Risk Protection and Redistribution in the Design of Social Insurance

Rory McGee and Martin O'Connell[*]

July 2025

Abstract

We study how heterogeneity in employment risk and income loss exposure shapes the social value of expanding unemployment insurance (UI) generosity. We estimate individual willingness-to-pay (WTP) for UI reforms, capturing both insurance value and cross-subsidization. WTP to expand U.S. UI peaks among lower income households, who are subsidized by others under the current system. We decompose the social value into gains from risk protection ($0.43 per benefit dollar) and redistribution through cross-subsidization ($0.10), net of incentive costs (-$0.71). Surplus redistribution raises the gains by 31%, split across (20%) and within (11%) income groups.

Keywords: social insurance, redistribution, risk protection, consumption, unemployment insurance

JEL classification: E21, E24, H23, H31, H50, I38, J64, J65

Acknowledgments: This project was supported by the ESRC through the ESRC-funded Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (grant reference ES/M010147/1). All errors and omissions remained the responsibility of the authors. We are grateful to Naoki Aizawa, Richard Audoly, Audra Bowlus, Richard Blundell, Eric French, Lance Lochner, Lee Lockwood, Costas Meghir, Davide Melcangi, Ryan Michaels, Corina Mommaerts, Luigi Pistaferri, Johannes Spinnewijn, Isaac Sorkin, and Todd Stinebrickner for their comments and insightful suggestions. We also thank seminar participants at various institutions and conferences.

[*]McGee: Western University and the Institute for Fiscal Studies. O'Connell: University of Wisconsin-Madison and the Institute for Fiscal Studies. **Correspondence**: rmcgee4@uwo.ca and moconnell9@wisc.edu

1. Introduction

A central role of government is to insure individuals against adverse shocks. In the United States, social insurance programs protect households against a wide range of risks, including unemployment, disability, and poor health. Roughly 36% of the federal government's budget—amounting to $2.4 trillion—is devoted to such programs.[1]

Designing effective programs involves a fundamental trade-off between providing insurance and minimizing behavioral distortions. The canonical framework for analyzing this trade-off uses a small set of empirical "sufficient statistics" to characterize welfare-relevant elasticities and consumption responses (e.g., Baily 1978; Chetty 2006; Chetty and Finkelstein 2013). Under certain conditions, the welfare gain from insurance can be inferred from the average consumption gap between "good" and "bad" states of the world, while the efficiency cost is summarized by a single behavioral elasticity.

While the benchmark approach offers powerful insights, it relies on average responses and abstracts from heterogeneity in both the value of insurance and the cost of behavioral distortions. In reality, individuals differ widely in their exposure to shocks, ability to smooth consumption (Blundell, Pistaferri, and Preston 2008), and responsiveness to incentives (Chetty 2008; Attanasio, Levell, Low, et al. 2018). Moreover, because premiums are rarely fully risk-adjusted, social insurance programs often redistribute through implicit cross-subsidization. These differences imply that the welfare consequences of reforms to program generosity can vary substantially across individuals.

This paper studies the value of changing the generosity of unemployment insurance (UI) in the United States, with particular emphasis on how heterogeneity in employment risk and exposure to income loss while unemployed shapes the welfare impact of reforms. We provide novel evidence on individual-level willingness-to-pay (WTP) for UI. A key innovation is that we estimate counterfactual consumption during unemployment for workers who are never observed unemployed in our data, enabling us to construct worker-level WTP. We use these estimates to decompose the welfare effects of reform into components reflecting the social value of risk protection, redistribution through cross-subsidization, and fiscal externalities.

We use data from the Panel Study of Income Dynamics (PSID) spanning 1999–2019. This nationally representative panel tracks individuals over time and

[1] We sum Social Security retirement, disability, and survivor insurance; unemployment insurance; and Medicare insurance for those age 65 and older. This share has remained relatively stable over time: Feldstein (2005) reports a value of 37% for 2003.

provides detailed information on labor market outcomes, demographics, and household consumption expenditures. To motivate our analysis, we begin by documenting descriptive evidence that workers who experience an unemployment spell have lower equivalized household consumption, individual earnings, and household income *prior* to job loss, relative to those who remain continuously employed. Moreover, even conditional on income, workers who go on to experience unemployment exhibit lower *in-work* consumption. These patterns highlight the redistributive role of unemployment insurance—both across and within income groups.

Our worker-level measure of WTP for UI reform depends on two components. The first is their *personalized insurance value*, which captures the individual's exposure to employment risk—specifically, the value they place on a marginal transfer of resources from states in which they are employed to those in which they are not. The second is their *net fiscal position* with respect to the program reform—that is, the extent to which they are, ex ante, net contributors or beneficiaries. The insurance value varies across individuals depending on the severity of income loss in unemployment and their capacity for self-insurance. The net fiscal position reflects the gap between an individual's risk and the implicit premium embedded in non-individualized taxes and benefits. When implicit contributions are not actuarially fair, this gap gives rise to redistribution through cross-subsidization.

To measure the insurance component of WTP, we build on the "consumption-based" approach to valuing social insurance, originally developed by Gruber (1997). We extend it by providing new evidence on the full distribution of unemployment-induced consumption declines across workers.

To do so, we apply clustering regression techniques, introduced by Lewis, Melcangi, and Pilossoph (2024) to estimate heterogeneous marginal propensities to consume (see also Boehm, Fize, and Jaravel 2025). These methods allow us to uncover heterogeneity in consumption responses across latent groups without requiring strong assumptions about which observable characteristics drive this variation. Crucially, the framework allows us to recover counterfactual consumption declines for workers we never observe experiencing unemployment. We estimate an average unemployment-induced consumption drop of 11.7%, broadly in line with prior studies (e.g., Gruber 1997; Hendren 2017; Kroft and Notowidigdo 2016). However, we also find substantial heterogeneity reflecting differences in the capacity for self-insurance: for example, the decline at the 95th percentile is roughly twice as large as the median decline.

The cross-subsidization component of WTP is shaped by the interaction between individual-level employment risk and the structure of the reform under consideration. We measure employment risk using estimates of how unemployment incidence varies across workers with different demographic and socioeconomic characteristics (see also, Anderson and Meyer 2006). We compare a natural expansion of the U.S. UI system—which replaces a fraction of lost earnings up to a cap—with alternative reforms that expand entitlement either by a flat amount or in proportion to lost earnings with no cap.

We find that WTP to expand the generosity of the current U.S. UI system is highly heterogeneous. Median WTP declines across the household income distribution, turning negative around the middle. This pattern arises despite all workers valuing actuarially fair insurance positively. The decline reflects the fact that, on average, workers above the median income level cross-subsidize others. Low-income workers benefit from the system due to higher employment risk—an effect strong enough to outweigh the link between benefit entitlements and earnings. A flat benefit reform further tilts the WTP distribution in favor of lower-income workers, whereas a proportional expansion dampens the income gradient by increasing entitlements more for higher earners.

We also find substantial heterogeneity in WTP conditional on income. For example, the standard deviation of WTP for expanding the current system among workers in the top income decile is more than half as large as the full range of median WTP across the income distribution. These reform-specific patterns of winners and losers are central to the social value of UI reforms.

We next consider the social value of a budget-balanced expansion in UI generosity. Since this requires inter-personal comparisons, our evaluation depends on the specification of social preferences. In our baseline implementation, we adopt welfare weights that value cross-income surplus redistribution at the marginal cost to the government of achieving equivalent redistribution through the income tax system (Hendren 2020), and place greater value on within-income-group surplus redistribution for individuals with lower long-run equivalized household consumption, which we treat as measuring lifetime income and economic wellbeing (Poterba 1989; Meyer and Sullivan 2003)

We show that the social value of expanding UI can be decomposed into three distinct components. The first is the value of risk protection, which takes the form of a risk- and welfare-weighted average of individual insurance values, scaled by an aggregate welfare weight. The individual weights reflect heterogeneity in the marginal social value of surplus among program beneficiaries, while the aggregate weight captures their average marginal value relative to those who fund the program. We estimate that the risk protection value of a budget-balanced expansion of the current UI program is $0.43 per dollar of increased benefit expenditure.

The second component captures the social value of cross-subsidization: redistributive gains that arise independently of insurance value. Absent in representative-agent frameworks, this component can generate welfare gains (or losses) even when individuals are fully insured through other means. For the current UI program, we estimate that the redistributive value from cross-subsidization associated with increased generosity is $0.10 per dollar of increased benefit expenditure. Combining these first two terms, the total social value of resource reallocation is $0.54, of which approximately 69% reflects efficiency gains from additional insurance ($0.37), 20% stems from cross-income surplus redistribution ($0.11), and 11% from within-income redistribution ($0.06).

The third component reflects the fiscal externality associated with behavioral responses to program generosity. In the canonical model, this is typically summarized by a single elasticity. In our setting, by contrast, it depends on the covariance between individuals' risk levels, their net fiscal position (i.e., the gap between taxes paid and benefits received across states), and their behavioral elasticities, which we calibrate based on evidence in Chetty (2008). We find that the direct fiscal externality from increased UI payments is $0.38 per dollar of increased benefit expenditure, with an additional cost of $0.33 due to lower income tax receipts. Thus, the total incentive cost arising from these sources exceeds the gains from risk protection and redistribution.

Our measure of the social value of changing UI generosity depends critically on the structure of the reform. Under a flat expansion—where UI entitlements increase uniformly across workers—the value of resource reallocation rises to $0.63, as this form of expansion directs a larger share of resources toward households that are less well-off and more exposed to employment risk. In this case, the welfare gains from increased generosity outweigh the associated incentive costs. Differences across reform types underscore that the social value of UI is likely to vary substantially across countries, due to substantial differences in program design (Spinnewijn 2020). The standard implementation effectively evaluates a flat reform under social preferences that are indifferent to how surplus is distributed, yielding a much lower estimated social value of resource reallocation of just $0.38. We show that abstracting from surplus redistribution across heterogeneous agents meaningfully alters the conclusions of UI policy analysis.

We contribute to a literature that uses a sufficient statistics approach to study the insurance-incentive trade-off in the design of UI (e.g., Schmieder, Von Wachter, and Bender 2012; Kolsrud, Landais, Nilsson, et al. 2018; Landais and Spinnewijn

2021).[2] We also build on recent work that measures the value of social transfers, including disability insurance (Deshpande and Lockwood 2022), public pensions (Kolsrud, Landais, Reck, et al. 2024), and welfare programs (Rafkin, Solomon, and Soltas 2023).

Our work also relates to a literature that examines how correlation between risk and ability shapes the optimal design of income taxes (e.g., Blomqvist and Horn 1984; Rochet 1991; Cremer and Pestieau 1996; Boadway, Leite-Monteiro, Marchand, et al. 2006), as well as research on how involuntary unemployment (Kroft, Kucko, Lehmann, et al. 2020) and the structure of non-linear UI benefit schedules (Ferey 2022) influence optimal tax formulas. We complement this body of work by empirically measuring the extent to which current UI systems redistribute surplus across heterogeneous workers—both across and within income groups.

The rest of this paper is structured as follows. In Section 2 we outline our measure of individual willingness-to-pay for increased UI generosity. Section 3 introduces our data and presents preliminary evidence on negative selection among the unemployed. Section 4 introduces our empirical approach for estimating willingness-to-pay and present results. Section 5 outlines a normative framework for valuing a change in UI generosity, along with our quantitative analysis. A final section concludes.

2. Individual's Value of a Reform

In this section, we derive an expression for individual-level willingness-to-pay for expanding the generosity of a social insurance program. We present a simple framework that captures the salient features unemployment insurance, as well as other programs such as workers' compensation, disability insurance, health insurance, and insurance against natural disaster. In Appendix A.1, we show that our results generalize to a much richer dynamic environment (Chetty 2006).

We build on the canonical set-up (see Chetty and Finkelstein 2013) by considering a *heterogeneous* population of individuals (agents) exposed to risk. Individual-level willingness-to-pay is a key input for evaluating the social value of reform, which depends both on its aggregation across individuals and on the resulting incentive costs (see Section 5).

[2] A complementary strand of the literature employs dynamic structural models to analyze UI design (e.g., Acemoglu and Shimer 1999; Lentz and Tranaes 2005; Krusell, Mukoyama, and Şahin 2010) and highlights UI policy's redistributive implications (e.g., Audoly 2024; Haan and Prowse 2024).

2.1. Theory

Agent's problem. There is a unit continuum of agents, indexed by i. Each agent faces uncertainty over two states: high (h) and low (l), for instance, indicating being employed and unemployed. Individuals potentially differ in their preferences, incomes, and the costs associated with mitigating risk. We denote income in the high and low state by y_i and z_i, respectively. The government provides a benefit $\mathcal{B}(y_i)$ to those in the low state and levies a tax $\mathcal{T}(y_i)$ on individuals in the high state. We allow the tax and benefit schedules to be arbitrary functions of high state incomes.[3] Consumption in the high state is therefore $c_i^h = y_i - \mathcal{T}(y_i)$, while consumption in the low state is $c_i^l = z_i + \mathcal{B}(y_i)$. Let $u_i^s(c)$ denote state-specific utility from consumption c; we assume $u_i^s(.)$ is increasing, concave and twice continuously differentiable. The difference in high and low state consumption is $\Delta c_i = c_i^h - c_i^l$. For the majority of agents, we expect $\Delta c_i > 0$.

Individuals can control the probability of being in the high state by undertaking actions, which we model through a scalar e and refer to as effort. We normalize the units of effort so that it equals the probability of being in the high state (and therefore $e \in [0, 1]$). Effort is costly, captured by $\psi_i(.)$, which is increasing and convex. The agent solves:

$$V_i = \max_e e u_i^h(c_i^h) + (1 - e) u_i^l(c_i^l) - \psi_i(e), \tag{1}$$

where V_i is the maximized value of their expected utility. Agent i's optimal effort choice, e_i, satisfies the first-order condition: $u_i^h(c_i^h) - u_i^l(c_i^l) = \psi_i'(e_i)$; the individual chooses the level of effort that equates the marginal benefit of exerting additional effort, equal to the difference in utilities in the two states, with the marginal cost of exerting additional effort.

Heterogeneity. The individual index i captures arbitrary heterogeneity arising from various sources. First, differences in earning ability—due, for example, to variation in innate ability, human capital, skills, or labor market opportunities—are reflected in high-state incomes, y_i. Low-state income, z_i, additionally reflects an individual's capacity to self-insure, for example through private savings or spousal insurance. As a result, differences in incomes across states reflect both the direct financial loss from entering the low state and individuals' self-insurance capacity. Combined with the influence of the tax and benefit system, these factors drive heterogeneity in the consumption gap, Δc_i.

[3] In our empirical implementation, we account for the role of family composition in the U.S. tax-and-transfer system. We omit this here to simplify exposition.

Additionally, heterogeneity in attitudes toward work and employment opportunities (in the unemployment or disability insurance context) or differences in the cost of abatement investments (in the health or natural disaster insurance context) are reflected in $\psi_i(.)$.

Finally, variation in the marginal utility of consumption (for example, due to household composition) and attitudes toward risk is captured by individual-specific utilities, $u_i^s(.)$. This approach also allows for variation in state-dependence across individuals, which may reflect differences in the opportunity cost of time or the availability of consumption substitutes.

Each of these factors influences the agent's optimal effort choice and therefore leads to heterogeneity in risk levels, e_i.

Arbitrary reforms. We consider the effects of an arbitrary marginal reform to the benefit and tax schedules.[4] We parameterize the reform by $(db, d\tau)$, such that the impact on state-specific consumption is given by:

$$\begin{aligned} dc_i^l &= \phi_{\mathcal{B}}(y_i)db \\ dc_i^h &= -\phi_{\mathcal{T}}(y_i)d\tau, \end{aligned} \tag{2}$$

where $\phi_{\mathcal{B}}(y_i)$ and $\phi_{\mathcal{T}}(y_i)$ are positive functions of high-state income y_i, determining the marginal effect of the reform on consumption in the low and high states, respectively. While these functions allow for reforms that vary arbitrarily by income, in practice reforms are typically simple functions of income—for example, a flat benefit increase ($\phi_{\mathcal{B}} = 1$) or proportional benefit increase ($\phi_{\mathcal{B}} = y_i$), funded by a flat tax adjustment ($\phi_{\mathcal{T}} = 1$). The structure of the reform is central to determining individual willingness-to-pay.

Willingness-to-pay. To define the individual-level willingness-to-pay (WTP) for a social insurance expansion, we consider the reform $d\theta \equiv (db = 1, d\tau = \frac{\int_i (1-e_i)\phi_{\mathcal{B}}(y_i)di}{\int_i e_i \phi_{\mathcal{T}}(y_i)di})$. This reform raises benefits and adjusts the tax schedule just enough to offset the resulting mechanical budgetary cost.

Let $\frac{dV_i}{d\theta}$ denote the impact of this reform on individual i's expected utility. We

[4]Focusing on marginal reforms has the advantage that, by the envelope theorem, the impact on individual utility is fully captured by the direct effect on consumption and is unaffected by behavioral responses. As a result, willingness-to-pay can be characterized without a full specification of the individual's optimization problem, and WTP in our simple framework aligns with that in a richer dynamic model (see Appendix A.1). The expression can also be interpreted as a first-order approximation to a non-marginal reform, or used to construct higher-order approximations by integrating along the path of a larger reform (e.g., see Kleven 2021).

define their willingness-to-pay for this reform as:

$$\begin{aligned}\mathrm{WTP}_i \equiv & \frac{1}{(1-e_i)\phi_{\mathcal{B}}(y_i)} \frac{1}{u_i^{h\prime}(c_i^h)} \frac{dV_i}{d\theta} \\ = & \underbrace{\left(\frac{u_i^{l\prime}(c_i^l) - u_i^{h\prime}(c_i^h)}{u_i^{h\prime}(c_i^h)}\right)}_{\text{insurance}_i} + \underbrace{\left(1 - \frac{e_i\phi_{\mathcal{T}}(y_i)}{(1-e_i)\phi_{\mathcal{B}}(y_i)} \bigg/ \frac{\int_{i'} e_{i'}\phi_{\mathcal{T}}(y_{i'})di'}{\int_{i'}(1-e_{i'})\phi_{\mathcal{B}}(y_{i'})di'}\right)}_{\text{cross-subsidization}_i}. \end{aligned} \tag{3}$$

The scaling in the first line expresses the utility gain from the reform per $1 increase in expected benefit payments, relative to the marginal value of an additional dollar in the high state.

The second line decomposes WTP into two conceptually distinct components. The insurance component captures the value to the individual of transferring a marginal unit of consumption from the high to the low state, valued at their own marginal utilities. This reflects the individual-specific insurance value of social insurance, and is likely to be positive for most individuals.

The cross-subsidization component captures how the individual's net expected transfer under the reform compares to that under an actuarially fair adjustment. Under actuarial fairness, the individual would face an expected tax increase of $e_i\phi_{\mathcal{T}}(y_i)$ to fund an expected benefit increase of $(1-e_i)\phi_{\mathcal{B}}(y_i)$. In practice, however, the actual reform may depart from actuarial fairness for most individuals. The cross-subsidization term reflects this deviation: it is positive for individuals whose actuarially fair price is higher than the implicit pooled price embedded in the actual reform, meaning they are net beneficiaries, and negative for those whose fair price is lower, implying they are net contributors. In other words, it measures whether the individual gains or loses from redistribution relative to a fair-pricing benchmark.

2.2. Discussion

Our willingness-to-pay decomposition relates to the expression in Finkelstein, Hendren, and Luttmer (2019), who show that the WTP for a marginal expansion of Medicaid can be decomposed into a "pure insurance" term, capturing the covariance between marginal utility and insured medical spending across health states, and a "pure transfer" term, representing the expected reduction in medical spending due to a lower out-of-pocket price. A key difference in our setting is that the pattern of transfers, or cross-subsidization, arises from departures from actuarial fairness embedded in the structure of the tax and benefit reform.

An alternative interpretation of our willingness-to-pay decomposition is that the cross-subsidization component reflects the ex post expected value of the reform, while the insurance component captures the extent to which the *ex ante* value of the reform to the individual exceeds its ex post value. This premium arises because the reform reallocates resources from the high state, where marginal utility is relatively low, to the low state, where marginal utility is higher. This interpretation aligns with Lieber and Lockwood's (2019) decomposition of the individual-level value of an in-kind transfer.[5]

2.3. Implementation

The individual-level insurance term depends on the gap in marginal utility of consumption between the low and high states. In our baseline implementation, we assume that marginal utility is state independent and that individuals share a common local curvature of utility at their high-state consumption level.

ASSUMPTION 1. *Marginal utility is state independent (i.e., $u_i^{s'}(.) = u'(.)$ for $s \in \{l, h\}$). and relative risk aversion at high-state consumption is homogeneous across individuals (i.e., $-\frac{u_i''(c_i^h)c_i^h}{u_i'(c_i^h)} = \gamma$).*

Below we discuss the implications of relaxing both state-independence and the assumption of common relative risk aversion. Following Gruber (1997), we adopt a consumption-based approach to measuring the marginal utility gap. Under assumption (1), a second-order (quadratic) approximation to the utility function yields:

$$\frac{u'(c_i^l) - u'(c^h)}{u'(c_i^h)} \approx \gamma \frac{\Delta c_i}{c_i^h}, \tag{4}$$

This re-expresses the value of insurance to an individual as the product of their percentage consumption decline between the high and low state and the coefficient of relative risk aversion.[6] Our empirical strategy estimates the unconditional distribution of consumption declines across the population.

[5] Both Finkelstein, Hendren, and Luttmer (2019) and Lieber and Lockwood (2019) consider settings with potentially many state of the world and convert utility to monetary terms by scaling by the average marginal utility. In contrast, because our empirical application focuses on unemployment insurance, we focus on two states of the world and scale by the marginal utility in the high-income (employment) state.

[6] Alternative approaches to deriving empirical measures of the marginal utility gap, implemented in the UI context, include decomposing unemployment responses into substitution and income effects (Chetty 2008), measuring reservation wage responses (Shimer and Werning 2008) and measuring gaps in the marginal propensity to consume across states (Landais and Spinnewijn 2021). Each of these methods could be used in place of the consumption-based approach to implement equation (3) (or our subsequent normative analysis).

Cross-subsidization depends both on state-specific probabilities and the structure of the reform under consideration. We estimate how unemployment probabilities vary across individuals, and consider a range of different realistic reforms.

3. Data and Setting

Our empirical analysis focuses on unemployment insurance (UI). In the United States, while the federal government provides guidelines for UI eligibility and benefit generosity, UI is administered at the state level with states determining benefit duration, generosity and eligibility (see Von Wachter 2019 for a discussion). Benefit payments, and entitlement, are computed using workers' (quarterly) employment history and earnings. The modal duration offered by states is six months outside of federal recession extensions. Eligible workers receive payments in proportion to their past earnings up to a maximum threshold.

The U.S. UI system is funded through payroll taxes that are legally incident on employers. The federally mandated minimum for taxable earnings is $7,000 per year. While most states apply a higher threshold, in practice the tax often functions as a head tax on the number of workers rather than a proportional earnings tax. A distinctive feature of the funding of U.S. UI is that tax rates vary across employers due to partial experience rating (see Guo and Johnston 2021 for discussion). Recent evidence suggests that firms are unable to pass the firm-specific component of UI taxes through to wages (Guo 2024). As a result, in our empirical quantification of UI expansion, we assume that only the common component of UI taxes is incident on wages.[7]

Sample. We use data from the Panel Study of Income Dynamics (PSID), a nationally representative survey of the U.S. population. The PSID has surveyed households since 1968 collecting information on demographics as well as labor market outcomes including unemployment, earnings, and wages. We use data from 1999-2019, the biennial so-called "new" PSID, which includes comprehensive information on consumption expenditures.[8]

We drop households from our sample where the reference person does not

[7] Guo (2024) shows the firm-specific component instead induces employment responses. See Spaziani (2024) for analysis of how experience rating alters the risk borne by employers. Anderson and Meyer (2000) similarly find that firm-specific components are largely incident on firms, though they emphasize that employment responses also effect layoffs and subsequent UI claims (and their denials).

[8] Before 1999 the PSID was collected annually. However, it included limited data on consumption outside of food expenditures. Both Gruber (1997) and Hendren (2017), for example, use this food expenditure data in studies of unemployment insurance.

provide any information on their educational attainment, and we focus on a sample where the reference person is aged between 23 and 60. This excludes people from our sample whose labor supply is likely driven by education and retirement decisions. To proxy for UI program eligibility, we omit households where the reference person never reports being employed, as well as periods immediately following a spell out of the labor force. This yields 55,671 observations (15,270 families observed for an average of 3.7 waves) of which 3,331 observations correspond to the reference person being unemployed (2,211 families for an average of 1.5 waves). For families with changes to the number of adults (such as divorce, marriage, or death) we construct panel identifiers for demographically stable units and omit the year of the change. We focus on the labor market status of the reference person (in some places referring to them as worker).

Key variables. In our main analysis, we define unemployment based on the employment status reported by the worker at the time of their survey interview. We measure worker earnings using reported wages when employed, while we define household income as the sum of worker and spousal wages. For each worker, we also construct a measure of permanent income based on worker-level fixed effect in a log wage regression, conditional on life-cycle wage growth. We use a comprehensive measure of expenditures on non-durable consumption and services, excluding spending on health care and housing.[9] To account for differences in household composition, we equivalize consumption expenditures using the square root of household size. We convert all financial variables to 2019 dollars. Summary statistics, details on sample selection, and the construction of key variables including our measure of worker permanent income are provided in Appendix B.

Preliminary evidence. In the next section, we document heterogeneity in willingness-to-pay for UI. In Section 5, we undertake a normative analysis, comparing a welfare-weighted aggregation of willingness-to-pay to the program's incentive costs. To motivate this analysis, we begin by presenting descriptive evidence on the extent of negative selection among the unemployed. On average, consumption among the unemployed is about three-quarters that of the employed. This reflects both negative selection—individuals who experience unemployment tend to have lower consumption levels even while working—and direct

[9]There are some differences in how the components of consumption are measured over time, including expansion of measured items and changes to the aggregation or disaggregation of certain categories. Our analysis uses the broadest set of consumption categories available for our measure in each wave. We include survey-wave specific fixed effects in our econometric strategy to account for changes in the coverage of our consumption measure.

consumption declines associated with unemployment. In this section, we focus on measuring the extent of negative selection (presenting estimates of unemployment-induced consumption declines in the next section).

Specifically, we examine the relationship between individuals who experience unemployment and measures of available resources during periods of employment. By focusing only on periods of employment, this analysis avoids mechanically induced reductions in earnings and associated consumption responses during unemployment spells.

To operationalize this, we estimate the following regression:

$$\ln y_{i,t} = \delta D_i + \beta X_{i,t} + u_{i,t}, \tag{5}$$

where $y_{i,t}$ is a measure of household resources, D_i is an indicator for whether worker (i.e., reference person) i has ever experienced an observed unemployment spell,[10] and $X_{i,t}$ is a vector of controls that include a cubic function of the worker's age, as well as indicators for household composition and year effects.[11] The coefficient of interest is δ, which captures the difference in household resources associated with having experienced unemployment.

Figure 1 reports estimates of δ and its 95% confidence interval for four measures of resources: household consumption, the reference person's labor income, household labor income, and worker permanent income. For each outcome, we report estimates without controls, with controls, and excluding post-unemployment periods to eliminate the effects of potential labor market scarring.[12] The figure clearly shows that households experiencing unemployment have lower consumption, individual income, and household income *while in work*, as well as lower permanent income, compared to households whose reference person is never unemployed.

To the extent that social insurance programs reallocate resources across income groups, this redistribution affects their social value. However, redistribution across income levels can be achieved, or offset, through adjustment to the tax and transfer system. In contrast, redistribution across households with the same income level but different levels of consumption is more difficult to achieve through tax and transfer policies. Moreover, because consumption may better reflect permanent income, redistribution based on consumption differences may be especially valuable. The

[10]Specifically, $D_i = 1$ if worker i reports an unemployment spell, including any occurring between survey waves.

[11]In Appendix D.1, we show that the presence of negative selection is robust to including controls for education and race in addition to household structure.

[12]For permanent income, differences reflect variation in the composition of workers observed before and after unemployment.

green marker in Figure 1 shows that households experiencing unemployment, even conditional on a flexible function of household income, have lower consumption than those that do not.

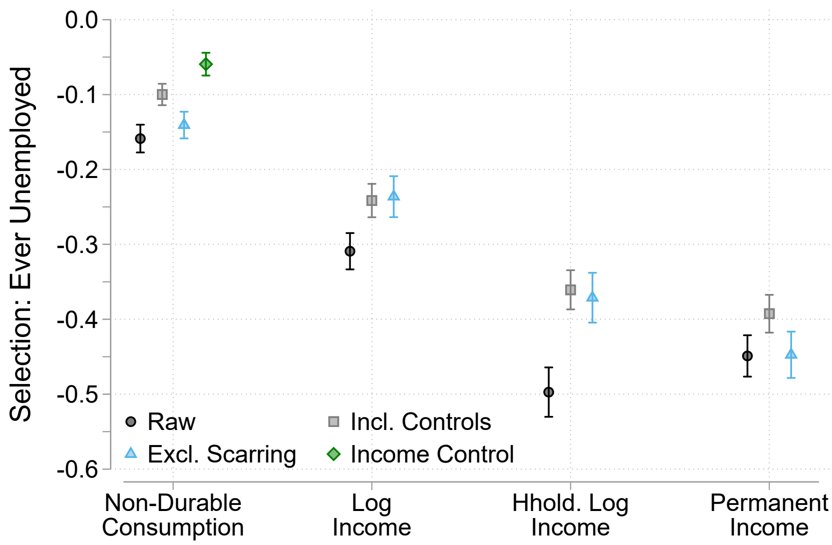

FIGURE 1. Negative Selection Among the Unemployed

Notes: Authors' calculation from the PSID data. All regressions drop periods of unemployment. The no scarring sample conditions on periods before first observed unemployment. Raw specification includes no controls. All other specifications additionally control for a cubic function of age, indicator variables for family size and whether they are married, and a set of year specific indicators. Income controls additionally controls for a cubic function of household income.

4. Measuring Willingness-to-Pay

In this section, we first outline how we measure heterogeneity in unemployment-induced consumption declines and employment risk, which, along with the coefficient of relative risk aversion determine worker risk exposure. In each case, we detail our estimation strategy, present results, and compare them to findings in the existing literature. We then use these estimates to compute individual willingness-to-pay for UI reform.

4.1. Consumption Declines

Method. The standard approach to recovering consumption declines (e.g., Gruber 1997) estimates an average effect of unemployment, conditional on being employed in the preceding period ($t-2$ in our case), using the following specification

$$\begin{aligned}\Delta^{FD}_{i,t} &\equiv \ln\left(c_{i,t}\right) - \ln\left(c_{i,t-2}\right) \\ &= \delta_0 + \delta_1 U_{i,t} + \beta X_{i,t} + \varepsilon_{i,t},\end{aligned} \tag{6}$$

where $U_{i,t}$ is an indicator variable denoting whether the worker is unemployed in period t[13] and $X_{i,t}$ are demographic controls. This is specified in first differences to allow for arbitrary permanent individual-level heterogeneity in consumption levels. The parameter of interest, δ_1, measures the average proportional consumption decline that results from becoming unemployed. The identification assumption underpinning this strategy is that the trend in consumption for the employed acts as a valid counterfactual for the evolution of in-work consumption of the unemployment had they remained employed.

To estimate the distribution of consumption declines, we specify a finite mixture approximation to the true distribution (e.g., Heckman and Singer 1984; Keane and Wolpin 1997). Our approach follows Lewis, Melcangi, and Pilossoph (2024), who estimate the distribution of marginal propensities to consume out of tax rebates. We assume each worker belongs to one of a finite number of latent types and augment the specification in equation (6) as follows:

$$\Delta^{FD}_{i,t} = \sum_{k\in K} \mathbb{1}\left[k(i)=k\right]\left(\delta^k_0 + \delta^k_1 U_{i,t}\right) + \beta X_{i,t} + \varepsilon_{i,t}. \tag{7}$$

We include group-specific intercepts δ^k_0, so that δ^k_1 can be interpreted as the consumption decline upon unemployment for group k. We assume that, conditional on group, controls, and unemployment status, the error term is mean zero.

This estimating equation has parallels with the reduced-form in Patterson (2023), who uses unemployment shocks to instrument for income declines. However, an important difference is that we model heterogeneity across unobserved latent types. To implement this, we use a Gaussian mixture linear regression (see Quandt 1972), which assumes that the errors are normally distributed with group-specific variances.

The key identifying assumption in equation (7) is that, conditional on group, the evolution of consumption for the employed acts as a valid counterfactual for the unobserved trend in consumption had the unemployed remained employed. This assumption mirrors that underlying equation (6), but is applied within group. Latent types are identified by grouping households with similar distributions of

[13] As in Gruber (1997) and Hendren (2017) we measure unemployment at the time of the interview.

consumption growth in each state. We estimate group membership jointly with the other parameters; see Appendix C for details.

We recover estimates of the group-specific parameters, the common effect of covariates, the group-specific variance of the regression residuals, and the unconditional probability of group assignment (π^k for each group k) by maximum likelihood estimation.

To obtain worker-level predictions of consumption declines, we compute mean squared error minimizing "posterior" probabilities of group assignment, $\hat{\pi}_i^k$, at the individual level, which depend on our parameter estimates and the individuals' sequence of observed regressors and outcome variables. We use these to simulate the value of unemployment-induced consumption declines within sample with $\Delta c_i / c_i$ as $\sum_{k \in k} \hat{\pi}_i^k \hat{\delta}_1^k$. These results are conditional on the specified number of groups G which we select using the Bayesian Information Criterion.

As discussed in Lewis, Melcangi, and Pilossoph (2024), this mixture regression can be viewed as a form of clustering regression which jointly (i) groups households together that have similar latent consumption responses to unemployment and employment and (ii) provides estimates of the consumption decline within these groups. Thus, our approach to estimating counterfactual consumption declines for those without observed unemployment spells can be viewed as a form of matching-based identification strategy. Intuitively, the approach can be thought of as entailing the following steps. First, we group together households observed in both states based on their consumption growth when employed and their consumption decline when unemployed. We then 'match' households who are always employed to those who experience similar consumption growth when employed, conditional on observables. Then, we assign them the matched household's observed consumption decline for their counterparts. In practice, we simultaneously estimate probabilistic group assignment and parameters, with the posterior weight, $\hat{\pi}_i^k$, serving as a (convex) imputation weight.[14]

Estimates. Table 1 reports our estimates from the finite mixture approach alongside estimates from the homogeneous specification (equation (6)). The homogeneous specification yields an estimate of the average consumption decline of 12.9% at the onset of unemployment.

[14]Note that if two groups experience similar consumption growth when employed, but different consumption declines (e.g., due to varying abilities to self-insure), our imputation assigns a household with similar consumption growth (who is only observed employed) a consumption decline that is a weighted average of these groups. As we show in Appendix C, these weights depend not only on the average consumption growth, but also its variability, which corresponds to the pass through of shocks to consumption (Blundell, Pistaferri, and Preston 2008). In principle, latent consumption risk may vary over time, and our approach could be extended to allow for a regime-switching model.

	Homogenous	Mixture Model Group 1	Group 2	Group 3
$\mathbb{1}\left[U_{i,t}=1\right]$	-0.129^{***}	-0.257^{***}	-0.129^{***}	-0.060^{***}
	(0.009)	(0.015)	(0.006)	(0.006)
Share ($\hat{\pi}^g$)		0.117^{***}	0.486^{***}	0.396^{***}
		(0.028)	(0.018)	(0.021)
Controls				
Age	*Yes*	*Yes*	*Yes*	*Yes*
Year	*Yes*	*Yes*	*Yes*	*Yes*
N	37, 260	37, 260	37, 260	37, 260

TABLE 1. Heterogeneity in Consumption Declines ($\Delta c/c$)

* $p < 0.10$, ** $p < 0.5$, *** $p < 0.01$. The first column reports the average consumption decline estimated under the assumption of homogeneous consumption drops (equation (6)). The remaining columns report the consumption declines, δ_1^k, and population shares, π^k, in our Gaussian mixture linear regression framework (equation (7)). In the mixture model standard errors account for uncertainty over both group assignment and parameter estimates conditional on group. In addition to the unemployment indicator, we control for a cubic polynomial in age, a series of year dummies, and the log of the change in family size. Results in Table D.1 provide heteroskedastic robust standard errors for the homogeneous estimate.

The implied average consumption decline from our mixture approach, aggregating the group-specific consumption decline δ_1^k, weighted by π^k, is slightly smaller at 11.7%.[15] This comparison of averages, however, ignores the large degree of heterogeneity in the exposure to unemployment-induced consumption declines that we find. The group-specific consumption declines we estimate range from 25.7% to 6.0%. The probability mass of the group with the smallest decline, 6.0%, is almost 40%. This group is relatively well insured against employment risk. A second group, with probability mass of 49%, has a consumption drop that equals the homogeneous estimate. The remaining probability mass is accounted for by a group with a much larger consumption decline of 25.7%

Figure 2 shows the distribution of posterior predicted consumption declines

[15]Part of this discrepancy is the average consumption declines represent subtly different estimands. In the homogeneous specification $\widehat{\delta}_1$ is an estimate of the average consumption decline among the unemployed (or the average treatment on the treated). In contrast, the unconditional estimate in our mixture model estimates the effect in the population (or the average treatment effect) which is possible due to the additional assumptions on the data generating process we impose. When we aggregate estimates of the individual consumption decline by ex-ante unemployment probabilities, we obtain a drop of 12.3%, which is very similar to the homogeneous estimate.

across worker, further emphasizing the large cross-sectional heterogeneity in risk exposure. Using the standard homogeneous estimator, we have a mass point at the estimate (the dashed vertical line). Instead, we find that 51% of households exhibit consumption drops that are smaller than the standard homogeneous estimate, but there is a long right tail of households with larger drops. This variation reflects broad differences in the ability of households to self-insure employment risk.

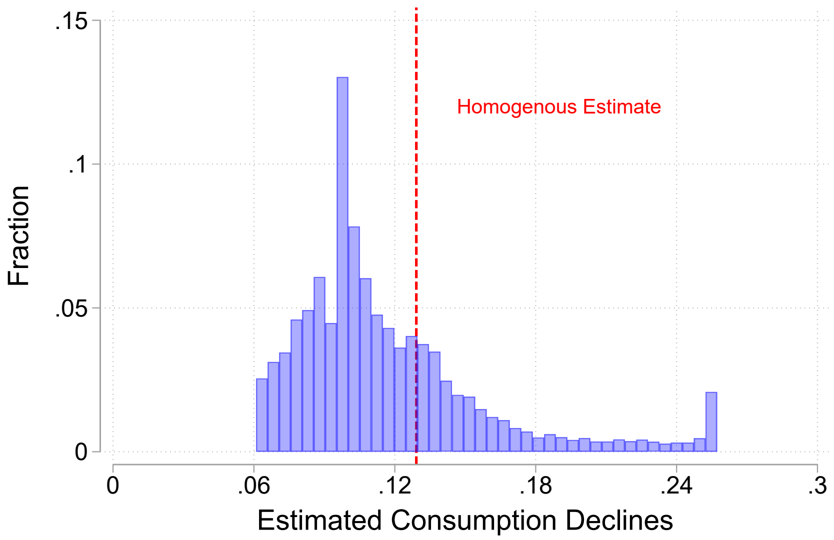

FIGURE 2. Distribution of Estimated Consumption Declines

Notes: Authors' estimates from the PSID data. The histogram (light blue bars) plots the worker-level proportional consumption declines constructed using the Gaussian mixture linear regression-estimated parameters and individuals' posterior probabilities for each group, $\hat{\pi}_i^k$. For each household we compute the posterior-weighted effect of unemployment across the discrete group-specific unemployment consumption declines. The sample is defined as in the text. The Bayesian Information Criterion selects K = 3. The homogeneous estimate (red dashed vertical line) is estimated imposing K = 1 in our baseline specification.

Discussion. Our homogeneous estimate is slightly larger than those reported by the existing literature (e.g., Gruber 1997 and Hendren 2017), which find declines in food expenditures of between 7 and 10%. There are three differences between our approach and those in the existing literature that can together account for this.

First, we measure consumption declines by comparing consumption when unemployed with in-work consumption two years prior. Hendren (2017) documents a statistically significant consumption decline of 2.7% in the year preceding unemployment; combining this with existing one-year horizon estimates reconciles the majority of the gap. Second, we use a broader measure of consumption expenditures. In Appendix D.2, we show that our results are robust to using a wide range of alternative consumption measures. Third, our sample differs from those used in prior work and is drawn from different survey waves. Using the PSID, East and Kuka (2015) document an increasing trend in the average decline in food expenditures following unemployment, which they attribute to factors other than changes in the survey design.[16]

Our finding of substantial heterogeneity in exposure to unemployment-induced consumption declines is consistent with a large body of evidence on how households respond to job-loss. Browning and Crossley (2001) find that the marginal propensity to consume out of UI benefits is largest for those without liquid savings. Using direct evidence on binding credit constraints, Crossley and Low (2014) find that 5% of job-losers face binding credit constraints and experience particularly large, welfare-reducing, consumption decline at job loss. They also document smaller, yet substantial, declines among households that are likely to be unconstrained. Ganong and Noel (2019) also document systematic variation in consumption responses by liquid assets, using de-identified banking data. Recent evidence from linked banking and administrative records in Denmark (Andersen, Jensen, Johannesen, et al. 2023) shows that households offset income losses from job-loss primarily by decumulating liquid savings, but also through added worker effects. Finally, both Patterson (2023) and Colarieti, Mei, and Stantcheva (2024) document *marginal* propensities to consume out of unemployment income losses of between 0.5 to 0.6. When workers face different income declines (e.g., married vs. single workers), even constant marginal propensities to consume imply heterogeneous consumption declines. Moreover, Patterson (2023) shows evidence of heterogeneity in these marginal effects by observables.

The key novelty in our analysis is that we construct worker-level estimates of unemployment-induced consumption declines without imposing assumptions about which observables drive this variation. In Appendix D.3, we examine how these estimated consumption declines correlate with proxies for a household's capacity to self-insure. We find larger consumption declines among households less likely to benefit from added worker effects (those who do not cohabit) and with low liquid assets.

[16]Similar to our analysis, East and Kuka (2015) find only modest differences across consumption measures. They attribute rising declines to a time trend rather than temporal sampling over the business cycle. In addition, our sample includes fewer married white workers compared to earlier work, as shown by comparing summary statistic we report in Table B.1 of Appendix B and Table 1 of Kroft and Notowidigdo (2016), who reproduce the estimate from Gruber (1997). We also explore the sensitivity of our results to alternative definitions of unemployment in the two-year panel and find our results are robust.

4.2. Employment Risk

Method. Patterns of program cross-subsidization depend on worker-level risk weights. To estimate these, we use a statistical model of unemployment that recovers rich demographic-specific employment probabilities. This approach infers how risk varies across individuals based on observed realizations of unemployment, sidestepping the need to specify a full model of workers' effort choices.

We specify a logistic regression,[17] where the probability individual i is unemployed at time t depends on the following observed characteristics: a cubic in age, indicators for gender marital status, a full set of education and race indicators, and a cubic in a measure of worker's permanent income.[18]

This yields an estimated employment probability, $\hat{e}_{i,t}$ for each individual and time period. While a non-parametric approach could in principle estimate $e(x) = Pr[U = 0|X]$ more flexibly, the moderate sample size of the PSID motivates our use of a parametric specification.

Estimates. We report the estimated coefficients in Table D.3 of Appendix D.4, and average partial effects with 95% confidence intervals in Figure 3. The largest effects are associated with our measure of permanent income—a one standard deviation increase corresponds to a 4 percentage point decrease in the probability of unemployment.

The results show systematic differences in unemployment rates across demographic groups. These differences are best interpreted as the effect of a given variable for two workers with similar wage histories, as permanent income is included as a control. Younger workers are more likely to be unemployed: a ten-year decrease in age is associated with a 1.4 percentage point increase in the probability of unemployment. We find a modest educational gradient: individuals who did not complete high school are approximately 1 percentage point more likely to be unemployed than those with higher educational attainment. This small effect is largely attributable to the direct role of permanent income. Even after controlling for education and permanent income, we find a substantial racial gap: Black individuals are 1.5 percentage points more likely to be unemployed than white

[17]An alternative approach uses idiosyncratic and subjective probabilities of job loss (Hendren 2017) and job finding (Spinnewijn 2015; Mueller, Spinnewijn, and Topa 2021). These measures are not directly available our dataset, and using them would require imputing probabilities based on similar demographic variables.

[18]In Appendix D.4, we show that results are nearly identical when we model employment risk directly as a function of estimated consumption declines. This suggests that other factors driving latent heterogeneity in consumption responses do not capture differences in employment risk. Including time effects improves overall fit but does not materially affect the relative risk ranking across demographic groups; in other words, relative employment risks are stable over time.

individuals. Conditional on other covariates, we find small effects for gender.

While the simple search model outlined in Section 2 rationalizes these differences through individuals' effort choices, it is important to emphasize that these disparities likely reflect broader factors. These include differences in market thickness, equilibrium congestion externalities, skill-biased labor demand, or outright discrimination—all of which are captured in our model through differences in the cost of exerting effort (i.e. in $\psi_i(.)$).

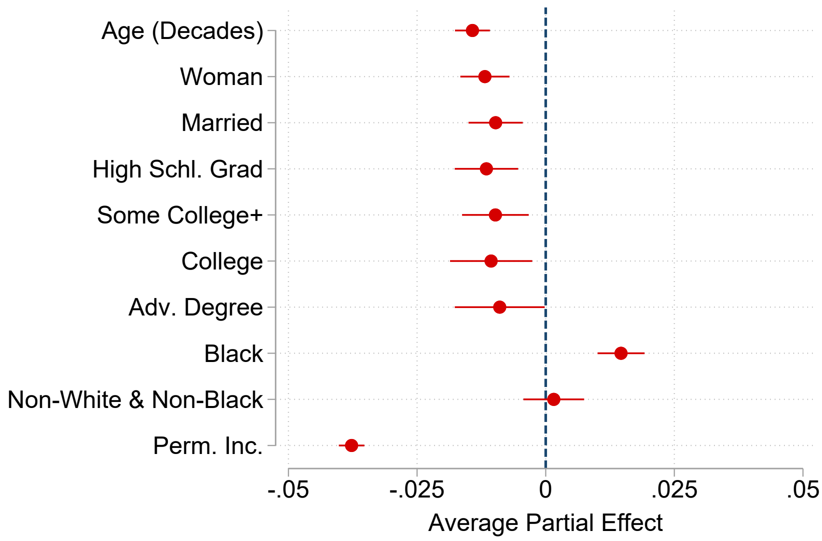

FIGURE 3. Effect of Household Characteristics on Predicted Probability of Unemployment

Notes: Authors' estimates from the PSID data. We estimate the logistic regression described in the text using employment status as the dependent variable. We calculate the marginal effects of each covariate on the unemployment probability using parameter estimates and integrate over the empirical distribution of household characteristics. The center of each circle corresponds to the point-estimate of the average partial effect (APE) and the horizontal lines span the 95% confidence intervals. The red dashed line indicates the origin. N=52,996 and standard errors are clustered at the household level. We report point-estimates for the APE and the underlying coefficients along with their standard errors in Table D.3 of Appendix D.4.

These patterns are consistent with findings documented in the broader literature. Prior research highlights a positive unemployment rate differential between young and old workers (e.g., Choi, Janiak, and Villena-Roldán 2015); a decline in unemployment rates with increased human capital (e.g., Ashenfelter and Ham 1979; Nickell 1979; Cairó and Cajner 2018); and persistent racial disparities in employment outcomes (e.g., Lang and Lehmann 2012) and callback rates (e.g., Bertrand and Mullainathan 2004). While we find a modest employment gap by gender, it is important to note that our estimates pertain to women who are the reference person in the PSID—a group that is disproportionately single. As prior work shows, participation gaps for unmarried men and women are relatively small, (e.g., Borella, De Nardi, and Yang 2023).

4.3. Risk Aversion and Reforms

Risk aversion. The insurance value of UI expansion to an individual worker is captured by the product of their consumption decline and coefficient of relative risk aversion. Based on evidence from a meta-survey of the intertemporal elasticity of substitution (Havránek 2015), we choose a baseline value for the coefficient of relative risk aversion, γ, of 3.

Reforms. We consider three types of expansion in UI generosity:

(a) Proportional-capped reform: $\phi_{\mathcal{B}}(y) = \min\{y, \kappa\}$ where κ is an earning cap

(b) Flat reform: $\phi_{\mathcal{B}}(y) = 1$

(c) Proportional reform: $\phi_{\mathcal{B}}(y) = y$

All reforms are funded by a flat tax increase, $\phi_{\mathcal{T}}(y) = 1$.

Reform (a) is designed to reflect the most natural expansion of the current U.S. UI system, which offers proportional UI benefits up to an earning cap—approximately at the 41st percentile of the earnings distribution—and is funded by a payroll tax that, in most states, functions as an employment tax. Reform (b) considers a flat benefit expansion for all workers, aligning with the standard reform considered in the UI literature. Reform (c) represents a proportional expansion of UI benefits with no cap, meaning benefits rise in proportion to income for all workers. The interaction between the shape of these reforms and the distribution of employment risk determine the extent of cross-subsidization inherent in each policy.

4.4. Distribution of Willingness-to-Pay

In Figure 4, we summarize the distribution of willingness-to-pay by reform type. Panel (A) shows the insurance component of WTP, which is common across reforms and always positive. Panels (B)-(D) show total WTP for each reform. Each panel plots how the median, interquartile range, and interdecile range of worker-level WTP vary across percentiles of the household income distribution.

Panel (A) shows that WTP for actuarially-fair insurance declines with household income and displays substantial variation conditional on income. Among households in the bottom income decile, the median WTP is \$0.38 per dollar of benefit expansion—over 20% higher than the median among households in the top decile. This difference is even more pronounced when comparing the upper tail of the WTP distribution within each group: the 90th percentile of WTP in the bottom income decile is \$0.68, 55% higher than that in the top income decile.

Panel (B) focuses on a proportional-capped reform, where WTP reflects both the value of insurance and cross-subsidization embedded in the reform. The wider vertical scale, relative to panel (A), highlights that much of the WTP variation by household income stems from cross-subsidization. Median WTP declines with income—from \$0.87 per dollar in the bottom income decile to -\$3.66 in the top. Even conditional on income, heterogeneity remains substantial; for example, among workers in the top income decile, the standard deviation in WTP is \$2.60.

Although all reforms share a common insurance component, they differ in their patterns of cross-subsidization and therefore in their resulting WTP distributions. Panel (C) shows that a flat reform tilts WTP in favor of workers in lower-income households: median WTP is \$1.10 per dollar in the bottom income decile falling to -\$6.40 in the top. Conversely, a proportional reform (panel (D)) tilts gains towards workers in higher-income households.

Under the flat reform, cross-subsidization arises from differences in employment risk: workers with below-average risk subsidize those with higher risk. As employment risk tends to be lower among higher-income households, WTP declines steeply with income. Under the proportional reform, cross-subsidization reflects both employment risk and worker earnings. Conditional on employment risk, higher earners gain more from the expansion, yet the overall decline in WTP with household income suggests that declining risk dominates the earnings effect. WTP for a proportional-capped reform resembles the proportional reform at the lower end of the income distribution and the flat reform at the upper end, reflecting the hybrid structure of the policy.

Individual workers' preference ordering over the three reforms depends on both their employment risk and earnings.[19] These reforms create distinct winners and losers and, as we discuss in the next section, differ in their associated fiscal externalities. Ranking their social value therefore requires aggregating willingness-to-pay and accounting for behavioral responses.

Sensitivity to preference specification. The insurance component of WTP depends on our calibrated value of the coefficient of relative risk-aversion and the assumption of state-independent utility. In Appendix E.2, we reproduce our WTP results under alternative assumptions.

Varying the coefficient of relative risk-aversion scales the insurance component

[19] Appendix E.1 documents the covariance of individual WTP across different reforms.

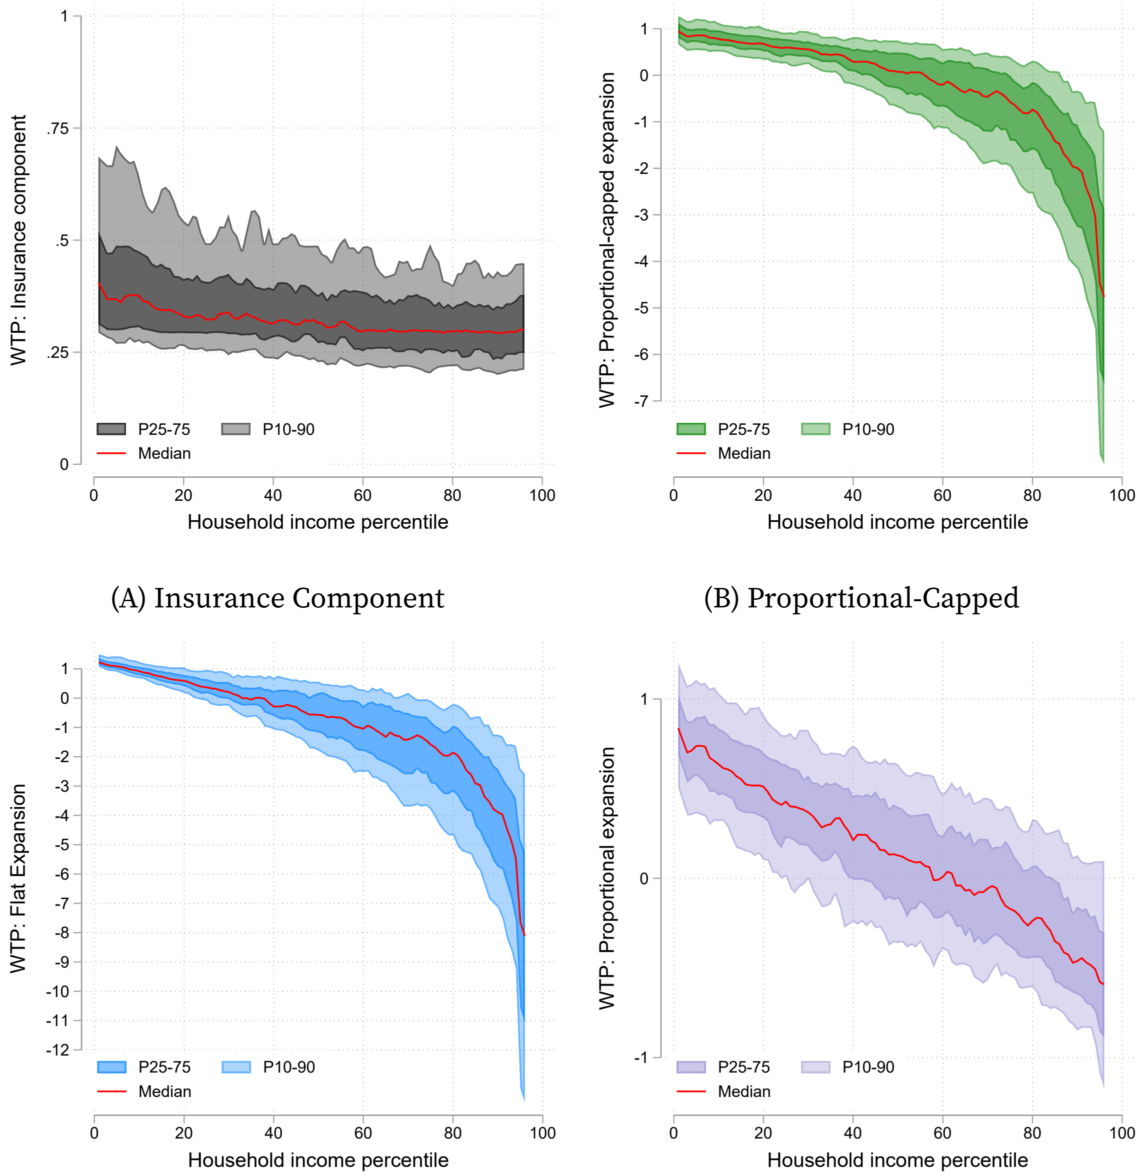

(A) Insurance Component

(B) Proportional-Capped

(C) Flat

(D) Proportional

FIGURE 4. Willingness-to-Pay for UI Expansion by Income and Reform Type
Notes: Authors' estimates from the PSID data. Each panel summarizes the distribution of workers willingness-to-pay for a UI expansion—reporting the median, as well as 10th, 25th, 75th, 90th percentiles. We compute worker-level WTP using equation (3), combining estimates of individual consumption declines and panel average employment risk, with worker incomes winsorized at \$10,000 and \$200,000. Panel (A) shows results for an actuarially fair reform at the worker level. Panel (B) shows results for a proportional-capped reform reflecting the current U.S. UI structure. Panels (C) and (D) show results for flat and proportional expansions, respectively. Quantiles are smoothed using a uniform rolling window of ± 1 percentiles.

proportionally but does not alter the shape of its distribution. We also allow for state-specific consumption prices, capturing the idea that unemployed individuals allocate more time to searching for lower prices. Assuming prices are 1.5% lower for the unemployed—based on evidence from Kaplan and Menzio (2015) and Campos and Reggio (2020)—modestly reduces the insurance component of WTP.

As emphasized by Andrews and Miller (2013), heterogeneity in risk aversion affects the distribution of insurance values to the extent it covaries with consumption declines. Unfortunately, direct evidence on such heterogeneity is limited.[20] However, as Figure 4 shows, any such effect will have much smaller impact on the distribution of WTP for UI reform than on the insurance competent itself. This is because cross-subsidization plays a dominant role in shaping WTP distributions under realistic reforms.

5. Social Value of Reform

In this section, we characterize the social value of expanding a publicly-funded insurance program. In the canonical setting, policy trades off the insurance benefit to a representative agent against the cost of distorting incentives, captured by a single elasticity. We show that with heterogeneous individuals, the single representative insurance value is replaced by a weighted sum of individual willingness-to-pay—which we decompose into the social value of risk protection and cross-subsidization based redistribution. Moreover, incentive costs depend on the covariance between individual-level behavioral responses, risk weights, and state-specific net tax liabilities.

5.1. Theory

The aggregate value of a social insurance reform depends both on its fiscal cost and on the weights assigned to the reform's impact on each individual.

To characterize these components, we define risk-weighted expectations in each state. For any variable x_i, we denote its risk-weighted expectation in the low and high state, respectively, as:

$$\begin{aligned}\mathbb{E}^l[x] &\equiv \int_i \left(\frac{(1-e_i)\phi_{\mathcal{B}}(y_i)}{\int_{i'}(1-e_{i'})\phi_{\mathcal{B}}(y_{i'})di'}\right) x_i\, di \\ \mathbb{E}^h[x] &\equiv \int_i \left(\frac{e_i\phi_{\mathcal{T}}(y_i)}{\int_{i'} e_{i'}\phi_{\mathcal{T}}(y_{i'})di'}\right) x_i\, di.\end{aligned} \tag{8}$$

These expectations assign risk-weights proportional to each individual's mechanical contribution to the fiscal impact of the reform: in the low state, this corresponds to

[20] The 1996 PSID survey wave includes survey-elicited measures of individual risk preferences (see Kimball, Sahm, and Shapiro (2009) for more details). These exhibit low covariance with our estimated consumption drops, and therefore imply little change in WTP distributions.

the expected benefit transfer, and in the high state, to the expected tax increase.

Budgetary impact. We focus on budget-neutral expansions of social insurance. This implies the policymaker faces the resource constraint $\int_i e_i \mathcal{T}(y_i)\, di = \int_i (1-e_i)\mathcal{B}(y_i)\, di + \bar{G}$, where $\bar{G}$ denotes an exogenous revenue requirement. The schedules $\mathcal{T}(.)$ and $\mathcal{B}(.)$ should be interpreted as net taxes levied in each state.

Budget-neutrality implies that the reform $d\theta = (db, d\tau)$, parameterized in equation (2), much satisfy:

$$\frac{d\tau}{db} = \underbrace{\frac{\int_i (1-e_i)\phi_{\mathcal{B}}(y_i)\, di}{\int_i e_i \phi_{\mathcal{T}}(y_i)\, di}}_{\text{Mechanical effect}} \left(1 + \underbrace{\frac{1}{\int_i (1-e_i)\phi_{\mathcal{B}}(y_i)\, di} \left(\int_i \frac{d(1-e_i)}{d\theta} (\mathcal{T}(y_i) + \mathcal{B}(y_i))\, di \right)}_{\text{Fiscal externality (FE)}} \right). \tag{9}$$

The first term captures the tax increase required to cover the mechanical revenue cost of the benefit expansion, holding behavior fixed. The second term reflects the additional tax adjustment needed to maintain budget balance due to behavioral responses—the fiscal externality due to moral hazard. This arises because higher benefits reduce the marginal return to exerting effort, thereby increasing the likelihood that individuals enter the low state.

Note that while the mechanical effect depends on the design of the reform—via the functions $\phi_{\mathcal{B}}(.)$ and $\phi_{\mathcal{T}}(.)$—the fiscal externality depends on the pre-reform structure of the tax and benefit schedules $\mathcal{T}(.)$ and $\mathcal{B}(.)$.

We can express the fiscal externality term as:

$$\text{FE} = \mathbb{E}^l \left[\epsilon^{(1-e,b)} \frac{\mathcal{T}(y) + \mathcal{B}(y)}{\mathcal{B}(y)} \right], \tag{10}$$

where $\epsilon_i^{1-e,b} \equiv \frac{1}{\phi_{\mathcal{B}}(y_i)} \frac{d(1-e_i)}{db} \frac{\mathcal{B}(y_i)}{1-e_i}$ is the individual-level elasticity of the low-state probability with respect to a benefit increase $\phi_{\mathcal{B}}(y)\, db$, financed by the balanced-budget tax adjustment. Equation (10) makes clear that, in general, the fiscal cost arising from moral hazard depends on the covariance between individual effort elasticities and both risk and tax and benefit liabilities.

Aggregate value. To capture the aggregate value of the reform, we define a weighted sum of expected utilities:

$$W(\theta) = \int_i \omega_i V_i(\theta) di, \tag{11}$$

where $\omega_i \geq 0$ is the Pareto weight assigned to individual i. As we discuss below, the inclusion of Pareto weights allows this function to accommodate a broad range of ethical positions regarding the redistributive value of social insurance.

We focus on the aggregate value of a marginal benefit adjustment, parameterized by equations (2) and (9). Analogous to our definition of individual-level willingness-to-pay, we scale the aggregate utility change $\frac{dW}{d\theta}$ to express it per $1 balanced-budget increase in the benefit bill relative to the aggregate utility value of a $1 decrease in the tax bill. This yields the condition:

$$\frac{dW^{MM}}{d\theta} = \underbrace{\frac{\mathbb{E}^l[\omega u^{l'}(c^l)] - \mathbb{E}^h[\omega u^{h'}(c^h)]}{\mathbb{E}^h[\omega u^{h'}(c^h)]}}_{\text{Resource reallocation (RR)}} - \text{FE} \tag{12}$$

The first term captures the gain of expanding social insurance that arises from shifting resources from high to low states of the world. It is measured by the percentage gap in the risk- and Pareto-weighted expected marginal utility of consumption between the low and high states. The overall aggregate value of the reform nets off the fiscal externality.

Insurance-redistribution decomposition. We define social marginal welfare weights as $g_i \equiv \omega_i u_i^{h'}(c_i^h)$. The following proposition decomposes the resource reallocation term in equation (12) into the societal value placed on insurance and redistributing resources across individuals.

PROPOSITION 1. *The aggregate value of resource reallocation from a marginal, budget-neutral rise in benefits satisfies:*

$$\text{RR} = \underbrace{\bar{\lambda}}_{\substack{\text{aggregate} \\ \text{welfare weight}}} \mathbb{E}^l \left[\underbrace{\lambda}_{\substack{\text{individual} \\ \text{welfare weight}}} \left(\underbrace{\frac{u_i^{l'}(c_i^l) - u_i^{h'}(c_i^h)}{u_i^{h'}(c_i^h)}}_{\text{insurance}} + \underbrace{1 - \frac{e_i \phi_{\mathcal{T}}(y_i)}{(1-e_i)\phi_{\mathcal{B}}(y_i)} \Big/ \frac{\int_{i'} e_{i'} \phi_{\mathcal{T}}(y_{i'})\, di'}{\int_{i'} (1-e_{i'})\phi_{\mathcal{B}}(y_{i'})\, di'}}_{\text{cross-subsidisation}} \right) \right] \tag{13}$$

where $\bar{\lambda} \equiv \frac{\mathbb{E}^l[g]}{\mathbb{E}^h[g]}$ is the ratio of risk-weighted social marginal welfare weights across the two states, and $\lambda_i \equiv \frac{g_i}{\mathbb{E}^l[g]}$ is the welfare weight of individual i, normalized such that the low-state risk-weighted average is 1. This can be equivalently expressed as:

$$\text{RR} = \underbrace{\bar{\lambda} \mathbb{E}^l \left[\lambda \left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} \right) \right]}_{\text{risk-protection}} + \underbrace{(\bar{\lambda} - 1)}_{\text{redistribution}} \tag{14}$$

Proof: See Appendix F.1.

Equation (13) expresses the welfare gain from social insurance expansion via resource reallocation as a weighted average of individuals' willingness-to-pay (or surplus), modulated by individual and aggregate social marginal welfare weights. The individual weight λ_i is scaled such that its low-state risk-weighted average is 1 (i.e., $\mathbb{E}^l[\lambda] = 1$). It reflects the social value of increasing individual i's surplus by \$1, relative to the average value among program beneficiaries. The aggregate weight $\bar{\lambda}$ captures the relative weight placed on individuals in the low state compared to those in the high state. It represents the average social value of providing the set of program beneficiaries with an additional \$1 of surplus, relative to the average value of doing so for program funders. This connects closely to the advantageous or adverse selection of beneficiaries relative to funders we document above. Importantly, the aggregate weight is reform-specific, through the expectation weightings, capturing how the benefit and tax changes vary with income.

Equation (14) expresses the welfare gain in terms of two components: (1) the social value of insurance, which we refer to as risk protection, and (2) the social value of the cross-subsidization embedded in the implicit pricing of social insurance, which we refer to as redistribution. These redistributive effects are captured by a single sufficient statistic—the reform-specific aggregate welfare weight $\bar{\lambda}$.[21]

These equations highlight two distinct channels through which social insurance reform affects aggregate welfare. First, reform alters the degree of risk protection provided to individuals. All else equal, the value of this protection is higher when there is a stronger positive correlation between individual insurance values and welfare weights. Moreover, its value increases with the relative weight society places on program beneficiaries compared to funders, as captured by the aggregate welfare weight $\bar{\lambda}$. Second, even in the absence of individual-level demand for insurance, the reform redistributes resources across individuals. If individuals in the low state have higher welfare weights than those in the high (i.e., $\bar{\lambda} > 0$), expanding the program delivers redistributive value. This effect arises through implicit cross-subsidization, as tax and benefit adjustments typically deviate from actuarially fair terms.

Importantly, because the risk-weighting in expectations depends on the design of the reform (via $\phi_{\mathcal{B}}$ and $\phi_{\mathcal{T}}$) both the aggregate insurance and redistributive effects

[21] Like our willingness-to-pay expression (equation (3)), the decompositions in equations (13) and (14) focus on marginal reforms. This allows us to apply the envelope theorem and ensures consistency with the analogous conditions in a rich dynamic model (see Appendix F.2). The decomposition is general in that it applies both to marginal reforms around the current system—our focus in this paper—and to reforms evaluated at the optimum. As such, it can be combined with a structural model to help characterize optimal policy.

are reform-specific.

5.2. Implementation

We maintain Assumption 1, which implies state-independent marginal utilities with common local curvature. In addition to individual-level insurance values and risk weights, we require two additional ingredients.

Welfare weights. The social marginal welfare weights, g_i, encode social preferences over interpersonal utility comparisons. As highlighted by Saez and Stantcheva (2016), these weights accommodate a wide range of normative views about how to compare welfare across individuals. In our baseline implementation, we use the following specification:

$$g_i = g_{y_i} \times \tilde{g}_i \tag{15}$$

where:

$$g_{y_i} \text{ is the inverse-optimum weight computed in Hendren (2020)}$$
$$\tilde{g}_i \propto (c_i^h)^{-\gamma} \quad \text{and} \quad \mathbb{E}[\tilde{g}|y] = 1$$

Under this specification, the social value of transferring \$1 of surplus between individuals at different income levels is determined by the inverse-optimum weights, g_y. These weights rationalize indifference to modifications of the prevailing income tax and transfer system and can be interpreted as the revealed social preference for redistribution across income groups.[22] Alternatively, they can be interpreted as the marginal cost to the government of providing \$1 of surplus to individuals of income y via an adjustment to the tax and transfer system. Thus, using these weights to value cross-income redistribution reflects the fact that equivalent transfers could, in principle, be implemented through the tax and transfer system.

Social insurance reform may also redistribute surplus among individuals with the same level of income. Our welfare weight specification values this within-income group surplus redistribution according to $\tilde{g}_i \propto (c_i^h)^{-\gamma}$.[23] This formulation implies that, conditional on their current income, the social value of providing an additional dollar of surplus to an individual declines with consumption. This is motivated by

[22]As in practice most married households file income taxes jointly, we map inverse optimum weights to household income.

[23]Welfare weights can also be expressed in terms of the Pareto weights ω_i and consumption marginal utilities $u_i^h(c_i^h)$. Under Assumption 1, we approximate $u_i^h(c_i^h) \approx (c_i^h)^{-\gamma}$, so setting $\omega_i = \frac{g_{y_i}}{\mathbb{E}[(c_i^h)^{-\gamma}|y_i]}$ yields the weights in equation (15).

empirical evidence that consumption is a reliable proxy for both lifetime income and economic well-being (e.g., Poterba 1989; Meyer and Sullivan 2003, respectively). We measure c_i^h using the panel average in-work consumption, equivalized by household size, to account for difference in needs due to household composition, and winsorize the measure at the 1*st* and 99*th* percentiles. While the income tax and transfer system redistributes across current income levels, social insurance reform may deliver additional redistributive value by directing resources to individuals with low consumption-levels within income groups.

Fiscal externality. The fiscal externality associated with social insurance reform (equation (10)) depends on the behavioral response elasticity, $\epsilon_i^{1-e,b}$, and state-contingent net taxes and benefits $(\mathcal{T}(y_i), \mathcal{B}(y_i))$.

We calibrate the elasticity based on evidence in Chetty (2008), who reports unemployment duration elasticities separately by quartiles of the wealth distribution.[24] The average elasticity reported in Chetty (2008) aligns closely with the median value reported in the broader literature (see Schmieder and Von Wachter 2016).

We compute state-contingent net taxes and benefits schedules using NBER TAXSIM, assuming married households file jointly (see Appendix G.1 for details). This means our measure of the incentive costs of reform take account of the fiscal impact of reduced worker search effort through both higher UI payments and the resulting decline income tax revenues.

Discussion. Our characterization of the welfare gains from resource reallocation parallels recent work on the social value of transfer programs. For example, Kolsrud, Landais, Reck, et al. (2024) examine the effects of shifting resources between an early and late retiree, showing that differences in their consumption levels and welfare weights captures the combined insurance and redistributive effects of reforms. A key difference in our setting is that we model individual-level heterogeneity across the full population of workers, so the social value of reform depends on the joint distribution of insurance values and welfare weights—along with employment risk, which is less important in the retirement context. As we show below, this is empirically relevant in our setting. In this respect, our approach also connects to recent work that uses consumption data from the PSID to evaluate

[24]The unemployment duration elasticity is the dynamic analogue of our static elasticity $\epsilon_i^{1-e,b}$ (see Appendix F.2). The estimates in Chetty (2008) range from -0.642 in the lowest wealth quartile to 0.016 for the highest (Table 2, pp 204). In our data, we assign each household the corresponding elasticity based on its wealth quartile. We adjust for the exhaustion of UI benefits after 6 months (see Appendix G.1 for details and Schmieder and Von Wachter (2016) for discussion).

the targeting value of disability insurance (Deshpande and Lockwood 2022) and welfare program take-up (Rafkin, Solomon, and Soltas 2023). In Appendix F.3 we relate our normative analysis to the marginal value of public funds framework.

5.3. Results

Baseline. Row (a) of Table 2 presents our baseline estimate of the social value of increasing UI generosity. These estimates are based on a budget-balanced expansion of the current U.S. UI program, which raises benefits proportionally to the minimum of a worker's earnings or an earnings cap, funded by a flat increase in payroll taxes. Columns (1) and (2) report the social value derived from risk protection and redistribution, respectively, while column (3) shows the total resource reallocation effect—the sum of these two components. We estimate that the risk protection value of a $1 budget-balanced expansion of the UI program is $0.43 and the social value of cross-subsidization is $0.10, yielding a total resource reallocation benefit of $0.54.

These gains must be weighed against the associated fiscal externalities. Column (4) reports the direct fiscal externality from increasing UI benefits, column (5) includes the additional fiscal externality due to reduced income tax receipts, and column (6) presents the total fiscal cost. We find that accounting for the broader tax and transfer system increases the fiscal externality from $0.38 (excluding these broader effects) to $0.71 (including them). Overall, the incentive costs associated with the expansion outweigh the resource reallocation gains.

Reform structure. Rows (b) and (c) examine two alternative reforms: a flat benefit expansion (independent of earnings) and an uncapped proportional expansion. Compared to the baseline, the flat expansion is more progressive, allocating a larger share of resources to individuals with higher welfare weights. This leads to increased social value from both risk protection and cross-subsidization based redistribution. Overall, the resource reallocation gain from the flat reform is $0.63.

In contrast, the proportional reform is more regressive than the baseline, directing a larger share of resources to workers with earnings above the cap. This reduces the value of risk protection and redistribution, resulting in a lower resource reallocation value of $0.40 for this reform.

The overall social value of the flat expansion is further enhanced by a reduction in the associated fiscal externality compared to the baseline. This divergence reflects two off-setting effects. On the one hand, the flat reform directs a larger share of the benefit-expansion gains toward lower earners, who tend to exhibit somewhat stronger behavioral responses. This results in a modest increase in the direct fiscal externality. However, this effect is more than offset by lower income tax liabilities associated with lower earners, which reduces the indirect fiscal externality. In contrast, the proportional reform lowers the direct fiscal externality relative to the baseline by shifting incentives to reduce effort towards higher earners who are less responsive. However, this is off-set by an increase in the indirect fiscal externality due to their larger tax contributions.

These differences across reform structures underscore the importance of incidence in determining the social value of a given policy change.

Social preferences. Rows (d)–(f) illustrate how the social value of the baseline reform varies depending on the weighting placed on individual willingness-to-pay. Money-metric weights ($g_i = 1$) represent the case where social preferences are indifferent to the distribution of surplus across individuals. One rationale for this approach is given by the Kaldor-Hicks Compensation Principle: if it is theoretically possible to combine the reform with individual-specific lump-sum taxes and transfers to achieve a Pareto improvement, the reform is considered desirable. Since this criterion eliminates any social value attached to redistributive gains, the resource reallocation value falls to \$0.37, attributable entirely to risk protection. A key drawback of this criterion is that it relies on a lump-sum redistribution scheme that is not practically implementable.

Row (e) reports results using inverse-optimum welfare weights, while assuming that social preferences are indifferent to the allocation of surplus conditional on income—that is, setting the within-income welfare weights $\tilde{g}_i = 1$. Hendren (2020) motivates inverse-optimum weights as an incentive-compatible extension of the Kaldor-Hicks Compensation Principle, which accounts for the distortionary costs of achieving cross-income redistribution through adjustments to the income tax system. Comparing these results to those obtained using money-metric weights highlights the value of cross-income group surplus redistribution achieved by UI expansion. Specifically, it increases the resource reallocation benefit of the reform by \$0.11, with over two-thirds of this gain attributable to the direct redistributive value arising from cross-subsidisation.

Comparing row (e) with the baseline results in row (a) reveals the additional social value of UI expansion that arises from within-income group surplus reallocation—specifically, between individuals with different levels of long-run in-work consumption. This within-income group effect further increases the social value of UI expansion arising from resource reallocation by \$0.06, split evenly between risk protection and the redistributive value of cross-subsidization. The fact that UI provides insurance against consumption risk—beyond the immediate shock of becoming unemployed—echoes findings from studies of disability insurance (Deshpande and Lockwood 2022) and self-selection into welfare programs among the eligible population (Rafkin, Solomon, and Soltas 2023).

Of the overall value of resource reallocation of \$0.54 in row (a), \$0.37 (69%) stems from efficiency gains, \$0.11 (20%) from cross-income surplus redistribution, and the remaining \$0.06 (11%) from within-income surplus redistribution.

Our use of inverse-optimum weights to value cross-income group redistribution can be motivated either by the extended Kaldor-Hicks Compensation Principle or by interpreting these weights as reflecting society's revealed preferences for redistribution. One potential objection to the first rationale is that, in practice, social insurance reforms are rarely accompanied by the complex income tax adjustments required to implement the compensation scheme. The second rationale may be challenged on normative grounds, as these weights may encode preferences for redistribution that are either too strong or too weak. In row (f), we present results based on utilitarian social welfare weights, defined as $g_i = (c_i^h)^{-\gamma}$. These weights reflect stronger preferences for cross-income redistribution than inverse-optimum weights, leading to an increase in both the value of risk protection and the redistributive benefit.

While the precise social value of surplus reallocation from UI reform ultimately depends on the policymaker's chosen welfare weighting scheme, this analysis demonstrates that its value will exceed that implied by a simple unweighted aggregation of money-metric surplus under broad conditions. Specifically, as long as the marginal surplus of (i) lower-income households is valued more than that of higher-income households, (ii) lower-consumption households is valued more than that of higher-consumption households, or (iii) some combination of both, the social value of the reform will be higher.

Alternative implementations.

Standard implementation. The penultimate row of Table 2 reports results for a flat reform evaluated using money-metric welfare weights, approximating the standard Baily-Chetty implementation. The corresponding social value of risk protection is estimated at 0.38. The difference relative to row (b)—which uses baseline welfare weights—reflects the distinction between valuing risk protection as a simple employment-risk-weighted average of individual insurance benefits and applying welfare weights that incorporate the social value of surplus redistribution.

Overall, the comparison between our approach and the standard implementation highlights the importance of incorporating individual-level heterogeneity—and its interaction with the structure of reforms—when evaluating both the welfare gains to workers and the incentive costs borne by the government budget.

Average implementation. The final row of Table 2 presents results from an alternative implementation of a flat reform that adapts the approach of Kolsrud, Landais, Reck, et al. (2024), originally applied to retirement benefit reform. This method evaluates equation (14) using *average* consumption among the employed and unemployed populations. It captures redistribution between these groups but abstracts from surplus dispersion within them. In this "average implementation," we also value the fiscal externality using average elasticity, tax liabilities and benefit entitlements of the unemployed. We provide further details in Appendix G.2.

This average implementation yields a resource reallocation gain of \$0.54—much closer to our estimate for the flat reform (\$0.63) than to the standard implementation (\$0.38). However, by abstracting from the positive covariance between employment risk and welfare weights within the unemployed population, it somewhat understates the social value of resource reallocation. It also overstates the indirect component of the fiscal externality by ignoring the correlation between employment probabilities, behavioral responses, and individuals' state-specific tax positions. The combined effect is a net social gain of -\$0.38 under the average implementation, compared with approximately zero under the full implementation.

Overall, the average implementation provides a substantially more accurate picture of the gains from resource reallocation than the standard implementation. However, in our context, it introduces some error, and, importantly for our purposes, is unable to compare the effects of different reforms that generate varying incidence across the unemployed population.

Cross-sectional implementation. The results in Table 2 exploit the panel dimension of the PSID. Panel data offer two main advantages over cross-sectional data. First, they enable us to estimate worker-level insurance values and WTP, which facilitates a decomposition of the social gains from expanding the generosity of UI into risk protection and redistribution based on cross-subsidization. Second, longitudinal information allow us to measure consumption and earnings using panel averages, helping to limit the influence of measurement error. However, if a researcher only has access to cross-sectional data, it is still possible to compute the social value of resource reallocation by directly implementing the social marginal utility gap in equation (12).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Resource reallocation			Fiscal externality			Total effect
	Risk protection	Redistribution	Total (1)+(2)	Direct	Indirect	Total (4)+(5)	(3)-(6)
(a) Baseline	0.43	0.10	0.54	0.38	0.33	0.71	-0.17
	[0.35, 0.53]	[0.09, 0.12]	[0.45, 0.64]	[0.38, 0.39]	[0.31, 0.34]	[0.70, 0.72]	[-0.26, -0.07]
Reform structure							
(b) Flat	0.47	0.16	0.63	0.40	0.19	0.59	0.04
	[0.37, 0.58]	[0.14, 0.18]	[0.53, 0.73]	[0.40, 0.40]	[0.17, 0.21]	[0.57, 0.61]	[-0.06, 0.15]
(c) Proportional	0.40	0.03	0.43	0.36	0.37	0.73	-0.30
	[0.32, 0.49]	[0.01, 0.06]	[0.34, 0.52]	[0.35, 0.37]	[0.35, 0.38]	[0.71, 0.74]	[-0.38, -0.21]
Social preferences							
(d) Money-metric	0.37	0	0.37	=(a)	=(a)	=(a)	-0.34
			[0.30, 0.44]				[-0.41, -0.27]
(e) Inverse-optimum	0.40	0.08	0.48	=(a)	=(a)	=(a)	-0.23
	[0.33, 0.48]	[0.07, 0.08]	[0.40, 0.55]				[-0.31, -0.15]
(f) Utilitarian	0.56	0.38	0.94	=(a)	=(a)	=(a)	0.23
	[0.44, 0.70]	[0.34, 0.42]	[0.82, 1.08]				[0.10, 0.38]
Alternative implementations							
(g) Standard implementation	0.38	-	0.38	0.40	-	0.40	-0.02
(flat reform+money-metric)			[0.31, 0.46]			[0.40, 0.40]	[-0.09, 0.06]
(h) Average implementation	0.43	0.12	0.54	0.40	0.40	0.80	-0.38
	[0.35, 0.52]	[0.11, 0.12]	[0.46, 0.63]	[0.40, 0.40]	[0.39, 0.42]	[0.79, 0.82]	[-0.46, -0.29]

TABLE 2. Risk Protection, Redistribution and Fiscal Externality from UI Expansion
Notes: Authors' estimates based on PSID data. Column (7) reports the total social welfare impact of a UI expansion, decomposed in columns (1)–(6). Row (a) corresponds to an expansion of the current U.S. system under baseline social preferences. Other rows present results under alternative reform structures or social preferences. All values are expressed in dollars per $1 increase in benefit expenditure. 95% confidence intervals, based on household-clustered bootstrap samples, are reported in square brackets.

6. Conclusion

In this paper, we study the positive and normative implications of heterogeneity in employment risk and income loss exposure for the design of unemployment insurance. We find substantial variation in workers' willingness-to-pay for UI, driven by two key factors: (i) differences in the value workers place on insurance, captured by our worker-level estimates of unemployment-induced consumption declines, and (ii) program cross-subsidization, arising from heterogeneity in employment risk and the incidence of reforms. The correlation between this heterogeneity and household wellbeing—reflected in social welfare weights—shapes the social value of expanding UI. We show that workers who experience unemployment shocks are negatively selected, both in terms of household income and, conditional on income, in terms of consumption while employed. As a result, UI expansion provides redistributive value, with surplus gains tilted toward less well-off households.

These social gains must be weighed against the distortionary costs of reform. We show that, after accounting for losses from reduced income tax payments, the net social value of a natural expansion of the current U.S. UI system is negative. In contrast, the net social value of a flat expansion—which delivers more resources to lower-income households—is approximately zero. Our analysis omits other fiscal externalities, such as the impact of unemployment on disability insurance take-up or future tax payments following re-employment. Quantifying these channels remains an important direction for future work.

Our approach relies on marginal analysis, which has the advantage of sidestepping the need to estimate a full structural model. While it provides insight into the merits of local reforms, it is not designed to characterize optimal policy. A promising avenue for future research is to combine our decomposition with a structural model to assess optimal UI reform.

References

Acemoglu and Shimer (1999). "Efficient unemployment insurance". *Journal of political Economy* 107.5, pp. 893–928.

Andersen, Jensen, Johannesen, Kreiner, Leth-Petersen, and Sheridan (2023). "How do households respond to job loss? lessons from multiple high-frequency datasets". *American Economic Journal: Applied Economics* 15.4, pp. 1–29.

Anderson and Meyer (2000). "The effects of the unemployment insurance payroll tax on wages, employment, claims and denials". *Journal of Public Economics* 78.1-2, pp. 81–106.

— (2006). "Unemployment insurance tax burdens and benefits: Funding family leave and reforming the payroll tax". *National Tax Journal* 59.1, pp. 77–95.

Andrews and Miller (2013). "Optimal Social Insurance with Heterogeneity". *mimeo* 78.1-2, p. 26.

Ashenfelter and Ham (1979). "Education, unemployment, and earnings". *Journal of Political Economy* 87.5, Part 2, S99–S116.

Attanasio, Levell, Low, and Sánchez-Marcos (2018). "Aggregating elasticities: intensive and extensive margins of women's labor supply". *Econometrica* 86.6, pp. 2049–2082.

Audoly (2024). “Self-Employment and Labor Market Risks”. *International Economic Review* 86.6, pp. 2049–2082.
Baily (1978). “Some Aspects of Optimal Unemployment Insurance”. *Journal of Public Economics* 10.6, p. 24.
Bertrand and Mullainathan (2004). “Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination”. *American Economic Review* 94.4, pp. 991–1013.
Blomqvist and Horn (1984). “Public Health Insurance and Optimal Income Taxation”. *Journal of Public Economics* 24.3, pp. 353–371.
Blundell, Pistaferri, and Preston (2008). “Consumption Inequality and Partial Insurance”. *American Economic Review* 98.5, pp. 1887–1921.
Boadway, Leite-Monteiro, Marchand, and Pestieau (2006). “Social Insurance and Redistribution with Moral Hazard and Adverse Selection”. *Scandinavian Journal of Economics* 108.2, pp. 279–298.
Boehm, Fize, and Jaravel (2025). “Five facts about MPCs: Evidence from a randomized experiment”. *American Economic Review* 115.1, pp. 1–42.
Borella, De Nardi, and Yang (2023). “Are marriage-related taxes and social security benefits holding back female labour supply?”. *The Review of Economic Studies* 90.1, pp. 102–131.
Browning and Crossley (2001). “Unemployment Insurance Benefit Levels and Consumption Changes”. *Journal of Public Economics* 80.1, pp. 1–23.
Cairó and Cajner (2018). “Human capital and unemployment dynamics: Why more educated workers enjoy greater employment stability”. *Economic Journal* 128.609, pp. 652–682.
Campos and Reggio (2020). “Do the Unemployed Pay Lower Prices? A Reassessment of the Value of Unemployment Insurance”. *Journal of the European Economic Association* 128.609, p. 47.
Chetty (2006). “A General Formula for the Optimal Level of Social Insurance”. *Journal of Public Economics* 90.10-11, pp. 1879–1901.
— (2008). “Moral Hazard versus Liquidity and Optimal Unemployment Insurance”. *Journal of Political Economy* 116.2, pp. 173–234.
Chetty and Finkelstein (2013). “Chapter 3 - Social Insurance: Connecting Theory to Data”. In: *Handbook of Public Economics*. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez. Vol. 5. Handbook of Public Economics, Vol. 5 10-11. Elsevier, pp. 111–193.
Choi, Janiak, and Villena-Roldán (2015). “Unemployment, participation and worker flows over the life-cycle”. *Economic Journal*. Handbook of Public Economics, Vol. 5 125.589. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1705–1733.
Colarieti, Mei, and Stantcheva (2024). *The how and why of household reactions to income shocks*. Tech. rep. 589. National Bureau of Economic Research, pp. 1705–1733.
Cremer and Pestieau (1996). “Redistributive Taxation and Social Insurance”. *International Tax and Public Finance*. Handbook of Public Economics, Vol. 5 3.3. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 281–295.
Crossley and Low (2014). “Job loss, credit constraints, and consumption growth”. *Review of Economics and Statistics*. Handbook of Public Economics, Vol. 5 96.5.

Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 876–884.
Deshpande and Lockwood (2022). “Beyond health: Nonhealth risk and the value of disability insurance”. *Econometrica*. Handbook of Public Economics, Vol. 5 90.4. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1781–1810.
East and Kuka (2015). “Reexamining the consumption smoothing benefits of Unemployment Insurance”. *Journal of Public Economics*. Handbook of Public Economics, Vol. 5 132.4. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 32–50.
Feldstein (2005). “Rethinking social insurance”. *American Economic Review*. Handbook of Public Economics, Vol. 5 95.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1–24.
Ferey (2022). “Redistribution and unemployment insurance”. *mimeo*. Handbook of Public Economics, Vol. 5 95.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1–24.
Finkelstein, Hendren, and Luttmer (2019). “The value of medicaid: Interpreting results from the oregon health insurance experiment”. *Journal of Political Economy*. Handbook of Public Economics, Vol. 5 127.6. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 2836–2874.
Ganong and Noel (2019). “Consumer Spending during Unemployment: Positive and Normative Implications”. *American Economic Review*. Handbook of Public Economics, Vol. 5 109.7. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 2383–2424.
Gruber (1997). “The Consumption Smoothing Benefits of Unemployment Insurance”. *American Economic Review*. Handbook of Public Economics, Vol. 5 87.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 192–205.
Guo (2024). “Payroll Tax Incidence: Evidence from Unemployment Insurance”. *Journal of Public Economics*. Handbook of Public Economics, Vol. 5 239.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, p. 105209.
Guo and Johnston (2021). “The Finance of Unemployment Compensation and Its Consequences”. *Public Finance Review*. Handbook of Public Economics, Vol. 5 49.3. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 392–434.
Haan and Prowse (2024). “The Heterogeneous Effects of Social Assistance and Unemployment Insurance: Evidence from a Life Cycle Model of Family Labor Supply and Savings”. *American Economic Journal: Macroeconomics*. Handbook of Public Economics, Vol. 5 16.2. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 127–181.
Havránek (2015). “Measuring intertemporal substitution: The importance of method choices and selective reporting”. *Journal of the European Economic Association*. Handbook of Public Economics, Vol. 5 13.6. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1180–1204.
Heckman and Singer (1984). “A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data”. *Econometrica*. Handbook of Public Economics, Vol. 5 52.2. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 271–320.

Hendren (2017). "Knowledge of Future Job Loss and Implications for Unemployment Insurance". *American Economic Review*. Handbook of Public Economics, Vol. 5 107.7. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1778–1823.
— (2020). "Measuring economic efficiency using inverse-optimum weights". *Journal of public Economics*. Handbook of Public Economics, Vol. 5 187.7. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, p. 104198.
Kaplan and Menzio (2015). "The morphology of price dispersion". *International Economic Review*. Handbook of Public Economics, Vol. 5 56.4. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1165–1206.
Keane and Wolpin (1997). "The career decisions of young men". *Journal of Political Economy*. Handbook of Public Economics, Vol. 5 105.3. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 473–522.
Kimball, Sahm, and Shapiro (2009). "Risk preferences in the PSID: Individual imputations and family covariation". *American Economic Review*. Handbook of Public Economics, Vol. 5 99.2. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 363–368.
Kleven (2021). "Sufficient statistics revisited". *Annual Review of Economics*. Handbook of Public Economics, Vol. 5 13.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 515–538.
Kolsrud, Landais, Nilsson, and Spinnewijn (2018). "The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden". *American Economic Review*. Handbook of Public Economics, Vol. 5 108.4-5. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 985–1033.
Kolsrud, Landais, Reck, and Spinnewijn (2024). "Retirement consumption and pension design". *American Economic Review*. Handbook of Public Economics, Vol. 5 114.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 89–133.
Kroft, Kucko, Lehmann, and Schmieder (2020). "Optimal income taxation with unemployment and wage responses: A sufficient statistics approach". *American Economic Journal: Economic Policy*. Handbook of Public Economics, Vol. 5 12.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 254–292.
Kroft and Notowidigdo (2016). "Should Unemployment Insurance Vary with the Unemployment Rate? Theory and Evidence". *Review of Economic Studies*. Handbook of Public Economics, Vol. 5 83.3. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1092–1124.
Krusell, Mukoyama, and Şahin (2010). "Labour-market matching with precautionary savings and aggregate fluctuations". *Review of Economic Studies*. Handbook of Public Economics, Vol. 5 77.4. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1477–1507.
Landais and Spinnewijn (2021). "The value of unemployment insurance". *The Review of Economic Studies*. Handbook of Public Economics, Vol. 5 88.6. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 3041–3085.
Lang and Lehmann (2012). "Racial discrimination in the labor market: Theory and empirics". *Journal of Economic Literature*. Handbook of Public Economics, Vol. 5

50.4. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 959–1006.
Lentz and Tranaes (2005). "Job Search and Savings: Wealth Effects and Duration Dependence". *Journal of Labor Economics*. Handbook of Public Economics, Vol. 5 23.3. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 959–1006.
Lewis, Melcangi, and Pilossoph (2024). *Latent heterogeneity in the marginal propensity to consume*. Tech. rep. 3. National Bureau of Economic Research, pp. 959–1006.
Lieber and Lockwood (2019). "Targeting with In-Kind Transfers: Evidence from Medicaid Home Care". *American Economic Review*. Handbook of Public Economics, Vol. 5 109.4. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1461–1485.
Meyer and Sullivan (2003). "Measuring the Well-Being of the Poor Using Income and Consumption". *Journal of Human Resources*. Handbook of Public Economics, Vol. 5 109.4. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 1180–1220.
Mueller, Spinnewijn, and Topa (2021). "Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias". *American Economic Review*. Handbook of Public Economics, Vol. 5 111.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 324–363.
Nickell (1979). "Education and lifetime patterns of unemployment". *Journal of Political Economy*. Handbook of Public Economics, Vol. 5 87.5, Part 2. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, S117–S131.
Patterson (2023). "The matching multiplier and the amplification of recessions". *American Economic Review*. Handbook of Public Economics, Vol. 5 113.4. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 982–1012.
Poterba (1989). "Lifetime Incidence and the Distributional Burden of Excise Taxes". *American Economic Review*. Handbook of Public Economics, Vol. 5 79.2. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 325–330.
Quandt (1972). "A new approach to estimating switching regressions". *Journal of American Statistical Association*. Handbook of Public Economics, Vol. 5 67.338. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 306–310.
Rafkin, Solomon, and Soltas (2023). "Self-targeting in US transfer programs". *Available at SSRN 4495537*. Handbook of Public Economics, Vol. 5 67.338. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 306–310.
Rochet (1991). "Incentives, Redistribution and Social Insurance". *Geneva Papers on Risk and Insurance Theory*. Handbook of Public Economics, Vol. 5 16.2. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 143–165.
Saez and Stantcheva (2016). "Generalized Social Marginal Welfare Weights for Optimal Tax Theory". *American Economic Review*. Handbook of Public Economics, Vol. 5 106.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 24–45.
Schmieder, Von Wachter, and Bender (2012). "The Effects of Extended Unemployment Insurance Over the Business Cycle: Evidence from Regression Discontinuity Estimates Over 20 Years". *Quarterly Journal of Economics*.

Handbook of Public Economics, Vol. 5 127.2. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 701–752.

Schmieder and Von Wachter (2016). “The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation”. *Annual Review of Economics*. Handbook of Public Economics, Vol. 5 8.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 547–581.

Shimer and Werning (2008). “Reservation Wages and Unemployment Insurance”. *Quarterly Journal of Economics*. Handbook of Public Economics, Vol. 5 112.3. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 547–581.

Spaziani (2024). “Optimal Unemployment Insurance Financing: Theory and Evidence from Two US States”. *mimeo*. Handbook of Public Economics, Vol. 5 112.3. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 547–581.

Spinnewijn (2015). “Unemployed but optimistic: Optimal insurance design with biased beliefs”. *Journal of the European Economic Association*. Handbook of Public Economics, Vol. 5 13.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 130–167.

— (2020). “The trade-off between insurance and incentives in differentiated unemployment policies”. *Fiscal Studies*. Handbook of Public Economics, Vol. 5 41.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 101–127.

Von Wachter (2019). “Unemployment insurance reform”. *The ANNALS of American Academy of Political and Social Science*. Handbook of Public Economics, Vol. 5 686.1. Ed. by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, pp. 121–146.

Online Appendix

A. Willingness-to-Pay: Theory

A.1. General model

Here we extend the static model from Section 2 to a richer dynamic setting, following the framework in Chetty (2006) and derive an expression for willingness-to-pay for social insurance reform. As in equation (3), WTP decomposes into components reflecting insurance and cross-subsidization. The difference is that, in the dynamic model, these components are evaluated as expectations over the agent's lifetime.

We consider a setting in which time is continuous and agents live $t \in [0, 1]$. Let $\varphi_{i,t}$ denote a state variable containing all relevant information up to time t in determining the agent's time t state status (i.e., whether they are in the high or low state) and behavior. $\varphi_{i,t}$ has unconditional distribution $F_{i,t}(\varphi_{i,t})$ at $t = 0$. Assume $F_{i,t}$ is smooth with maximal support φ for all (i, t).

Let $c_{i,t}(\varphi_{i,t})$ denote agent i's time t state-contingent consumption. Let $x_{i,t}(\varphi_{i,t})$ denote M other choices the agent makes (for instance, different dimensions of effort, actions to self-insure like borrowing from family, spousal labor supply decisions and so on). Denote the agent's flow utility function $u_i^s(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))$ for $s \in \{l, h\}$. Let $\xi_{i,t}(\varphi_{i,t})$ denote the state the agent is in at time t and given state variable $\varphi_{i,t}$. If $\xi = 1$ if the agent is in the high state, if $\xi = 0$ the agent is the low state.

Denote the full program of agent i's state-contingent choices:

$$c_i = \{c_{i,t}(\varphi_{i,t})\}_{t\in[0,1],\, \varphi_{i,t}\in\varphi},$$
$$x_i = \{x_{i,t}(\varphi_{i,t})\}_{t\in[0,1],\, \varphi_{i,t}\in\varphi}.$$

When in the high state the agent earns $y_i - \mathcal{T}(y_i)$ and when in the low state they receive benefits $\mathcal{B}(y_i)$. The agent can also earn additional income $f_{i,t}(x_{i,t}(\varphi_{i,t}))$. They face the flow budget constraint:

$$\dot{A}_{i,t}(\varphi_{i,t}) = \xi_{i,t}(\varphi_{i,t})\left(y_i - \mathcal{T}(y_i)\right) + (1 - \xi_{i,t}(\varphi_{i,t}))\mathcal{B}(y_i) + f_{i,t}(x_{i,t}(\varphi_{i,t})) - c_{i,t}(\varphi_{i,t})$$

with terminal condition: $A_{i1}(\varphi_{i,1}) > \bar{A}_i$ for all $\varphi_{i,1}$. They also face N additional constraints in each state $\varphi_{i,t}$ at each time t:

$$g_{it}^n(c_{it}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t}) \geq 0$$

The agent's problem is to choose the program (c_i, x_i) to solve:

$$\max \int_t \int_{\varphi_{i,t}} \left(\xi_{i,t}(\varphi_{i,t})u_i^h(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t})) + (1 - \xi_{i,t}(\varphi_{i,t}))u_i^l(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))\right) dF_{i,t}(\varphi_{i,t})dt$$

$$+\int_t\int_{\varphi_{i,t}} \lambda^A_{it}(\varphi_{i,t})\left(\xi_{i,t}(\varphi_{i,t})\left(y_i-\mathcal{T}(y_i)\right)+(1-\xi_{i,t}(\varphi_{i,t}))\mathcal{B}(y_i)+f_{i,t}(x_{i,t}(\varphi_{i,t}))-c_{i,t}(\varphi_{i,t})\right)dF_{i,t}(\varphi_{i,t})dt$$
$$+\int_{\varphi_{i,1}} \lambda^A_{i1}(\varphi_{i,1})(A_{i1}(\varphi_{i1})-\bar{A}_i)dF_{i,1}(\varphi_{i,1})$$
$$+\sum_{n=1}^{N}\int_t\int_{\varphi_{i,t}} \lambda^n_{it}(\varphi_{i,t})g^n_{it}(c_{it}(\varphi_{i,t}),x(\varphi_{i,t});\varphi_{i,t})dF_{i,t}(\varphi_{i,t})dt$$

Denote the maximum function of this problem $V_i(\theta)$, written as a function of the policy parameters θ, which parameterizes modifications to the tax and benefit schedules. We assume the following regularity conditions:

ASSUMPTION 2 (Regularity Conditions). *Assume*

i. Total lifetime utility is smooth, increasing and strictly quasiconvex in (c, x)

ii. The choices (c, x) *that satisfy the constraints are convex*

iii. $V_i(\theta)$ *is differentiable*

where (i) and (ii) ensures the agent's problem has a unique solution and (iii) ensures the envelope theorem applies.

Consider the reform, $d\theta$, to the benefit and tax schedules parameterized by:

$$\begin{aligned}\frac{d\mathcal{B}(y_i)}{d\theta} &= \phi_{\mathcal{B}}(y_i)\\ \frac{d\mathcal{T}(y_i)}{d\theta} &= \phi_{\mathcal{T}}(y_i)\times\frac{\int_{i'}\int_t\int_{\varphi_{i',t}}\phi_{\mathcal{B}}(y_{i'})\left(1-\xi_{i',t}(\varphi_{i',t})\right)dF_{i',t}(\varphi_{i',t})dtdi'}{\int_{i'}\int_t\int_{\varphi_{i',t}}\phi_{\mathcal{T}}(y_{i'})\xi_{i',t}(\varphi_{i',t})dF_{i',t}(\varphi_{i,t})dtdi'}\end{aligned}$$

This reform is budget-neutral in the absence of any behavioral responses.

Assume that the constraints $g^n_{i,t}$ satisfy the regularity conditions:

$$\begin{aligned}\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}),x(\varphi_{i,t});\varphi_{i,t})}{\partial\mathcal{B}(y_i)} &= -(1-\xi_{i,t}(\varphi_{i,t}))\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}),x(\varphi_{i,t});\varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}\\ \frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}),x(\varphi_{i,t});\varphi_{i,t})}{\partial\mathcal{T}(y_i)} &= \xi_{i,t}(\varphi_{i,t})\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}),x(\varphi_{i,t});\varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}\\ \frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}),x(\varphi_{i,t});\varphi_{i,t})}{\partial c_{i,s}(\varphi_{i,s})} &= 0 \quad \text{if } t\neq s\end{aligned}$$

for $n = 1, \ldots, N$. See Chetty (2006) for a demonstration of the mildness of these conditions and an example of when they do not hold.

The impact of the reform $d\theta$ on individual i's expected utility is:

$$\frac{dV_i}{d\theta} = -\int_t\int_{\varphi_{i,t}}\left(\xi_{i,t}(\varphi_{i,t})\lambda^A_{i,t}(\varphi_{i,t})-\sum_{n=1}^{N}\lambda^n_{i,t}(\varphi_{i,t})\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}),x(\varphi_{i,t});\varphi_{i,t})}{\partial\mathcal{T}(y_i)}\right)dF_{i,t}(\varphi_{i,t})dt\frac{d\mathcal{T}(y_i)}{d\theta}$$
$$+\int_t\int_{\varphi_{i,t}}\left((1-\xi_{i,t}(\varphi_{i,t}))\lambda^A_{i,t}(\varphi_{i,t})+\sum_{n=1}^{N}\lambda^n_{i,t}(\varphi_{i,t})\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}),x(\varphi_{i,t});\varphi_{i,t})}{\partial\mathcal{B}(y_i)}\right)dF_{i,t}(\varphi_{i,t})dt\frac{d\mathcal{B}(y_i)}{d\theta}$$

Under the third assumption on the constraints, agent i's optimal consumption choice satisfies:

$$\frac{\partial u_i^s(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} = \lambda_{i,t}^A(\varphi_{i,t}) - \sum_{n=1}^{N} \lambda_{i,t}^n(\varphi_{i,t}) \frac{\partial g_{i,t}^n(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}$$

for all t and $\varphi_{i,t}$. The first two assumptions imply for all t and $\varphi_{i,t}$:

$$\sum_{n=1}^{N} \lambda_{i,t}^n(\varphi_{i,t}) \frac{\partial g_{i,t}^n(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial \mathcal{T}(y_i)} = \sum_{n=1}^{N} \lambda_{i,t}^n(\varphi_{i,t}) \xi_{i,t}(\varphi_{i,t}) \frac{\partial g_{i,t}^n(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}$$
$$\sum_{n=1}^{N} \lambda_{i,t}^n(\varphi_{i,t}) \frac{\partial g_{i,t}^n(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial \mathcal{B}(y_i)} = -\sum_{n=1}^{N} \lambda_{i,t}^n(\varphi_{i,t})(1 - \xi_{i,t}(\varphi_{i,t})) \frac{\partial g_{i,t}^n(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}$$

Hence, we can re-write $dV_i/d\theta$:

$$\begin{aligned} \frac{dV_i}{d\theta} = &-\int_t \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i) \xi_{i,t}(\varphi_{i,t}) \frac{\partial u_i^h(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t}) dt \\ &\times \frac{\int_{i'} \int_t \int_{\varphi_{i',t}} \phi_{\mathcal{B}}(y_{i'}) \left(1 - \xi_{i',t}(\varphi_{i',t})\right) dF_{i',t}(\varphi_{i',t}) dt di'}{\int_{i'} \int_t \int_{\varphi_{i',t}} \phi_{\mathcal{T}}(y_{i',t}) \xi_{i',t}(\varphi_{i',t}) dF_{i',t}(\varphi_{i,t}) dt di'} \\ &+ \int_t \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)(1 - \xi_{i,t}(\varphi_{i,t})) \frac{\partial u_i^l(c_{i,t}(\varphi_{i,t})), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t}) dt \end{aligned}$$

We define willingness-to-pay, scaling $\frac{dV_i}{d\theta}$ as follows:

$$WTP_i \equiv \frac{dV_i}{d\theta} \times \frac{\frac{1}{\int_t \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)(1 - \xi_{i,t}(\varphi_{i,t})) dF_{i,t}(\varphi_{i,t}) dt}}{\frac{\int_t \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i) \xi_{i,t}(\varphi_{i,t}) \frac{\partial u_i^h(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t}) dt}{\int_t \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i) \xi_{i,t}(\varphi_{i,t}) dF_{i,t}(\varphi_{i,t}) dt}}$$

The first factor expresses the utility gains per \$1 increase in lifetime benefit payments and the second factor converts to money-metric terms by comparing to the utility value of an unfunded tax cut (also expressed per \$1 worth of lifetime tax reduction).

Define the individual's state-specific average marginal utility of consumption:

$$\widetilde{\mathbb{E}}_i^h \left[\frac{\partial u^h(c, x)}{\partial c} \right] = \int_t \int_{\varphi_{i,t}} \frac{\phi_{\mathcal{T}}(y_i) \xi_{i,t}(\varphi_{i,t})}{\int_{t'} \int_{\varphi_{i,t'}} \phi_{\mathcal{T}}(y_i) \xi_{i,t'}(\varphi_{i,t'}) dF_{i,t'}(\varphi_{i,t'}) dt'} \frac{\partial u_i^h(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t}) dt$$
$$\widetilde{\mathbb{E}}_i^l \left[\frac{\partial u^l(c, x)}{\partial c} \right] = \int_t \int_{\varphi_{i,t}} \frac{\phi_{\mathcal{B}}(y_i)(1 - \xi_{i,t}(\varphi_{i,t}))}{\int_{t'} \int_{\varphi_{i,t'}} \phi_{\mathcal{B}}(y_i)(1 - \xi_{i,t'}(\varphi_{i,t'})) dF_{i,t'}(\varphi_{i,t'}) dt'} \frac{\partial u_i^l(c_{i,t}(\varphi_{i,t})), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t}) dt$$

We can then express willingness-to-pay for the reform:

$$WTP_i = \frac{\widetilde{\mathbb{E}}_i^l\left[\frac{\partial u^l(c,x)}{\partial c}\right] - \widetilde{\mathbb{E}}_i^h\left[\frac{\partial u^h(c,x)}{\partial c}\right]}{\widetilde{\mathbb{E}}_i^h\left[\frac{\partial u^h(c,x)}{\partial c}\right]} +$$
$$= 1 - \frac{\int_t \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t})dF_{i,t}(\varphi_{i,t})dt}{\int_t \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)(1-\xi_{i,t}(\varphi_{i,t}))dF_{i,t}(\varphi_{i,t})dt} \Bigg/ \frac{\int_{i'} \int_t \int_{\varphi_{i',t}} \phi_{\mathcal{T}}(y_{i'})\xi_{i',t}(\varphi_{i',t})dF_{i',t}(\varphi_{i,t})dtdi'}{\int_{i'} \int_t \int_{\varphi_{i',t}} \phi_{\mathcal{B}}(y_{i'})\left(1-\xi_{i',t}(\varphi_{i',t})\right)dF_{i',t}(\varphi_{i',t})dtdi'},$$

where the first expression on the right-hand side captures the lifetime value of insurance and the second expression captures the lifetime value of program cross-subsidization.

A.2. State-dependent utility

Consider the simple framework outlined in Section 2. Suppose the worker has state-dependent utility, which we capture through the marginal utility shifter ρ, defined such that $u_i^{h'}(c) = u_i'(c)$ and $u_i^{l'}(c) = \rho u_i'(c)$.

In this case, the individual's WTP is:

$$\mathrm{WTP}_i = \left(\frac{\rho u_i'(c_i^l) - u_i'(c_i^h)}{u_i'(c_i^h)}\right) + \left(1 - \frac{e_i \phi_{\mathcal{T}}(y_i)}{(1-e_i)\phi_{\mathcal{B}}(y_i)} \Bigg/ \frac{\int_{i'} e_{i'}\phi_{\mathcal{T}}(y_{i'})di'}{\int_{i'}(1-e_{i'})\phi_{\mathcal{B}}(y_{i'})di'}\right). \quad \text{(A.1)}$$

The state-dependence term enters the insurance component of WTP, by scaling marginal utility in the low state. In this case, the consumption-based implementation uses the approximation:

$$\frac{\rho u_i'(c_i^l) - u_i'(c_i^h)}{u_i'(c_i^h)} \approx (\rho - 1) + \rho \frac{\Delta c_i}{c_i^h}$$

The state-dependence term captures factors that cause the same level of consumption expenditure to yield different marginal utilities across states. For example, leisure time may vary by state and act as a complement to or substitute for consumption. There may also be state-specific expenditures (e.g., work-related costs) that do not directly generate utility (Browning and Crossley 2001). Additionally, the extent to which individuals combine consumption expenditures with home production may differ across states (Aguiar and Hurst 2005).

The state-dependence parameter can also capture situations in which individuals face state-specific consumption prices (see Campos and Reggio 2020). Suppose the consumption price in the high state is ρ times the price in the low state, and utility is otherwise state independent. In this case, WTP takes the same form as in equation (A.1). The consumption-based implementation also leads to the same state-dependence-adjusted formula. However, it is necessary to correct observed expenditure changes to account for the state-specific price difference. Let x_i^s denote state-specific expenditures, which are observed in the data. These are related to consumption through:

$$\frac{\Delta c_i}{c_i^h} = \frac{x_i^h/\rho - x_i^l}{x_i^h/\rho} = \rho\left(\frac{\Delta x_i}{x_i^h}\right) - (\rho - 1)$$

B. Description of PSID Data

We use the Panel Study of Income Dynamics to estimate the sufficient statistics that allow us to quantify the welfare effects of unemployment insurance expansion. We use data from 1999-2019, the biennial so-called "new" PSID, which includes high quality information on consumption expenditures and asset holdings. For our baseline sample we focus on non-immigrant households.[25] We include both single and married households. We do not select the sample on the basis of the reference person's gender.[26] Our sample has non-missing information on key demographics (age, education, and state of residence). We focus only on a sample of those where the reference person is aged between 23 and 60 and is participating in the labor force. Thus, we exclude observations where the reference person is retired, permanently disabled and neither working or looking for work, on sick (or maternity) leave or temporarily laid off, as well as those in education, who are a homemaker or in prison. In addition, to proxy for UI program eligibility we exclude those who never report employment, as well as periods immediately following a period out of the labor force. We include in our sample periods when the reference person is either in employment or unemployment, and we do not condition our sample on the employment status of their spouse. This is because we are explicitly interested in the employment risk facing households, including any self-insurance that may occur through spousal labor supply (Blundell, Pistaferri, and Saporta-Eksten 2016).

To reduce the influence of measurement error, we also drop observations with extremely high asset values, as well as observations that exhibit large fluctuations in key outcomes of interest. Following, Blundell, Pistaferri, and Saporta-Eksten (2016) we exclude households with net worth greater than $20 million to limit the impact of outliers and drop observations reporting extreme jumps (the bottom and top 0.25^{th} percentile) in wages, income and consumption to limit the role of measurement error. We do not use data displaying extreme "jumps" from one period to the next as we view this as most likely due to measurement error. A "jump" is defined as an extremely positive (negative) change from $t-2$ to t, followed by an extreme negative (positive) change from t to $t+2$. Formally, for each variable x, we construct the biennial log difference $\Delta_2 \ln(x_t)$, and drop the relevant variables for observation in the bottom 0.25 percent of the product $\Delta_2 \ln(x_t)\Delta_2 \ln(x_{t-2})$. Furthermore, in our analysis of permanent income we do not use earnings and wage data when the implied hourly wage is below one-half

[25] This is an additional sample frame beginning from the late 1990s. We do, however, include the Survey of Economic Opportunity households which oversamples low income families. Results using survey weights are similar.

[26] Previously, the PSID referred to the reference person as the household head.

of the state minimum wage. Table B.1 provides summary statistics for our sample of interest.

To measure education we create five categories based on completed years of education: those completing less than high school, those who completed high school, those who complete high school and some college (including those who dropout of four year degrees, and those who attain a community college degree or other diploma), those who complete four years of college, and, finally, those who complete further study. Our measure of race is derived from the self-reported race of the household reference person, from which we define three race categories (White, Black, and other).

The majority of our empirical analysis focuses on the relationship between employment risk and both levels and growth rates in consumption. The PSID measures disaggregated consumption across a number of different categories of household expenditure which are designed to cover approximately 70% of aggregate consumption. Respondents can also indicate the period the expenditure covers. We convert all expenditures to the annual level (e.g., we multiply weekly expenditures by 52 and monthly expenditures by 12) and treat missing values in the consumption subcategories as zeros. We focus on expenditure categories that are measured consistently across survey waves. Our primary consumption measure of interest captures a range of utility-relevant expenditures comprising non-durable purchases plus service. We show, however, robustness of our empirical results to a variety of alternative consumption measures (including food consumption as in Gruber (1997) and Hendren (2017)) which we describe here. Consumption in the PSID is measured at the household level, to account for differences in household size we equivalize using the square root of the household size.

Baseline: Non-Durable Expenditure including Services. Our baseline measure includes a broad range of non-durable expenditures and services, as long as those services do not include an investment or durable component (for example, vehicle maintenance). To build our baseline consumption series we first construct a food expenditure series by summing food purchased to be consumed at home, food purchased to be consumed away from home, and those food purchases covered by the Supplemental Nutrition Assistance Program. Inclusion of food expenditures covered by the Supplemental Nutrition Assistance Program (formerly known as the Food Stamp Program and colloquially known as food stamps) is important because they are a relevant source of financing expenditures for low-income households, while Low and Pistaferri (2015) show that food stamps can act as substitutes for social insurance.

We then construct a household expenditure series for services without a durable component. We sum spending on home and auto-mobile insurance, utilities, parking costs and other direct transportation costs (such as bus fare and payments to taxis) that do not correspond to maintenance for a vehicle, as well as child care costs.

Finally, we combine the aggregated food expenditure series with household spending on gasoline expenses and the household expenditure series on services without a durable component.

Food Expenditure. To construct a series for food expenditure, which excludes the other components of our baseline measure, we sum food purchased to be consumed at home, food purchased to be consumed away from home, and those food purchases covered by food the Supplemental Nutrition Assistance Program.

Services Expenditure. To construct a series for broad services, inclusive of those that may relate to durables or have an investment component, we combine the services measured in our baseline expenditure with a set of additional spending categories. These additional categories include healthcare related spending (out-of-pocket payments including for hospital and nursing home stays, doctor visits, and prescription drugs as well as insurance premia), vehicle repairs, and payments for educational services or schooling costs such as school tuition.

Total (Non-Housing) Expenditure. We combine our baseline measure with the additional services described in the preceding paragraph (Services Expenditure) to produce a household-level series for total non-housing expenditures. As the PSID consumption categories are not designed to have full coverage, the name total expenditure is a misnomer. Furthermore, we continue to exclude the purchase of durables, such as vehicles, and memory goods (Hai, Krueger, and Postlewaite 2020), such as vacations, that have durable-like properties.

Total (Including Housing) Expenditure. Our final consumption measure incorporates a measure of housing services. We sum the total non-housing expenditures with the consumption value of housing services. For renters, we use reported rental expenditures. For homeowners, we approximate the rental equivalent flow of housing services as a 6 percent yield on the house price (Poterba and Sinai 2008).

We report summary statistics for our sample in Table B.1.

B.1. Constructing a Measure of Permanent Income

We include a measure of worker's permanent income in our statistical model for employment risk. Here we describe how we construct that measure.

We define permanent income as the predicted time-invariant individual component from the following log wage regression

$$\ln w_{i,t} = \beta_0 + \beta_1 age_{i,t} + beta_2 age_{i,t}^2 + \beta_3 age_{i,t}^3 + \mu_i + \eta_t + u_{i,t},$$

where the dependent variable is the annual employment earnings of the reference person. We control for a third-order polynomial in age, and year fixed effects. We estimate this regression using only observations in which the worker is employed.

	All	Employed	Unemployed
Ref. Person's Age (years)	39.57	39.69	37.56
	(10.31)	(10.29)	(10.31)
Married	0.53	0.55	0.30
	(0.50)	(0.50)	(0.46)
Number of Children	1.03	1.02	1.15
	(1.21)	(1.20)	(1.37)
Share White	0.55	0.56	0.32
	(0.50)	(0.50)	(0.47)
Share Black	0.13	0.13	0.13
	(0.33)	(0.33)	(0.33)
Ref. Person's Schooling (years)	13.78	13.84	12.76
	(2.45)	(2.44)	(2.41)
Unemployed	0.06	0.00	1.00
	(0.24)	(0.00)	(0.00)
Non-Durable Expenditure	20494.70	20852.78	14868.23
	(10760.47)	(10769.96)	(8,884.20)
Food Expenditure	9,596.98	9,741.15	7,331.58
	(5,750.15)	(5,771.84)	(4,866.71)
Services Expenditure	16226.51	16584.40	10603.04
	(13272.99)	(13344.30)	(10618.33)
Total Expenditure (Non-Housing)	25719.53	26211.44	17990.13
	(16316.70)	(16356.78)	(13496.92)
Total Expenditure (Incl. Housing)	38476.78	39289.33	25709.24
	(24918.80)	(24984.72)	(19899.40)
N	55671	52340	3331
Waves	3.65	3.45	1.51
Unique Households	15270	15188	2211

TABLE B.1. Summary Statistics for Our Sample of Interest
Computed from the 1999-2019 waves of the PSID. Tables shows and means and standard deviations in parenthesis. We denominate all dollar values using 2019 prices.

We then normalize the estimated individual fixed effects by their standard-deviation to construct our measure of permanent income:

$$\overline{Y}_i = \frac{\mu_i}{\sigma_\mu}.$$

This normalization provides an interpretable scale for permanent income. We use the resulting z-score in our estimation of employment risk, so the coefficient on permanent income captures the effect of a one standard deviation increase in permanent income.

C. Maximum Likelihood Estimator for Latent Heterogeneity

In the main text we specify group-specific heterogeneity as (equation (7) above):

$$\Delta_{i,t}^{FD} = \sum_{k \in K} \mathbb{1}\left[k(i) = k\right] \left(\delta_0^k + \delta_1^k U_{i,t}\right) + \beta X_{i,t} + \varepsilon_{i,t}, \tag{C.1}$$

and assume that $\varepsilon_{i,t} \sim N(0, \sigma^2_{g(i)})$. In other words, we make a parametric restriction on the latent group specific density of errors. As discussed by Lewis, Melcangi, and Pilossoph (2024), under this parametric restriction, identification of latent heterogeneity in consumption growth and the consumption response to the onset of unemployment exploits two complementarity sources of information: i) panel data information on persistent differences in households' consumption growth over time, and (ii) cross-sectional restrictions on the errors (which lack first-order autoregressive group structure under the Gaussian assumption) that can be distinguished from group-specific variation in the conditional mean. It is because we have relatively short panels for each household that lead us to leverage both sources of identification.

Let D collect the vector of individual-specific indicators defined by the partition $k(i)$ and $\mathbb{X}$ collect the covariates and unemployment indicators. Then, were the assignment to groups known with certainty, the complete-data likelihood is given by

$$L(\Delta^{FD}, \mathbb{X}, D; \chi) = \prod_{i=1}^{N} \prod_{t=1}^{T} \prod_{k=1}^{K} \left(\pi^k\right)^{d_{i,k} \times o_{i,t}} f\left(\Delta_{i,t}^{FD} | \underbrace{\delta_0^k + \delta_1^k U_{i,t} + \beta X_{i,t}}_{\mu^k(U_{i,t}, X_{i,t})}, (\sigma^g)^2\right)^{d_{i,k} \times o_{i,t}} \tag{C.2}$$

where $o_{i,t}$ is an indicator for whether the household is observed at time t, $f(\cdot|\mu, \sigma^2)$ is the density for the normal distribution with mean μ and variance σ^2. χ collects the parameters of interest: the unconditional probability of belonging to each group ($\pi^k = E[d_{i,k}]$), the (group-specific) parameters in the linear regression (δ^k and β), and the variance (σ^k).

As the assignment is not observed by the econometrician ex ante, we maximize the expected log-likelihood instead:

$$E_{D|\Delta^{FD}, \mathbb{X}}[\ln L(\Delta^{FD}, \mathbb{X}, D; \chi)] = \sum_{i=1}^{N} \sum_{t=1}^{T} o_{i,t} \sum_{g=1}^{G} \pi_i^g \left(\ln(\pi^k) + \ln f\left(\Delta_{i,t}^{FD} | \mu^k(U_{i,t}, X_{i,t}), (\sigma^k)^2\right)\right), \tag{C.3}$$

where

$$\pi_i^k = Pr(d_{i,k} = 1 | \Delta_i^{FD}, \mathbb{X}_i) = \frac{\pi^k \prod_t f\left(\Delta_{i,t}^{FD} | \mu^g(U_{i,t}, X_{i,t}), (\sigma^k)^2\right)}{\sum_{k'=1}^{K} \pi^{k'} \prod_t f\left(\Delta_{i,t}^{FD} | \mu^h(U_{i,t}, X_{i,t}), (\sigma^{k'})^2\right)} \tag{C.4}$$

are posterior weights which capture the econometrician's ex post uncertainty over group assignment. We do not explicitly include other covariates in the conditioning set, instead these are valid posteriors conditional on the outcome and the regressors included in the linear model. We use the value of this posterior probability at our estimated parameters to construct the analogue $\hat{\pi}_i^k$ which we use to simulate the value of our sufficient statistics within sample with $\Delta c_i / c_i$ as $\sum_{k \in K} \hat{\pi}_i^k \hat{\delta}_1^k$.

As Lewis, Melcangi, and Pilossoph (2024) highlight, the likelihood in (C.3) admits a sequential estimation procedure using the Expectation-Maximization (E-M) algorithm Dempster, Laird, and Rubin (1977). We implement this estimation procedure using the R package `flexmix`. Due to the local convergence properties of the E-M algorithm we initialize the algorithm from 2000 different starting values to account for the possibility of local optima and select the estimates providing the largest value of the likelihood.

Inference on these objects is based on the Fisher Information Matrix.

D. Additional Empirical Results

D.1. Robustness of Negative Selection Estimates to Alternative Controls

In our baseline empirical test of negative selection, we condition on a limited set of household characteristics. We exclude education and race controls, as redistributive policies are typically not designed with the explicit purpose of solely redistributing within these groups. Here, we assess the robustness of our findings to the inclusion of additional controls. Figure D.1 shows that average consumption differences remain robust when these controls are added. The right panel, which focuses on *within-group* differences, shows some attenuated effects of negative selection, reflecting the fact that we are controlling for other household characteristics relevant to consumption and labor market outcomes. Nevertheless, the quantitative impact of including these additional controls is modest.

D.2. Robustness of Consumption Decline Estimates to Alternative Consumption Measures

In addition to measuring consumption using our baseline expenditure measures, we also consider robustness to alternative consumption series. Table D.1 shows that estimates of the average consumption decline are similar across these alternative measure and Figure D.2 shows how the distribution of consumption declines differs between our baseline measure and a measure of food expenditure as well as reporting the confidence interval for our estimate of the distribution.

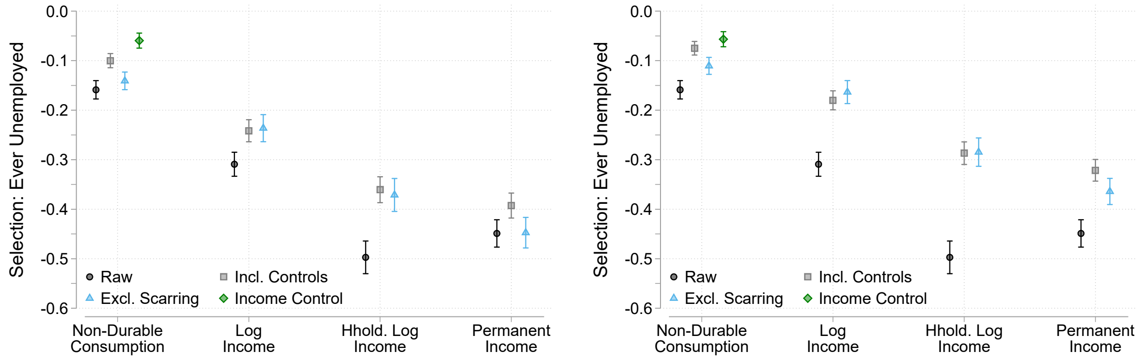

(A) Baseline

(B) Additional Controls

FIGURE D.1. Negative Selection Among the Unemployed

Notes: Authors' calculation from the PSID data. All regressions drop periods of unemployment. The no scarring sample conditions on periods before first observed unemployment. Raw specification includes no controls. In the left panel, all other specifications additionally control for a cubic function of age, indicator variables for whether they are married, a set of indicators for family size, and a set of year specific indicators. In the right panel we additionally control for gender, and a set of indicator variables capturing education status and race.

	Baseline	Alternative Measures			
	Non-durable	Food Expenditure	Services	Total (Non-Housing)	Total (Incl. Housing)
$\mathbb{1}\left[U_{i,t}=1\right]$	-0.129^{***}	-0.145^{***}	-0.143^{***}	-0.129^{***}	-0.150^{***}
	(0.012)	(0.017)	(0.018)	(0.014)	(0.014)
Controls					
Age	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Year	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
Family Size	*Yes*	*Yes*	*Yes*	*Yes*	*Yes*
N	37, 260	36, 997	37, 246	37, 289	37, 295

TABLE D.1. Consumption Declines by Alternative Consumption Measures ($\Delta c/c$)

* $p < 0.10$, ** $p < 0.5$, *** $p < 0.01$. The first column reports the average consumption decline estimated (under the assumption of homogeneous consumption drops) for our baseline measure of consumption. The remaining columns report the consumption declines for alternative measures of consumption constructed from reported expenditures in the PSID data. Standard errors are clustered at the household-level. The number of observations across regressions differs due to item non-response.

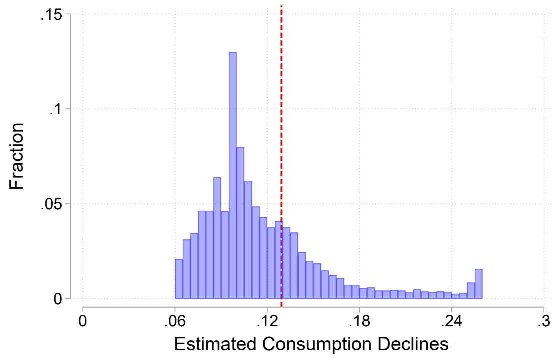

(A) Baseline (Non-Durable)

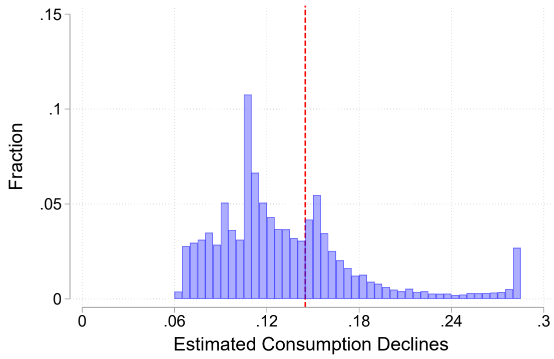

(B) Food Expenditures

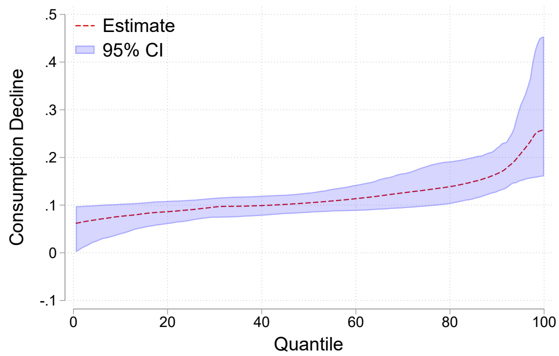

(C) Baseline (Non-Durable)

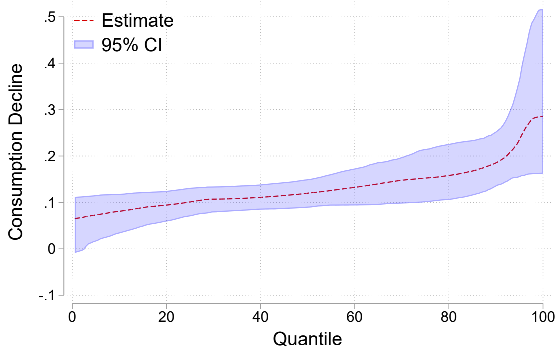

(D) Food Expenditures

FIGURE D.2. Distribution of Estimated Consumption Declines for Alternative Consumption Series

Notes: Authors' estimates from the PSID data. The left panel reports our baseline estimates for our broad measure of non-durable consumption and the right panel reports results based on an alternative measure of consumption, food expenditures. The histogram (light blue bars) plots the individual proportional consumption declines constructed using the Gaussian mixture linear Regression-estimated parameters and individuals' posterior probabilities for each group, $\hat{\pi}_i^k$. For each household we compute the posterior-weighted effect of unemployment across the discrete group-specific unemployment consumption declines. The sample is defined as in the text. The Bayesian Information Criterion selects K = 3. The homogeneous estimate (red dashed vertical line) is estimated imposing K = 1 in our baseline specification. The second row shows our point estimates of each quantile along with the 95% confidence interval constructed from our bootstrap replications.

D.3. Observable Heterogeneity in Consumption Declines

Our approach to estimating unemployment-induced consumption declines avoids the need to specify observable determinants of its heterogeneity ex ante. We can, however, correlate our predicted declines with observables ex post. In Table D.2 we report the results of this exercise in order to understand observable determinants of heterogeneity in consumption declines.

We focus on proxies of a household's capacity to self-insure measured in periods of employment. This enables us to assess evidence for whether our estimated consumption declines are higher for those less able to self-insure.[27] To proxy for the availability of private savings, by which households may smooth consumption, we include indicators for each quartile of the distribution of liquid wealth and an indicator for whether or not the household owns their home. To construct liquid wealth, we follow Carroll and Samwick (1997) and sum the value of cash (including checking accounts) and savings, stocks owned outside of a retirement account, and bonds and government treasuries. These are the most liquid assets and can be used to smooth shocks in the short run. We additionally include indicators for quartiles of earned income, which captures both earning capacity and the size of the income loss a household will experience in unemployment. Finally, we include an indicator for marriage and cohabitation, which proxies for the additional insurance provided by the added work effect or spousal insurance.

We report the coefficient estimates from a regression of consumption declines (column 1) as well as an estimate of the conditional heteroskedasticity (column 2), which captures the extent to which the variability of consumption declines is impacted by observables. We find that, conditional on earning quartile, households with a larger ability to self insure (i.e., higher liquid assets, cohabiting, and home owning) have smaller consumption declines. Moreover, these factors also act to lower the dispersion of these consumption declines. These effects are economically and statistically significant.

We view this as both reinforcing the mechanisms highlighted by the existing literature, and highlighting the key role played by heterogeneity. Yet, we also find that these measures explain only a relatively small fraction of the variation in outcomes: an R^2 of 0.11. There are two interpretations of this finding. First, there is considerable latent heterogeneity that drives differences in the consumption exposure to unemployment (e.g., heterogeneous beliefs, preferences, or stochastic processes that lead to the optimality of different consumption behavior). Second, the proxies of a household's capacity to self insure are relatively weak (e.g., because liquidity is measured during employment and two years before job–loss). We conjecture that both are important, but note our approach is robust to both factors. This highlights a key advantage of the flexible approach we take to estimating the distribution of consumption declines.

D.4. Modeling the Correlation between Employment Risk and Consumption Declines

In Table D.3 we report the underlying estimates that correspond to Figure 3

In our main empirical implementation we model employment risk as a function of observables. In this subsection we show that directly modeling the correlation between employment risk and our estimate of individual-specific consumption drops yields similar estimates of employment risk. To do so, we

[27] As Lewis, Melcangi, and Pilossoph (2024) note in the context of marginal propensities to consume, a key advantage of using latent types to capture heterogeneity is that it implicitly allows for various observables to be included jointly, without loss of statistical power from the interaction of successively smaller groups in an interacted consumption decline specification.

	Consumption Decline	log(Variance)
Second Quartile Liquidity	-0.009^{***}	-0.253^{***}
	(0.001)	(0.044)
Third Quartile Liquidity	-0.013^{***}	-0.381^{***}
	(0.001)	(0.047)
Top Quartile Liquidity	-0.011^{***}	-0.211^{***}
	(0.001)	(0.051)
Second Quartile Earnings	-0.007^{***}	-0.171^{***}
	(0.001)	(0.046)
Third Quartile Earnings	-0.010^{***}	-0.234^{***}
	(0.001)	(0.050)
Top Quartile Earnings	-0.008^{***}	-0.158^{***}
	(0.001)	(0.056)
Home Owner	-0.013^{***}	-0.265^{***}
	(0.001)	(0.042)
Married	-0.011^{***}	-0.111^{**}
	(0.001)	(0.046)
N	47,735	47,735

TABLE D.2. Consumption Declines ($\Delta c/c$) and Proxies of Access to Self-Insurance
* $p < 0.10$, ** $p < 0.5$, *** $p < 0.01$. For each household we compute the posterior-weighted effect of unemployment across the discrete group-specific unemployment consumption declines. This is constructed using regression-estimated parameters and individuals' posterior probabilities for each group $\hat{\pi}_i^K$. We use the sample as defined as in the main text and consider periods of employment. The first column reports how individual consumption declines predicted by our Gaussian Linear Mixture Model vary with measures of ability to self-insure. The second column reports estimates of the conditional heteroskedasticity. As discussed above we define liquidity using the definition of "Very Liquid Assets" given by Carroll and Samwick (1997). Standard errors are clustered at the household level.

add to our model of employment risk (described in Section 4.2) a cubic function of the idiosyncratic consumption exposure.

Figure D.3 reports the results of this exercise. Panel (A) shows that the implied average partial effects in the extended model are very similar for all regressors included in the baseline model. These effects are slightly attenuated towards zero as the additional measure of consumption exposure absorbs some of their explanatory power. However, the changes in effect sizes are economically insignificant. The effect of consumption exposure is statistically significant, but small. On average, an increase in the consumption exposure of 10 percentage points is associated with an increase in employment risk of 1.6 percentage points.

Panel (B) displays the average employment risk in our extended model as a function of our baseline estimates. There are no clear or systematic deviations in the predictions across models: values lie on or close to the 45-degree line and the predictions have a positive correlation of almost 1. Given we estimate the

	Coefficient Estimates	Average Partial Effect
Age (Decades)	−2.013*	−0.014***
	(1.058)	(0.002)
Age^2	0.350	
	(0.273)	
Age^3	−0.021	
	(0.023)	
Woman	−0.284***	−0.012***
	(0.060)	(0.002)
High Schl. Grad	−0.257***	−0.012***
	(0.068)	(0.003)
Some College+	−0.214***	−0.010***
	(0.071)	(0.003)
College	−0.235**	−0.011***
	(0.092)	(0.004)
Adv. Degree	−0.194*	−0.009**
	(0.099)	(0.004)
Black	0.339***	0.015***
	(0.054)	(0.002)
Non-White & Non-Black	0.041	0.002
	(0.078)	(0.003)
Perm. Inc.	−1.048***	−0.038***
	(0.048)	(0.001)
PI^2	0.017	
	(0.014)	
PI^3	0.038***	
	(0.004)	
Constant	0.477	
	(1.315)	
N	52, 996	52, 996

TABLE D.3. Logit Estimation of e_i

Notes: Table reports the point estimates and standard errors for the coefficients of our estimated model of employment risk described in Section 4.2. In addition, we report the average partial effects and their corresponding standard errors which we display in Figure 3 in the main text.

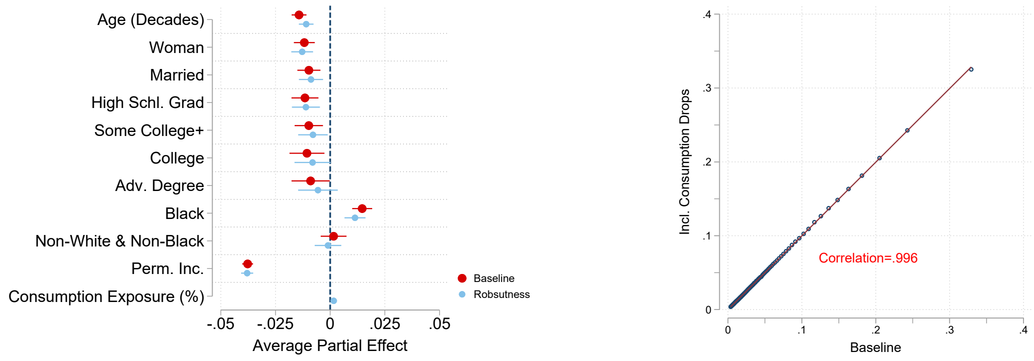

(A) Average Partial Effects

(B) Individual Predictions

FIGURE D.3. Correlation Between Consumption Exposure and Employment Risk
Notes: Authors' calculation from the PSID data. Our baseline measures of risk and consumption declines are estimated as we describe above and correspond to the model results summarized in Figures 2 and 3. Our robustness check includes a cubic function of the idiosyncratic consumption exposure in our model of employment risk. We use binned scatter plots with 100 bins in the right panel.

consumption declines in an alternative estimation procedure which complicates the problem of inference, we choose to use the simpler model of employment risk in our baseline analysis.

Figure 1 shows how both the estimated consumption declines and labor market risks vary with in-work consumption. Here we provide additional evidence on the link between employment risk and consumption declines. We regress our individual-level estimates of consumption declines on the same demographic and education information that we use to model employment risk. This directly illustrates the covariance between weights and consumption declines in our risk-weighted expectation.

Figure D.4 reports the average partial effects from this exercise (panel A), alongside our baseline estimates of average partial effects for employment risk (panel B, repeating Figure 3 in the main text). We find that factors that are systematically associated with smaller consumption declines are also associated with smaller probabilities of becoming unemployed and vice-versa. Consistent with the results shown above in Appendix D.3, we find a larger effect of marriage on consumption declines relative to employment risk, and a smaller permanent income gradient.

E. Willingness-to-Pay: Additional Results

E.1. Worker Preferences Across Alternative Reforms

Figure E.1 illustrates how WTP for UI expansions covaries across different reform types. We begin by dividing the sample into 20 equally sized groups based on WTP under a proportional-capped reform. Within each of these groups, we further divide individuals into 20 equally sized bins based on WTP for an alternative expansion. We then plot the joint distribution by showing

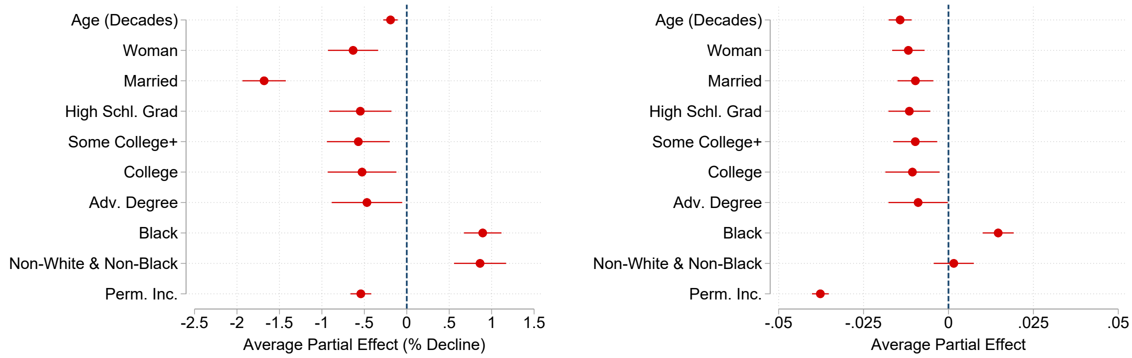

(A) Consumption Decline

(B) Employment Risk

FIGURE D.4. Consumption Exposure, Employment Risk and Demographics
Notes: Authors' calculation from the PSID data. Our baseline measures of risk and consumption declines are estimated as we describe above and correspond to the model results summarized in Figures 3 and 2. We predict consumption declines as a function of the same variables we included in our estimated employment risk and report the average partial effects for both.

average WTP across these bins for each reform. Panel (A) compares WTP under the baseline proportional-capped reform with WTP under a flat expansion. In the upper-right quadrant, we observe individuals that contribute relatively little to financing either reform. For these individuals WTP values converge to the 45-degree line, as, at this point, their contribution to reform financing is negligible and their reform valuations are therefore largely independent of its structure. Panel (B) shows a similar pattern when comparing the proportional-capped expansion with a proportional expansion.

By displaying the joint distribution, we can partition the WTP space into four mutually exclusive and exhaustive quadrants. The "Northeast" and "Southwest" quadrants contain individuals with consistent evaluation of both reforms: those in the Northeast derive positive surplus from either expansion, while those in the Southwest experience negative surplus under both. Two types of individuals fall in the Northeast quadrant: (i) those who in net terms contribute to funding the reform, but whose insurance value outweighs cross-subsidization effects, and (ii) those who benefit from cross-subsidization under both reforms—typically workers with high employment risk for whom the pooled reform price is below their actuarially fair price. Conversely, individuals in the Southwest quadrant find the pooled price more expensive than an actuarially fair individual price. This group disproportionately includes high-income individuals, who tend to face lower employment risk.

The differences across reform types highlight that workers may disagree about the desirability of UI expansions. Workers in the Southeast quadrant favor an expansion to the proportional-capped system but oppose a flat (Panel A) or proportional (Panel B) expansion. Conversely, those in the Northwest quadrant oppose a proportional-capped expansion, but favor a flat (Panel (A)) or proportional (panel (B)) expansion.

Additionally, we can classify individuals based on their relative ranking of alternative reforms. Those above the 45-degree line in Panel (A) prefer a flat expansion to a proportional-capped expansion. Similarly, those above the 45-degree line in panel (B) favor an increase in replacement rates to an expansion of the current proportional-capped system. Overall, 81% of workers prefer expanding the proportional-capped system to a uniform benefit increase, and 55% prefer it to an increase in the replacement rate.

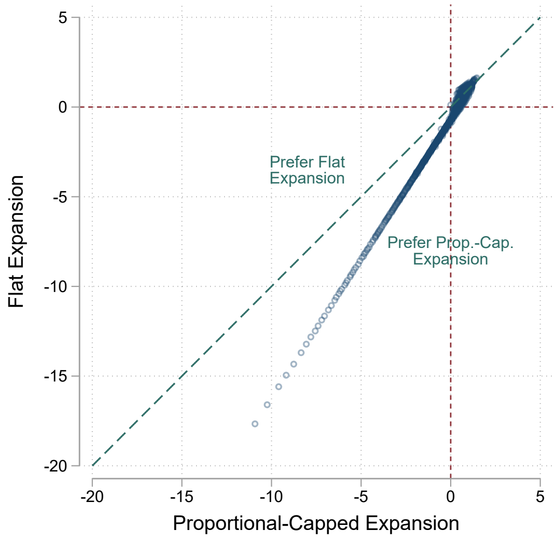

(A) Flat

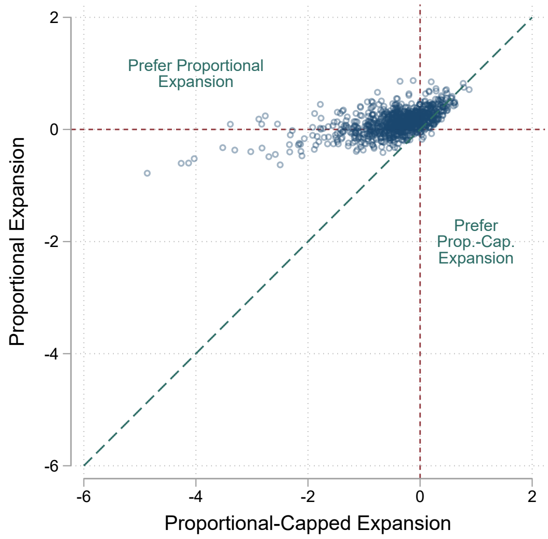

(B) Proportional

FIGURE E.1. Individual Preferences over Reforms

Notes: Authors' estimates from the PSID data. We compare the proportional-capped system corresponding to current U.S. UI policy against a flat expansion (Panel (A)) and proportional (Panel (B)) reform.

E.2. Sensitivity to Preference Specification

Our willingness-to-pay results in Figure 4 are based on a coefficient of risk aversion of 3 and assume state-independent utility. Figure E.2 presents sensitivity analyses exploring the robustness of the implied WTP distributions to these assumptions. Since both risk-aversion and state dependence affect only the insurance component of WTP, we focus on that component.

Panel (A) reproduces the baseline distribution of the insurance value of WTP, conditional on household income (Figure 4(A)). Panel (B) introduces state-dependent utility: specifically, we assume that unemployed households face consumption prices 1.5% lower than those of employed households, following evidence from Kaplan and Menzio (2015) and Campos and Reggio (2020).[28] This assumption affects WTP in two offsetting ways (see Appendix A.2). First, lower prices raise the marginal value of a dollar when unemployed, increasing insurance value. Second, because observed expenditure understates true consumption for the unemployed, measured expenditure-based

[28] Kaplan and Menzio (2015) find that married U.S. households with one unemployed member pay 1.5% lower prices; Campos and Reggio (2020) reports similar findings for Spanish households.

consumption drops overstate insurance values. Adjusting to correct the second effect dominates, modestly lowering the WTP distribution relative to the baseline.

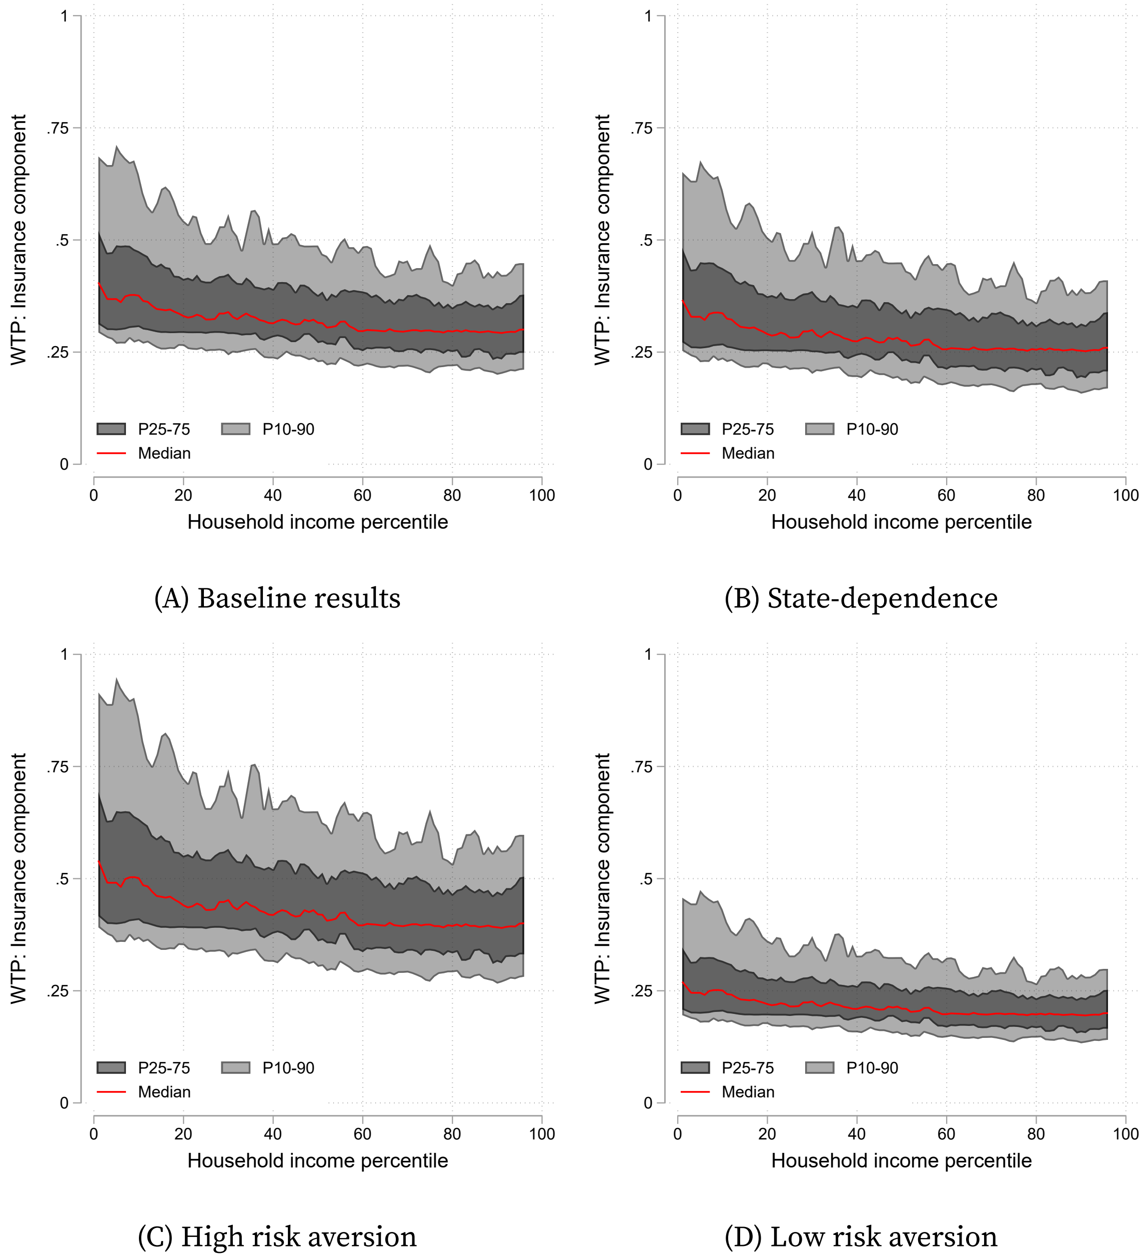

(A) Baseline results

(B) State-dependence

(C) High risk aversion

(D) Low risk aversion

FIGURE E.2. Willingness-to-Pay for Actuarially Fair UI Expansion by Income

Notes: Authors' estimates from the PSID data. Each panel summarizes the distribution of workers willingness-to-pay for an actuarially fair UI expansion (i.e,, the insurance component of overall WTP)—reporting the median, as well as 10th, 25th, 75th, 90th percentiles. We compute worker-level WTP using estimates of individual consumption declines. Panel (A) repeats our baseline results (Figure 4(A)). Panel (B) shows results assuming the price of consumption is 1.5% lower when unemployed. Panels (C) and (D) show results when the coefficient of relative risk aversion is 4 and 2, respectively. Quantiles are smoothed using a uniform rolling window of ± 1 percentiles.

Panels (C) and (D) show the WTP distribution assuming coefficients of relative risk aversions of 4 and 2, respectively. As expected, the distribution scales proportionately with risk aversion. However, because the distribution of WTP for realistic reforms is shaped primarily by cross-subsidization patterns (as shown in Figure 4), the overall distribution remains relatively insensitive to the calibrated level of risk aversion.

F. Normative Framework: Theory

F.1. Section 5.1 Derivations

Social welfare is given by

$$W(\theta) = \int_i \omega_i V_i(\theta) di$$

where individual expected utility, V_i, is defined in equation (1), and indexed by the reform parameter θ.

We consider the reform parameterized by equation (2). Its impact on social welfare is:

$$\frac{dW}{d\theta} = \int_i (1-e_i)\phi_{\mathcal{B}}(y_i)\omega_i u_i^{l\prime}(c_i^l) di - \int_i e_i \phi_{\mathcal{T}}(y_i)\omega_i u_i^{h\prime}(c_i^h) di \frac{d\tau}{db}$$

We focus on a budget-balanced reform, which implies that $\frac{d\tau}{db}$ is given by equation (9), repeated here:

$$\frac{d\tau}{db} = \frac{\int_i (1-e_i)\phi_{\mathcal{B}}(y_i)\, di}{\int_i e_i \phi_{\mathcal{T}}(y_i)\, di} \left(1 + \frac{1}{\int_i (1-e_i)\phi_{\mathcal{B}}(y_i)\, di} \left(\int_i \frac{d(1-e_i)}{d\theta} (\mathcal{T}(y_i) + \mathcal{B}(y_i))\, di\right)\right).$$

Note, we can re-express the fiscal externality (i.e., the portion of the tax adjustment needed to cover behavioral responses):

$$\begin{aligned}
\text{FE} \equiv & \frac{1}{\int_i (1-e_i)\phi_{\mathcal{B}}(y_i)\, di} \left(\int_i \frac{d(1-e_i)}{d\theta} (\mathcal{T}(y_i) + \mathcal{B}(y_i))\, di\right) \\
= & \int_i \frac{(1-e_i)\phi_{\mathcal{B}}(y_i)}{\int_{i'} (1-e_{i'})\phi_{\mathcal{B}}(y_{i'})\, di'} \frac{1}{\phi_{\mathcal{B}}(y_i)} \frac{\mathcal{B}(y_i)}{1-e_i} \frac{d(1-e_i)}{d\theta} \frac{\mathcal{T}(y_i) + \mathcal{B}(y_i)}{\mathcal{B}(y_i)} di \\
= & \mathbb{E}^l \left[\epsilon^{(1-e,b)} \frac{\mathcal{T}(y) + \mathcal{B}(y)}{\mathcal{B}(y)}\right]
\end{aligned}$$

where the final line corresponds to equation (10), and uses the definition of the risk-weighted expectation (equation (8)) and the individual behavioral elasticity $\epsilon_i^{1-e,b} \equiv \frac{1}{\phi_{\mathcal{B}}(y_i)} \frac{d(1-e_i)}{db} \frac{\mathcal{B}(y_i)}{1-e_i}$. Note, in the definition of this elasticity db should be understood as the budget-neutral benefit change.

Using equations (8) and (9), we can then re-write $dW/d\theta$:

$$\frac{dW}{d\theta} = \int_i (1-e_i)\phi_{\mathcal{B}}(y_i)\, di \left(\mathbb{E}^l \left[\omega_i u_i^{l\prime}(c_i^l)\right] - \mathbb{E}^h \left[\omega_i u_i^{h\prime}(c_i^h)\right] \times (1 + \text{FE})\right)$$

Define the money-metric social welfare impact as:

$$\frac{dW^{MM}}{d\theta} = \frac{dW/d\theta}{\int_i (1-e_i)\phi_{\mathcal{B}}(y_i)\,di} \Big/ \frac{\partial W/\partial\tau}{\int_i e_i \phi_{\mathcal{T}}(y_i)\,di}$$

Division by $\int_i (1-e_i)\phi_{\mathcal{B}}(y_i)\,di$ expresses the change in social welfare per \$ expansion in the benefit bill. The scaling by the inverse of $\frac{\partial W/\partial\tau}{\int_i e_i\phi_{\mathcal{T}}(y_i)\,di}$ expresses the social welfare change in money-metric terms, relative to the value of an unfunded tax cut. Note, $\frac{\partial W}{\partial\tau} = \int_i e_i\phi_{\mathcal{T}}(y_i)\,di\mathbb{E}^h\left[\omega_i u_i^{h\prime}(c_i^h)\right]$. As our implementation sets $\phi_{\mathcal{T}}(y_i)=1$, this is equivalent to the effects of a lump-sum increase in high-state incomes.

Hence, we obtain equation (12), repeated here:

$$\frac{dW^{MM}}{d\theta} = \frac{\mathbb{E}^l[\omega u^{l\prime}(c^l)] - \mathbb{E}^h[\omega u^{h\prime}(c^h)]}{\mathbb{E}^h[\omega u^{h\prime}(c^h)]} - \text{FE}$$

Note that:

$$\begin{aligned}
\mathbb{E}^l[\omega u^{l\prime}(c^l)] - \mathbb{E}^h[\omega u^{h\prime}(c^h)] &= \mathbb{E}^l\left[\omega\left(u^{l\prime}(c^l) - \frac{e_i\phi_{\mathcal{T}}(y_i)}{(1-e_i)\phi_{\mathcal{B}}(y_i)} \Big/ \frac{\int_{i'} e_{i'}\phi_{\mathcal{T}}(y_{i'})di'}{\int_{i'}(1-e_{i'})\phi_{\mathcal{B}}(y_{i'})di'} u^{h\prime}(c^h)\right)\right] \\
&= \mathbb{E}^l\left[\omega u^{h\prime}(c^h)\left(\frac{u^{l\prime}(c^l)-u^{h\prime}(c^h)}{u^{h\prime}(c^h)} + \left(1 - \frac{e_i\phi_{\mathcal{T}}(y_i)}{(1-e_i)\phi_{\mathcal{B}}(y_i)} \Big/ \frac{\int_{i'} e_{i'}\phi_{\mathcal{T}}(y_{i'})di'}{\int_{i'}(1-e_{i'})\phi_{\mathcal{B}}(y_{i'})di'}\right)\right)\right] \\
&= \mathbb{E}^l[\omega u^{h\prime}(c^h)]\mathbb{E}^l\left[\lambda\left(\frac{u^{l\prime}(c^l)-u^{h\prime}(c^h)}{u^{h\prime}(c^h)} + \left(1 - \frac{e_i\phi_{\mathcal{T}}(y_i)}{(1-e_i)\phi_{\mathcal{B}}(y_i)} \Big/ \frac{\int_{i'} e_{i'}\phi_{\mathcal{T}}(y_{i'})di'}{\int_{i'}(1-e_{i'})\phi_{\mathcal{B}}(y_{i'})di'}\right)\right)\right]
\end{aligned}$$

where $\lambda_i \equiv \frac{\omega_i u_i^{h\prime}(c^h)}{\mathbb{E}^l[\omega u^{h\prime}(c^h)]}$. Hence we obtain the decomposition in equation (13),

$$\frac{dW^{MM}}{d\theta} = \bar{\lambda}\mathbb{E}^l\left[\lambda\left(\frac{u^{l\prime}(c^l)-u^{h\prime}(c^h)}{u^{h\prime}(c^h)} + 1 - \frac{e_i\phi_{\mathcal{T}}(y_i)}{(1-e_i)\phi_{\mathcal{B}}(y_i)} \Big/ \frac{\int_{i'} e_{i'}\phi_{\mathcal{T}}(y_{i'})di'}{\int_{i'}(1-e_{i'})\phi_{\mathcal{B}}(y_{i'})di'}\right)\right] - \text{FE},$$

where $\bar{\lambda} \equiv \frac{\mathbb{E}^l[\omega u^{h\prime}(c^h)]}{\mathbb{E}^h[\omega u^{h\prime}(c^h)]}$. Note that:

$$\mathbb{E}^l\left[\lambda\left(1 - \frac{e_i\phi_{\mathcal{T}}(y_i)}{(1-e_i)\phi_{\mathcal{B}}(y_i)} \Big/ \frac{\int_{i'} e_{i'}\phi_{\mathcal{T}}(y_{i'})di'}{\int_{i'}(1-e_{i'})\phi_{\mathcal{B}}(y_{i'})di'}\right)\right] = 1 - \mathbb{E}^h[\lambda] = 1 - \frac{1}{\bar{\lambda}}$$

hence we have the decomposition in equation (14):

$$\frac{dW^{MM}}{d\theta} = \bar{\lambda}\mathbb{E}^l\left[\lambda\left(\frac{u^{l\prime}(c^l)-u^{h\prime}(c^h)}{u^{h\prime}(c^h)}\right)\right] + (\bar{\lambda}-1) - \text{FE}.$$

F.2. Extension to General Model

Under the general model outlined in Appendix A.1, the impact of the reform on social welfare is given by:

$$\frac{dW}{d\theta} = -\int_i \omega_i \int_t \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t}) \frac{\partial u_i^h(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t})dtdi \times \frac{d\tau}{db}$$
$$+ \int_i \omega_i \int_t \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)(1-\xi_{i,t}(\varphi_{i,t})) \frac{\partial u_i^l(c_{i,t}(\varphi_{i,t})), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t})dtdi,$$

where budget-balance requires:

$$\frac{d\tau}{db} = \frac{\int_i \int_t \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)\left(1-\xi_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t})dtdi}{\int_i \int_t \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t})dF_{i,t}(\varphi_{i,t})dtdi} \times \left(1 + \frac{1}{\int_i \int_t \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)\left(1-\xi_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t})dtdi} \times \right.$$
$$\left. \int_i \int_t \int_{\varphi_{i,t}} \frac{d(1-\xi_{i,t}(\varphi_{i,t}))}{d\theta}(\mathcal{T}(y_i) + \mathcal{B}(y_i))dF_{i,t}(\varphi_{i,t})dtdi \right)$$
$$= \frac{\int_i \int_t \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)\left(1-\xi_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t})dtdi}{\int_i \int_t \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t})dF_{i,t}(\varphi_{i,t})dtdi} \times \left(1 + \underbrace{\mathbb{E}^l\left[\epsilon \times \frac{\mathcal{T}(y_i) + \mathcal{B}(y_i)}{\mathcal{B}(y_i)}\right]}_{\text{FE}}\right).$$

ϵ_i is given by:

$$\epsilon_i = \int_t \int_{\varphi_{i,t}} \frac{d(1-\xi_{i,t}(\varphi_{i,t}))}{db} dF_{i,t}(\varphi_{i,t})dt \frac{1}{\phi_B(y_i)} \frac{B(y_t)}{\int_t \int_{\varphi_{i,t}} (1-\xi_{i,t}(\varphi_{i,t}))dF_{i,t}(\varphi_{i,t})dt}$$

and the low-state risk-weighted expectation over any i-specific variable x_i is given by:

$$\mathbb{E}^l\,[x] \equiv \int_i \left(\frac{\int_t \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)(1-\xi_{i,t}(\varphi_{i,t}))dF_{i,t}(\varphi_{i,t})dt}{\int_{i'} \int_t \int_{\varphi_{i',t}} \phi_{\mathcal{B}}(y_{i'})(1-\xi_{i',t}(\varphi_{i',t}))dF_{i',t}(\varphi_{i',t})dtdi'} \right) x_i di$$

with the high-state risk-weighted expectation defined analogously (see equation (8)).

Under the standard assumption that when a worker finds a job they remain employed thereafter, $D_i \equiv \int_t \int_{\varphi_{i,t}} (1-\xi_{i,t}(\varphi_{i,t}))dF_{i,t}(\varphi_{i,t})dt$ is the worker's expected unemployment duration, ϵ_i is the elasticity of worker i's unemployment duration with respect to the benefit change $\phi_i(y_i)db$ they face under the budget-balanced reform, and the low risk-weighted expectation weights by individuals' expected benefit rise, which depends on their expected duration and benefit adjustment, $\phi_i(y_i)$.

To express the social welfare impact per \$ of increased benefit expenditure in and money-metric terms, relative to an unfunded tax cut, we rescale $dW/d\theta$ according to:

$$\frac{dW^{MM}}{d\theta} = \frac{dW/d\theta}{\int_i \int_t \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)\left(1-\xi_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t})dtdi} \Big/ \frac{\partial W/\partial\tau}{\int_i \int_t \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t})dF_{i,t}(\varphi_{i,t})dtdi}.$$

After a couple of lines of algebra analogous to those in the preceding subsection, we obtain:

$$\frac{dW^{MM}}{d\theta} = \frac{\mathbb{E}^l\left[\frac{\partial u^l(c,x)}{\partial c}\right] - \mathbb{E}^h\left[\frac{\partial u^h(c,x)}{\partial c}\right]}{\mathbb{E}^h\left[\frac{\partial u^h(c,x)}{\partial c}\right]} - \text{FE}$$

F.3. Relationship with Marginal Value of Public Funds

Our normative analysis considers a reform to UI generosity that raises benefits financed by a flat increase in payroll taxation. We focus on this combined reform because the structure of the U.S. UI system links UI benefits to a hypothecated tax. However, it is instructive to consider the benefit and tax components separately. Doing so clarifies the relationship between our analysis and the marginal value of public funds (MVPF) framework (see Hendren and Sprung-Keyser 2020).

The MVPF for an unfunded benefit rise is:

$$\text{MVPF}^{\mathcal{B}} = \frac{\mathbb{E}^l[\text{WTP}^{\mathcal{B}}]}{1+\text{FE}^{\mathcal{B}}}$$

where the willingness-to-pay is:

$$\text{WTP}_i^{\mathcal{B}} = \left(\frac{u_i^{l\prime}(c_i^l) - u_i^{h\prime}(c_i^h)}{u_i^{h\prime}(c_i^h)} + 1\right)$$

and $\text{FE}^{\mathcal{B}}$ denotes the associated fiscal externality from behavioral responses.

Willingness-to-pay to pay to avoid a flat tax increase is 1, and the MVPF for the tax rise is:

$$\text{MVPF}^{\mathcal{T}} = \frac{1}{1+\text{FE}^{\mathcal{T}}},$$

where $\text{FE}^{\mathcal{T}}$ is the fiscal externality associated with the payroll tax increase.

To assess whether the combined reform is socially valuable, we can compare the MVPF of the benefit expansion to the MVPF of the tax increases, both valued at the incidence-weighted average marginal social welfare weights of affected groups. The combined reform raises welfare if:

$$\mathbb{E}^l\left[g\left(\frac{\text{WTP}^{\mathcal{B}}}{\mathbb{E}^l[\text{WTP}^{\mathcal{B}}]}\right)\right]\frac{\mathbb{E}^l[\text{WTP}^{\mathcal{B}}]}{1+\text{FE}^{\mathcal{B}}} > \mathbb{E}^h[g]\frac{1}{1+\text{FE}^{\mathcal{T}}}$$

$$\frac{\mathbb{E}^l[g]}{\mathbb{E}^h[g]}\mathbb{E}^l\left[\frac{g}{\mathbb{E}^l[g]}\text{WTP}^{\mathcal{B}}\right] - \frac{1+\text{FE}^{\mathcal{B}}}{1+\text{FE}^{\mathcal{T}}} > 0$$

$$\frac{\mathbb{E}^l[g]}{\mathbb{E}^h[g]}\mathbb{E}^l\left[\frac{g}{\mathbb{E}^l[g]}\left(\frac{u^{l\prime}(c^l)-u^{h\prime}(c^h)}{u^{h\prime}(c^h)}\right)\right]+\left(\frac{\mathbb{E}^l[g]}{\mathbb{E}^h[g]}-1\right)-\left(\frac{1+\mathrm{FE}^{\mathcal{B}}}{1+\mathrm{FE}^{\mathcal{T}}}-1\right)>0.$$

Recognizing that $\bar{\lambda}=\frac{\mathbb{E}^l[g]}{\mathbb{E}^h[g]}$, $\lambda_i=\frac{g}{\mathbb{E}^l[g]}$ and that the fiscal externality of the combined reform in the main text is equivalent to $\mathrm{FE}=\frac{1+\mathrm{FE}^{\mathcal{B}}}{1+\mathrm{FE}^{\mathcal{T}}}-1$, makes clear that the left-hand side of the inequality corresponds exactly to the condition in equations (12) and (14).[29]

G. Normative Framework: Implementation

G.1. Measuring Fiscal Externality

Our fiscal externality term decomposes as:

$$\mathrm{FE}=\mathbb{E}^l\left[\epsilon^{(1-e,b)}\right]+\mathbb{E}^l\left[\epsilon^{(1-e,b)}\frac{\mathcal{T}(y)}{\mathcal{B}(y)}\right], \tag{G.1}$$

where we refer to the first component as the direct fiscal externality from UI expansion, arising from higher benefit payments. The second component captures the indirect effect, which operates through changes in income tax payments.

The first component depends on risk weights (in the expectation), which we construct using our employment risk model and observed data on worker wages, along with the behavioral elasticity with respect to UI benefits. We draw on elasticity estimates from Chetty (2008), who reports how this elasticity varies across wealth quartiles (see Table 2, p. 204). We assign each household in our data the corresponding elasticity based on its wealth quartile. Following Chetty (2008), we adjust these elasticity to account for benefit exhaustion after six months. Specifically, we scale them by the ratio of average benefit and unemployment duration (15.8/24.3) (see also discussion in Schmieder and Von Wachter, 2016).

The indirect component of the fiscal externality depends on risk weights, behavioral elasticities, and, additionally, the state-specific tax position of workers. We compute these tax positions based using the NBER TAXSIM (version 32) U.S. tax and transfer simulator.

Specifically, for all observations in which the household's reference person is employed, we compute household annual tax liability, accounting for spousal earnings in married households (assuming joint filing) and household composition. This yields an estimate of the household's tax liability in the high state, denoted by T_i^h.

To construct the net tax liability in the low state, we assume the reference person experiences an unemployed spell of average duration and receives unemployment benefits equal to the lower of 50% of wages or the UI earnings

[29] In practice, it is likely that $\mathrm{FE}^{\mathcal{B}}\approx\mathrm{FE}$ and $\mathrm{FE}^{\mathcal{T}}\approx 0$, since the fiscal externality of the combined reform is likely primarily driven by the benefit expansion rather than the small funding-side payroll tax increase. However, the above equivalence does not rely on this being the case.

cap (approximately 41% of average wages). Denote UI payments (ui_i). We then compute the household's tax liability, in low state net states, denoted by T_i^l.

The empirical counterpart of $\frac{\mathcal{T}(y)}{\mathcal{B}(y)}$ is given by:

$$\frac{T_i^h - T_i^l}{ui_i},$$

which reflects the difference in state-specific income tax liabilities, scaled by unemployment benefit payments.

G.2. Average Implementation

Our full implementation of the flat reform entails evaluating the expression:

$$\frac{\mathbb{E}^l[g_y\tilde{g}]}{\mathbb{E}^h[g_y\tilde{g}]}\mathbb{E}^l\left[\frac{g_y\tilde{g}}{\mathbb{E}^l[g_y\tilde{g}]}\gamma\frac{\Delta c}{c^h}\right]+\left(\frac{\mathbb{E}^l[g_y\tilde{g}]}{\mathbb{E}^h[g_y\tilde{g}]}-1\right)-\left(\mathbb{E}^l\left[\epsilon^{(1-e,b)}\right]+\mathbb{E}^l\left[\epsilon^{(1-e,b)}\frac{\mathcal{T}(y)}{\mathcal{B}(y)}\right]\right),$$

which is obtained by combining equations (4), (12), (14), (15) and (G.1). Recall, that the welfare weights include a cross-income (inverse optimum) component g_y and a within-income-group component $\tilde{g}_i = \frac{(c_i^h)^{-\gamma}}{\mathbb{E}[(c_i^h)^{-\gamma}|y]}$

In the average implementation, we treat the employed and unemployed populations as single agents with average values, $\mathbb{E}^s[x]$ for $s \in \{h, l\}$, for the income-based welfare weight (g_y), consumption (c_i^h), consumption drop ($\Delta c/c_h$), benefit entitlement ($\mathcal{B}$), and tax liability ($\mathcal{T}$). This leads to the simplified expression:

$$\frac{\mathbb{E}^l[g_y]}{\mathbb{E}^h[g_y]}\times\gamma\mathbb{E}^l\left[\frac{\Delta c}{c^h}\right]+\left(\frac{\mathbb{E}^h[g_y]}{\mathbb{E}^l[g_y]}-1\right)-\left(\mathbb{E}^l\left[\epsilon^{(1-e,b)}\right]+\mathbb{E}^l\left[\epsilon^{(1-e,b)}\right]\frac{\mathbb{E}^l\left[\mathcal{T}(y)\right]}{\mathbb{E}^l\left[\mathcal{B}(y)\right]}\right)$$

The average implementation captures differences in average welfare weights between the employed and unemployed populations, but it ignores within-group heterogeneity. As a result, it is invariant to the type of reform considered, since reform differences stem from variation in their incidence across the unemployed population.

References

Aguiar and Hurst (2005). "Consumption versus expenditure". *Journal of Political Economy* 113.5, pp. 919–948.

Blundell, Pistaferri, and Saporta-Eksten (2016). "Consumption inequality and family labor supply". *American Economic Review* 106.2, pp. 387–435.

Browning and Crossley (2001). "Unemployment Insurance Benefit Levels and Consumption Changes". *Journal of Public Economics* 80.1, pp. 1–23.

Campos and Reggio (2020). "Do the Unemployed Pay Lower Prices? A Reassessment of the Value of Unemployment Insurance". *Journal of the European Economic Association* 80.1, p. 47.

Carroll and Samwick (1997). “The nature of precautionary wealth”. *Journal of Monetary Economics* 40.1, pp. 41–71.
Chetty (2006). “A General Formula for the Optimal Level of Social Insurance”. *Journal of Public Economics* 90.10-11, pp. 1879–1901.
— (2008). “Moral Hazard versus Liquidity and Optimal Unemployment Insurance”. *Journal of Political Economy* 116.2, pp. 173–234.
Dempster, Laird, and Rubin (1977). “Maximum likelihood from incomplete data via the EM algorithm”. *Journal of the Royal Statistical Society: series B (methodological)* 39.1, pp. 1–22.
Gruber (1997). “The Consumption Smoothing Benefits of Unemployment Insurance”. *American Economic Review* 87.1, pp. 192–205.
Hai, Krueger, and Postlewaite (2020). “On the welfare cost of consumption fluctuations in the presence of memorable goods”. *Quantitative Economics* 11.4, pp. 1177–1214.
Hendren (2017). “Knowledge of Future Job Loss and Implications for Unemployment Insurance”. *American Economic Review* 107.7, pp. 1778–1823.
Hendren and Sprung-Keyser (2020). “A unified welfare analysis of government policies”. *Quarterly journal of economics* 135.3, pp. 1209–1318.
Kaplan and Menzio (2015). “The morphology of price dispersion”. *International Economic Review* 56.4, pp. 1165–1206.
Lewis, Melcangi, and Pilossoph (2024). *Latent heterogeneity in the marginal propensity to consume*. Tech. rep. 4. National Bureau of Economic Research, pp. 1165–1206.
Low and Pistaferri (2015). “Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off”. *American Economic Review* 105.10, pp. 2986–3029.
Poterba and Sinai (2008). “Tax expenditures for owner-occupied housing: Deductions for property taxes and mortgage interest and the exclusion of imputed rental income”. *American Economic Review* 98.2, pp. 84–89.
Schmieder and Von Wachter (2016). “The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation”. *Annual Review of Economics* 8.1, pp. 547–581.