.018	0.0575
s.e.	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Weight	0.471	0.202	0.470	0.483	0.456	0.365	0.486	0.508	0.2771
s.e.	(0.024)	(0.071)	(0.052)	(0.035)	(0.030)	(0.032)	(0.034)	(0.037)	(0.0341)
T(p)	4.360	8.492	4.228	4.325	4.442	5.217	4.223	5.736	4.7392
s.e.	0.203	4.194	4.576	0.306	0.279	0.417	2.181	3.008	0.5227
3rd Mass									
Location	0.641	0.639	0.642	0.639	0.643	0.626	0.649	0.641	0.6420
s.e.	(0.004)	(0.004)	(0.028)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.0055)
Weight	0.082	0.086	0.081	0.081	0.083	0.105	0.061	0.070	0.1105
s.e.	(0.002)	(0.002)	(0.006)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.0040)
Controls									
Demographics	X	X	X	X	X	X	X	X	X
Job Characteristics	X		X	X	X	X	X	X	X
Health Characteristics			X						
Num of Obs.	26,640	26,640	22,831	11,134	15,506	13,320	13,320	17,850	8,790
Num of HHs	3,467	3,467	3,180	2,255	3,231	2,916	2,259	2,952	2,437


Notes: This table presents estimates of the distribution of private information about unemployment risk, P . Column (1) reports the baseline specification. Columns (2) uses only demographic controls; Column (3) uses demographic, job characteristics, and health characteristics. Columns (4)-(9) report results for the baseline specification on various subsamples including below and above age 55 (Columns 4 and 5), above and below-median wage earners (Columns 6 and 7) and above and below 5 years of job tenure. The $F(p)$ estimates report the location and mass given to each point mass, evaluated at the mean $q = \Pr\{U=1\} = 0.031$. For example, in the baseline specification, the results estimate a point mass at 0.001, 0.031, and 0.641 with weights 0.446, 0.471 and 0.082. The values of $T(p)$ represent the markup that individuals at this location in the distribution would have to be willing to pay to cover the pooled cost of worse risks. All parameter estimates are constructed using maximum likelihood. Because of the non-convexity of the optimization program, I assess the robustness to 1000 initial starting values. All standard errors are constructed using bootstrap re-sampling using 1000 re-samples at the household level.

ONLINE APPENDIX FIGURE I: “First Stage” Impact of Unemployment on Beliefs



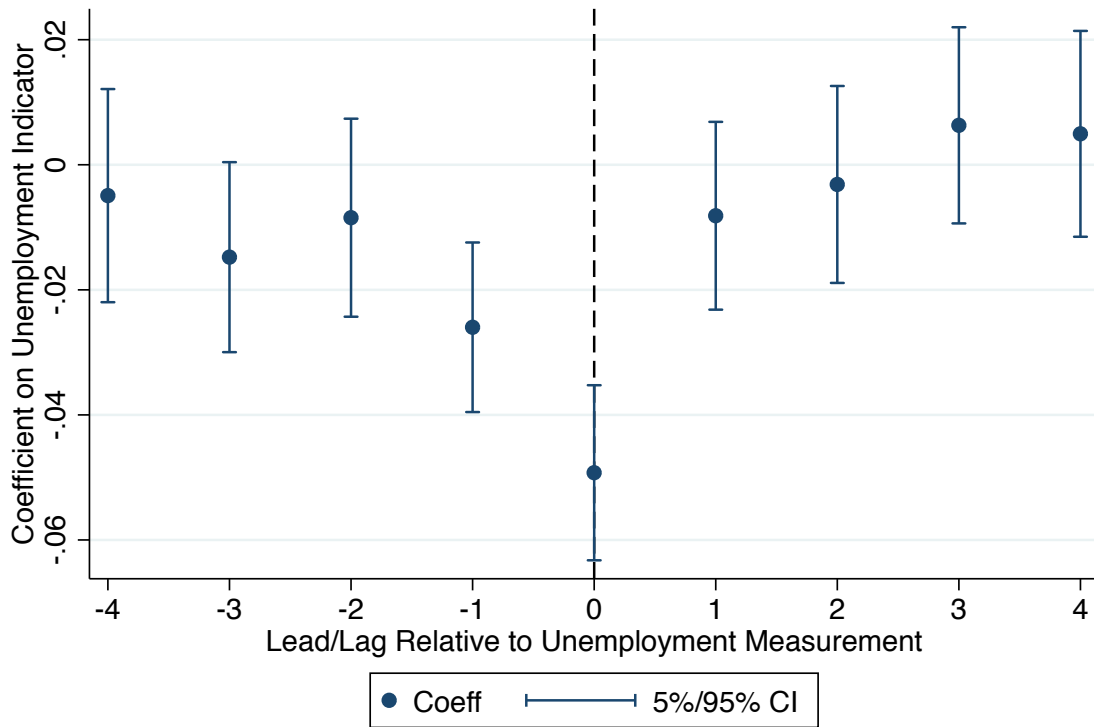
Notes: This figure presents the estimated coefficients of a regression of the elicitation (elicited in year t) on unemployment indicators in year $t + j$ for $j = 1, \dots, 8$. To construct the unemployment indicators for each year $t + j$, I construct an indicator for involuntary job loss in any survey wave (occurring every 2 years). I then use the data on when the job loss occurred to assign the job loss to either the first or second year in between the survey waves. Because of the survey design, this definition potentially misses some instances of involuntary separation that occur in back-to-back years in between survey waves. To the extent to which such transitions occur, the even-numbered years in the Figure are measured with greater measurement error. The figure presents estimated 5/95% confidence intervals using standard errors clustered at the household level.

ONLINE APPENDIX FIGURE II: Impact of Unemployment on Total Consumption Expenditure (2-year intervals)




Notes: This figure presents the estimated coefficients of a regression of leads and lags of log household consumption expenditure on an indicator for unemployment. The figure replicates the sample and specification in Figure IV (Panel B) by replacing the dependent variable with log total consumption expenditure on a sample beginning in 1999, surveyed every two years. I restrict the sample to household heads who are employed in $t - 2$ or $t - 4$. Following the baseline specification, the sample is restricted to observations with less than a threefold change in consumption expenditures. Note that after 1999, the PSID asks a broader set of consumption questions but is conducted only every two years, which prevents analyzing total 1-year interval responses to unemployment.

ONLINE APPENDIX FIGURE III: Impact of Job Loss on Consumption



Notes: This figure re-constructs the analysis in Figure IV using job loss instead of unemployment. I define job loss as an indicator for being laid off or fired from the job held in the previous wave of the survey. The figure present coefficients from separate regressions of leads and lags of the log change in food expenditure on an indicator of job loss, along with controls for year indicators and a cubic in age. Sample is restricted to household heads who are employed in years $t - 1$ and $t - 2$.

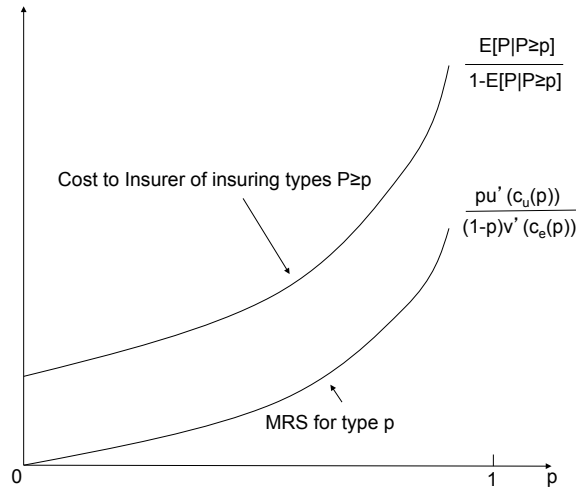
ONLINE APPENDIX FIGURE IV: Household Income Pattern Around Unemployment



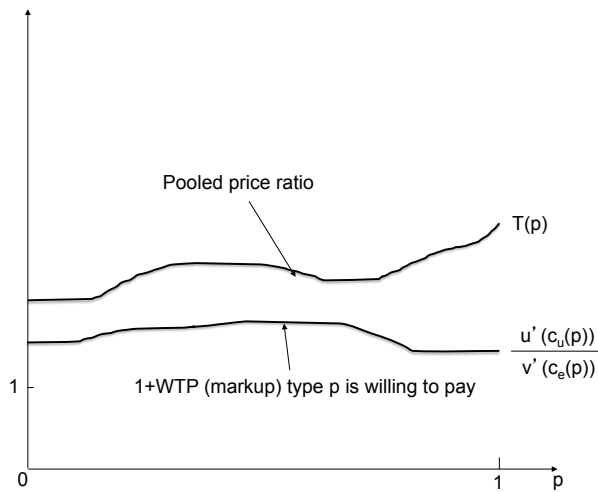
Notes: This figure presents the estimated coefficients of a regression of leads and lags of log household income on an indicator for unemployment. The figure replicates the sample and specification in Figure IV by replacing the dependent variable with log household income as opposed to the change in log food expenditure. I restrict the sample to household heads who are not unemployed in $t - 1$ or $t - 2$.

ONLINE APPENDIX FIGURE V: Illustration of No Trade Condition Using Demand and Average Cost Curves

A. MRS versus Average Cost

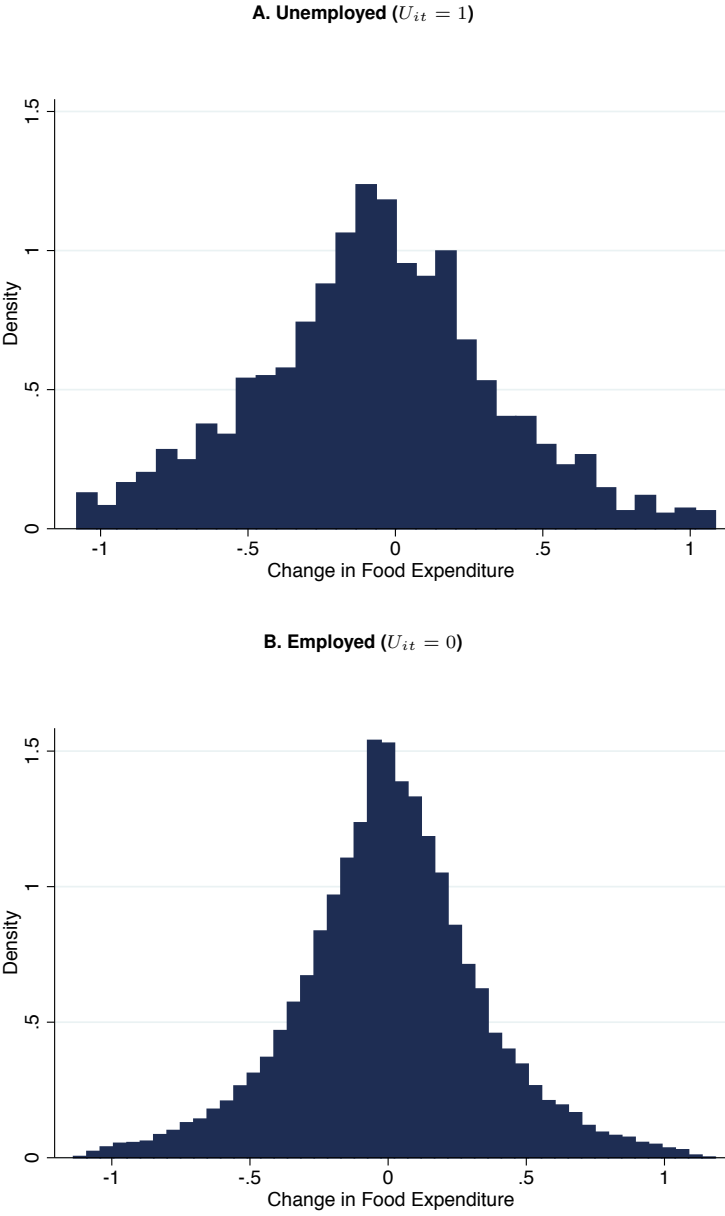


B. Pooled Price Ratio versus 1 + Willingness to pay (Markup)



Notes: This figure illustrates the no trade condition using the marginal and average cost curves as in Einav, Finkelstein and Cullen (2010). Panel A presents an illustrative example for the demand for a contract that pays \$1 in the event of becoming unemployed. The willingness to pay out of income if employed is given by the marginal rate of substitution, $\frac{p}{1-p} \frac{u'(c_u(p))}{v'(c_e(p))}$. Under a standard single crossing condition, all types $P \geq p$ would also purchase the insurance policy (see text for discussion of multi-dimensional heterogeneity). Therefore, the cost to the insurer of the contract is given by the average likelihood that the payment is made, $E[P|P \geq p]$ relative to the likelihood the payment is received, $1 - E[P|P \geq p]$, $\frac{E[P|P \geq p]}{1 - E[P|P \geq p]}$. Panel B normalizes by $\frac{1-p}{p}$ to illustrate the empirical approach that compares the pooled price ratio, $T(p) = \frac{1-p}{p} \frac{E[P|P \geq p]}{1 - E[P|P \geq p]}$, to one plus the markup individuals are willing to pay for insurance, $\frac{u'(c_u(p))}{v'(c_e(p))}$. The empirical results suggest the willingness to pay lies below the pooled price ratio, as depicted in Panels A and B.

ONLINE APPENDIX FIGURE VI: Distribution of Reported Log Food Expenditure Growth, Δ_{it}^* , for Employed and Unemployed



Notes: This figure plots histograms of Δ_{it}^* for those who are employed, $U_{it} = 0$, and unemployed, $U_{it} = 1$. Δ_{it}^* is defined as the residual from a regression of $\log(c_{it}) - \log(c_{it-1})$ on an age cubic and year dummies, X . I restrict the sample to household heads who are employed in $t - 1$ and $t - 2$. Following the baseline specification, the sample is restricted to observations with less than a threefold change in consumption expenditures.