# Risk Protection and Redistribution in the Design of Social Insurance

Rory McGee and Martin O'Connell*

April 2026

**Abstract**

We study how heterogeneity in unemployment risk and exposure to income loss 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 for expanding U.S. UI is highest among lower-income households, who are subsidized by others under the current system. For a $1 increase in UI benefit expenditure, social resource-reallocation gains are $0.48 from risk protection and $0.19 from cross-subsidization, while incentive costs are $0.75. Of this $0.67 resource-reallocation gain, 57% reflects efficiency gains, 19% cross-income redistribution and 24% within-income redistribution.

**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, Hamish Low, 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 households 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 percent of federal spending—about $2.4 trillion—is devoted to such programs.[1]

Designing effective programs involves a fundamental trade-off between providing insurance and limiting behavioral distortions. The canonical sufficient-statistic approach summarizes this trade-off using a small set of empirical moments—consumption changes that capture the welfare gain from insurance and behavioral elasticities that summarize incentive costs (e.g., Baily 1978; Chetty 2006; Chetty and Finkelstein 2013). In standard implementations, the welfare gain from insurance is measured using the average consumption gap between "good" and "bad" states of the world, while the associated incentive cost is summarized by a single behavioral elasticity.

While the benchmark approach offers powerful insights, it abstracts from heterogeneity in both the value of insurance and the cost of behavioral distortions. In practice, individuals differ widely in their exposure to shocks, ability to smooth consumption (Blundell et al. 2008), and responsiveness to incentives (Chetty 2008; Attanasio et al. 2018). Moreover, because taxes and benefits are rarely fully risk-adjusted, social insurance programs often involve implicit cross-subsidization. As a result, evaluating changes in program generosity using average responses alone may miss the welfare consequences of the way policy redistributes resources across individuals.

This paper studies the value of expanding unemployment insurance (UI) generosity in the United States, emphasizing how heterogeneity in unemployment risk and unemployment-induced consumption declines shapes the social value of reforms. We provide new evidence on individual willingness-to-pay (WTP) for UI expansions and recover the distribution of WTP across workers. Using these estimates, we decompose reform impacts into (i) the social value of resource reallocation—through risk protection and cross-subsidization—and (ii) incentive costs captured by fiscal externalities.

We use data from the Panel Study of Income Dynamics (PSID) from 1999–2019. This nationally representative panel provides detailed information on labor market outcomes, demographics, and household consumption expenditures. To motivate our analysis, we document that workers who experience an unemployment spell

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

have lower equivalized household consumption, earnings, and household income even *before* job loss than those who remain continuously employed. Moreover, even conditional on income, workers who later become unemployed exhibit lower *in-work* consumption. These patterns highlight the redistributive role of UI—both across and within income groups.

Our worker-level measure of WTP for UI reform has two components. The first is the individual's *personalized insurance value*: the value of a marginal transfer of resources from the employment state to the unemployment state. The second is their *net fiscal position*: the extent to which they are, ex ante, net contributors to or beneficiaries of the reform. Insurance values vary with the severity of income loss in unemployment and with self-insurance capacity. Net fiscal positions reflect the gap between actuarially fair contributions to the reform and the implicit premium embedded in non-individualized taxes and benefits, giving rise to cross-subsidization in the form of expected net transfers across individuals.

To measure the insurance component of WTP, we extend the "consumption-based" approach to valuing social insurance developed by Gruber (1997) in two ways. First, we provide new evidence on the distribution of unemployment-induced consumption declines across workers, rather than focusing only on average effects. Second, we allow for unobserved heterogeneity in both unemployment risk and consumption responses, as well as correlation between the two.

We adopt the finite-mixture approach of Lewis et al. (2026) (see also Boehm et al. 2025), which uncovers heterogeneity in consumption responses across latent groups without requiring strong assumptions about which observable characteristics drive that variation. We combine this model for consumption responses with a model of unemployment risk in a joint finite-mixture framework. Unemployment risk is estimated using a finite mixture of logistic regressions that allow for demographic and socioeconomic characteristics to shape baseline unemployment rates (see also Anderson and Meyer 2006). The joint structure yields, for each worker, a posterior distribution over latent risk–consumption types, which we use to construct both unemployment-induced consumption declines and unemployment probabilities—even for workers we never observe experiencing unemployment.

We estimate an average unemployment-induced consumption decline of about 13.1%, broadly in line with prior studies (e.g., Gruber 1997; Hendren 2017; Kroft and Notowidigdo 2016), but find substantial heterogeneity reflecting differences in self-insurance capacity: for example, the 95th-percentile decline is roughly twice the median. We also recover large and systematic differences in unemployment risk by earnings, age, education, and race. Finally, workers who are more likely to become unemployed also tend to experience larger consumption declines conditional on unemployment, so many workers face both high unemployment risk and large consumption losses.

The cross-subsidization component of WTP depends on how reform design interacts with heterogeneity in unemployment risk. Using the estimated joint distribution of risk and consumption types, together with workers' earnings, we compute individual risk weights implied by alternative UI reforms. We compare a natural expansion of the U.S. UI system—which replaces a fraction of lost earnings up to a cap—with reforms that expand entitlement either by a flat amount or in proportion to lost earnings with no cap.

We find that WTP for expanding the generosity of the current U.S. UI system is highly heterogeneous. It is positive at the bottom of the income distribution but declines steeply with household income, becoming negative just above the bottom quintile. This pattern arises even though all workers value actuarially fair insurance. The decline reflects that, for most workers, the expected net tax burden of the reform exceeds its insurance value: on average, workers above the bottom quintile are net contributors who cross-subsidize others. Low-income workers benefit because they face higher unemployment risk—an effect strong enough to outweigh the link between benefit entitlements and earnings. A flat benefit reform tilts the WTP distribution further toward 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 within income groups. For example, within the top income decile, the standard deviation of WTP for expanding the current system exceeds half the gap between median WTP at the bottom and top of the income distribution. These reform-specific patterns of winners and losers—both across and within income groups—shape the social value of alternative UI reforms.

We next consider the social value of a budget-balanced expansion of UI generosity. Since this requires interpersonal comparisons, the evaluation depends on social preferences. In our baseline implementation, we use welfare weights that (i) value redistribution across income groups at the marginal cost of achieving equivalent redistribution through the income tax system (Hendren 2020), and (ii) place greater value on within-income-group redistribution toward individuals with lower long-run equivalized household consumption, a natural measure of lifetime income and economic well-being (Poterba 1989; Meyer and Sullivan 2003). This formulation reflects the idea that broad redistribution across income groups can, in principle, be implemented through the income-tax-and-transfer system. By contrast, UI also reallocates resources within income groups, toward households that face high unemployment risk and have low in-work consumption or limited self-insurance capacity—dimensions that are difficult to observe and target in the tax code. Our framework is designed to distinguish the redistributive value of UI along these margins from the component that could instead be achieved through changes to the income tax schedule.

We show that the social value of expanding UI can be decomposed into three components. The first is risk protection: a risk- and welfare-weighted average of individual insurance values, scaled by an aggregate welfare weight. The individual weights capture heterogeneity in the marginal social value of surplus among program beneficiaries, while the aggregate weight captures the average marginal value of transferring surplus to beneficiaries relative to the workers who finance the reform. We estimate that the risk-protection value from a budget-balanced expansion of the current UI program is $0.48 per dollar of increased benefit expenditure.

The second component captures the social value of cross-subsidization, which arises from heterogeneity in net incidence rather than from insurance value itself. In representative-agent frameworks, this component is absent; in our setting, it can generate welfare gains (or losses) even when individuals are fully insured through other means. For the current UI program, we estimate that the cross-subsidization component associated with increased generosity is $0.19 per dollar of increased benefit expenditure. Together, these two terms imply total resource-reallocation gains of $0.67. To understand what drives this value, we compare our baseline results to two benchmarks that sequentially remove within- and across-income variation in individual welfare weights. This decomposition implies that about 57% of the resource-reallocation gain reflects efficiency improvements from better insurance, 19% reflects redistribution across income groups, and 24% reflects redistribution within income groups.

The third component is the fiscal externality arising from behavioral responses to program generosity. In the canonical model, this incentive cost is summarized by an aggregate elasticity. In our setting, it depends on how behavioral elasticities covary with individuals' risk and net fiscal positions (i.e., the gap between taxes paid and benefits received across states). We calibrate these heterogeneous elasticities using evidence from Chetty (2008). We estimate a direct fiscal externality of $0.27 per dollar of increased benefit expenditure from higher UI payments induced by behavioral responses, plus an additional $0.48 from reduced income tax receipts. Thus, the total incentive cost of $0.75 per dollar slightly exceeds the $0.67 resource-reallocation gains from risk protection and cross-subsidization.

Our measure of the social value of changing UI generosity depends critically on reform design. Under a flat expansion—where UI entitlements increase uniformly across workers—the value of resource reallocation rises to \$0.78 per dollar of increased benefit expenditure, because this reform directs a larger share of resources toward households that are less well-off and more exposed to unemployment risk. In this case, the welfare gains from increased generosity outweigh the associated incentive costs. By contrast, the standard sufficient-statistics implementation corresponds to a flat reform under social preferences that are indifferent to how surplus is distributed, yielding a much lower estimated resource-reallocation value of \$0.39. We therefore show that abstracting from redistribution across heterogeneous individuals meaningfully alters the conclusions of UI policy analysis. More broadly, these differences across reforms highlight that the social value of UI depends on program design and is therefore likely to vary across countries (Spinnewijn 2020).

We contribute to a literature that uses sufficient-statistics approaches to study the insurance-incentive trade-off in UI design (e.g., Schmieder et al. 2012; Kolsrud et al. 2018; Landais and Spinnewijn 2021).[2] Our contribution is to extend the sufficient-statistics framework for UI beyond the representative-agent setting by combining microdata on consumption and unemployment risk to recover individual-level WTP and to show how heterogeneity shapes the decomposition of the social value of UI reforms across workers. We also build on recent work that measures the value of social transfers, including disability insurance (Deshpande and Lockwood 2022), public pensions (Kolsrud et al. 2024), and welfare programs (Rafkin et al. 2023).

Our work also relates to a literature that examines how the correlation between risk and ability shapes the joint design of income taxes and social insurance (e.g., Blomqvist and Horn 1984; Rochet 1991; Cremer and Pestieau 1996; Boadway et al. 2006). It also connects to research on how involuntary unemployment (Kroft et al. 2020) and the structure of non-linear UI benefit schedules (Ferey 2022) influence optimal tax design. This work emphasizes that redistribution across income or ability types can, in principle, be implemented through the non-linear income tax and transfer system, while social insurance instruments such as unemployment insurance play a distinct role by providing state-contingent transfers within types and across employment states. 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 the paper is organized as follows. Section 2 outlines our measure of individual willingness-to-pay for increased UI generosity. Section 3 introduces the

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

data and presents preliminary evidence on negative selection into unemployment. Section 4 describes our empirical approach and presents the estimates. Section 5 develops the normative framework and reports the quantitative welfare analysis. Section 6 concludes.

## 2. Individual Value of a Reform

In this section, we derive an expression for individual willingness-to-pay (WTP) for an increase in the generosity of a social insurance program. We present a simple framework that captures the salient features of unemployment insurance, as well as other programs such as workers' compensation, disability insurance, health insurance, and insurance against natural disasters. The sufficient-statistics expressions we derive extend to a much richer dynamic environment (Chetty 2006; see Appendix A.1).

We build on the canonical sufficient-statistics setup (see Chetty and Finkelstein 2013) by allowing for heterogeneous individuals facing risk. Individual WTP is a key input for evaluating the social value of reform, which in our framework depends both on how WTP is aggregated across individuals and on the incentive costs generated by behavioral responses (see Section 5).

### 2.1. Theory

***Individual's problem.*** There is a unit continuum of individuals indexed by $i$. Each individual faces uncertainty over two states: high ($h$) and low ($l$), for instance corresponding to being employed and unemployed. Individuals may differ in their preferences, incomes, and the costs associated with mitigating risk. Let $y_i$ and $z_i$ denote the individual's resources in the high and low state, respectively. Taxes and benefits are indexed by high-state income $y_i$: 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.[3]

Consumption in the high state is $c_i^h = y_i - \mathcal{T}(y_i)$, 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 $u_i^s(\cdot)$ is increasing, concave, and twice continuously differentiable. Let $\Delta c_i \equiv c_i^h - c_i^l$ denote the gap in consumption between the high and low states. For most individuals, 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 effort

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

so that $e$ equals the probability of being in the high state (and therefore $e \in [0, 1]$). Effort entails cost $\psi_i(e)$, where $\psi_i(\cdot)$ is increasing and convex. The individual solves:

$$V_i = \max_{e \in [0,1]} \; e\, u_i^h(c_i^h) + (1 - e)\, u_i^l(c_i^l) - \psi_i(e), \tag{1}$$

where $V_i$ is the maximized value of their expected utility. Individual $i$'s optimal effort choice, $e_i$, satisfies the first-order condition: $u_i^h(c_i^h) - u_i^l(c_i^l) = \psi_i'(e_i)$, i.e., the individual chooses effort so that the utility gain from the high state relative to the low state equals the marginal cost of effort.

***Heterogeneity.*** The index $i$ captures arbitrary heterogeneity from several sources. First, differences in earning ability—for example due to variation in innate ability, human capital, skills, or labor market opportunities—are reflected in $y_i$, the individual's high-state resources. Low-state resources $z_i$ additionally reflect the individual's capacity to self-insure, for example through private savings or spousal insurance. As a result, variation in resources across states reflect both the direct financial loss from entering the low state and differences in self-insurance capacity. Together with the tax and benefit system, these factors generate heterogeneity in the consumption gap $\Delta c_i$.

Second, heterogeneity in attitudes toward work and employment opportunities (in the unemployment or disability insurance context), or in the cost of abatement investments (in health or natural-disaster insurance contexts), is captured by the effort-cost function $\psi_i(\cdot)$.

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

Each of these factors affects the individual's optimal effort choice and therefore generates 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 state-contingent consumption possibilities shift according to:

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

i.e., the consumption possibilities in the low and high states change by $\phi_{\mathcal{B}}(y_i)db$ and $-\phi_{\mathcal{T}}(y_i)d\tau$, respectively, where $\phi_{\mathcal{B}}(y_i)$ and $\phi_{\mathcal{T}}(y_i)$ are positive functions of high-state income $y_i$. While these functions allow for reforms that vary arbitrarily with income, in practice reforms are typically simple functions—for example, a flat benefit increase ($\phi_{\mathcal{B}} = 1$) or a 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 individual willingness-to-pay because it governs both who receives additional resources in each state and who finances them.

***Willingness-to-pay.*** To define the individual-level WTP for an increase in social insurance generosity, we consider the reform

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

This reform raises benefits and adjusts the tax schedule just enough to offset the resulting mechanical budgetary cost (holding behavior fixed).

Let $\frac{dV_i}{d\theta}$ denote the impact of this reform on individual $i$'s expected utility. We define individual $i$'s willingness-to-pay for this reform as:

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

[4] Focusing on marginal reforms has the advantage that, by the envelope theorem, the first-order impact of the reform on expected utility is fully captured by the change in the state-contingent budget constraint, evaluated at the pre-reform optimum. Throughout, when we write the effect of the reform in terms of $dc_i^l$ and $dc_i^h$, this should be understood as the induced change in state-contingent consumption possibilities rather than as the policy directly choosing consumption. 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).

The scaling in the first line expresses the utility gain from the reform per $1 increase in individual expected benefit payments, measured in units of the marginal utility 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 using their 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, tax adjustments would be individualized so that the individual's expected tax increase, $e_i \, \phi_{\mathcal{T}}(y_i)\, d\tau$, finances their expected benefit increase, $(1 - e_i)\, \phi_{\mathcal{B}}(y_i)\, db$. In practice, however, the actual reform typically departs from actuarial fairness. 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 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 is a net beneficiary or net contributor relative to a fair-pricing benchmark.

***Discussion.*** Our willingness-to-pay decomposition is related to Finkelstein et al. (2019), who show that 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 (cross-subsidization) arises from departures from actuarial fairness embedded in the structure of the tax and benefit reform.

An alternative interpretation of our decomposition is that the cross-subsidization component corresponds to the expected net transfer from the reform when it is valued at a common marginal utility across states, while the insurance component captures the additional value that arises because marginal utility is higher in the low state than in the high state. In this sense, the insurance term measures how much the *ex ante*, risk-adjusted value of the reform exceeds its ex post value valued at a common marginal utility. This interpretation aligns with Lieber and Lockwood's

(2019) decomposition of the individual-level value of an in-kind transfer.[5]

### 2.2. 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., ${u^{s}}'_i(\cdot) = u'_i(\cdot)$ for $s \in \{l, h\}$) and relative risk aversion at high-state consumption is homogeneous across individuals (i.e., $-\frac{u''_i(c^h_i)c^h_i}{u'_i(c^h_i)} = \gamma$).*

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 Taylor approximation to the utility function around $c^h_i$ yields:

$$\frac{u'_i(c^l_i) - u'_i(c^h_i)}{u'_i(c^h_i)} \approx \gamma \frac{\Delta c_i}{c^h_i}, \tag{4}$$

so that the individual insurance component of WTP can be expressed as the product of the percentage consumption decline between the high and low states and the coefficient of relative risk aversion at $c^h_i$. Our empirical strategy estimates the unconditional distribution of these consumption declines across the population.

Cross-subsidization depends both on state probabilities and on the structure of the reform under consideration. In our application, we estimate how unemployment probabilities vary across individuals and consider a range of realistic reforms that differ in how benefit and tax adjustments vary with income.

***Discussion.*** Our implementation adopts a sufficient-statistics approach, using estimates of individual insurance value and unemployment risk to construct WTP. We estimate heterogeneous consumption responses to job loss using panel data on earnings and nondurable consumption, building on Gruber (1997) and recent applications in UI (e.g., Kolsrud et al. 2018; Ganong and Noel 2019) and other social

[5] Both Finkelstein et al. (2019) and Lieber and Lockwood (2019) consider settings with potentially many states 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.

insurance programs (e.g., Meyer and Mok 2019; Kolsrud et al. 2024). Under Assumption 1, these estimates map into the insurance component of WTP. We implement this using a finite-mixture regression that recovers a discrete distribution of unemployment-induced consumption declines across latent types.

We estimate unemployment risk using a finite mixture of logistic regressions for unemployment incidence as a function of demographics and a measure of permanent income constructed from a multi-year earnings history. We recover latent "risk–consumption" types, and the correlation between them, from a joint mixture model for unemployment and consumption dynamics. We interpret the implied unemployment probabilities as summarizing heterogeneity in exposure to unemployment risk in our sample, rather than as a full structural model of the joint dynamics of wages, assets, and search.

Other approaches to valuing social insurance programs—based on search responses (Chetty 2008; Shimer and Werning 2008), differences in marginal propensities to consume across states, or revealed preference using supplemental UI purchases (Landais and Spinnewijn 2021)—recover insurance value under different assumptions and institutional and data requirements. In principle, any such method for recovering the insurance value of UI could be used in place of, or alongside, our consumption-based approach to implement equation (3) and the subsequent normative analysis. In this paper, we focus on the consumption-based method because it can be applied consistently in the U.S. setting we study.

## 3. Data and Setting

Our empirical analysis focuses on unemployment insurance (UI). In the United States, the federal government sets broad guidelines on UI eligibility and benefit generosity, but programs are administered at the state level, with states determining benefit duration, generosity, and eligibility criteria (see Von Wachter 2019 for a discussion). Benefit levels and entitlement are computed from workers' employment and earnings histories. Outside of temporary federal extensions during recessions, most states offer benefits for up to 26 weeks. Eligible workers receive payments in proportion to past earnings, up to a cap.

The U.S. UI system is financed through payroll taxes that are statutorily levied 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 more like a per-worker charge than a proportional earnings tax. A distinctive feature of U.S. UI finance is that tax rates vary across employers due to partial experience rating (see Guo and Johnston 2021 for discussion). Recent evidence suggests that firms do not pass the firm-specific component of UI taxes through to wages (Guo 2024). Accordingly, in our quantification of a UI expansion, we assume that only the common component of UI taxes is incident on wages.[6]

***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 and labor market outcomes including unemployment, earnings, and wages. We use data from the 1999–2019 biennial PSID waves, which includes comprehensive information on consumption expenditures.[7]

We drop households for which the reference person does not report educational attainment and restrict the sample to reference persons aged 23–60, excluding individuals whose labor supply is likely driven by education or retirement decisions. To proxy for UI eligibility, we omit households where the reference person never reports being employed, as well as observations immediately following a spell out of the labor force. This yields 55,257 observations (14,995 families observed for an average of 3.7 waves). Of these, 3,283 observations correspond to the reference person being unemployed (2,178 families for an average of 1.5 waves). For families that experience changes in the number of adults (e.g., divorce, marriage, or death), we construct panel identifiers for demographically stable units and omit the year of the change. Throughout, we focus on the labor market status of the reference person, whom we refer to as the "worker."

***Key variables.*** In our main analysis, we define current unemployment based on the employment status reported by the worker at the time of the survey interview. We measure worker earnings using reported labor earnings when employed, and define household labor income as the sum of worker and spousal earnings. For each worker, we also construct a measure of permanent income based on the worker fixed effect in a log earnings regression, conditional on life-cycle earning growth.

We use a comprehensive measure of expenditures on nondurable consumption

[6] 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 affect layoffs and subsequent UI claims (and denials).

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

and services, excluding spending on health care and housing.[8] 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 WTP for UI. In Section 5, we undertake a normative analysis, comparing a welfare-weighted aggregation of WTP to the program's incentive costs. To motivate this analysis, we begin by presenting descriptive evidence on 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 the direct 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 how household resources during periods of employment differ between individuals who ever experience unemployment and those who never do. By restricting attention to observations in which the worker is employed, this analysis avoids mechanically induced reductions in earnings and associated consumption responses during unemployment.

Restricting the sample to observations in which the worker is employed, we estimate:

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

where $y_{i,t}$ is a measure of household resources; $D_i$ is an indicator for whether worker $i$ (the reference person) ever experiences an observed unemployment spell;[9] and $X_{i,t}$ includes a cubic in the worker's age, household composition indicators, and year fixed effects; in some specifications, it additionally includes a cubic in household income[10] The coefficient of interest $\delta$ captures the difference in household resources

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

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

[10]The presence of negative selection is robust to additionally controlling for education and race (see Appendix D.1).

associated with ever experiencing unemployment.

Figure 1 reports estimates of δ and associated 95% confidence intervals for four measures of resources: household nondurable consumption, the reference person's labor income, household labor income, and worker permanent income. For each outcome, we report three baseline specifications: (i) without controls, (ii) with controls, and (iii) excluding post-unemployment periods to eliminate potential labor-market scarring.[11] Across these specifications, households whose reference person ever experiences unemployment have lower consumption, individual income, and household income *while employed*, as well as lower permanent income, than households whose reference person is never unemployed.

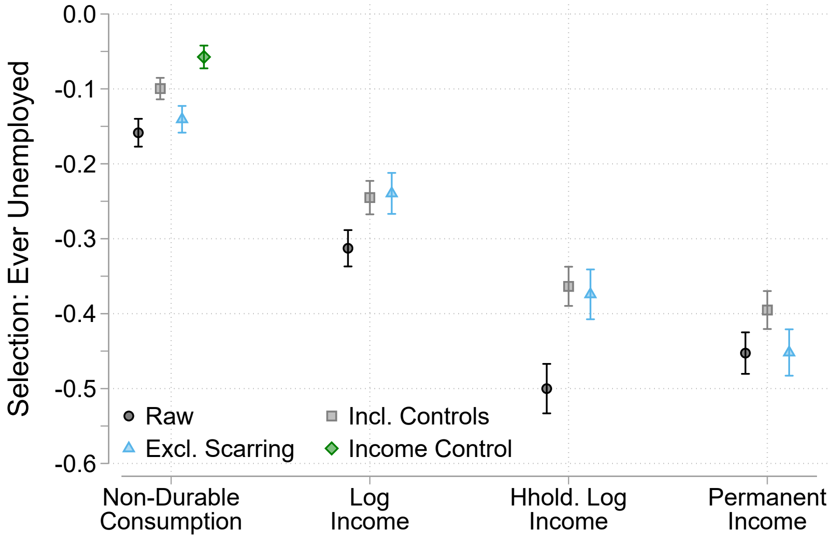


FIGURE 1. Negative Selection Among the Unemployed

Notes: All regressions drop periods during which the reference person is unemployed. For all outcomes, we report the raw, controls, and excl. scarring specifications. The raw specification includes no controls. The remaining specifications control for a cubic in age, indicators for family size and marital status, and year fixed effects. The excl. scarring specification further restricts the sample to periods before the first observed unemployment spell. For nondurable consumption only, we also report an "income controls" specification that additionally controls for a cubic in household income.

To the extent that social insurance programs reallocate resources across income groups, redistribution shapes their social value. However, redistribution across income levels can, in principle, be achieved or offset through adjustments to the tax-and-transfer system. By contrast, reallocating resources within income groups—across households with the same current income but different levels of consumption—is much harder to implement through standard tax-and-transfer

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

policies. Moreover, because consumption may better reflect permanent income and longer-run economic well-being, redistribution based on consumption differences may be especially valuable. For household nondurable consumption, the income-controls specification (green markers in Figure 1) shows that even conditional on a flexible function of household income, households that experience unemployment have lower consumption than those that do not.

## 4. Measuring Willingness-to-Pay

In this section, we outline how we measure heterogeneity in unemployment risk and unemployment-induced consumption declines, which, along with the coefficient of relative risk aversion, determine workers' exposure to income risk. 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. Empirical Framework

We allow for discrete unobserved heterogeneity in both unemployment risk and unemployment-induced consumption declines. Specifically, we assume that each worker $i$ belongs to one of $H$ unemployment-risk types and one of $K$ consumption types, giving rise to $H \times K$ latent joint types. Risk types capture differences in unemployment hazards, while consumption types capture differences in the magnitude of the consumption decline following job loss. We estimate this structure using a finite-mixture framework and recover the joint distribution of risk and consumption types by maximum likelihood. For expositional clarity, we first describe the unemployment-risk block and then the consumption block; the joint likelihood and estimation algorithm are detailed in Appendix C. In our baseline specification we set $H = 2$ and $K = 3$, yielding six joint types; we select $(H, K)$ using the Bayesian Information Criterion.

Our use of a finite-mixture approach is guided by the need to construct worker-level measures of WTP, which depend on worker-level estimates of unemployment-induced consumption declines—including for workers never observed unemployed—and on their correlation with unemployment risk. The joint mixture model assigns each worker posterior probabilities over latent risk–consumption types based on their history of consumption, employment and observables, which we use to construct individual unemployment-induced consumption declines and to characterize the joint distribution of unemployment risk and consumption drops across the population.

***Unemployment risk.*** Patterns of program cross-subsidization depend on worker-level risk weights. To estimate these, we use a statistical model of unemployment incidence that allows for both observable and latent heterogeneity in risk. We model unemployment risk as a function of a cubic in age, indicators for gender, marital status, education, and race, and a cubic in permanent income, along with discrete unobserved risk types.[12]

Specifically, we assume that each worker belongs to one of $H$ unemployment-risk types $h \in \{1, \ldots, H\}$ with population share $\pi(h)$. Conditional on type $h$, unemployment at time $t$ follows a logit model

$$\Pr\left(U_{i,t} = 1 \mid h(i) = h, X_{i,t}^{U}\right) = g\left(\eta_0^h + X_{i,t}^{U\prime}\gamma\right),$$

where $g(\cdot)$ is the logistic function, $\eta_0^h$ is a type-specific intercept, and $\gamma$ is a vector of coefficients on observables $X_{i,t}^{U}$. We assume that, conditional on risk type $h$ and observables $X_{i,t}^{U}$, unemployment shocks are independent across time and across workers. Identification of risk types uses systematic differences in unemployment incidence across workers with similar observables.

***Consumption declines.*** The standard approach to recovering consumption declines (following 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:

$$\begin{aligned}\Delta_{i,t}^{FD} &\equiv \ln c_{i,t} - \ln c_{i,t-2} \\ &= \delta_0 + \delta_1 U_{i,t} + X_{i,t}^{C\prime}\beta + \varepsilon_{i,t},\end{aligned} \tag{6}$$

where $U_{i,t}$ is an indicator variable denoting whether the worker is unemployed in period $t$ and $X_{i,t}^{C}$ are demographic controls. This first-difference specification allows for arbitrary permanent individual heterogeneity in consumption levels. The parameter of interest, $\delta_1$, measures the average proportional consumption decline that results from becoming unemployed. The identification assumption underpinning this strategy is absent unemployment, workers who become unemployed would have experienced the same consumption growth as those who remain employed.

We estimate the *distribution* of consumption declines using a finite-mixture

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

approximation to the underlying distribution (e.g., Heckman and Singer 1984; Keane and Wolpin 1997). We build on Lewis et al. (2026), by adopting their finite-mixture regression for consumption responses and embedding it in a joint mixture model with unemployment risk. We assume each worker belongs to one of a finite number of latent types and augment the specification in equation (6) as follows:

$$\Delta^{FD}_{i,t} = \sum_{k=1}^{K} \mathbb{1}\left[k(i) = k\right]\left(\delta_0^k + \delta_1^k U_{i,t}\right) + X^{C\,\prime}_{i,t}\beta + \varepsilon_{i,t}. \tag{7}$$

We include group-specific intercepts $\delta_0^k$, so that $\delta_1^k$ can be interpreted as the consumption decline associated with unemployment for group $k$. We assume that, conditional on consumption type $k$, observables, and unemployment status, shocks $\varepsilon_{i,t}$ are independent over time and satisfy $\mathbb{E}[\varepsilon_{i,t} \mid k(i) = k, X^C_{i,t}, U_{i,t}] = 0$, with type-specific variance $\sigma_k^2$.

To implement equation (7), we use a Gaussian mixture linear regression (Quandt 1972), which assumes normally distributed errors $\varepsilon_{i,t}$. The key identifying assumption is that, conditional on latent type $k$, the evolution of consumption for employed workers provides a valid counterfactual for the trend in consumption that unemployed workers would have experienced had they remained employed. Intuitively, latent types are identified by grouping households with similar distributions of consumption growth in each state.

As discussed in Lewis et al. (2026), this mixture regression can be viewed as a form of clustering regression which jointly (i) groups households together with similar latent consumption dynamics during employment and unemployment and (ii) estimates the consumption decline within each group. This provides a useful matching interpretation for how we impute counterfactual consumption declines to workers never observed unemployed. Intuitively, the approach resembles the following steps. First, we group households observed in both states based on their consumption growth when employed and their consumption decline when unemployed. We then "match" always-employed households to those with similar consumption growth when employed, conditional on observables, and assign them the matched households' consumption decline. In practice, we estimate group assignments and parameters jointly. The marginal posterior probability of each consumption type $k$, $p_i(k) \equiv \sum_{h=1}^{H} p_i(k,h)$ serves as a convex imputation weight.[13]

***Estimation.*** Estimation proceeds in two stages. In the first stage, we estimate the parameters of the unemployment-risk mixture,

$$\chi_U = \{\pi(h), \eta_0^h, \gamma\}_{h=1}^{H}$$

by maximum likelihood using workers' employment status histories. In the second stage, conditional on these first-stage estimates, we estimate the consumption-decline parameters and the joint type probabilities using an expectation–maximization algorithm. Let

$$\chi_C = \left\{\delta_0^k, \delta_1^k, \sigma_k^2, \beta\right\}_{k=1}^{K}$$

denote the consumption-block parameters, and let $\pi(h,k)$ denote the joint population share of consumption type $k$ and risk type $h$. The joint type probabilities $\pi(h,k)$ are identified from the joint distribution of unemployment histories and consumption growth paths, conditional on observables. Intuitively, when workers who experience frequent unemployment spells also exhibit large consumption declines, the model assigns higher posterior weight to high-risk/high-decline joint types.

For each worker, we then compute posterior probabilities $\hat{p}_i(k,h)$ of belonging to each latent type. We use these posteriors to construct worker-level employment probabilities and unemployment-induced consumption declines,

$$\hat{e}_{i,t} = 1 - \sum_{h=1}^{H} \sum_{k=1}^{K} \hat{p}_i(k,h)\ \Pr\left(U_{i,t} = 1 \mid h(i) = h, X_{i,t}^U; \hat{\chi}_U\right), \quad \widehat{(\Delta c/c)}_i = \sum_{h=1}^{H} \sum_{k=1}^{K} \hat{p}_i(k,h)\, \hat{\delta}_1^k.$$

Appendix C describes the likelihood and algorithm in detail.

### 4.2. Estimates

#### Unemployment risk

Table 1 reports estimates from our finite-mixture logit model. We report the type-specific probabilities of unemployment, averaged over the empirical distribution of

[13] 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 while employed) a consumption decline that is a weighted average of these groups. 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 et al. 2008). In principle, latent consumption risk may vary over time, and our approach could be extended to allow for a regime-switching model.

covariates, along with the estimated population share of each type. The estimates reveal substantial latent heterogeneity in unemployment risk. Roughly three quarters of workers belong to the low-risk group with an average unemployment probability of about 2.6%, while the remaining quarter form a high-risk group facing an average unemployment probability of close to 14%.

TABLE 1. Predicted Probabilities of Unemployment and Type Shares

| | Overall | Group 1 | Group 2 |
|---|---|---|---|
| $\bar{P}\left(U_{i,t}=1\right)$ | 0.053 | 0.026 | 0.140 |
| | (0.001) | (0.005) | (0.023) |
| Share ($\widehat{\pi}(h)$) | | 0.759 | 0.241 |
| | | (0.066) | (0.066) |

Notes: $\bar{P}(U_{i,t}=1)$ denotes the sample-averaged predicted probability of unemployment, computed from the finite-mixture logit model and averaged over the empirical distribution of covariates. $\widehat{\pi}(h)$ denotes the estimated population shares of each unemployment-risk type. The estimation sample contains $N = 50{,}380$ person-wave observations. Bootstrap standard errors based on 400 household-level replications are shown in parentheses. Covariates include a cubic in age, indicators for gender, marital status, education, and race, and a cubic in permanent income.

We report average partial effects with 95% confidence intervals in Figure 2. The largest average partial effects are associated with our measure of permanent income: a one-standard deviation increase is associated with a 5.8 percentage-point decrease in the unemployment probability.

The results show systematic differences in unemployment rates across demographic groups. These differences are best interpreted as the effect of a given variable comparing two workers with similar earnings histories, since permanent income is included as a control. Younger workers are more likely to be unemployed: being ten years younger is associated with a 1.1 percentage point increase in the unemployment probability. 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 2.1 percentage points more likely to be unemployed than White 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 enter our model through differences in the cost of exerting effort (i.e., in $\psi_i(\cdot)$).

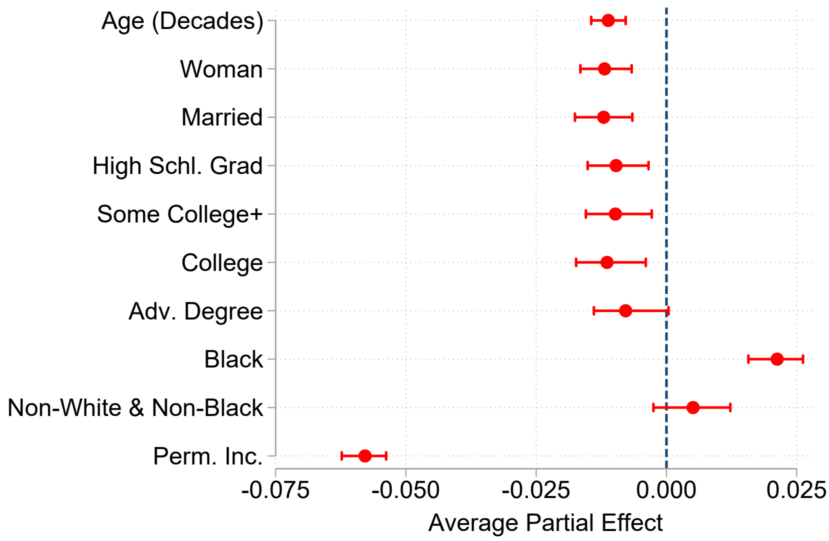


FIGURE 2. Effect of Household Characteristics on Predicted Probability of Unemployment

Notes: We estimate the finite mixture logit model described in the text using unemployment status as the dependent variable, and compute marginal effects of each covariate on unemployment probability by integrating over the empirical distribution of household characteristics and the estimated distribution of risk types. Each market corresponds to the point estimate of the average partial effect (APE), and the horizontal lines span the 95% confidence intervals. The dashed line indicates zero. The estimation sample contains $N$ = 50,380 person-wave observations, and bootstrap standard errors are based on 400 household-level replications.

While the contribution of unobserved heterogeneity is substantial, the estimated average partial effects line up with patterns documented in the broader literature. Prior work finds than unemployment rates are higher among younger workers (e.g., Choi et al. 2015); decline with higher education and other measures of human capital (e.g., Ashenfelter and Ham 1979; Nickell 1979; Cairó and Cajner 2018); and exhibit persistent racial disparities in employment outcomes (e.g., Lang and Lehmann 2012) and callback rates (e.g., Bertrand and Mullainathan 2004). We also find only a modest unemployment gap by gender. However, our estimates pertain to women who are the PSID reference person—a group that is disproportionately single. Consistent with our small gender gap, prior work finds that labor force participation gaps between unmarried men and women are small, (e.g., Borella et al. 2023).

### Consumption declines

Table 2 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 13.1% associated with unemployment.

| | | Mixture Model | | |
|---|---|---|---|---|
| | Homogeneous | Group 1 | Group 2 | Group 3 |
| $\mathbb{1}\left[U_{i,t} = 1\right]$ | −0.131 | −0.263 | −0.113 | −0.069 |
| | (0.012) | (0.067) | (0.026) | (0.025) |
| Share ($\widehat{\pi}(k)$) | | 0.122 | 0.494 | 0.384 |
| | | (0.076) | (0.098) | (0.073) |

TABLE 2. Heterogeneity in Consumption Declines ($\Delta c/c$)

Notes: The first column reports the estimated average consumption decline from the homogeneous specification (equation (6)). The remaining columns report the consumption declines, $\delta_1^k$, and marginal population shares, $\pi(k) = \sum_{h=1}^{H} \pi(h, k)$, from our Gaussian mixture linear regression framework (equation (7)). Standard errors are shown in the parentheses; in the mixture model, they account for uncertainty over both group assignment and parameter estimates conditional on group. The estimation sample contains $N$ = 50,380 person-wave observations, and bootstrap standard errors are based on 400 household-level replications. Covariates include a cubic in age, indicators for year, and the log of the change in family size.

The implied average consumption decline from our mixture approach, aggregating the group-specific consumption declines $\delta_1^k$ weighted by their marginal population shares $\pi(k)$, is slightly smaller at 11.4%.[14] This comparison of averages, however, ignores the large degree of heterogeneity in exposure to unemployment-induced consumption declines that we find. The group-specific consumption declines we estimate range from 26.3% to 6.9%. The probability mass of the group with the smallest decline, 6.9%, exceeds 38%. This group is relatively well insured against unemployment risk. A second group, with probability mass of 49%, has a consumption decline close to the homogeneous estimate. The remaining probability mass is accounted for by a group with a much larger consumption decline of 26.3%.

[14]This discrepancy arises because the two averages correspond to different estimands. In the homogeneous specification, $\widehat{\delta}_1$ estimates the average consumption decline among the unemployed (the average treatment on the treated). In contrast, aggregating the mixture model's group-specific declines using the marginal type shares estimates the population average effect (the average treatment effect), which is possible under the additional assumptions on the data-generating process implied by the mixture structure. When we aggregate individual consumption declines using ex ante unemployment probabilities, we obtain a decline of 13.0%, very close to the homogeneous estimate.

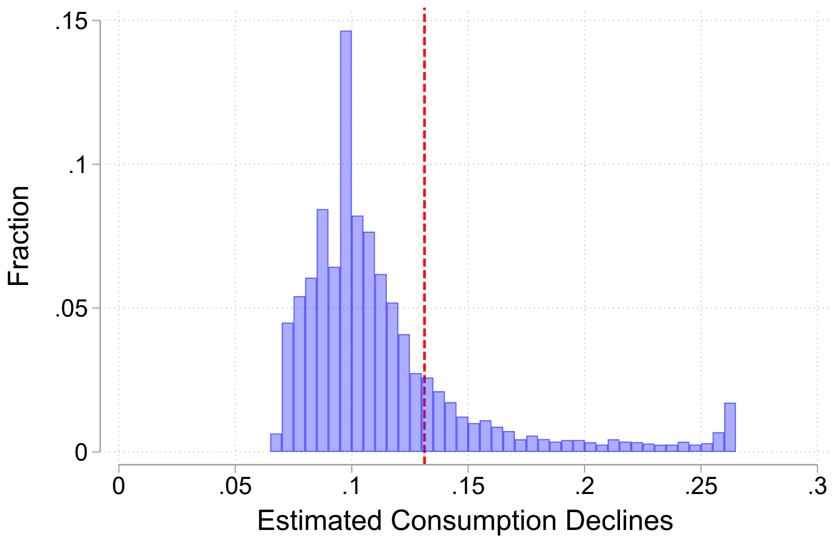


FIGURE 3. Distribution of Estimated Consumption Declines

Notes: The histogram plots worker-level unemployment-induced consumption declines implied by the Gaussian mixture linear regression, constructed by combining estimated parameters with each worker's posterior probabilities of belonging to each consumption type, $p_i(k) = \sum_{h=1}^{H} p_i(h,k)$. For each household, we compute the posterior-weighted average of the type-specific unemployment coefficients $\delta_1^k$. The estimation sample contains $N = 50{,}380$ person-wave observations. The homogeneous estimate (dashed vertical line) is obtained by imposing $K = 1$ in our baseline specification. Appendix D.2 reports bootstrap confidence intervals for quantiles of the worker-level consumption declines.

Figure 3 shows the distribution of posterior-implied consumption declines across workers, further emphasizing substantial cross-sectional heterogeneity in exposure to unemployment risk. Using the homogeneous specification, all households are assigned the same predicted decline (the dashed vertical line). By contrast, we find that 81% of households exhibit consumption declines smaller than this estimate, while a long right tail experiences much larger declines. This variation reflects broad differences in households' ability to self-insure against unemployment risk.

***Comparison to existing homogeneous estimates.*** 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 between 7 and 10%. There are three differences between our approach and prior work that account for this.

First, we measure consumption declines by comparing consumption when unemployed with in-work consumption two years earlier. 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. Our results are robust to alternative consumption measures (see Appendix D.3). Third, our sample differs from those used in prior work and is drawn from different years. 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.[15]

***Threats to identification.*** We assess the key identifying assumption that, conditional on type, consumption growth for workers who remain employed provides a valid counterfactual for workers who become unemployed. Relative to a homogeneous specification, conditioning on latent type may help mitigate endogenous selection into unemployment by making the counterfactual comparison within groups of workers who are more comparable in persistent consumption dynamics and unemployment risk. Following Hendren (2017), we estimate lead–lag specifications that replace the contemporaneous unemployment indicator with indicators for unemployment $k$ periods away; the homogeneous and type-specific profiles show little evidence of differential pre-trends prior to job loss (Appendix D.4). In addition, we exploit PSID information on reasons for job separation to construct an indicator for separations due to plant closure or firm shutdown, and use this as an instrument for our baseline unemployment indicator. The resulting estimate is similar to our baseline and not statistically distinguishable from it (Appendix D.5), suggesting that selection into unemployment and violations of parallel trends are unlikely to be major drivers of our estimated consumption declines. As an additional validation of the mixture specification, we compare the distribution of consumption drops implied by the finite-mixture model for unemployed workers in a training sample to quantile-regression estimates in an independent holdout sample; the two align closely across most of the distribution, with differences only in the extreme lower tail (Appendix D.6).

***Robustness to PSID design.*** A limitation of the PSID's biennial structure is that it may under-represent short unemployment spells and the associated consumption responses. In our baseline specification, we classify a worker as unemployed at time $t$ if they are employed at $t-2$ but unemployed at $t$, so that $\Delta_{i,t}^{FD}$ compares pre- and

[15] Similar to our analysis, East and Kuka (2015) find 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 (see Appendix Table B.1 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.

post-unemployment consumption. Unemployment spells that start and end between these interviews appear as employment at both dates and therefore do not enter our "unemployed" group, tilting the sample toward longer spells for which consumption declines may be larger. Using post-2003 questions on whether the worker experienced any unemployment in the previous 12 or 24 months, we compare our homogeneous estimate to specifications that use alternative sample restrictions based on reported unemployment timing (see Appendix D.7). The point estimates are similar, with samples that exclude more reported unemployment outside the interview window yielding slightly smaller (but statistically indistinguishable) consumption drops. This suggests that any bias from the biennial sampling frame is modest and, if anything, would lead us to slightly overstate the insurance value of UI.

***Relation to evidence on heterogeneous consumption responses.*** Our finding of substantial heterogeneity in 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 declines at job loss. They also document smaller, yet substantial, declines among households that are likely to be unconstrained. Ganong and Noel (2019) 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 et al. 2023) shows that households offset income losses from job loss primarily by decumulating liquid savings, but also through added worker effects.

Recent work provides complementary evidence on the magnitude and heterogeneity of consumption responses during unemployment. Patterson (2023) and Colarieti et al. (2024) estimate average marginal propensities to consume out of unemployment income shocks between 0.5 and 0.6 and document variation in these responses across households. Combining our homogeneous estimate of the average consumption decline with our estimate of the corresponding earnings loss implies an MPC of about 0.6. Thus, while our framework focuses unemployment-induced consumption declines rather than directly on MPCs, the implied average responsiveness of consumption to income losses is similar to that found in these recent studies.

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. We examine how these estimated consumption declines correlate with proxies for a household's capacity to self-insure and find larger consumption declines among households less likely to benefit from added worker effects (those who do not live with a partner) and among those with low liquid assets (see Appendix D.8).

#### Correlation between risk and consumption types

Our joint-type framework allows us to study how unemployment risk and unemployment-induced consumption declines covary in the population. Table 3 reports the estimated joint distribution of risk and consumption types, $\pi(h, k)$, together with the conditional share of each consumption type within each risk type. We find that the high-risk type accounts for about one-quarter of workers and is almost twice as likely to belong to the high-drop consumption group as the low-risk type (19% vs. 10%), and substantially less likely to be in the low-drop group (27% vs. 42%). This pattern implies that workers who face a higher probability of unemployment also tend to experience larger consumption declines when unemployment occurs, amplifying ex ante welfare losses from labor market risk. Below, we explore the implications of this correlation for the distribution of willingness-to-pay for UI reforms.

| | | Consumption type ($k$) | | |
|---|---|---|---|---|
| | Risk type ($h$) | High drop | Medium drop | Low drop |
| $\hat{\pi}(h, k)$ | High | 0.045 | 0.131 | 0.064 |
| | | [0.022, 0.104] | [0.029, 0.212] | [0.031, 0.181] |
| | Low | 0.076 | 0.363 | 0.320 |
| | | [0.048, 0.348] | [0.072, 0.448] | [0.229, 0.432] |
| $\hat{\pi}(k \mid h)$ | High | 0.189 | 0.545 | 0.266 |
| | | [0.128, 0.483] | [0.157, 0.624] | [0.160, 0.608] |
| | Low | 0.100 | 0.478 | 0.421 |
| | | [0.065, 0.454] | [0.095, 0.564] | [0.312, 0.551] |

TABLE 3. Joint and Conditional Distribution of Risk and Consumption Types

Notes: The table reports estimates of the joint and conditional type probabilities in the joint risk–consumption mixture model. The entries labeled $\hat{\pi}(h, k)$ give the estimated joint population share of risk type $h$ and consumption type $k$. The entries labeled $\hat{\pi}(k \mid h)$ give the conditional share of each consumption type within risk type $h$, i.e. $\hat{\pi}(k \mid h) = \hat{\pi}(h, k) / \sum_{k'} \hat{\pi}(h, k')$. 95% confidence intervals are based on 400 household-level replications are shown in brackets.

### 4.3. Risk Aversion and Reforms

***Risk aversion.*** In our framework, the insurance value of a UI expansion to an individual worker is captured by the product of their consumption decline in unemployment and their coefficient of relative risk aversion (see equation (4)). Based on evidence from a meta-survey of the intertemporal elasticity of substitution (Havránek 2015), we set our baseline coefficient of relative risk aversion to $\gamma = 3$.

***Reforms.*** We consider three expansions of UI generosity:

(a) Proportional–capped reform: $\phi_{\mathcal{B}}(y) = \min\{y, \kappa\}$, where $\kappa$ is an earnings 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 a natural expansion of the current U.S. UI system, which offers proportional UI benefits up to an earnings 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 earnings for all workers. The interaction between the reform design and the distribution of unemployment risk determines 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. These WTP estimates are defined for reforms that adjust the financing tax to offset the mechanical cost; Section 5 incorporates the additional financing required by behavioral responses. Panel (A) shows the insurance component of WTP, which is common across reforms and positive for all workers. Panels (B)–(D) show total WTP for each reform. Each panel plots how the median, interquartile range, and interdecile range vary across percentiles of the household income distribution.

Panel (A) shows that WTP for actuarially fair insurance declines with household income and exhibits substantial within-income variation. Among households in the bottom income decile, the median WTP is \$0.36 per dollar of benefit expansion—over 20% higher than the median among households in the top decile. This difference is even more pronounced in the upper tail: the 90th percentile of WTP in the bottom income decile is $0.68, over 65% higher than that in the top income decile.

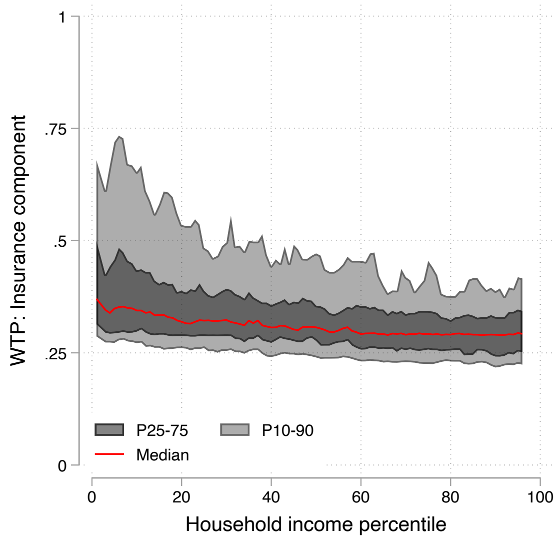


(A) Insurance Component

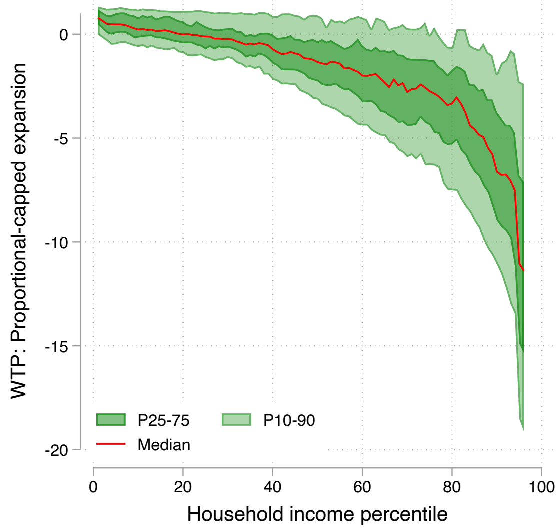


(B) Proportional-Capped

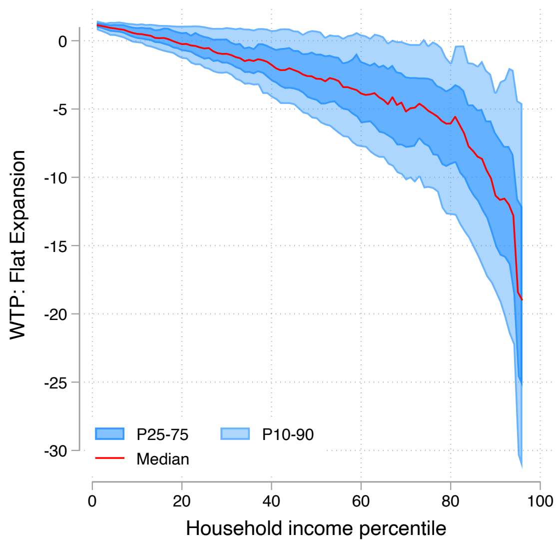


(C) Flat

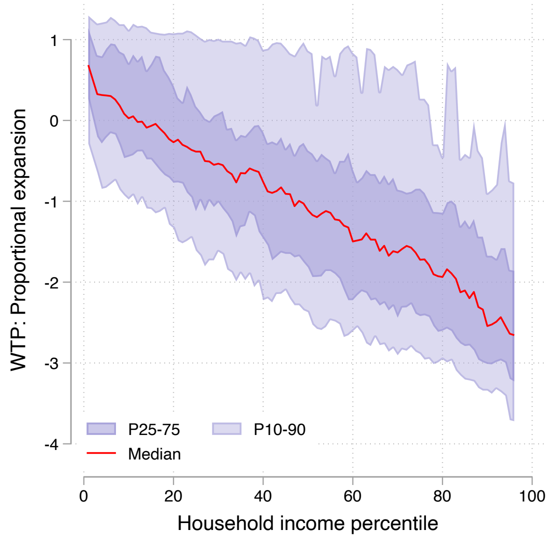


(D) Proportional

FIGURE 4. Willingness-to-Pay for UI Expansion by Income and Reform Type

Notes: Each panel summarizes the distribution of workers' willingness-to-pay (WTP) for a UI expansion, reporting the median and the 10th, 25th, 75th, and 90th percentiles. We compute worker-level WTP using equation (3), combining estimates of individual consumption declines and panel-average unemployment risk; worker incomes winsorized at $10,000 and $200,000. Panel (A) shows the insurance component (WTP for actuarially insurance). 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. Appendix E.1 documents the covariance of individual WTP across different reforms.

Panel (B) focuses on a proportional–capped reform, where WTP reflects both the value of insurance and the cross-subsidization embedded in the reform. The wider vertical scale relative to panel (A) highlights that much of the WTP variation by household income reflects cross-subsidization. Median WTP declines with income—from \$0.54 per dollar in the bottom income decile to -\$9.25 per dollar 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 \$5.64.

While the insurance component is common across reforms, 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 \$0.95 per dollar in the bottom income decile, falling to -\$15.54 in the top. Conversely, a proportional reform (panel (D)) tilts gains toward workers in higher-income households.

Under the flat reform, cross-subsidization arises solely from differences in unemployment risk: workers with below-average risk subsidize those with higher risk. As unemployment risk tends to be lower among higher-income households, WTP declines steeply with income. Under the proportional reform, cross-subsidization reflects both unemployment risk and earnings. Conditional on unemployment risk, higher earners gain more from the expansion, but 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.

An individual worker's preference ordering over the three reforms depends on both unemployment risk and earnings. 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 calibration of the coefficient of relative risk aversion and on the assumption of state-independent utility.

Varying the coefficient of relative risk aversion scales the insurance component 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 (see Appendix E.2).

As emphasized by Andrews and Miller (2013), heterogeneity in risk aversion affects the distribution of insurance values only insofar as it covaries with consumption declines. Direct evidence on such heterogeneity is limited.[16] However, as Figure 4 shows, variation in risk aversion has a much smaller effect on the distribution of total WTP for UI reform than on the distribution of the insurance component, because under realistic reforms the cross-subsidization term plays a dominant role in shaping WTP.

## 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 design trades off insurance gains for a representative agent against incentive costs, summarized by a single elasticity. We show that with heterogeneous individuals, the single representative-agent insurance term is replaced by a welfare-weighted sum of individual willingness-to-pay, which we decompose into the social value of risk protection and cross-subsidization. Moreover, incentive costs depend on how individual behavioral responses covary with risk weights and state-specific net tax liabilities.

### 5.1. Theory

The aggregate value of a social insurance reform depends on its fiscal cost and on how the reform's effects are weighted across individuals.

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:

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

By construction, the weights in each state sum to one, so $\mathbb{E}^l[x]$ and $\mathbb{E}^h[x]$ are state-specific averages under weights proportional to each individual's mechanical contribution to the reform's fiscal impact: in the low state, expected benefit payments $(1-e_i)\,\phi_{\mathcal{B}}(y_i)$, and in the high state, expected tax payments $e_i\,\phi_{\mathcal{T}}(y_i)$.

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

***Budgetary impact.*** We focus on budget-neutral expansions of social insurance. The policymaker faces the budget 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 $\mathcal{B}(.)$ summarize the overall tax-and-transfer system in each state: $\mathcal{T}(y_i)$ denotes the net payment to the government in the high state, and $\mathcal{B}(y_i)$ the net transfer received in the low state.

Budget-neutrality implies that the reform $d\theta = (db, d\tau)$, parameterized in equation (2), must 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 cost of the benefit expansion, holding behavior fixed. The second term reflects the additional tax adjustment needed to maintain budget balance in response to behavioral changes—the fiscal externality. This arises because higher benefits reduce the marginal return to exerting effort, increasing the likelihood that individuals are in the low state. The government then both forgoes taxes $\mathcal{T}(y_i)$ and pays benefits $\mathcal{B}(y_i)$, raising the required tax adjustment.

Note that the mechanical effect depends on reform design—through $\phi_{\mathcal{B}}(.)$ and $\phi_{\mathcal{T}}(.)$—whereas the fiscal externality depends on pre-reform 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 $\phi_{\mathcal{B}}(y_i)\, db$, financed by the balanced-budget tax adjustment. Equation (10) makes clear that, in general, the fiscal cost from behavioral responses depends on the covariance between individual effort elasticities and both (i) the risk weights and (ii) state-specific 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 willingness-to-pay, we convert the change in social welfare into money-metric terms by scaling the aggregate utility change $\frac{dW}{d\theta}$ per \$1 balanced-budget increase in benefit expenditures by the aggregate utility value of a \$1 decrease in the tax bill. Denoting this money-metric welfare effect by $\frac{dW^{MM}}{d\theta}$, we obtain:

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

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

***Insurance–cross-subsidization decomposition.*** We define social marginal welfare weights as $g_i \equiv \omega_i u_i^{h\prime}(c_i^h)$. The following proposition decomposes the resource-reallocation term in equation (12) into the social value of insurance—the within-individual, actuarially fair transfer of resources between states—and cross-subsidization—the across-individual net expected transfer embedded in the reform.

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

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

*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, i.e.,* $\mathbb{E}^l[\lambda] = 1$*. This can be equivalently expressed as:*

$$\mathrm{RR} = \underbrace{\bar{\lambda}\mathbb{E}^l\left[\lambda\left(\frac{u^{l\prime}(c^l) - u^{h\prime}(c^h)}{u^{h\prime}(c^h)}\right)\right]}_{\substack{\text{risk protection} \\ \text{(within-individual)}}} + \underbrace{(\bar{\lambda} - 1)}_{\substack{\text{cross-subsidization} \\ \text{(across-individuals)}}} \tag{14}$$

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, modulated by individual and aggregate social marginal welfare weights. The individual weight $\lambda_i$ is normalized so 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—that is, the average social value of providing an additional \$1 of surplus to program beneficiaries relative to funders. This connects to our descriptive evidence showing households who experience unemployment have systematically lower pre-unemployment consumption and earnings, implying that reform beneficiaries differ from funders in economically meaningful ways. Importantly, $\bar{\lambda}$ is reform-specific because the risk weights depend on $\phi_{\mathcal{B}}(\cdot)$ and $\phi_{\mathcal{T}}(\cdot)$, capturing how benefit and tax changes vary with income.

Equation (14) expresses the welfare gain from the reform 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. This across-individual redistribution component is captured by a single sufficient statistic—the reform-specific aggregate welfare weight $\bar{\lambda}$.[17]

These equations highlight two distinct channels through which social insurance reform affects aggregate welfare. First, a reform alters the degree of risk protection provided to individuals. All else equal, risk protection is more valuable when individual insurance values are more positively correlated with 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 $\bar{\lambda}$.

[17] 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.

Second, even absent individual demand for insurance, the reform reallocates net resources across individuals through cross-subsidization. If individuals in the low state have higher welfare weights than those in the high state (i.e., $\bar{\lambda} > 1$), expanding the program generates positive value purely from this reallocation. This cross-subsidization arises because the implicit pricing embedded in the tax and benefit adjustments typically deviates from actuarially fair terms.

Importantly, because the risk weighting in expectations depends on reform design (via $\phi_{\mathcal{B}}$ and $\phi_{\mathcal{T}}$), both the risk protection and cross-subsidization components are reform specific.

### 5.2. Implementation

We maintain Assumption 1, which implies state-independent marginal utility 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:

$$g_i \equiv g_{y_i} \tilde{g}_i, \qquad \tilde{g}_i \equiv (c_i^h)^{-\gamma} \Big/ \mathbb{E}[(c^h)^{-\gamma} \mid y_i], \tag{15}$$

where $g_{y_i}$ is the inverse-optimum weight computed in Hendren (2020), and $c_i^h$ denotes long-run in-work (equivalized) consumption. By construction, $\mathbb{E}[\tilde{g} \mid y_i] = 1$, so the inverse-optimum weights $g_{y_i}$ determine the *average* welfare weight placed on each income group, while $\tilde{g}_i$ reweights *within* an income group toward individuals with lower long-run consumption.[18]

Under this specification, the social value of transferring \$1 of surplus between individuals at different income levels is governed by the inverse-optimum weights $g_y$. These weights rationalize the existing income-tax-and-transfer system as locally optimal: small, revenue-neutral modifications have approximately zero first-order effect on social welfare. They can be interpreted either as revealed social preferences for redistribution across income groups or as the marginal cost to the government of providing \$1 of surplus to individuals of income $y$ through an adjustment to the

[18] Equivalently, we can write $g_i = \omega_i u_i'(c_i^h)$ in terms of Pareto weights $\omega_i$ and marginal utilities $u_i'(\cdot)$. Approximating $u_i'(c_i^h) \approx (c_i^h)^{-\gamma}$ and choosing $\omega_i = \frac{g_{y_i}}{\mathbb{E}[(c^h)^{-\gamma} \mid y_i]}$ implies the social marginal welfare weights in equation (15), $g_i = g_{y_i} (c_i^h)^{-\gamma} / \mathbb{E}[(c^h)^{-\gamma} \mid y_i]$.

tax-and-transfer system. 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. As most married households file income taxes jointly, we evaluate $g_y$ at household income.

Social insurance reform may also reallocate surplus among individuals with the same level of income. Our welfare-weight specification values this within-income-group surplus redistribution according to $\tilde{g}_i \propto (c_i^h)^{-\gamma}$, implying that, conditional on current income, the social value of providing an additional dollar of surplus to an individual declines with their consumption level. This is motivated by empirical evidence that consumption is a reliable proxy for both lifetime income and economic well-being (e.g., Poterba 1989; Meyer and Sullivan 2003). We measure $c_i^h$ using panel-average in-work consumption, equivalized by household size to account for differences in needs due to household composition, and winsorize this measure at the 1st and 99th 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 on state-specific net taxes and benefits, $\mathcal{T}(y_i)$ and $\mathcal{B}(y_i)$.

We calibrate the elasticity using evidence from Chetty (2008), who reports unemployment-duration elasticities separately by quartiles of the wealth distribution.[19] 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 tax and benefit schedules using NBER TAXSIM, assuming married households file jointly (see Appendix G.1 for details). This means our measure of the fiscal externality captures the impact of behavioral responses through both higher UI payments and changes in income tax revenues.

[19] Chetty (2008) estimates elasticities of unemployment exit hazards with respect to UI benefits that range from −0.642 in the lowest wealth quartile to 0.016 in the highest (Table 2, p. 204). Following Chetty (2008), we map these hazard elasticities into unemployment-duration elasticities by using the fact that, under a constant-hazard approximation, the elasticity of duration is the negative of the elasticity of the exit hazard. This duration elasticity is the dynamic counterpart of our static elasticity $\epsilon_i^{(1-e,b)}$ (see Appendix F.2). In the PSID, we assign each household the corresponding elasticity based on its wealth quartile. Following Schmieder and Von Wachter (2016), we account for the exhaustion of UI benefits after six months when mapping duration elasticities into fiscal costs (see Appendix G.1)

***Discussion.*** Our characterization of the welfare gains from resource reallocation parallels recent work on the social value of transfer programs. For example, Kolsrud et al. (2024) examine reforms that shift resources between an early and a late retiree and show that differences in their consumption levels and welfare weights capture the combined insurance and redistributive effects. A key difference in our setting is that we model heterogeneity across the entire distribution of workers, so the social value of reform depends on the joint distribution of insurance values and welfare weights, together with unemployment risk, which is central in our setting. In this respect, our approach also connects to recent work using PSID consumption data to evaluate the targeting value of disability insurance (Deshpande and Lockwood 2022) and welfare programs (Rafkin et al. 2023).

### 5.3. Results

#### Baseline

Row (a) of Table 4 presents our baseline estimate of the social value of increasing UI generosity.[20] These estimates are based on a budget-neutral expansion of the current U.S. UI program that raises benefits in proportion to the minimum of a worker's earnings and an earnings cap, funded by a flat increase in payroll taxes. Columns (1) and (2) report the social value derived from risk protection and cross-subsidization, 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-neutral expansion of the UI program is \$0.48 and the social value of cross-subsidization is \$0.19, yielding a total resource-reallocation gain of \$0.67 per \$1 of increased benefit expenditure.

These gains must be weighed against the associated fiscal externalities. Column (4) reports the direct fiscal externality from higher UI payments, column (5) reports the additional fiscal externality due to reduced income tax receipts, and column (6) presents the total. We find that accounting for the broader tax-and-transfer system increases the fiscal externality from \$0.27 (excluding income-tax revenue effects) to \$0.75 (including them). Overall, the fiscal externality (\$0.75) slightly exceed the resource-reallocation gains (\$0.67), implying a modest negative net welfare effect under our baseline welfare weights.

[20] Our welfare results are broadly stable to alternative choices of the number of latent types, and rescaling the insurance component of WTP leaves the relative ranking of the baseline, flat, and proportional reforms essentially unchanged (see Appendix H).

### Reform structure

Rows (b) and (c) examine two alternative reforms: a flat benefit expansion (independent of earnings) and a proportional expansion with a cap. Relative to the baseline reform, the flat expansion is more progressive, allocating a larger share of additional resources to individuals with higher welfare weights. This raises the social value of both risk protection and cross-subsidization; the total resource-reallocation gain from the flat reform is $0.78 per $1 of increased benefit expenditure.

By 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 social value of risk protection and cross-subsidization, yielding a lower resource-reallocation gain of $0.61 for this reform.

The overall social value of the flat expansion is further enhanced by a lower fiscal externality than the baseline. This reduction reflects two offsetting forces. On the one hand, the flat reform directs a larger share of the benefit expansion toward lower earners, who tend to exhibit somewhat stronger behavioral responses; this modestly increases the direct fiscal externality. On the other hand, lower earners face smaller income-tax liabilities, which reduces the indirect fiscal externality linked to lost tax revenue—and this effect dominates. For the proportional reform, the pattern is reversed: shifting the incentives to reduce effort toward higher earners who are less responsive lowers the direct fiscal externality relative to the baseline, but this is more than offset by a higher indirect fiscal externality due to their larger tax contributions.

Taken together, these differences across reform structures underscore the importance of incidence in determining both the distributional benefits of reform and the magnitude of the associated fiscal externalities.

### Social preferences

Row (a) reports our baseline results under the welfare weights in equation (15); rows (d)–(f) illustrate how the social value of the reform varies depending on the weighting placed on individual willingness-to-pay.

***Money-metric weights.*** Under money-metric marginal social welfare weights ($g_i$ = 1), social preferences are indifferent to how surplus is distributed across individuals. One rationale for this approach is 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 is distribution-neutral, cross-subsidization contributes zero by construction, and risk protection is valued using common welfare weights, yielding a resource-reallocation value of \$0.38 (row (d)). A key drawback of this criterion is that it relies on a lump-sum redistribution scheme that is not feasible in practice.

Money-metric weights also clarify the connection between the budget-neutral reform that we focus on and the "unpackaged" components: (i) an unfunded benefit increase and (ii) a tax change. UI beneficiaries' average WTP for an *unfunded* UI expansion, per mechanical \$1 increase in benefits, is \$1.38—the direct \$1 plus \$0.38 in money-metric risk-protection value (Table 4, row (d)).

To unpack the budget-neutral reform into an unfunded benefit increase and a financing tax change, we attribute the fiscal externality to the benefit expansion. Then each \$1 of higher benefits costs the government \$1.75 once behavior is accounted for (the \$1 mechanical cost plus \$0.75 through both UI payments and the broader tax system.). The implied marginal value of public funds (MVPF) for an unfunded UI benefit expansion is therefore \$1.38/\$1.75 $\approx$ 0.79. Under the same unpacking, an unfunded payroll tax reduction has a mechanical cost of \$1, generates no fiscal externality, and raises recipients' surplus by exactly \$1, so its MVPF is 1. With money-metric welfare weights, this comparison implies that an additional dollar of government revenue is more valuable when used to cut payroll taxes than to expand UI benefits. When welfare weights value surplus redistribution, the joint reform can be unpacked analogously into appropriately welfare-weighted MVPFs for UI benefits and for the financing tax (Appendix F.3).

***Inverse-optimum weights.*** Row (e) uses inverse-optimum welfare weights and assumes that social preferences are indifferent to the distribution of surplus *within* income groups—i.e., it sets 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 that accounts for the distortionary costs of achieving cross-income redistribution through the income tax system. Under these weights, the social value of resource reallocation is \$0.51.

***Cross- versus within-income redistribution.*** Comparing money-metric weights, inverse-optimum weights, and the baseline specification quantifies the social value of surplus redistribution.

Comparing row (d) with row (e) isolates the value of *cross*-income surplus redistribution delivered by a UI expansion. Moving from money-metric to inverse-optimum weights increases the resource-reallocation gain by \$0.13, with over two-thirds of this increase coming from the cross-subsidization component and the remainder from changes in the welfare-weighted insurance term.

Comparing row (e) with the baseline results in row (a) reveals the additional social value of *within*-income-group surplus reallocation—toward individuals with lower long-run in-work consumption. This within-income effect further increases the resource-reallocation gain by $0.16, with around two-thirds of the increase arising from the cross-subsidization component. This pattern reflects that, conditional on current income, lower long-run consumption households are more likely to become unemployed and thus receive a larger share of UI transfers, consistent with evidence on targeting and selection in disability insurance (Deshpande and Lockwood 2022) and welfare programs (Rafkin et al. 2023).

Of the overall resource-reallocation gain of $0.67 in row (a), we attribute $0.38 (57%) to efficiency gains from risk protection, $0.13 (19%) to cross-income surplus redistribution, and the remaining $0.16 (24%) to within-income surplus redistribution.

***Utilitarian weights.*** Our use of inverse-optimum weights to value cross-income 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 income-tax adjustments required to implement the implied 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 therefore present results based on utilitarian social welfare weights, $g_i = (c_i^h)^{-\gamma}$. These weights place greater value on cross-income redistribution than inverse-optimum weights, increasing both the value of risk protection and the cross-subsidization component of resource reallocation.

While the precise social value of surplus reallocation from UI reform ultimately depends on the policymaker's chosen welfare-weighting scheme, this analysis shows that it will generally exceed the value implied by money-metric weights. Specifically, as long as marginal surplus is valued more for (i) lower-income than higher-income households, (ii) lower-consumption than higher-consumption households, or (iii) some combination of both, the social value of the reform exceeds that implied by money-metric weights.

### Alternative implementations

***Standard implementation.*** The penultimate row of Table 4 reports results for a flat reform evaluated using money-metric welfare weights, approximating the standard

Baily–Chetty sufficient-statistics implementation. The implied risk protection value is \$0.39 per dollar of increased benefit expenditure, reflecting the use of unemployment-risk-weighted averages rather than welfare weights that place additional value on surplus redistribution (as in row (b)). The resulting net welfare effect is positive (approximately \$0.11), reflecting the lower fiscal externality under this implementation.

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

***Average implementation.*** The final row of Table 4 presents results from an alternative implementation of a flat reform that adapts the approach of Kolsrud et al. (2024) developed in the context of retirement benefit reform. This method evaluates equation (14) using *average* consumption levels among the employed and unemployed populations. It captures redistribution between employment states but abstracts from surplus dispersion within each state. In this "average implementation," we likewise value the fiscal externality using the 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 \$0.57 per \$1 of increased benefit expenditure—much closer to our full individual-level estimate for a flat reform (\$0.78) than the standard implementation (\$0.39). However, by collapsing the unemployed to a representative individual, it abstracts from heterogeneity in expected benefit receipt within the unemployed population—specifically, the positive covariance between welfare weights and expected benefit payments—and therefore understates the social value of resource reallocation. It also overstates the indirect fiscal externality by ignoring how unemployment risk and behavioral responses covary with individuals' state-specific tax liabilities. The combined effect is a net social value of -\$0.30 under the average implementation, compared with +\$0.26 under the full implementation.

Overall, the average implementation more accurately approximates resource-reallocation gains than the standard implementation. However, in our context, it introduces non-trivial error and, importantly for our purposes, cannot compare reforms generating heterogeneous incidence within unemployed workers.

***Cross-sectional implementation.*** The results in Table 4 use the panel dimension of the PSID. Panel data offer two main advantages over cross-sectional data. First, they enable us to estimate individual-level insurance values and WTP, which facilitates a decomposition of the social gains from expanding UI generosity into within-individual risk protection and across-individual cross-subsidization. Second, longitudinal information allows us to measure consumption and earnings using within-household panel averages, helping to limit the influence of measurement error. However, when only cross-sectional data are available, 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 | | | Overall effect |
| | Risk protection | Cross-subsidization | Total (1)+(2) | Direct | Indirect | Total (4)+(5) | (3)-(6) |
| (a) Baseline | 0.48 | 0.19 | 0.67 | 0.27 | 0.48 | 0.75 | -0.07 |
| | [0.39, 0.59] | [0.15, 0.22] | [0.56, 0.78] | [0.26, 0.27] | [0.45, 0.51] | [0.72, 0.78] | [-0.18, 0.04] |
| *Reform structure* | | | | | | | |
| (b) Flat | 0.52 | 0.25 | 0.78 | 0.28 | 0.23 | 0.51 | 0.26 |
| | [0.42, 0.64] | [0.22, 0.29] | [0.65, 0.89] | [0.28, 0.28] | [0.19, 0.28] | [0.47, 0.56] | [0.13, 0.38] |
| (c) Proportional | 0.46 | 0.15 | 0.61 | 0.25 | 0.56 | 0.82 | -0.21 |
| | [0.37, 0.56] | [0.11, 0.18] | [0.50, 0.70] | [0.25, 0.26] | [0.53, 0.60] | [0.79, 0.85] | [-0.33, -0.11] |
| *Social preferences* | | | | | | | |
| (d) Money-metric | 0.38 | 0 | 0.38 | =(a) | =(a) | =(a) | -0.36 |
| | | | [0.31, 0.46] | | | | [-0.44, -0.28] |
| (e) Inverse-optimum | 0.42 | 0.09 | 0.51 | =(a) | =(a) | =(a) | -0.24 |
| | [0.34, 0.51] | [0.08, 0.09] | [0.43, 0.59] | | | | [-0.33, -0.15] |
| (f) Utilitarian | 0.64 | 0.53 | 1.17 | =(a) | =(a) | =(a) | 0.43 |
| | [0.51, 0.80] | [0.46, 0.61] | [1.01, 1.34] | | | | [0.25, 0.60] |
| *Alternative implementations* | | | | | | | |
| (g) Standard implementation | 0.39 | - | 0.39 | 0.28 | - | 0.28 | 0.11 |
| *(flat reform+money-metric)* | | | [0.32, 0.48] | | | [0.28, 0.28] | [0.04, 0.20] |
| (h) Average implementation | 0.44 | 0.13 | 0.57 | 0.28 | 0.58 | 0.86 | -0.30 |
| | [0.36, 0.53] | [0.12, 0.13] | [0.48, 0.66] | [0.28, 0.28] | [0.55, 0.62] | [0.83, 0.90] | [-0.39, -0.20] |

TABLE 4. Risk Protection, Cross-subsidization, and the Fiscal Externality from UI Expansion

Notes: Column (7) reports the overall social welfare impact of a UI expansion, decomposed in columns (1)–(6). Row (a) corresponds to an expansion of the current U.S. UI system under baseline social preferences; other rows report results for alternative reform structures or social preferences. All entries are expressed in dollars per $1 increase in benefit expenditure. 95% confidence intervals are based on 400 household-level bootstrap replications and are reported in square brackets.

***Discussion.*** Our analysis advances the literature on the design of social insurance by showing how individual-level heterogeneity and reform structure shape insurance value, implicit cross-subsidization, and incentive cost—and, through these channels, the social value of UI reforms. Our framework also has several limitations that are important to keep in mind when interpreting the quantitative results.

First, our baseline calculations abstract from incomplete eligibility and take-up by treating all workers in our sample as eligible for UI and as taking it up when unemployed. In practice, many low-income workers either do not qualify for benefits or fail to claim them conditional on eligibility. Our framework can accommodate this by introducing an indicator $L_i \in \{0, 1\}$ for "eligible and takes up UI if unemployed," which scales the marginal benefit adjustment $\phi_{\mathcal{B}}(y_i)$ in the resource-reallocation term, so that only those with $L_i = 1$ receive benefits when unemployed. Using recent evidence on income gradients in UI take-up (Lachowska et al., 2025), we recompute the baseline resource-reallocation term allowing $L_i$ to vary by income and find that it falls modestly, from \$0.67 to about \$0.64 per dollar of increased benefit expenditure. The reduction reflects lower take-up at the bottom of the income distribution, partly offset by relatively low take-up at the top. Our framework can also be used to study reforms that explicitly expand coverage or take-up among currently uncovered low-income workers, but we leave a full exploration of such policies to future work.

Second, our treatment of fiscal externalities focuses on the impact of UI generosity on employment and earnings, and hence on UI payments and income-tax revenues, holding other programs fixed. In reality, UI interacts with the broader safety net. More generous UI may reduce reliance on other transfer programs (see, e.g., Mueller et al. (2016) on UI and Social Security Disability Insurance), generating additional fiscal savings that are not captured in our baseline fiscal externality. At the same time, UI reforms may affect other behavioral margins that could amplify or offset these effects (see, e.g., Anderson and Meyer (1997) on UI generosity and take-up). We view a full accounting of these additional fiscal externalities as an important avenue for future research building on our framework.

## 6. Conclusion

We study how heterogeneity in unemployment risk and income losses shapes both individual demand for unemployment insurance (UI) and the social value of expanding UI generosity. We estimate the distribution of workers' willingness-to-pay for UI reforms and show that it reflects two forces: heterogeneous insurance value, captured by unemployment-induced consumption declines, and cross-subsidization arising from pooled financing and heterogeneous unemployment risk. Because workers who experience unemployment are negatively selected—by income and, conditional on income, by in-work consumption—UI expansions tend to redistribute surplus toward less well-off households through both risk protection and cross-subsidization.

We weigh these gains against distortionary costs. Under our baseline welfare weights, a natural proportional–capped expansion of the current U.S. UI system has net value close to zero once reduced income-tax revenues are accounted for, while a flat expansion that targets more resources toward lower-income households yields positive net value.

Our approach is marginal and avoids estimating a full structural model. An important direction for future work is to combine our decomposition with a structural framework to assess optimal UI design.

## References


Acemoglu, D. and R. Shimer (1999). Efficient unemployment insurance. *Journal of political Economy 107*(5), 893–928.

Andersen, A. L., A. S. Jensen, N. Johannesen, C. T. Kreiner, S. Leth-Petersen, and A. Sheridan (2023). How do households respond to job loss? lessons from multiple high-frequency datasets. *American Economic Journal: Applied Economics 15*(4), 1–29.

Anderson, P. M. and B. D. Meyer (1997). Unemployment insurance takeup rates and the after-tax value of benefits. *Quarterly Journal of Economics 112*(3), 913–937.

Anderson, P. M. and B. D. Meyer (2000). The effects of the unemployment insurance payroll tax on wages, employment, claims and denials. *Journal of Public Economics 78*(1-2), 81–106.

Anderson, P. M. and B. D. Meyer (2006). Unemployment insurance tax burdens and benefits: Funding family leave and reforming the payroll tax. *National Tax Journal 59*(1), 77–95.

Andrews, I. and C. Miller (2013). Optimal Social Insurance with Heterogeneity. *mimeo*.

Ashenfelter, O. and J. Ham (1979). Education, unemployment, and earnings. *Journal of Political Economy 87*(5, Part 2), S99–S116.

Attanasio, O., P. Levell, H. Low, and V. Sánchez-Marcos (2018). Aggregating elasticities: intensive and extensive margins of women's labor supply. *Econometrica 86*(6), 2049–2082.

Audoly, R. (2024). Self-employment and labor market risks. *International Economic Review*.

Baily, N. (1978). Some aspects of optimal unemployment insurance. *Journal of Public Economics 10*, 24.

Bertrand, M. and S. Mullainathan (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review 94*(4), 991–1013.

Blomqvist, Å. and H. Horn (1984). Public health insurance and optimal income taxation. *Journal of Public Economics 24*(3), 353–371.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption Inequality and Partial Insurance. *American Economic Review 98*(5), 1887–1921.
Boadway, R., M. Leite-Monteiro, M. Marchand, and P. Pestieau (2006). Social Insurance and Redistribution with Moral Hazard and Adverse Selection. *Scandinavian Journal of Economics 108*(2), 279–298.
Boehm, J., E. Fize, and X. Jaravel (2025). Five facts about mpcs: Evidence from a randomized experiment. *American Economic Review 115*(1), 1–42.
Borella, M., M. De Nardi, and F. Yang (2023). Are marriage-related taxes and social security benefits holding back female labour supply? *The Review of Economic Studies 90*(1), 102–131.
Browning, M. and T. F. Crossley (2001). Unemployment insurance benefit levels and consumption changes. *Journal of Public Economics 80*(1), 1–23.
Cairó, I. and T. Cajner (2018). Human capital and unemployment dynamics: Why more educated workers enjoy greater employment stability. *Economic Journal 128*(609), 652–682.
Campos, R. and I. Reggio (2020). Do the Unemployed Pay Lower Prices? A Reassessment of the Value of Unemployment Insurance. *Journal of the European Economic Association*, 47.
Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics 90*(10-11), 1879–1901.
Chetty, R. (2008). Moral Hazard versus Liquidity and Optimal Unemployment Insurance. *Journal of Political Economy 116*(2), 173–234.
Chetty, R. and A. Finkelstein (2013). Chapter 3 - Social Insurance: Connecting Theory to Data. In A. J. Auerbach, R. Chetty, M. Feldstein, and E. Saez (Eds.), *Handbook of Public Economics*, Volume 5 of *Handbook of Public Economics, Vol. 5*, pp. 111–193.
Choi, S., A. Janiak, and B. Villena-Roldán (2015). Unemployment, participation and worker flows over the life-cycle. *Economic Journal 125*(589), 1705–1733.
Colarieti, R., P. Mei, and S. Stantcheva (2024). The how and why of household reactions to income shocks. Technical report, National Bureau of Economic Research.
Cremer, H. and P. Pestieau (1996). Redistributive taxation and social insurance. *International Tax and Public Finance 3*(3), 281–295.
Crossley, T. F. and H. W. Low (2014). Job loss, credit constraints, and consumption growth. *Review of Economics and Statistics 96*(5), 876–884.
Deshpande, M. and L. M. Lockwood (2022). Beyond health: Nonhealth risk and the value of disability insurance. *Econometrica 90*(4), 1781–1810.
East, C. N. and E. Kuka (2015). Reexamining the consumption smoothing benefits of unemployment insurance. *Journal of Public Economics 132*, 32–50.
Feldstein, M. (2005). Rethinking social insurance. *American Economic Review 95*(1), 1–24.
Ferey, A. (2022). Redistribution and unemployment insurance. *mimeo*.
Finkelstein, A., N. Hendren, and E. F. Luttmer (2019). The value of medicaid: Interpreting results from the oregon health insurance experiment. *Journal of Political Economy 127*(6), 2836–2874.
Ganong, P. and P. Noel (2019). Consumer Spending during Unemployment: Positive and Normative Implications. *American Economic Review 109*(7), 2383–2424.

Gruber, J. (1997). The Consumption Smoothing Benefits of Unemployment Insurance. *American Economic Review 87*(1), 192–205.
Guo, A. (2024). Payroll tax incidence: Evidence from unemployment insurance. *Journal of Public Economics 239*, 105209.
Guo, A. and A. C. Johnston (2021). The Finance of Unemployment Compensation and Its Consequences. *Public Finance Review 49*(3), 392–434.
Haan, P. and V. 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 16*(2), 127–181.
Havránek, T. (2015). Measuring intertemporal substitution: The importance of method choices and selective reporting. *Journal of the European Economic Association 13*(6), 1180–1204.
Heckman, J. and B. Singer (1984). A method for minimizing the impact of distributional assumptions in econometric models for duration data. *Econometrica 52*(2), 271–320.
Hendren, N. (2017). Knowledge of Future Job Loss and Implications for Unemployment Insurance. *American Economic Review 107*(7), 1778–1823.
Hendren, N. (2020). Measuring economic efficiency using inverse-optimum weights. *Journal of public Economics 187*, 104198.
Kaplan, G. and G. Menzio (2015). The morphology of price dispersion. *International Economic Review 56*(4), 1165–1206.
Keane, M. P. and K. I. Wolpin (1997). The career decisions of young men. *Journal of Political Economy 105*(3), 473–522.
Kimball, M. S., C. R. Sahm, and M. D. Shapiro (2009). Risk preferences in the psid: Individual imputations and family covariation. *American Economic Review 99*(2), 363–368.
Kleven, H. J. (2021). Sufficient statistics revisited. *Annual Review of Economics 13*(1), 515–538.
Kolsrud, J., C. Landais, P. Nilsson, and J. Spinnewijn (2018). The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden. *American Economic Review 108*(4-5), 985–1033.
Kolsrud, J., C. Landais, D. Reck, and J. Spinnewijn (2024). Retirement consumption and pension design. *American Economic Review 114*(1), 89–133.
Kroft, K., K. Kucko, E. Lehmann, and J. Schmieder (2020). Optimal income taxation with unemployment and wage responses: A sufficient statistics approach. *American Economic Journal: Economic Policy 12*(1), 254–292.
Kroft, K. and M. J. Notowidigdo (2016). Should Unemployment Insurance Vary with the Unemployment Rate? Theory and Evidence. *Review of Economic Studies 83*(3), 1092–1124.
Krusell, P., T. Mukoyama, and A. Şahin (2010). Labour-market matching with precautionary savings and aggregate fluctuations. *Review of Economic Studies 77*(4), 1477–1507.
Lachowska, M., I. Sorkin, and S. A. Woodbury (2025). Employers and unemployment insurance take-up. *American Economic Review 115*(8), 2529–2573.
Landais, C. and J. Spinnewijn (2021). The value of unemployment insurance. *Review of Economic Studies 88*(6), 3041–3085.

Lang, K. and J.-Y. K. Lehmann (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature 50*(4), 959–1006.

Lentz, R. and T. Tranaes (2005). Job Search and Savings: Wealth Effects and Duration Dependence. *Journal of Labor Economics 23*(3).

Lewis, D., D. Melcangi, and L. Pilossoph (2026). Latent heterogeneity in the marginal propensity to consume. *Review of Economic Studies*.

Lieber, E. M. and L. M. Lockwood (2019). Targeting with In-Kind Transfers: Evidence from Medicaid Home Care. *American Economic Review 109*(4), 1461–1485.

Meyer, B. D. and W. K. Mok (2019). Disability, earnings, income and consumption. *Journal of Public Economics 171*, 51–69.

Meyer, B. D. and J. X. Sullivan (2003). Measuring the well-being of the poor using income and consumption. *Journal of Human Resources*, 1180–1220.

Mueller, A. I., J. Rothstein, and T. M. Von Wachter (2016). Unemployment insurance and disability insurance in the great recession. *Journal of Labor Economics 34*(S1), S445–S475.

Mueller, A. I., J. Spinnewijn, and G. Topa (2021). Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias. *American Economic Review 111*(1), 324–363.

Nickell, S. (1979). Education and lifetime patterns of unemployment. *Journal of Political Economy 87*(5, Part 2), S117–S131.

Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review 113*(4), 982–1012.

Poterba, J. M. (1989). Lifetime Incidence and the Distributional Burden of Excise Taxes. *American Economic Review 79*(2), 325–330.

Quandt, R. E. (1972). A new approach to estimating switching regressions. *Journal of American Statistical Association 67*(338), 306–310.

Rafkin, C., A. Solomon, and E. Soltas (2023). Self-targeting in US transfer programs. *Available at SSRN 4495537*.

Rochet, J.-C. (1991). Incentives, Redistribution and Social Insurance. *Geneva Papers on Risk and Insurance Theory 16*(2), 143–165.

Saez, E. and S. Stantcheva (2016). Generalized Social Marginal Welfare Weights for Optimal Tax Theory. *American Economic Review 106*(1), 24–45.

Schmieder, J. F. and T. Von Wachter (2016). The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation. *Annual Review of Economics 8*(1), 547–581.

Schmieder, J. F., T. Von Wachter, and S. Bender (2012). The Effects of Extended Unemployment Insurance Over the Business Cycle: Evidence from Regression Discontinuity Estimates Over 20 Years. *Quarterly Journal of Economics 127*(2), 701–752.

Shimer, R. and I. Werning (2008). Reservation wages and unemployment insurance. *Quarterly Journal of Economics 112*(3).

Spaziani, S. (2024). Optimal Unemployment Insurance Financing: Theory and Evidence from Two US States. *mimeo*.

Spinnewijn, J. (2015). Unemployed but optimistic: Optimal insurance design with biased beliefs. *Journal of the European Economic Association 13*(1), 130–167.

Spinnewijn, J. (2020). The trade-off between insurance and incentives in differentiated unemployment policies. *Fiscal Studies 41*(1), 101–127.

Von Wachter, T. (2019). Unemployment insurance reform. *The ANNALS of American Academy of Political and Social Science 686*(1), 121–146.

# Appendix: For online publication

## Risk Protection and Redistribution in the Design of Social Insurance

Rory McGee and Martin O'Connell
April 2026

## 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), willingness-to-pay decomposes into components reflecting insurance and cross-subsidization. The difference is that, in the dynamic model, these components are evaluated as expectations over the individual's lifetime.

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

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

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

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

When in the high state the individual earns $y_i - \mathcal{T}(y_i)$ and when in the low state they receive benefits $\mathcal{B}(y_i)$. The individual 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: $A_{i,1}(\varphi_{i,1}) > \bar{A}_i$ for all $\varphi_{i,1}$. They also face $N$ additional constraints in each state $\varphi_{i,t}$ at each time $t$:

$$g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}); \varphi_{i,t}) \geq 0$$

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

$$\begin{aligned}
\max & \int_0^1 \int_{\varphi_{i,t}} \left(\xi_{i,t}(\varphi_{i,t})u^h_i(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t})) + (1 - \xi_{i,t}(\varphi_{i,t}))u^l_i(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))\right) dF_{i,t}(\varphi_{i,t})dt \\
& + \int_0^1 \int_{\varphi_{i,t}} \lambda^A_{it}(\varphi_{i,t})\left(\xi_{i,t}(\varphi_{i,t})\left(y_i - \mathcal{T}(y_i)\right) + (1 - \xi_{i,t}(\varphi_{i,t}))\mathcal{B}(y_i) + f_{i,t}(x_{i,t}(\varphi_{i,t})) - c_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t})dt \\
& + \int_{\varphi_{i,1}} \lambda^A_{i,1}(\varphi_{i,1})(A_{i1}(\varphi_{i,1}) - \bar{A}_i)dF_{i,1}(\varphi_{i,1}) \\
& + \sum_{n=1}^{N} \int_0^1 \int_{\varphi_{i,t}} \lambda^n_{it}(\varphi_{i,t})g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}); \varphi_{i,t})dF_{i,t}(\varphi_{i,t})dt
\end{aligned}$$

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

Assumption 1 (Regularity Conditions). *Assume*

*i. Total lifetime utility is smooth, increasing and strictly quasi-concave in $(c, x)$*

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

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

*where (i) and (ii) ensure the individual'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:

$$\begin{aligned}
\frac{d\mathcal{B}(y_i)}{d\theta} &= \phi_{\mathcal{B}}(y_i) \\
\frac{d\mathcal{T}(y_i)}{d\theta} &= \phi_{\mathcal{T}}(y_i) \times \frac{\int_{i'} \int_0^1 \int_{\varphi_{i',t}} \phi_{\mathcal{B}}(y_{i'})\left(1 - \xi_{i',t}(\varphi_{i',t})\right) dF_{i',t}(\varphi_{i',t})dtdi'}{\int_{i'} \int_0^1 \int_{\varphi_{i',t}} \phi_{\mathcal{T}}(y_{i'})\xi_{i',t}(\varphi_{i',t})dF_{i',t}(\varphi_{i',t})dtdi'}
\end{aligned}$$

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

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

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

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:

$$\begin{aligned}
\frac{dV_i}{d\theta} = &-\int_0^1 \int_{\varphi_{i,t}} \left( \xi_{i,t}(\varphi_{i,t})\lambda^A_{i,t}(\varphi_{i,t}) - \sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t}) \frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}); \varphi_{i,t})}{\partial \mathcal{T}(y_i)} \right) dF_{i,t}(\varphi_{i,t})dt \frac{d\mathcal{T}(y_i)}{d\theta} \\
&+ \int_0^1 \int_{\varphi_{i,t}} \left( (1-\xi_{i,t}(\varphi_{i,t}))\lambda^A_{i,t}(\varphi_{i,t}) + \sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t}) \frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}); \varphi_{i,t})}{\partial \mathcal{B}(y_i)} \right) dF_{i,t}(\varphi_{i,t})dt \frac{d\mathcal{B}(y_i)}{d\theta}
\end{aligned}$$

Under the regularity assumptions on the constraints, individual $i$'s optimal consumption choice satisfies:

$$\frac{\partial u^s_i(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} = \lambda^A_{i,t}(\varphi_{i,t}) - \sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t}) \frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}$$

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

$$\begin{aligned}
\sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t}) \frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}); \varphi_{i,t})}{\partial \mathcal{T}(y_i)} &= \sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t})\xi_{i,t}(\varphi_{i,t}) \frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})} \\
\sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t}) \frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}); \varphi_{i,t})}{\partial \mathcal{B}(y_i)} &= -\sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t})(1-\xi_{i,t}(\varphi_{i,t})) \frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}
\end{aligned}$$

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

$$\begin{aligned}
\frac{dV_i}{d\theta} = &-\int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t}) \frac{\partial u^h_i(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t})dt \\
&\times \frac{\int_{i'} \int_0^1 \int_{\varphi_{i',t}} \phi_{\mathcal{B}}(y_{i'}) \left(1-\xi_{i',t}(\varphi_{i',t})\right) dF_{i',t}(\varphi_{i',t})dtdi'}{\int_{i'} \int_0^1 \int_{\varphi_{i',t}} \phi_{\mathcal{T}}(y_{i'})\xi_{i',t}(\varphi_{i',t})dF_{i',t}(\varphi_{i',t})dtdi'} \\
&+ \int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)(1-\xi_{i,t}(\varphi_{i,t})) \frac{\partial u^l_i(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t})dt
\end{aligned}$$

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

$$\text{WTP}_i \equiv \frac{dV_i}{d\theta} \times \left[ \frac{\frac{1}{\int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)(1-\xi_{i,t}(\varphi_{i,t}))dF_{i,t}(\varphi_{i,t})dt}}{\frac{\int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t})\frac{\partial u_i^h(c_{i,t}(\varphi_{i,t}),x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})}dF_{i,t}(\varphi_{i,t})dt}{\int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t})dF_{i,t}(\varphi_{i,t})dt}} \right]$$

The scaling expresses the utility gains from the reform per \$1 increase in expected lifetime benefit payments, measured in units of the average of marginal utility 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:

$$\widetilde{\mathbb{E}}_i^h\left[\frac{\partial u^h(c,x)}{\partial c}\right] = \int_0^1 \int_{\varphi_{i,t}} \frac{\phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t})}{\int_{t'}\int_{\varphi_{i,t'}} \phi_{\mathcal{T}}(y_i)\xi_{i,t'}(\varphi_{i,t'})dF_{i,t'}(\varphi_{i,t'})dt'} \frac{\partial u_i^h(c_{i,t}(\varphi_{i,t}),x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t})dt$$

$$\widetilde{\mathbb{E}}_i^l\left[\frac{\partial u^l(c,x)}{\partial c}\right] = \int_0^1 \int_{\varphi_{i,t}} \frac{\phi_{\mathcal{B}}(y_i)(1-\xi_{i,t}(\varphi_{i,t}))}{\int_{t'}\int_{\varphi_{i,t'}} \phi_{\mathcal{B}}(y_i)(1-\xi_{i,t'}(\varphi_{i,t'}))dF_{i,t'}(\varphi_{i,t'})dt'} \frac{\partial u_i^l(c_{i,t}(\varphi_{i,t}),x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t})dt$$

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

$$\text{WTP}_i = \frac{\widetilde{\mathbb{E}}_i^l\left[\frac{\partial u^l(c,x)}{\partial c}\right] - \widetilde{\mathbb{E}}_i^h\left[\frac{\partial u^h(c,x)}{\partial c}\right]}{\widetilde{\mathbb{E}}_i^h\left[\frac{\partial u^h(c,x)}{\partial c}\right]} +$$
$$1 - \left(\frac{\int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t})dF_{i,t}(\varphi_{i,t})dt}{\int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)(1-\xi_{i,t}(\varphi_{i,t}))dF_{i,t}(\varphi_{i,t})dt}\right) \Big/ \left(\frac{\int_{i'}\int_0^1 \int_{\varphi_{i',t}} \phi_{\mathcal{T}}(y_{i'})\xi_{i',t}(\varphi_{i',t})dF_{i',t}(\varphi_{i',t})dtdi'}{\int_{i'}\int_0^1 \int_{\varphi_{i',t}} \phi_{\mathcal{B}}(y_{i'})\left(1-\xi_{i',t}(\varphi_{i',t})\right)dF_{i',t}(\varphi_{i',t})dtdi'}\right),$$

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 $\rho > 0$, defined such that $u_i^{h\prime}(c) = u_i'(c)$ and $u_i^{l\prime}(c) = \rho u_i'(c)$.

In this case, the individual's WTP is:

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

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 around $c_i^h$:

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

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 $\rho$ 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 quantify the welfare effects of unemployment insurance expansion. We use data from 1999–2019 biennial PSID waves, which include information on consumption expenditures and asset holdings. For our baseline sample, we focus on non-immigrant households.[1] We include both single and married households and do not select the sample on the basis of the reference person's gender.[2] Our sample has non-missing information on key demographics (age, education, and state of

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

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

residence).

We restrict attention to households in which the reference person is aged 23–60 and participating in the labor force. We therefore exclude observations in which the reference person is retired, permanently disabled and neither working nor looking for work, on sick (or maternity) leave, temporarily laid off, in education, a homemaker, or in prison. In addition, to proxy for UI program eligibility we exclude those who never report employment, as well as the first observation following a period out of the labor force. We retain observations in which the reference person is either employed or unemployed, and we do not condition on spousal employment status, since we are explicitly interested in household exposure to unemployment risk, including self-insurance through spousal labor supply (Blundell et al. 2016).

To reduce the influence of measurement error, we drop observations with extremely high asset values and observations that exhibit implausibly large fluctuations in key outcomes. Following Blundell et al. (2016), we exclude households with net worth greater than $20 million and drop observations with extreme "jumps" in earnings, household income, and consumption which we attribute to measurement error. A "jump" is defined as an extreme 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 compute the biennial log difference $\Delta_2 \ln(x_t)$ and drop the observation in the bottom 0.25 percent of the product $\Delta_2 \ln(x_t)\Delta_2 \ln(x_{t-2})$. Furthermore, in constructing permanent income we exclude earnings observations when the implied hourly wage is below one-half 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 schooling: less than high school, high school graduate, some college (including those who dropout of four-year degrees and those who attain a community college degree or other diploma), four-year college graduate; and postgraduate education. 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 unemployment risk and both levels and growth rates in consumption. The PSID measures disaggregated consumption across a number of different household expenditure categories designed to cover approximately 70% of aggregate consumption. Respondents also report the reference period for the expenditure item. We convert all expenditures to annual values (e.g., multiplying weekly expenditures by 52 and monthly expenditures by 12) and treat missing values in subcategories as zeros. We focus on expenditure categories measured consistently across survey waves. Our primary consumption measure captures a range of utility-relevant expenditures comprising nondurable purchases plus service. We also show that our results are robust to alternative consumption measures (including food consumption as in Gruber (1997) and Hendren (2017)), which we describe below. 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: Nondurable Expenditure including Services.*** Our baseline measure includes a broad range of nondurable expenditures and services, excluding categories with an investment or durable component (e.g., vehicle maintenance). To build our baseline consumption series, we first construct a food-expenditure series by summing food consumed at home, food consumed away from home, and food purchased using Supplemental Nutrition Assistance Program (SNAP). Including SNAP-financed food expenditures (formerly the Food Stamp Program) is important because they are a relevant source of consumption financing for low-income households, and Low and Pistaferri (2015) show that food stamps can act as substitutes for social insurance.

We then construct a service-expenditure series excluding durables. We sum spending on home and automobile insurance, utilities, parking, and other direct transportation costs (e.g., bus fare and taxi payments) that do not correspond to vehicle maintenance, and child care.

Finally, we combine the aggregated food series with gasoline expenditures and the services series without a durable component.

***Food Expenditure.*** To construct a food-expenditure series that excludes other components of our baseline measure, we sum food consumed at home, food be consumed away from home, and food purchased using SNAP benefits.

***Services Expenditure.*** To construct a series for broad services, including categories 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 health-care-related spending (out-of-pocket payments including for hospital and nursing home stays, doctor visits, and prescription drugs as well as insurance premiums), 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 above to produce a household-level series for total non-housing expenditures. As the PSID consumption categories are not designed to have full coverage, the term "total" is a misnomer. We continue to exclude durable purchases (e.g., vehicles) and memory goods (Hai et al. 2020), such as vacations, that have durable-like properties.

***Total (Including Housing) Expenditure.*** Our final consumption measure incorporates housing services by adding the consumption value of housing services to the total non-housing expenditures. 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 unemployment risk. Here we describe how we construct that measure.

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

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

where $w_{i,t}$ denotes annual labor earnings of the reference person. We control for a third-order polynomial in age and year fixed effects.

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

| | All | Employed | Unemployed |
|---|---|---|---|
| Ref. Person's Age (years) | 39.55 | 39.67 | 37.60 |
| | (10.29) | (10.28) | (10.30) |
| Married | 0.54 | 0.55 | 0.30 |
| | (0.50) | (0.50) | (0.46) |
| Number of Children | 1.03 | 1.02 | 1.15 |
| | (1.21) | (1.20) | (1.38) |
| Share White | 0.54 | 0.56 | 0.32 |
| | (0.50) | (0.50) | (0.47) |
| Share Black | 0.13 | 0.13 | 0.12 |
| | (0.33) | (0.33) | (0.33) |
| Ref. Person's Schooling (years) | 13.79 | 13.85 | 12.77 |
| | (2.45) | (2.44) | (2.41) |
| Unemployed | 0.06 | 0.00 | 1.00 |
| | (0.24) | (0.00) | (0.00) |
| Non-Durable Expenditure ($1,000s) | 20.49 | 20.85 | 14.90 |
| | (10.73) | (10.74) | (8.88) |
| Food Expenditure ($1,000s) | 9.60 | 9.74 | 7.35 |
| | (5.74) | (5.76) | (4.87) |
| Services Expenditure ($1,000s) | 16.24 | 16.59 | 10.65 |
| | (13.28) | (13.35) | (10.66) |
| Total Expenditure (Non-Housing, $1,000s) | 25.72 | 26.21 | 18.05 |
| | (16.30) | (16.34) | (13.52) |
| Total Expenditure (Incl. Housing, $1,000s) | 38.49 | 39.29 | 25.81 |
| | (24.92) | (24.98) | (19.97) |
| N | 55257 | 51974 | 3283 |
| Waves | 3.69 | 3.48 | 1.51 |
| Unique Households | 14995 | 14936 | 2178 |

TABLE B.1. Summary Statistics for Our Sample of Interest

Notes: Computed from the 1999–2019 waves of the PSID. Table reports means, with standard deviations in parenthesis. We deflate all dollar to 2019 dollars.

## C. Estimation of the Joint Risk–Consumption Mixture

This appendix describes the estimation of the joint latent-type model for unemployment risk and unemployment-induced consumption declines. We first present the specifications for the unemployment-risk and consumption-drop blocks, then derive the observed-data likelihood and the conditional consumption likelihood, and finally describe our estimation procedure.

### C.1. Model specification

We observe a panel of workers indexed by $i = 1, \ldots, N$ and survey waves indexed by $t = 1, \ldots, T_i$. Let $\Delta^{FD}_{i,t} \equiv \ln c_{i,t} - \ln c_{i,t-2}$ denote the two-year change in log consumption, and let $U_{i,t}$ denote an indicator for unemployment in period $t$. Let $X^B_{i,t}$ denote a vector of block-specific covariates, and let $o^B_{i,t}$ indicate whether $(i, t)$ is observed in our estimation sample for block $B \in \{U, C\}$ (unemployment-risk block $U$ and consumption block $C$).

Each worker belongs to one of $H$ unobserved *unemployment-risk types* and one of $K$ unobserved *consumption types*. We denote the risk type by $h(i) \in \{1, \ldots, H\}$ and the consumption type by $k(i) \in \{1, \ldots, K\}$. We sometimes use indicator notation

$$d^H_{i,h} \equiv \mathbb{1}[h(i) = h], \qquad d^K_{i,k} \equiv \mathbb{1}[k(i) = k].$$

***Unemployment risk.*** The unemployment-risk block corresponds to the specification in the main text: conditional on risk type $h$ and covariates $X^U_{i,t}$, unemployment in period $t$ follows a logit model

$$\Pr\left(U_{i,t} = 1 \mid h(i) = h, X^U_{i,t}\right) = g\left(\eta^h_0 + X^{U\prime}_{i,t}\gamma\right), \tag{C.1}$$

where $g(\cdot)$ is the logistic function, $\eta^h_0$ is a type-specific intercept, and $\gamma$ is a vector of common coefficients on $X^U_{i,t}$. We assume that, conditional on $(h(i), X^U_{i,1:T_i})$, unemployment shocks are independent over $t$ and across workers. Let $\pi(h)$ denote the population share of risk type $h$.

**Consumption declines**

The consumption-decline block follows the finite-mixture specification in the main text (equation (7)):

$$\Delta^{FD}_{i,t} = \sum_{k=1}^{K} \mathbb{1}[k(i) = k]\left(\delta^k_0 + \delta^k_1 U_{i,t}\right) + X^{C\prime}_{i,t}\beta + \varepsilon_{i,t}, \tag{C.2}$$

where $\delta_0^k$ and $\delta_1^k$ are type-specific intercepts and consumption decline, respectively, and $\beta$ collects the coefficients on observables. Conditional on consumption type $k$, observables, and unemployment status, we assume that:

$$\varepsilon_{i,t} \mid \left(k(i) = k, X_{i,t}^C, U_{i,t}\right) \sim N\left(0, \sigma_k^2\right),$$

and that the shocks $\varepsilon_{i,t}$ are independent over $t$ and across workers, conditional on $(k(i), X_{i,1:T_i}^C, U_{i,1:T_i})$.

***Independence across blocks.*** Finally, we assume that, conditional on the latent types and observables, the unemployment and consumption shocks are independent:

$$\begin{aligned}
u_{i,t} &\perp u_{i,s} && \mid X_{i,1:T_i}^U, h(i), \ \forall t \neq s, \\
\varepsilon_{i,t} &\perp \varepsilon_{i,s} && \mid X_{i,1:T_i}^C, k(i), \ \forall t \neq s, \\
(\varepsilon_{i,1:T_i}) &\perp (u_{i,1:T_i}) && \mid X_{i,1:T_i}^U, X_{i,1:T_i}^C, h(i), k(i),
\end{aligned}$$

where $u_{i,t}$ denotes the latent logistic error underlying equation (C.1). This assumption, combined with equations (C.1)–(C.2), implies that the joint likelihood factorizes into an unemployment-risk component and a consumption component, conditional on type.

### C.2. Likelihood

Let

$$\begin{aligned}
\chi_U &= \left\{\pi(h), \eta_0^h, \gamma\right\}_{h=1}^H, \\
\chi_C &= \left\{\delta_0^k, \delta_1^k, \sigma_k^2, \beta\right\}_{k=1}^K
\end{aligned}$$

denote the parameters of the unemployment-risk and consumption blocks, respectively. In addition, let $\pi(h, k)$ denote the joint population share of risk type $h$ and consumption type $k$, satisfying

$$\sum_{h=1}^{H} \sum_{k=1}^{K} \pi(h, k) = 1, \qquad \pi(h, k) \geq 0 \ \ \forall h, k,$$

with marginal shares $\pi(h) = \sum_{k=1}^K \pi(h, k)$ and $\pi(k) = \sum_{h=1}^H \pi(h, k)$.

***Likelihood contributions conditional on type.*** For a given worker $i$ and consumption type $k$, the conditional likelihood contribution of the consumption data is

$$\begin{aligned}\ell_i^C(k;\chi_C) &\equiv \ell_i^C(\Delta_i^{FD}, U_i, X_i^C \mid k(i)=k;\chi_C)\\ &= \prod_{t=1}^{T_i} f\Big(\Delta_{i,t}^{FD} \,\Big|\, \underbrace{\delta_0^k + \delta_1^k U_{i,t} + X_{i,t}^{C\prime}\beta}_{\mu^k(U_{i,t},X_{i,t}^C)}, \sigma_k^2\Big)^{o_{i,t}^C},\end{aligned} \tag{C.3}$$

where $f(\cdot \mid \mu, \sigma^2)$ denotes the normal density with mean $\mu$ and variance $\sigma^2$.

For a given worker $i$ and risk type $h$, the conditional likelihood contribution of the employment data is

$$\begin{aligned}\ell_i^U(h;\chi_U) &\equiv \ell_i^U(U_i, X_i^U \mid h(i)=h;\chi_U)\\ &= \prod_{t=1}^{T_i}\Big(g(\eta_0^h + X_{i,t}^{U\prime}\gamma)^{U_{i,t}}\big[1-g(\eta_0^h + X_{i,t}^{U\prime}\gamma)\big]^{1-U_{i,t}}\Big)^{o_{i,t}^U}.\end{aligned} \tag{C.4}$$

Conditional on $(h(i)=h, k(i)=k)$, the joint likelihood contribution for individual $i$ factorizes:

$$\ell_i(h,k;\chi_U,\chi_C) = \ell_i^U(h;\chi_U)\,\ell_i^C(k;\chi_C). \tag{C.5}$$

***Observed-data likelihood.*** Because the latent types are unobserved, the observed-data likelihood for individual $i$ sums over all $(h,k)$ pairs:

$$\begin{aligned}\ell_i(\pi;\chi_U,\chi_C) &\equiv p(\Delta_i^{FD}, U_i \mid X_i^C, X_i^U;\pi,\chi_U,\chi_C)\\ &= \sum_{h=1}^{H}\sum_{k=1}^{K}\pi(h,k)\,\ell_i^U(h;\chi_U)\,\ell_i^C(k;\chi_C).\end{aligned} \tag{C.6}$$

The sample log-likelihood is then

$$\mathcal{L}(\pi;\chi_U,\chi_C) = \sum_{i=1}^{N}\ln \ell_i(\pi;\chi_U,\chi_C).$$

***Conditional consumption likelihood given employment outcomes.*** For estimation, it is helpful to consider the conditional likelihood of the consumption process given the employment history and covariates:

$$\ell_i^{C|U}(\pi;\chi_U,\chi_C) \equiv p(\Delta_i^{FD} \mid U_i, X_i^C;\pi,\chi_U,\chi_C)$$

$$= \frac{\sum_{h=1}^{H}\sum_{k=1}^{K}\pi(h,k)\,\ell_i^U(h;\chi_U)\,\ell_i^C(k;\chi_C)}{\sum_{b=1}^{H}\sum_{a=1}^{K}\pi(b,a)\,\ell_i^U(b;\chi_U)}. \tag{C.7}$$

Define the individual-specific mixing weights

$$\alpha_i(k;U_i,X_i^U) \equiv \frac{\sum_{h=1}^{H}\pi(h,k)\,\ell_i^U(h;\chi_U)}{\sum_{b=1}^{H}\sum_{a=1}^{K}\pi(b,a)\,\ell_i^U(b;\chi_U)}, \tag{C.8}$$

so that equation (C.7) can be written as

$$\ell_i^{C|U}(\pi;\chi_U,\chi_C) = \sum_{k=1}^{K}\alpha_i(k;U_i,X_i^U)\,\ell_i^C(k;\chi_C). \tag{C.9}$$

In general, $\alpha_i(k;U_i,X_i^U)$ depends on the full employment history $U_i$ and the covariates $X_i^U$. Estimating the consumption model in isolation via

$$\prod_{i=1}^{N}\sum_{k=1}^{K}\pi(k)\,\ell_i^C(k;\chi_C)$$

implicitly imposes $\alpha_i(k;U_i,X_i) \equiv \pi(k)$ for all $i, k$, which holds only under additional restrictions (e.g., $h(i)$ independent of $k(i)$). The joint model (equation (C.6)) therefore provides a coherent way to allow for correlation between unobserved heterogeneity in unemployment risk and consumption declines.

### C.3. Estimation procedure

Estimation proceeds in two stages, as described in the main text. In the first stage, we estimate the unemployment-risk mixture and obtain $\hat{\chi}_U$ by maximizing the marginal likelihood of workers' unemployment histories. In the second stage, conditional on $\hat{\chi}_U$, we estimate the consumption parameters $\chi_C$ and the joint type probabilities $\pi(h,k)$ by maximum likelihood using an EM algorithm.[3]

***Stage 1: Unemployment-risk mixture.*** We estimate the parameters of the risk block, $\chi_U \equiv \{\pi(h), \eta_0^h, \gamma\}_{h=1}^{H}$, by maximizing

$$\hat{\chi}_U \in \arg\max_{\chi_U}\prod_{i=1}^{N}\sum_{h=1}^{H}\pi(h)\,\ell_i^U(h;\chi_U), \tag{C.10}$$

[3]As in Lewis et al. (2026), we initialize the EM algorithm from a large number of random starting values and retain the solution with the highest log-likelihood.

subject to $\sum_{h=1}^{H} \pi(h) = 1$ and $\pi(h) \geq 0$ for all $h$.

In the EM algorithm, the E-step computes the posterior risk-type probabilities

$$\begin{aligned}\tau_i(h) &\equiv \Pr\left(h(i) = h \mid U_i, X_i^U; \pi, \chi_U\right) \\ &= \frac{\pi(h)\, \ell_i^U\left(h; \chi_U\right)}{\sum_{b=1}^{H} \pi(b)\, \ell_i^U\left(b; \chi_U\right)}.\end{aligned} \tag{C.11}$$

The M-step updates the type probabilities and the logit parameters. The optimal mixture weights are

$$\pi^{\text{new}}(h) = \frac{1}{N} \sum_{i=1}^{N} \tau_i(h),$$

and the logit parameters solve

$$(\{\eta_0^h\}_{h=1}^{H}, \gamma)^{\text{new}} = \arg\max_{\{\eta_0^h\}_h, \gamma} \sum_{i=1}^{N} \sum_{h=1}^{H} \tau_i(h)\, \ln \ell_i^U\left(h; \chi_U\right),$$

which corresponds to a weighted logistic regression with type-specific intercepts and a common slope vector.

***Stage 2: Consumption mixture and joint type probabilities.*** Given $\hat{\chi}_U$ from Stage 1, we estimate $\chi_C$ and the joint type probabilities $\pi(h, k)$ by maximizing the plug-in likelihood

$$(\{\pi(h, k)\}_{h,k}, \chi_C) \in \arg\max_{\{\pi(h,k)\}_{h,k}, \chi_C} \prod_{i=1}^{N} \sum_{h=1}^{H} \sum_{k=1}^{K} \pi(h, k)\, \ell_i^U\left(h; \hat{\chi}_U\right) \ell_i^C\left(k; \chi_C\right), \tag{C.12}$$

subject to

$$\sum_{k=1}^{K} \pi(h, k) = \hat{\pi}(h) \quad \forall h, \qquad \sum_{h=1}^{H} \sum_{k=1}^{K} \pi(h, k) = 1, \qquad \pi(h, k) \geq 0 \ \ \forall h, k. \tag{C.13}$$

The constraints (C.13) ensure that the marginal risk-type shares in the joint distribution coincide with the Stage 1 estimates.

As a preliminary step, we estimate a consumption-only mixture that ignores correlation between $h(i)$ and $k(i)$:

$$(\{\tilde{\pi}(k)\}_{k=1}^{K}, \tilde{\chi}_C) \in \arg\max_{\{\pi(k)\}_k, \chi_C} \prod_{i=1}^{N} \sum_{k=1}^{K} \pi(k)\, \ell_i^C\left(k; \chi_C\right), \tag{C.14}$$

subject to $\sum_{k=1}^{K} \pi(k) = 1$ and $\pi(k) \geq 0$. These estimates are used to initialize estimation of the joint model:

$$\pi^{(0)}(h,k) = \hat{\pi}(h)\,\tilde{\pi}(k), \qquad \chi_C^{(0)} = \tilde{\chi}_C.$$

The EM algorithm for equation (C.12) proceeds as follows. In the E-step, for each individual we compute the posterior probability of each joint type:

$$\begin{aligned} p_i(h,k) &\equiv \Pr\left(k(i)=k, h(i)=h \mid \Delta_i^{FD}, U_i, X_i^C, X_i^U; \pi, \chi_C, \hat{\chi}_U\right) \\ &= \frac{\pi(h,k)\,\ell_i^U(h;\hat{\chi}_U)\,\ell_i^C(k;\chi_C)}{\sum_{b=1}^{H}\sum_{a=1}^{K}\pi(b,a)\,\ell_i^U(b;\hat{\chi}_U)\,\ell_i^C(a;\chi_C)}. \end{aligned} \tag{C.15}$$

In the M-step, we update the joint type probabilities and the consumption parameters. The constrained update for the joint type probabilities is

$$\pi(h,k)^{\text{new}} = \hat{\pi}(h)\cdot\frac{\sum_{i=1}^{N} p_i(h,k)}{\sum_{i=1}^{N}\sum_{a=1}^{K} p_i(h,a)}, \tag{C.16}$$

which enforce the constraints (C.13) exactly. The consumption parameters are updated by maximizing the expected complete-data log-likelihood,

$$\chi_C^{\text{new}} = \arg\max_{\chi_C} \sum_{i=1}^{N}\sum_{h=1}^{H}\sum_{k=1}^{K} p_i(h,k)\,\ln \ell_i^C(k;\chi_C), \tag{C.17}$$

which corresponds to estimating $K$ Gaussian regressions with type-specific parameters using weights

$$p_i^C(k) \equiv \sum_{h=1}^{H} p_i(h,k)$$

for each consumption type $k$.

We iterate the E- and M-steps until convergence in the log-likelihood (C.12). The resulting estimates $(\hat{\pi}(h,k), \hat{\chi}_U, \hat{\chi}_C)$ are used in the main text to construct worker-level unemployment probabilities and consumption declines as

$$\hat{e}_{i,t} = 1 - \sum_{h=1}^{H}\sum_{k=1}^{K} \hat{p}_i(h,k)\,\Pr\left(U_{i,t}=1 \mid h(i)=h, X_{i,t}^U; \hat{\chi}_U\right), \quad \widehat{(\Delta c/c)}_i \approx \widehat{\Delta \ln c_i} = \sum_{h=1}^{H}\sum_{k=1}^{K} \hat{p}_i(h,k)\,\hat{\delta}_1^k.$$

***Inference.*** We obtain standard errors for all parameters and welfare objects using a nonparametric bootstrap. Specifically, we draw 400 bootstrap samples by resampling households with replacement, re-estimate the full model in each sample, and compute the distribution of mixture parameters, type probabilities, and welfare statistics. We report bootstrap standard errors based on these 400 replications.

# 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, since redistributive policies are typically not designed with the explicit purpose of redistributing within these groups. Here, we assess the robustness to the inclusion of additional controls. Figure D.1 shows that average consumption differences remain similar when these controls are added. The right panel—which focuses on within-group gender/education/race comparisons—shows some attenuation, reflecting the fact that we control for additional household characteristics relevant to consumption and labor market outcomes. Nevertheless, the quantitative impact of including these additional controls is modest.

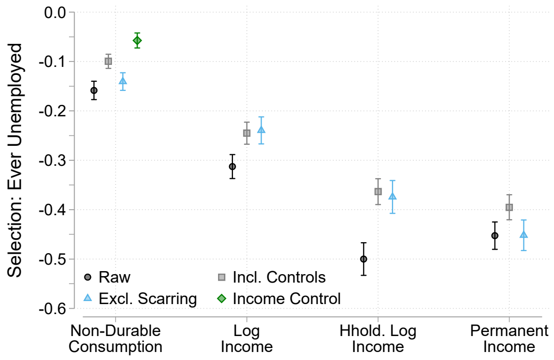


(A) Baseline

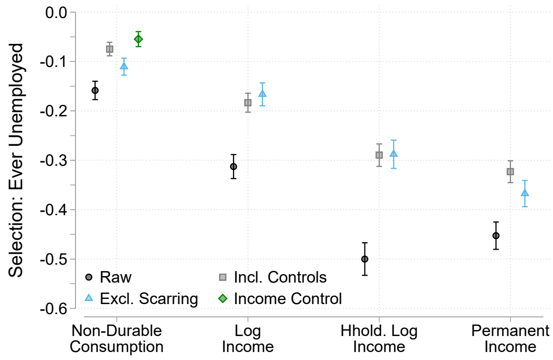


(B) Additional Controls

FIGURE D.1. Negative Selection Among the Unemployed

Notes: All regressions drop periods during which the reference person is unemployed. For all outcomes, we report the raw, controls, and excl. scarring specifications. The raw specification includes no controls. The remaining specifications control for a cubic in age, indicators for family size and marital status, and year fixed effects. Estimates in panel (B) also control for indicators for gender, education, and race. The excl. scarring specification further restricts the sample to periods before the first observed unemployment spell. For nondurable consumption only, we also report an "income controls" specification that additionally controls for a cubic in household income.

## D.2. Posterior-implied Consumption Decline Confidence Bands

Panel (A) of Figure D.2 plots the distribution of posterior-implied consumption declines across workers (replicating Figure 3) and the right panel plots quantiles of the distribution along with 95% confidence intervals.

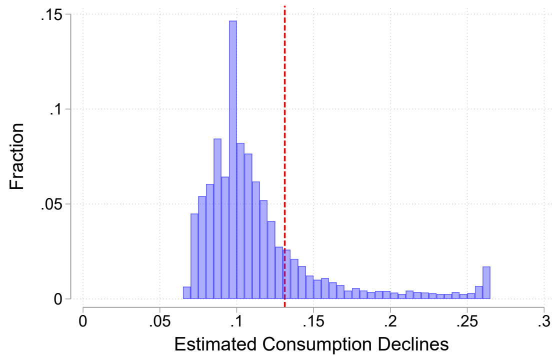


(A) Distribution

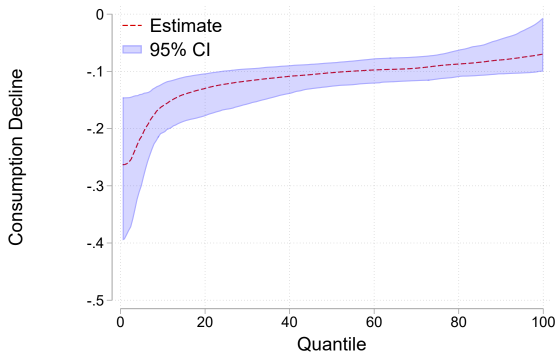


(B) Quantiles and CI

FIGURE D.2. Distribution of Estimated Consumption Declines

Notes: Panel (A) plots the distribution of posterior-implied unemployment-induced consumption declines under our baseline nondurable-consumption measure. Panel (B) reports the corresponding quantiles and 95% bootstrap confidence intervals. Worker-level declines are constructed from the Gaussian mixture linear regression by combining the estimated type-specific unemployment coefficients $\delta_1^k$ with each worker's posterior type probabilities $p_i(k)$. The sample is defined as in the text. The Bayesian Information Criterion selects $K = 3$. The homogeneous estimate (dashed vertical line in left panel) is obtained by imposing $K = 1$.

## D.3. Consumption Decline Estimates Under Alternative Consumption Measures

In addition to our baseline consumption measure, we assess robustness to alternative expenditure series. Table D.1 shows that estimated average consumption decline are similar across measure.

| | Baseline | Alternative Measures | | | |
|---|---|---|---|---|---|
| | Non-durable | Food Expenditure | Services | Total (Non-Housing) | Total (Incl. Housing) |
| $\mathbb{1}\left[U_{i,t} = 1\right]$ | $-0.131^{***}$ | $-0.148^{***}$ | $-0.145^{***}$ | $-0.129^{***}$ | $-0.149^{***}$ |
| | (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, 143 | 36, 880 | 37, 126 | 37, 167 | 37, 170 |

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

Notes: $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$. The first column reports the estimated average consumption decline from the homogeneous specification (imposing a common unemployment coefficient) for our baseline consumption measure. The remaining columns report estimated declines for alternative consumption measures constructed from PSID expenditures. Standard errors are clustered at the household level. The number of observations differs across regressions due to item non-response. Covariates include a cubic in age, indicators for year, and the log of the change in family size.

### D.4. Pre-trends and Anticipation Effects

To assess the identifying assumption that, conditional on type, consumption growth for workers who remain employed provides a valid counterfactual for those who become unemployed, we follow the lead–lag exercise in Hendren (2017). Specifically, we estimate variants of the homogeneous consumption-drop specification in equation (6) that replace the contemporaneous unemployment indicator $U_{i,t}$ (equal to one if the worker is employed at $t-2$ and unemployed at $t$) with leads and lags $U_{i,t-k}$ for $k \in \{-6, -4, -2, 0, 2, 4, 6\}$:

$$\Delta_{i,t}^{FD} \equiv \ln c_{i,t} - \ln c_{i,t-2} = \delta_0 + \delta_1 U_{i,t-k} + X_{i,t}^{C\,\prime}\beta + \varepsilon_{i,t}. \tag{D.1}$$

Negative values of $k$ capture "leads", while positive values correspond to "lags". We estimate this specification (i) in the homogeneous model and (ii) using posterior type probabilities from the joint mixture model as weights to obtain type-specific profiles.[4]

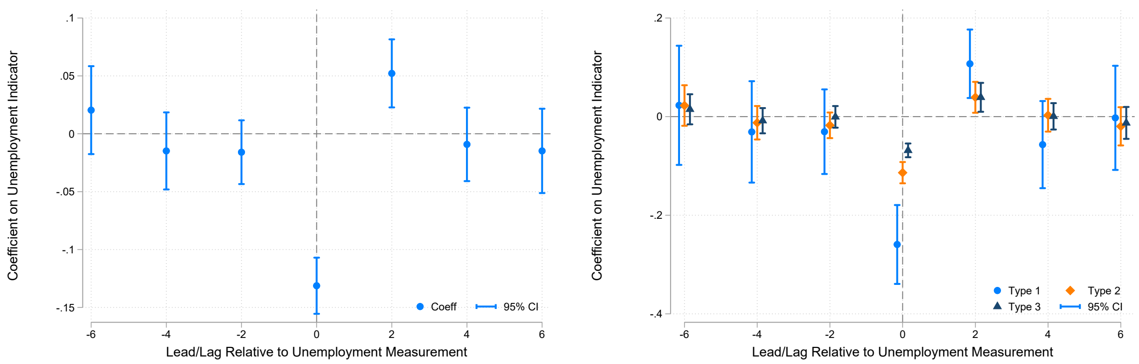


(A) Homogeneous estimates

(B) Type-specific estimates

FIGURE D.3. Impact of Unemployment on Consumption Growth

Notes: Each point plots the estimated effect of unemployment on two-year log consumption growth at a given event time $k$, with vertical bars showing 95% confidence intervals. Event time $k = 0$ corresponds to workers who are employed at $t-2$ and unemployed at $t$; negative (positive) values of $k$ are leads (lags) relative to this unemployment event. All specifications control for the same covariates as in the baseline homogeneous regression and restrict the sample to workers employed at $t-2$. Panel (A) reports homogeneous estimates from equation (D.1) with the unemployment indicator replaced by lead/lag indicators $U_{i,t-k}$. Panel (B) reports type-specific estimates obtained by re-estimating these regressions using posterior type probabilities from the joint mixture model as weights (treating the posteriors as given). Standard errors are clustered at the household level.

[4]We obtain type-specific profiles by re-estimating the lead–lag regressions with each observation weighted by the worker's posterior probability of belonging to each latent type (treating these posteriors as given). This "soft classification" approach is standard in finite-mixture settings and avoids assigning each worker to a single type. Moreover, the generalized least squares estimator corresponds to the M-step in our mixture estimation.

Figure D.3 plots the estimated coefficients and 95% confidence intervals. We find little evidence of differential consumption trends prior to job loss in either the homogeneous or type-specific estimates: the lead coefficients are close to zero and statistically indistinguishable from zero, and do not display a systematic pre-trend.. We observe a sizable decline in the period when unemployment occurs, and some "bounce back" in consumption two years after the spell, consistent with partial recovery. While Hendren (2017) documents an anticipatory decline in (one-year) food consumption one year prior to unemployment, our biennial consumption measure does not exhibit this effect. Overall, the lead–lag patterns provide evidence in favor of the parallel-trends assumption underlying our consumption-drop estimates.

### D.5. Alternative Unemployment Definitions

We examine robustness to alternative definitions of job loss and to an IV strategy using plant closures. The PSID asks respondents who leave a job to report the reason for separation. We use these responses to construct an indicator for separations due to plant closure or firm shutdown (a proxy for mass layoffs).

Table D.2 reports estimates from the homogeneous consumption-decline specification under several alternative treatments of unemployment. Column (1) reproduces the baseline specification using our standard unemployment indicator. Column (2) instruments this indicator with the plant-closure indicator. Columns (3)–(4) use alternative measures of job loss based on plant closures and broader "displacement" definitions that include workers who report being fired or laid off, following Stephens (2002) and Charles and Stephens (2004). The IV and displacement-based estimates of the consumption decline are slightly smaller in magnitude than the baseline, but imprecisely estimated, and none is statistically distinguishable from the baseline estimate. As in prior PSID work, we find only modest differences between plant closures and the broader displaced group. These results suggest that selection into unemployment along dimensions captured by separation reason is unlikely to be a first-order source of bias in our estimated consumption drops.

| | Unemployed at Interview | | Alternative Definitions | |
|---|---|---|---|---|
| | Baseline | IV | Plant Closed | Displaced |
| $\mathbb{1}\left[U_{i,t}=1\right]$ | $-0.129^{***}$ | $-0.100^{**}$ | | |
| | (0.012) | (0.047) | | |
| $\mathbb{1}\left[U_{i,t}=1\right]\times\mathbb{1}\left[Closed(t,t-2)_{i,t}=1\right]$ | | | $-0.095^{**}$ | |
| | | | (0.045) | |
| $\mathbb{1}\left[U_{i,t}=1\right]\times\mathbb{1}\left[Displaced(t,t-2)_{i,t}=1\right]$ | | | | $-0.132^{***}$ |
| | | | | (0.015) |
| N | 37, 260 | 37, 260 | 37, 260 | 37, 260 |

TABLE D.2. Consumption Declines by Alternative Displacement Measures

Notes: $^{*}$ $p<0.10$, $^{**}$ $p<0.05$, $^{***}$ $p<0.01$. Column (1) reports the estimated average consumption decline from the homogeneous specification using our baseline unemployment definition. Column (2) instruments this unemployment indicator with an indicator for plant closures. The Kleibergen-Paap F-statistic is $1.8\times10^{5}$. The remaining columns report estimated consumption declines using alternative job-loss definitions. Standard errors are clustered at the household level. Covariates include a cubic in age, indicators for year, and the log of the change in family size.

### D.6. Validation of the Mixture Model

To validate our finite-mixture specification, we compare the distribution of consumption drops it implies to a quantile-regression benchmark estimated on an independent holdout sample. We randomly split workers into two equally sized subsamples (assigning all observations for a given worker to the same subsample). In the first (training) subsample, we estimate the joint risk–consumption mixture model and recover posterior probabilities over consumption types; we then use these to construct the implied distribution of unemployment-induced consumption declines among unemployed workers. In the second (holdout) subsample, we estimate conditional quantile regressions for consumption growth to recover a flexible estimate of the distribution of consumption drops among unemployed workers, along with 95% confidence bands.

Figure D.4 plots percentiles of the consumption-decline distribution implied by the mixture model in the training sample, reweighted by ex ante unemployment probabilities so that the implied distribution is comparable to the ATT distribution targeted by quantile regression in the holdout sample. It also plots the corresponding quantile-regression estimates and confidence intervals from the holdout sample. Across most of the distribution, the mixture-based percentiles lie well within the 95% confidence bands of the quantile-regression benchmark and are typically close to their midpoints. The main discrepancy arises in the extreme lower tail (roughly the bottom decile), where the quantile regressions imply somewhat larger drops than the mixture model, though with wide confidence intervals. We interpret this exercise as providing reassuring external validation that the finite-mixture specification does a good job of recovering the distribution of consumption responses to unemployment.

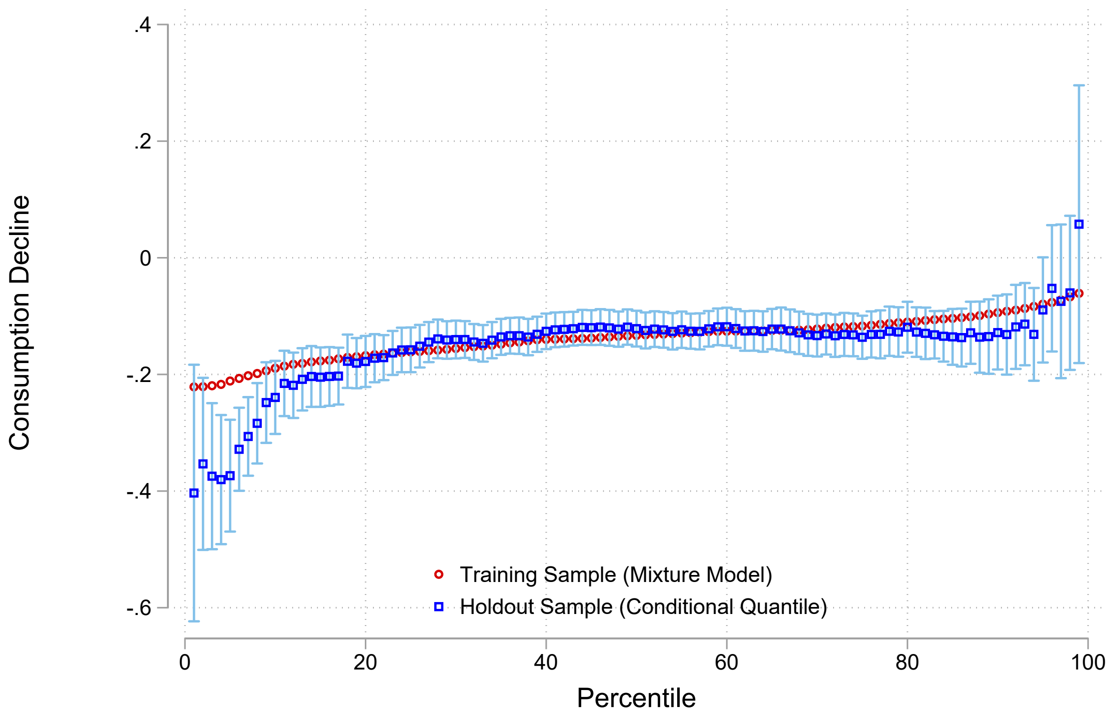


FIGURE D.4. Distribution of Consumption Declines: Prediction Sample vs Holdout

Notes: The red circles plot percentiles of the estimated consumption declines implied by the finite-mixture model, using a random training sample comprising 50% of observations, reweighted by ex ante unemployment probabilities. The blue squares plot the corresponding quantiles in the remaining holdout sample, estimated by quantile regression; shaded bands show 95% bootstrap confidence intervals. Consumption declines are defined as $\Delta^{FD}_{i,t} = \ln c_{i,t} - \ln c_{i,t-2}$.

### D.7. Robustness of Consumption Declines to Unemployment Spell Length

A potential concern with our consumption-based estimates is that the PSID's biennial structure may under-represent short unemployment spells and their associated consumption responses. In our baseline specification, we classify a worker as in unemployment at period $t$ when we observe them employed at $t-2$ but unemployed at $t$, so that $\Delta^{FD}_{i,t}$ compares pre- and post-unemployment consumption. Spells that begin and end between these two interview dates appear as employment at both $t-2$ and $t$ and therefore do not enter our "unemployed" group, tilting the sample toward longer spells for which measured consumption declines may be larger.

To assess the importance of this issue, we exploit additional information in the post–2003 PSID waves on unemployment between interviews. In these waves, respondents report whether they experienced any unemployment spell in the previous 12 months and in the preceding 12–24 months. We first replicate the homogeneous specification using only post–2003 waves (column 2 of Table D.3). We then re-estimate the same specification under alternative sample restrictions that drop observations reporting additional unemployment outside the contemporaneous interview state. Specifically, we consider three alternative restrictions based on the two unemployment-history indicators, excluding in turn: (i) observations with unemployment only in the preceding 12–24 months; (ii) observations with unemployment only in the previous 12 months; and (iii) observations with unemployment in either window. Columns 3–5 of Table D.3 report the resulting estimates.

| | Baseline | Post-2003 | No Spell Two Calendar Years Prior ($t-2$) | No Spell In Previous Calendar Year ($t-1$) | Neither |
|---|---|---|---|---|---|
| $\mathbb{1}\left[U_{i,t}=1\right]$ | $-0.131^{***}$ | $-0.133^{***}$ | $-0.132^{***}$ | $-0.117^{***}$ | $-0.110^{***}$ |
| | (0.012) | (0.013) | (0.014) | (0.017) | (0.018) |
| N | 37, 143 | 33, 449 | 31, 701 | 31, 459 | 30, 485 |

Table D.3. Consumption Declines under Alternative Unemployment Spell Definitions

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Column (1) reports the estimated average consumption decline from the homogeneous specification (equation (6)) in our full baseline sample. Column (2) repeats this specification using only post–2003 data, for which we observe whether the respondent experienced unemployment in the previous 12 months and in the preceding 12–24 months. The remaining columns report estimates from alternative specifications that impose sample restrictions based on these unemployment-history indicators, as described in the text. Standard errors are clustered at the household level. Covariates include a cubic in age, indicators for year, and the log of the change in family size.

The estimates in Table D.3 are similar across specifications. Restricting attention to shorter spells yields slightly smaller (but statistically indistinguishable) consumption drops relative to the post–2003 baseline. This pattern suggests that any bias from the biennial sampling frame is modest and, if anything, would lead us to slightly overstate the insurance value of UI in our baseline estimates.

### D.8. Observable Heterogeneity in Consumption Declines

Our approach to estimating unemployment-induced consumption declines does not require specifying observable determinants of heterogeneity ex ante. Ex post, however, we can relate our predicted declines to observables. Table D.4 reports the results of this exercise, which we use to understand how observable characteristics correlate with heterogeneity in consumption declines.

We focus on proxies for a household's capacity to self-insure, measured while employed. This enables us to assess whether predicted consumption declines are larger for households that appear less able to self-insure.[5] To proxy for the availability of private savings, we include indicators for each quartile of the liquid-wealth distribution and an indicator for home ownership. Following Carroll and Samwick (1997), we construct liquid wealth as the sum of cash (including checking accounts), savings, stocks held outside retirement account, and bonds and government treasuries—assets that are relatively liquid 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 losses during unemployment, and an indicator for marriage/cohabitation, which proxies for the additional insurance through added-work effect or spousal income effects.

Table D.4 reports coefficient estimates from a regression of predicted worker-level consumption declines on those observables. We find that, conditional on earnings quartile, households with greater apparent ability to self-insure (higher liquid assets, homeownership, and cohabitation) experience smaller consumption declines when unemployed. These magnitudes are economically and statistically significant. At the same time, these observables explain only a modest fraction of total variation in predicted declines (an $R^2$ of 0.11). This is consistent with two interpretations: (i) there is substantial latent heterogeneity in exposure to income risk and self-insurance capacity (e.g., differences in beliefs, preferences, or shock processes); and (ii) imperfect proxies for self-insurance capacity, especially because liquid assets are measured while employed and up to two years before job loss. We suspect both forces are important. Importantly for our purposes, our latent-type approach does not rely on these observables for identification and is therefore robust to omitted or mismeasured proxies—a key advantage of the flexible mixture specification.

[5] As Lewis et al. (2026) emphasize in the context of marginal propensities to consume, a key advantage of using latent types to capture heterogeneity is that it implicitly allows many observables to enter jointly, without the rapid loss of statistical power that arises from the fully interacted specifications.

| | Consumption Decline |
|---|---|
| Second Quartile Liquidity | 0.010*** |
| | (0.001) |
| Third Quartile Liquidity | 0.013*** |
| | (0.001) |
| Top Quartile Liquidity | 0.010*** |
| | (0.001) |
| Second Quartile Earnings | 0.006*** |
| | (0.001) |
| Third Quartile Earnings | 0.009*** |
| | (0.001) |
| Top Quartile Earnings | 0.007*** |
| | (0.001) |
| Home Owner | 0.013*** |
| | (0.001) |
| Married | 0.010*** |
| | (0.001) |
| N | 36, 942 |
| R-squared | 0.105 |

TABLE D.4. Consumption Declines ($\Delta c/c$) and Proxies of Access to Self-Insurance

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. For each household, we compute the posterior-weighted unemployment-induced consumption decline by averaging the type-specific declines $\hat{\delta}_1^K$ using the worker's posterior consumption-type probabilities $p_i^C(k)$. We then regress these predicted declines on proxies for access to self-insurance. The estimation sample is the same as in the main text, and covariates are measured during periods of employment. Liquid wealth is defined using the "very liquid assets" concept in Carroll and Samwick (1997). Standard errors clustered at the household level are reported in parentheses.

We also compare our mixture-based estimates to "interacted observable" specifications. Specifically, we estimate versions of equation (7) that partition the sample into terciles of age, liquid wealth, or permanent income, and interact unemployment with group indicators. Table D.5 reports the resulting group-specific consumption declines. The first column reproduces our baseline mixture estimates; the remaining columns show interacted specifications based on terciles of age, liquid wealth, permanent income, and average employed consumption, respectively. Within each column, we order the tercile groups from largest to smallest which corresponds to the order of the estimated consumption declines (largest to smallest).

The interacted specifications display intuitive patterns: consumption declines decrease with age, liquid wealth, and permanent income. However, even the grouping that generates the largest spread (age terciles) captures less heterogeneity than the mixture model, and, as Table D.4 shows, these observables jointly explain only a modest share of the variance in predicted declines. Moreover, the mixture-based estimates are generally more precisely estimated.[6] These findings aligns with Lewis et al. (2026), who show that latent-type methods can deliver substantial gains in statistical power relative to fully interacted observable-group approaches.

| | Mixture | Age | Liquid Wealth | Permanent Income |
|---|---|---|---|---|
| Group 1 | −0.259*** | −0.159*** | −0.146*** | −0.143*** |
| | (0.041) | (0.018) | (0.014) | (0.014) |
| Group 2 | −0.114*** | −0.152*** | −0.122*** | −0.114*** |
| | (0.011) | (0.018) | (0.020) | (0.020) |
| Group 3 | −0.068*** | −0.078*** | −0.090*** | −0.098*** |
| | (0.007) | (0.018) | (0.020) | (0.027) |

TABLE D.5. Interacted Observable Models for Consumption Declines ($\Delta c/c$) and Proxies of Access to Self-Insurance

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Each column reports estimated consumption declines for different groupings of households. Column (1) reproduces our baseline Gaussian mixture linear regression estimates. The remaining columns, partition the sample into terciles based on age, liquid wealth, or permanent income, and interact unemployment with these group indicators. Groups are ordered from largest to smallest consumption decline in each column. Robust standard errors are reported in parentheses; for the mixture model, condition on estimated posterior group assignments, as discussed in the text. Covariates include a cubic in age, indicators for year, and the log of the change in family size.

## 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 first partition the sample into 50 equally sized groups based on WTP under a proportional-capped reform. Within each group, we further partition individuals

[6] For the mixture model, we report standard errors that condition on posterior group assignment (i.e., treat posteriors as given), which is the closest analogue to assuming group membership is a known function of observables in the interacted specifications. These standard errors therefore do not reflect uncertainty in posterior group membership.

into 50 equally sized bins based on WTP under an alternative expansion. For each pair of bins, we plot the average WTP across under the two reforms.

Panel (A) plots worker-level WTP under the proportional-capped reform against WTP under a flat expansion. Panel (B) compares the proportional-capped and proportional expansion. Points along the 45-degree line correspond to workers whose valuations are similar under the two reforms; departures from the line reflect differences in cross-subsidization across reform designs.

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 have positive WTP for either expansion, while those in the Southwest experience negative WTP 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 this cross-subsidization effects, and (ii) those who benefit from cross-subsidization under both reforms—typically workers with high unemployment risk for whom the implicit pooled reform price is below their actuarially fair price. Conversely, individuals in the Southwest quadrant face an implied pooled price above their actuarially fair price under both reforms, and their insurance value is insufficient to offset this cost. This group disproportionately includes high-income individuals, who tend to face lower unemployment 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.

We can also 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. Those above the 45-degree line in panel (B) prefer an increase in replacement rates (a proportional expansion) to an expansion of the current proportional-capped system. Overall, 82% of workers prefer expanding the proportional-capped system to a uniform benefit increase, and 52% prefer it to an increase in the replacement rate.

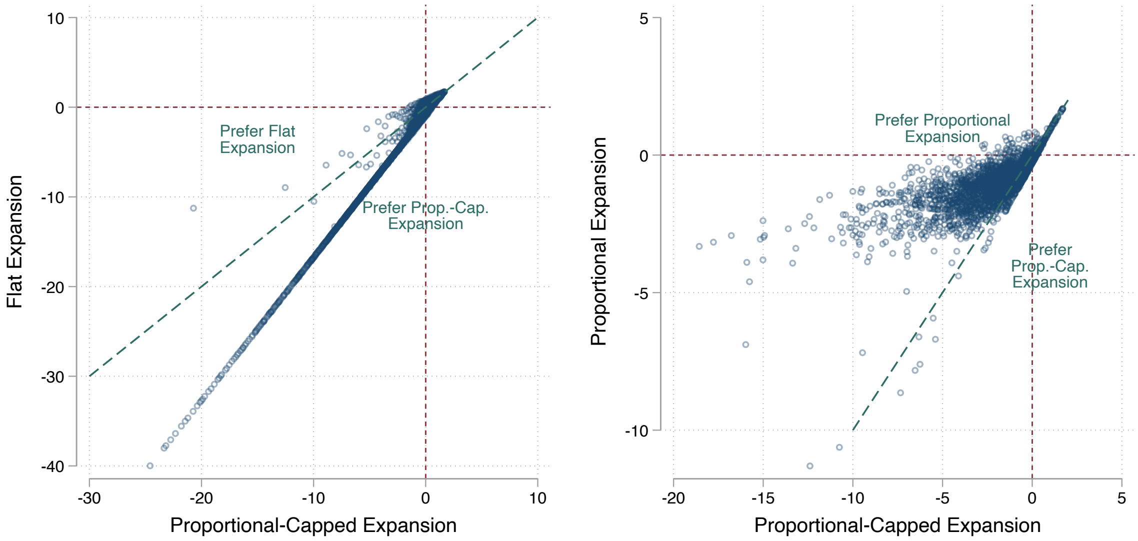


(A) Flat

(B) Proportional

FIGURE E.1. Individual Preferences over Reforms

Notes: We compare the proportional-capped system corresponding to current U.S. UI policy against a flat expansion (Panel (A)) and proportional expansion (Panel (B)).

## E.2. Sensitivity to Preference Specification

Our willingness-to-pay results in Figure 4 assume a coefficient of risk aversion of 3 and state-independent utility. Figure E.2 presents sensitivity analyses exploring robustness 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 component of WTP by household income (Figure 4(A)). Panel (B) introduces state dependence by assuming that consumption prices are 1.5% lower in unemployment, following Kaplan and Menzio (2015) and Campos and Reggio (2020).[7] This assumption affects the insurance component through two offsetting channels (see Appendix A.2). Lower prices increase the marginal value of a dollar in unemployment, increasing insurance value. However, because observed expenditure understates true consumption for the unemployed, measured expenditure-based consumption drops overstate insurance values. The latter effect dominates, modestly lowering the distribution relative to the baseline.

[7] 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.

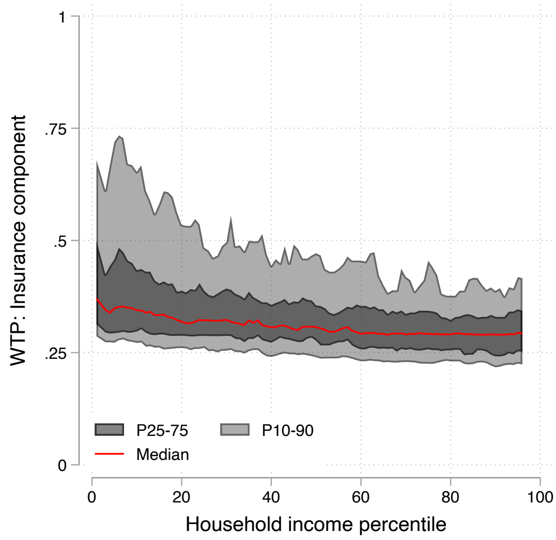


(A) Baseline results

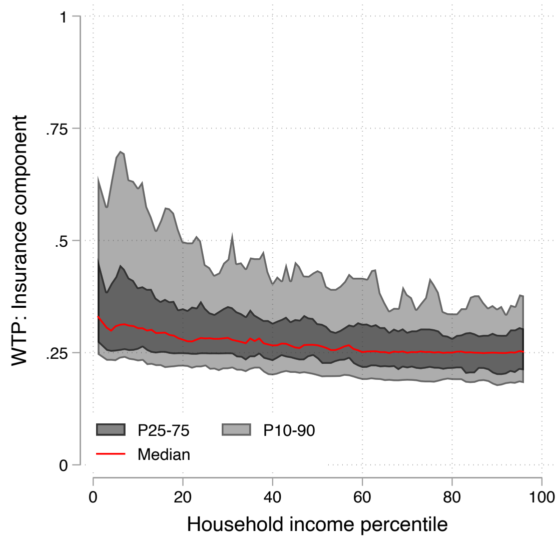


(B) State-dependence

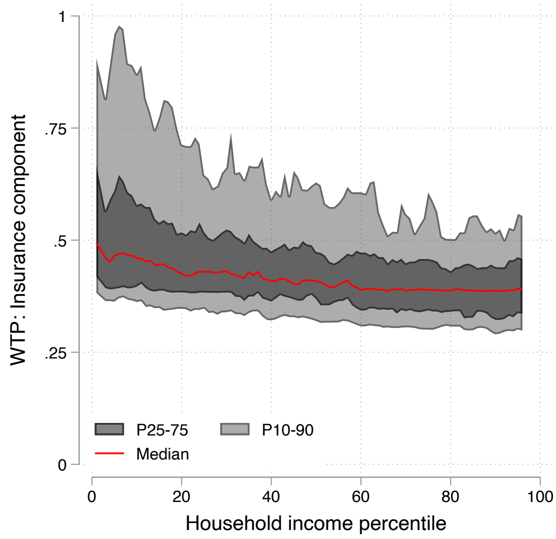


(C) High risk aversion

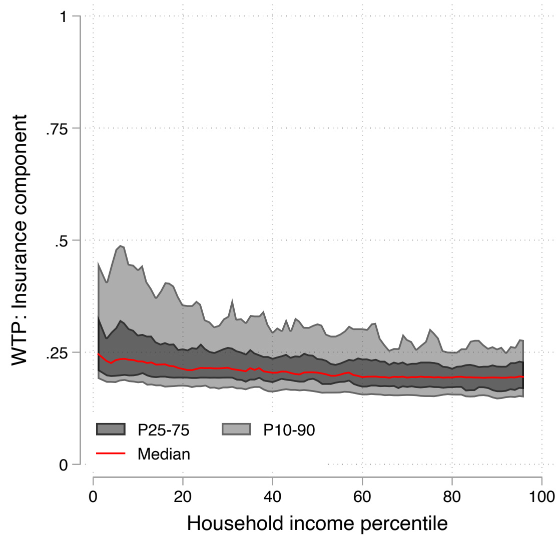


(D) Low risk aversion

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

Notes: Each panel summarizes the distribution of workers' WTP for an actuarially fair UI expansion (i.e,, the insurance component of overall WTP), reporting the median and the 10th, 25th, 75th, and 90th percentiles. We compute worker-level WTP using estimates of individual consumption declines. Panel (A) repeats our baseline results (Figure 4(A)). Panel (B) assumes the price of consumption is 1.5% lower in unemployment. Panels (C) and (D) assume coefficients of relative risk aversion of 4 and 2, respectively. Quantiles are smoothed using a uniform rolling window of $\pm 1$ percentiles.

Panels (C) and (D) show the distribution assuming coefficients of relative risk aversions of 4 and 2, respectively. As expected, the insurance component scales proportionately with risk aversion. However, because WTP for realistic reforms is shaped primarily by cross-subsidization (Figure 4), the overall distribution of total

WTP is 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 $\theta$.

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

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

We focus on a budget-neutral 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 that 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 reform path.

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

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

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 \$1 expansion in benefit expenditure. 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 that $\frac{\partial W}{\partial\tau} = \int_i e_i\phi_{\mathcal{T}}(y_i)\, di \mathbb{E}^h\left[\omega u^{h\prime}(c^h)\right]$. As our implementation sets $\phi_{\mathcal{T}}(y_i) = 1$, $\frac{\partial W}{\partial\tau}$ corresponds to the social value of a \$1 increase in high-state incomes.

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

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

To decompose the resource-reallocation term, note that:

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

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

$$\mathrm{RR} = \bar{\lambda}\mathbb{E}^l\left[\lambda\left(\frac{u^{l\prime}(c^l) - u^{h\prime}(c^h)}{u^{h\prime}(c^h)} + 1 - \frac{e\phi_{\mathcal{T}}(y)}{(1-e)\phi_{\mathcal{B}}(y)} \Big/ \frac{\int_{i'} e_{i'}\phi_{\mathcal{T}}(y_{i'})di'}{\int_{i'}(1-e_{i'})\phi_{\mathcal{B}}(y_{i'})di'}\right)\right],$$

where $\bar{\lambda} \equiv \frac{\mathbb{E}^l[\omega u^{h\prime}(c^h)]}{\mathbb{E}^h[\omega u^{h\prime}(c^h)]}$ and $\frac{dW^{MM}}{d\theta} = \mathrm{RR} - \mathrm{FE}$. Note that:

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

Hence we have the decomposition in equation (14):

$$\mathrm{RR} = \bar{\lambda}\mathbb{E}^l\left[\lambda\left(\frac{u^{l\prime}(c^l) - u^{h\prime}(c^h)}{u^{h\prime}(c^h)}\right)\right] + (\bar{\lambda} - 1).$$

### 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:

$$\frac{dW}{d\theta} = -\left(\frac{d\tau}{db}\right)\int_i \omega_i \int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t})\frac{\partial u_i^h(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t})\, dt\, di$$
$$+\int_i \omega_i \int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)(1-\xi_{i,t}(\varphi_{i,t}))\frac{\partial u_i^l(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t})\, dt\, di,$$

where budget-neutrality requires:

$$\frac{d\tau}{db} = \frac{\int_i\int_0^1\int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)\left(1-\xi_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t})\, dt\, di}{\int_i\int_0^1\int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t}) dF_{i,t}(\varphi_{i,t})\, dt\, di} \times \left(1 + \frac{1}{\int_i\int_0^1\int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)\left(1-\xi_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t})\, dt\, di}\right.$$
$$\left.\int_i\int_0^1\int_{\varphi_{i,t}} \frac{d(1-\xi_{i,t}(\varphi_{i,t}))}{d\theta}(\mathcal{T}(y_i)+\mathcal{B}(y_i))dF_{i,t}(\varphi_{i,t})\, dt\, di\right)$$
$$= \frac{\int_i\int_0^1\int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)\left(1-\xi_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t})\, dt\, di}{\int_i\int_0^1\int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t}) dF_{i,t}(\varphi_{i,t})\, dt\, di} \times \left(1 + \underbrace{\mathbb{E}^l\left[\epsilon \times \frac{\mathcal{T}(y_i)+\mathcal{B}(y_i)}{\mathcal{B}(y_i)}\right]}_{\text{FE}}\right).$$

$\epsilon_i$ is given by:

$$\epsilon_i = \int_0^1\int_{\varphi_{i,t}} \frac{d(1-\xi_{i,t}(\varphi_{i,t}))}{db} dF_{i,t}(\varphi_{i,t})\, dt \frac{1}{\phi_{\mathcal{B}}(y_i)} \frac{\mathcal{B}(y_i)}{\int_0^1\int_{\varphi_{i,t}}(1-\xi_{i,t}(\varphi_{i,t}))dF_{i,t}(\varphi_{i,t})\, dt}$$

and the low-state risk-weighted expectation over $x_i(\psi_{i,t})$ is given by:

$$\mathbb{E}^l[x] \equiv \int_i\int_0^1\int_{\varphi_{i,t}} \left(\frac{\phi_{\mathcal{B}}(y_i)(1-\xi_{i,t}(\varphi_{i,t}))}{\int_{i'}\int_0^1\int_{\varphi_{i',t}} \phi_{\mathcal{B}}(y_{i'})(1-\xi_{i',t}(\varphi_{i',t}))dF_{i',t}(\varphi_{i',t})dtdi'}\right) x_i(\psi_{i,t})dF_{i,t}(\varphi_{i,t})\, dt\, di$$

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, $D_i \equiv \int_0^1\int_{\varphi_{i,t}}(1-\xi_{i,t}(\varphi_{i,t}))dF_{i,t}(\varphi_{i,t})dt$ is the worker's expected unemployment duration, $\epsilon_i$ is the elasticity of worker $i$'s unemployment duration with respect to the benefit, where the derivative is taken along the budget-balanced reform path indexed by $\phi_{\mathcal{B}}(y_i)db$, and the low risk-weighted expectation weights individuals' by their expected benefit rise, which depends on both expected unemployment duration and benefit adjustment, $\phi_{\mathcal{B}}(y_i)$.

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

$$\frac{dW^{MM}}{d\theta} = \frac{dW/d\theta}{\int_i \int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{B}}(y_i)\left(1 - \xi_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t})\, dt\, di} \Bigg/ \frac{\partial W/\partial \tau}{\int_i \int_0^1 \int_{\varphi_{i,t}} \phi_{\mathcal{T}}(y_i)\xi_{i,t}(\varphi_{i,t}) dF_{i,t}(\varphi_{i,t})\, dt\, di}.$$

After algebra analogous to that in the preceding subsection, we obtain:

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

### F.3. Relationship with Marginal Value of Public Funds

Our normative analysis considers a reform to UI generosity that raises benefits and is 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 increase and financing tax change 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 increase is:

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

where the willingness-to-pay is:

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

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

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

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

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

To assess whether the combined reform is socially valuable, we compare the welfare-weighted MVPF of the benefit expansion to that of the financing tax increase. Weighting each MVPF by the incidence-weighted average marginal social welfare weights of affected groups, we require:

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

Rewriting this inequality, we obtain:

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

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

Recognizing that $\bar{\lambda} = \frac{\mathbb{E}^{l}[g]}{\mathbb{E}^{h}[g]}$, $\lambda_i = \frac{g_i}{\mathbb{E}^{l}[g]}$ and that the fiscal externality of the combined reform in the main text is $\text{FE} = \frac{1+\text{FE}^{\mathcal{B}}}{1+\text{FE}^{\mathcal{T}}} - 1$, shows that the left-hand side is exactly the welfare condition in equations (12) and (14).[8]

## G. Normative Framework: Implementation

### G.1. Measuring the Fiscal Externality

Our fiscal externality term decomposes as:

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

where $\mathcal{T}(y_i)$ denotes the net payment to the government in the high state, and $\mathcal{B}(y_i)$ the net transfer received in the low state. We refer to the first component as the direct fiscal externality from UI expansion, arising from higher benefit payments, and to the second component as the indirect effect operating through changes in income tax liabilities.

The first component depends on risk weights (in the expectation), which we construct using our unemployment risk model and observed data on worker earnings, 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.

To account for benefit exhaustion, we follow Schmieder and Von Wachter (2016),

[8] In practice, a natural approximation is $\text{FE}^{\mathcal{B}} \approx \text{FE}$ and $\text{FE}^{\mathcal{T}} \approx 0$, since the combined reform's fiscal externality is likely primarily driven by the benefit expansion rather than the small financing-side payroll tax increase. However, the equivalence above does not rely on this.

who show that under a constant-hazard approximation, the fiscal cost of a benefit increase depends on both total non-employment responses and the fraction of time spent receiving benefits. Let $D$ denote average non-employment duration and $D_B$ average benefit duration. Then the fiscal externality can be written as:

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

where

$$\xi \equiv 1-(1+sP)e^{-sP}$$

captures how benefit increases affect compensated duration relative to total duration, with $P$ denoting potential benefit duration and $s$ the unemployment exit hazard. We set $s = 1/D$ and use $D = 24.3$, $D_B = 15.8$ from Chetty (2008), and $P = 26$.

The indirect component depends on risk weights, behavioral elasticities, and the state-specific tax position of households. We compute these tax positions using the NBER TAXSIM (version 32) U.S. tax simulator, implementing 2019 federal and state tax schedules.

Specifically, for all observations in which the household's reference person is employed, we compute each household's 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 household's net tax liability in the low state, we assume the reference person experiences an unemployment spell of average duration and receives unemployment benefits equal to the lower of 50% of earnings or the UI earnings cap. Denote these UI payments by $ui_i$. We then compute the household's tax liability in the low state, denoted by $T_i^l$.

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

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

which captures the difference in state-specific income tax liabilities between employment and unemployment, scaled by UI 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 component $\tilde{g}_i = \frac{(c_i^h)^{-\gamma}}{\mathbb{E}[(c^h)^{-\gamma}|y_i]}$.

In the average implementation, we treat the employed and unemployed populations as representative individuals with model-implied average values. Specifically, for $s \in \{h, l\}$

$$\widetilde{\mathbb{E}}^h[x] = \int_i \left(\frac{e_i}{\int_{i'} e_{i'} di'}\right) x_i\, di \qquad \widetilde{\mathbb{E}}^l[x] = \int_i \left(\frac{1-e_i}{\int_{i'} (1-e_{i'}) di'}\right) x_i\, di$$

where $e_i$ is the worker's estimated high-state probability. We then evaluate the social welfare expression by replacing worker-level objects with $\widetilde{\mathbb{E}}^s[\cdot]$, for the income-based welfare weight ($g_y$), consumption drop ($\Delta c/c^h$), behavioral elasticity ($\epsilon^{(1-e,b)}$), benefit entitlement ($\mathcal{B}$), and tax liability ($\mathcal{T}$). This leads to the simplified expression:

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

The average implementation captures differences in average welfare weights between the employed and unemployed populations, but it abstracts from within-group heterogeneity. As a result, it is invariant across reform structure, since differences between reforms stem from heterogeneity in incidence within the low state.

## H. Additional Welfare Results

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| | Resource reallocation | | | Fiscal externality | | | Total effect |
| | Risk protection | Cross-subsidization | Total (1)+(2) | Direct | Indirect | Total (4)+(5) | (3)-(6) |
| H=2, K=3 (Baseline) | 0.483 | 0.190 | 0.673 | 0.268 | 0.479 | 0.747 | -0.074 |
| H=1, K=3 (-1 H) | 0.449 | 0.130 | 0.579 | 0.264 | 0.485 | 0.749 | -0.170 |
| H=3, K=3 (+1 H) | 0.478 | 0.192 | 0.670 | 0.262 | 0.522 | 0.784 | -0.114 |
| H=2, K=2 (-1 K) | 0.474 | 0.190 | 0.663 | 0.268 | 0.479 | 0.747 | -0.084 |
| H=2, K=4 (+1 K) | 0.542 | 0.191 | 0.732 | 0.268 | 0.479 | 0.747 | -0.015 |

TABLE H.1. Robustness of Social Value of Expanding UI Generosity by Number of Latent Types

Notes: Column (7) reports the total social welfare impact of a UI expansion, decomposed in columns (1)–(6). $H$ denotes the number of unemployment-risk types and $K$ the number of consumption types; the baseline results in the paper correspond to $H = 2$, $K = 3$. Entries corresponds to an expansion of the current U.S. UI system under baseline social preferences and are expressed in dollars per $1 increase in benefit expenditure.

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) |
|---|---|---|---|---|---|---|---|
| | Resource reallocation | | | Fiscal externality | | | Total effect |
| | Risk protection | Cross-subsidization | Total (1)+(2) | Direct | Indirect | Total (4)+(5) | (3)-(6) |
| Baseline | 0.483 | 0.190 | 0.673 | 0.268 | 0.479 | 0.747 | -0.074 |
| Age | 0.462 | 0.130 | 0.591 | 0.264 | 0.485 | 0.749 | -0.157 |
| Liquid Wealth | 0.448 | 0.130 | 0.578 | 0.264 | 0.485 | 0.749 | -0.171 |
| Permanent Income | 0.443 | 0.130 | 0.573 | 0.264 | 0.485 | 0.749 | -0.176 |

TABLE H.2. Robustness of Social Value of Expanding UI Generosity by Observable Only Heterogeneity

Notes: Column (7) reports the total social welfare impact of a UI expansion, decomposed in columns (1)–(6). Entries corresponds to an expansion of the current U.S. UI system under baseline social preferences and are expressed in dollars per $1 increase in benefit expenditure.

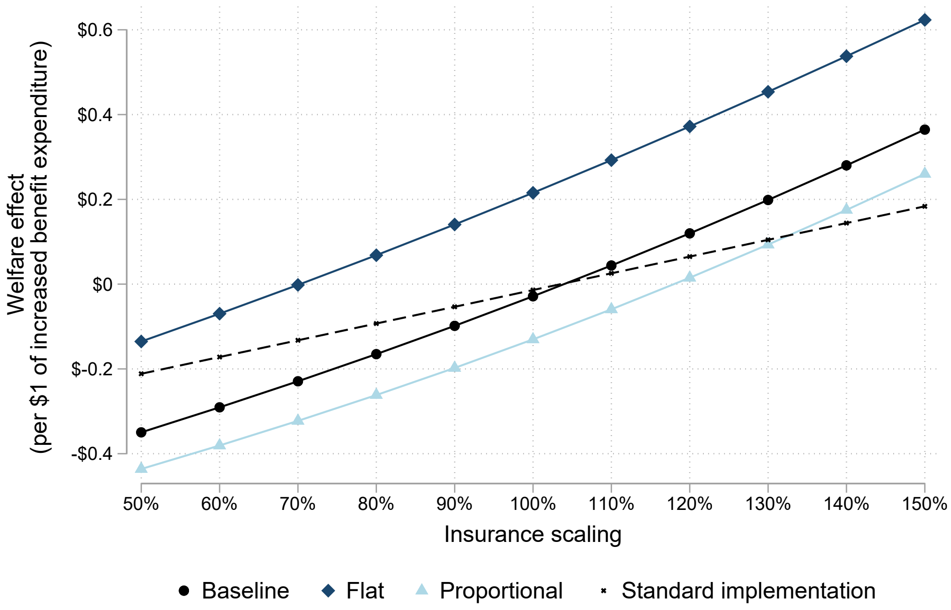


FIGURE H.1. Sensitivity of Welfare Effects to Insurance Value

Notes: Each line plots the net social value per $1 of additional UI expenditure as a function of a scaling factor $s$ applied to the individual-level insurance component of WTP, holding the cross-subsidization and fiscal-externality components fixed. Values of $s < 100\%$ correspond to proportionally smaller consumption drops (or lower risk aversion), and $s > 100\%$ to larger drops (or higher risk aversion). The figure reports results for the baseline reform (solid circles), a flat benefit expansion (squares), a proportional expansion (triangles), and the standard homogeneous implementation (dashed).

| | **(1)** | **(2)** | **(3)** |
|---|---|---|---|
| | Risk protection | Cross-subsidization | Total (1)+(2) |
| Baseline | 0.48 | 0.19 | 0.67 |
| Heterogeneous take-up | 0.47 | 0.17 | 0.63 |

TABLE H.3. Robustness of Resource-Reallocation Gain of Expanding UI Generosity to Heterogeneous Take-up

Notes: Row (1) reports our baseline results. Row (2) assumes take-up probabilities of 0.25, 0.35, 0.425, 0.4 and 0.3 for the income percentiles 0–15, 16–30, 31–80, 81–90 and 91–100, respectively (based on Lachowska et al. (2025; Figure 3D)). Entries corresponds to an expansion of the current U.S. UI system under baseline social preferences and are expressed in dollars per $1 increase in benefit expenditure.

## References


Aguiar, M. and E. Hurst (2005). Consumption versus expenditure. *Journal of Political Economy 113*(5), 919–948.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review 106*(2), 387–435.

Browning, M. and T. F. Crossley (2001). Unemployment insurance benefit levels and consumption changes. *Journal of Public Economics 80*(1), 1–23.

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

Carroll, C. D. and A. A. Samwick (1997). The nature of precautionary wealth. *Journal of Monetary Economics 40*(1), 41–71.

Charles, K. K. and M. Stephens, Jr (2004). Job displacement, disability, and divorce. *Journal of Labor Economics 22*(2), 489–522.

Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics 90*(10-11), 1879–1901.

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

Gruber, J. (1997). The Consumption Smoothing Benefits of Unemployment Insurance. *American Economic Review 87*(1), 192–205.

Hai, R., D. Krueger, and A. Postlewaite (2020). On the welfare cost of consumption fluctuations in the presence of memorable goods. *Quantitative Economics 11*(4), 1177–1214.

Hendren, N. (2017). Knowledge of Future Job Loss and Implications for Unemployment Insurance. *American Economic Review 107*(7), 1778–1823.

Hendren, N. and B. Sprung-Keyser (2020). A unified welfare analysis of government policies. *Quarterly journal of economics 135*(3), 1209–1318.

Kaplan, G. and G. Menzio (2015). The morphology of price dispersion. *International Economic Review 56*(4), 1165–1206.

Lachowska, M., I. Sorkin, and S. A. Woodbury (2025). Employers and unemployment insurance take-up. *American Economic Review 115*(8), 2529–2573.

Lewis, D., D. Melcangi, and L. Pilossoph (2026). Latent heterogeneity in the marginal propensity to consume. *Review of Economic Studies*.

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

Poterba, J. and T. 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), 84–89.

Schmieder, J. F. and T. Von Wachter (2016). The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation. *Annual Review of Economics 8*(1), 547–581.

Stephens, Jr, M. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics 20*(3), 504–537.