Risk Protection and Redistribution in the Design of Social Insurance

Rory McGee and Martin O'Connell*

January 2025

Abstract

We develop and empirically implement a framework to evaluate the societal value of social insurance, explicitly accounting for arbitrary individual-level heterogeneity. Our approach extends the classic Baily-Chetty formula to incorporate how differences in individuals' willingness-to-pay for insurance, program cross-subsidisation, and inequality-averse social preferences influence the social value of reforms. Applying this framework to U.S. Unemployment Insurance (UI), we provide novel evidence on the distribution of willingness-to-pay for UI and the societal value of the current system. Our results show that accounting for worker heterogeneity significantly increases the estimated value of a lump-sum UI expansion, enough to overturn standard policy recommendations found in the existing literature.

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, Davide Melcangi, Ryan Michaels, Corina Mommaerts, Luigi Pistaferri, 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

One of the primary functions of governments is to provide individuals with insurance against adverse shocks. These social insurance programs provide protection against a wide range of risks, including unemployment, health issues, and disability. In the U.S., approximately 36% of the federal government's budget is allocated to such programs.[1] At the heart of their design is a trade-off between providing insurance and minimising distortions to incentives. At the same time, risk levels vary widely across individuals due to differences in the costs and benefits of investing in risk abatement. And, since the implicit premiums for these programs are often embedded in the tax and benefit system and are not individualised, social insurance inherently redistributes across individuals with different risk levels.

In this paper we develop and empirically implement a framework to characterise the societal value of social insurance that explicitly accounts for heterogeneity across individuals, including in their risk level, risk exposure and living standards. We make two contributions to the existing literature. We extend the classic sufficient-statistic approach to valuing social insurance reforms (see Baily 1978, Chetty 2006) to explicitly capture how heterogeneity in individuals' willingness-to-pay for insurance, and inequality-averse social preferences, influence the social value of reforms. Our approach is straightforward to implement and can be carried out using the same datasets typically used for applying the standard formula. Second, we present new evidence on the distribution of willingness-to-pay for unemployment insurance and the social value of the U.S. Unemployment Insurance (UI) system. We show that accounting for worker heterogeneity increases the value of a lump-sum benefit expansion compared to estimates in the existing literature.

We begin by developing a model in which a population of agents is exposed to risk (e.g., unemployment or natural disaster) and must choose their effort level, representing costly actions that reduce the risk of entering the bad state (e.g., job search or abatement investment). A social planner decides on tax and benefit levels that transfer resources from the good to the bad state. In the canonical framework (see Chetty and Finkelstein 2013), policy balances the insurance value of the program for a representative agent against the costs of distorting effort incentives. We extend this framework by characterising the implications of agent heterogeneity – in incomes,

[1] As in Feldstein (2005), who finds a value of 37% for 2003, we sum Social Security retirement, disability, and survivor insurance; unemployment insurance; and Medicare insurance for those age 65 and older.

preferences and effort costs – on the value of social insurance.

In our setting, each agent's willingness-to-pay for an additional dollar of social insurance depends on two components: (1) their personalised insurance value, which reflects their risk exposure (i.e., the value they place on a marginal transfer of resources from the good to bad state), and (2) the extent to which they are, ex ante, net contributors to or recipients of the social program. The first component may vary significantly across individuals, depending on the magnitude of the financial loss in the bad state and their ability to self-insure. The second component arises if individual risk levels differ from the implicit social insurance premium, due to a lack of individualised tax and benefits.

We use a social welfare function, defined as a weighted sum of individuals' expected utilities, to characterise the aggregate value of the program. We show that the marginal social value of program expansion – and, implicitly, the optimal policy – is characterised by three sets of statistics. First, an average of individual-level insurance values, weighted by both risk and social marginal welfare weights. This extends the standard representative agent or average insurance value by accounting for how individual insurance values covary with the societal value placed on redistributing resources to them. Second, an aggregate welfare weight that captures how the average social marginal utility of consumption among program beneficiaries compares to program funders. Absent in the representative agent framework, this measure reflects the social value of the pure redistributive effects arising from program cross-subsidisation, which persist even if individuals are fully insured against the risk the social program aims to protect against. This aggregate welfare weight depends not only on societal preferences for equity but also on the structure of the program reform, thereby encoding the reform's progressivity. Third, a measure of the sensitivity of individuals' effort choices to changes in program generosity. This captures the moral hazard costs of the program and, in the case of a flat reform – a rise in a lump-sum benefit funded by a lump-sum tax rise – is determined by a single aggregate elasticity.

An advantage of the sufficient-statistic approach is its ability to evaluate marginal policy reforms without requiring a full specification of the individual's optimisation problem.[2] Our approach preserves this strength while relaxing the restrictive representative agent assumption embedded in the classic Baily-Chetty formula.

[2]For instance, we demonstrate our static framework extends with minimal modification to a richer dynamic setting (see Chetty 2006). A complementary strand of recent research uses dynamic structural models to quantify the welfare implications of social insurance reforms, highlighting the importance of redistributive effects. Examples include Audoly (2024) and Haan and Prowse (2024) in the context of UI, and Braun et al. (2017) and Arapakis et al. (2023) in the context of social insurance programs for seniors.

Although it requires estimating a broader set of statistics – for example, the distribution of individual insurance values rather than just the average – our empirical implementation demonstrates that this can readily be accomplished using the same datasets commonly employed for the standard analysis.

In our empirical analysis, we examine the value of expanding the U.S. UI system using data from the Panel Study of Income Dynamics (PSID) covering the period 1999–2019. This nationally representative dataset tracks individuals over time, and includes detailed measures of labour market outcomes, demographics, and consumption expenditures. We exploit the panel dimension of the data to estimate how household consumption responds when the primary earner becomes unemployed, building on the "consumption-based" approach to estimating insurance values first used in Gruber (1997). Additionally, we also use the panel consumption data as inputs for a schedule of social marginal welfare weights that capture societal redistributive preferences (Saez 2002). Our analysis makes two main empirical contributions.

First, we provide new evidence on the full distribution of unemployment-driven consumption declines. To do this we apply clustering regression techniques (Lewis et al. 2024), which have recently been used to estimate heterogeneity in marginal propensities to consume (see also Boehm et al. 2025). This method allows us to account for heterogeneity in consumption declines across latent groups without requiring a priori assumptions about which observable characteristics might correlate with these declines and to recover counterfactual consumption declines for those who we do not observe experiencing unemployment. We estimate an average unemployment-induced consumption drop of 11.6%, broadly consistent with existing estimates. However, we also find meaningful heterogeneity; for example, the decline at the 95th percentile is twice as large at the median decline.

This variation in consumption declines reflects heterogeneity along two dimensions. First, differences in the income loss at unemployment, which include the extent to which other social programs insure income. Second, differences in the ability to self-insure, for example, because of added worker effects or binding liquidity constraints. We show that, on average, each log point increase in equivalised household consumption when employed is associated with a 1.6 percentage point reduction in the size of unemployment-induced consumption declines.

We also provide novel evidence on the relationship between labour market risk, consumption levels, and unemployment-related consumption declines. We show that, cross-sectionally, a household's unemployment risk is systematically related to their

consumption level; an employed worker with equivalised household consumption at the 25th percentile faces unemployment risk that is 34% higher than that of a household at the 75th percentile. This suggests that the UI system has significant redistributive effects.

Our second main empirical contribution is to quantify the social welfare impact of expanding UI while accounting for heterogeneity in individual willingness-to-pay. We first examine a flat increase in UI benefits, funded by a flat tax increase. When social preferences exhibit zero inequality aversion – meaning the planner assigns equal weight to each individual's money-metric willingness-to-pay – the aggregate gains to households from the program expansion is approximately offset by the central estimates of the moral hazard costs of UI expansion (see Krueger and Meyer 2002, Schmieder and Von Wachter 2016, and Cohen and Ganong 2024). However, when social preferences exhibit inequality aversion, the social gains from expansion become positive. Under baseline social preference, commonly used in optimal tax analysis (e.g., Saez 2002, Allcott et al. 2019), the aggregate gains to households are 40% higher than those suggested in the existing literature.

A key channel captured in our framework, absent in the representative agent case, is that the welfare impacts of program expansion depend on the nature of the reform. The U.S. UI system currently replaces a fraction of lost earnings up to a cap. A natural reform involves increasing the replacement rate below the cap while providing an increase proportional to the cap for higher earners. This reform is more regressive than a flat benefit expansion, as higher-income households benefit more in cash terms. Under our baseline social preferences, we find that the aggregate benefit to households would only modestly exceed the central estimate of the moral hazard costs of UI. Additionally, we show that more progressive tax funding can substantially raise the societal value of UI expansion.

Our work contributes to the recent literature, which extends the sufficient-statistics approach to valuing social insurance by incorporating more realistic features. For example, existing work quantifies how differences in the composition of UI recipients over the duration of unemployment spells shape the optimal benefit profile (Schmieder et al. 2012, Kolsrud et al. 2018). We advance this literature by demonstrating how ex-ante heterogeneity among program participants interacts with the structure of reforms to shape their welfare impact. We use tools from the taxation literature to aggregate across heterogeneous agents into a social welfare metric (Piketty and Saez 2013) and to parameterise various reform types (e.g., Heathcote et al. 2017).

Our work also relates to a literature that studies how unemployment affects optimal income tax design. This includes theoretical results, which show that if individual risk levels correlate with earnings (a central mechanism in our analysis) accounting for this in income tax design can enable more efficient redistribution (see Blomqvist and Horn 1984; Rochet 1991; Cremer and Pestieau 1996; Boadway et al. 2006), and recent empirical work that quantifies the impact of involuntary unemployment on optimal tax formulae (e.g., Kroft et al. 2020). Ferey (2022) considers the optimal design of a tax and unemployment benefit schedules and demonstrates that allowing for non-linearity in the benefit schedule relaxes incentive-compatibility constraints within a Mirrleesian framework. We complement this work by demonstrating, both conceptually and quantitatively in a rich empirical setting, how accounting for the reallocation of resources across households can influence the value of social insurance reforms. We view this as of first-order importance, since it is common for governments to implement program reforms without reoptimising the entire tax and transfer and benefit system.[3]

The rest of this paper is structured as follows. In Section 2 we outline our social insurance design framework. In Section 3 we summarise the key features of the U.S. UI system and our dataset. Section 4 outlines our empirical approach for estimating the distribution of unemployment-induced consumption declines, and individual labour market risk. In Section 5 we present our quantitative analysis on the social welfare impacts of UI expansion. A final section concludes.

2. Social Insurance Design

In this section we describe an approach for quantifying the trade-off inherent in expanding a publicly funded insurance program. We present a simple framework designed to reflect the salient features of social insurance programs, including unemployment insurance, workers' compensation, disability insurance, health insurance, and insurance against natural disasters. We build on the canonical set-up (see Chetty and Finkelstein 2013) by considering social insurance design when there is a *heterogeneous* population of individuals (agents) exposed to risk. We focus here on a stylised static setting, and show in Appendix A.2 that our results generalise to a much richer dynamic environment (Chetty 2006). We include derivations of all equations in

[3]Our approach is related to that taken in Lockwood (2024), who shows that accounting for the interaction of a broad range of different risks is key to valuing the risk protection provided by health insurance. In our setting other risks both influence the size of unemployment-induced consumption declines, and households' consumption levels.

this section in Appendix A.

2.1. The Agent's Problem

There is a unit continuum of agents, indexed by i. Each agent faces uncertainty over two states: high (h) and low (l). Individuals potentially differ in their preferences, incomes, and the costs associated with mitigating risk. We denote income in the high and low state by y_i and z_i, respectively. We begin by focusing on a case in which the government provides a lump-sum benefit b to those in the low state and levies a lump-sum tax τ on individuals in the high state. Unlike in the standard representative agent setting, in our set-up the structure of the social insurance reform affects its social value. In the following subsection, we extend our framework to accommodate more flexible reforms. Consumption in the high state is therefore $c_i^h = y_i - \tau$, while consumption in the low state is $c_i^l = z_i + b$. $u_i^s(c)$ is state-specific utility from consumption c; we assume $u_i^s(.)$ is increasing, concave and twice continuously differentiable. The difference in high and low state consumption is $\Delta c_i = c_i^h - c_i^l$. For the majority of agents, we expect $\Delta c_i > 0$.

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

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

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

The individual index i captures arbitrary heterogeneity arising from various sources. First, differences in earning ability – due, for example, to variation in innate ability, human capital, skills, or labour market opportunities – are reflected in state-specific incomes, y_i and z_i. Individual-specific differences in incomes across states determine the direct financial loss of entering the low state, which, combined with variation in individuals' capacity to self-insure (for example, private savings or spousal insurance), drives heterogeneity in the consumption gap Δc_i. Additionally,

variation in the marginal utility of consumption and attitudes toward risk is captured by individual-specific utilities, $u_i^s(.)$. This approach also allows for variation in state-dependence across individuals, which may reflect differences in the opportunity cost of time or the availability of consumption substitutes. Finally, heterogeneity in attitudes toward work and employment opportunities (in the unemployment or disability insurance context) or differences in the cost of abatement investments (in the health or natural disaster insurance context) are reflected in $\psi_i(.)$. Each of these factors influences the agent's optimal effort choice and therefore leads to heterogeneity in risk levels, e_i.

Willingness-to-pay for social insurance expansion. To understand individual i's willingness-to-pay for a marginal expansion of social insurance it is instructive to consider the impact on their expected utility of a rise in the benefit $db = 1$, funded by a tax rise of $d\tau = \frac{1-\bar{e}}{\bar{e}}$, where $\bar{e} = \int_i e_i di$ denotes the expected share of individuals in the high state, and $1 - \bar{e}$ the expected share of individuals in the low state. This change would be revenue-neutral if program expansion engenders only a mechanical budgetary cost. The impact on agent i's expected utility is:

$$dV_i = (1 - e_i)u_i^{l\prime}(c_i^l) - e_i u_i^{h\prime}(c_i^h)\frac{1-\bar{e}}{\bar{e}}$$

To convert dV_i from "utils" to a money-metric cardinalisation we define: $\theta_i = \frac{dV_i}{(1-e_i)} \Big/ \frac{dV_i/dy_i}{e_i}$. θ_i measures the change in the individual's expected utility from program expansion expressed in terms of the gains from a \$ rise in high state income. We refer to θ_i as the individual's willingness-to-pay for a marginal expansion in social insurance (or, as shorthand, their willingness-to-pay). Note:

$$\theta_i = \underbrace{\left(\frac{u_i^{l\prime}(c_i^l) - u_i^{h\prime}(c_i^h)}{u_i^{h\prime}(c_i^h)}\right)}_{\text{insurance}} + \underbrace{\left(1 - \frac{e_i/\bar{e}}{(1-e_i)/(1-\bar{e})}\right)}_{\text{cross-subsidisation}}, \tag{2}$$

The first term in equation (2), which reflects the agent's risk exposure, captures the benefit to the individual of a marginal transfer of resources from the high to low state that is actuarially fair at the individual level, where the resources are valued at the marginal utility of consumption in each state. This captures the value to the individual of the additional insurance provided by a marginal social insurance expansion. The second term captures the net expected transfer to the individual from marginal program

expansion, which we refer to as the effect of cross-subsidisation. While it is likely that the marginal utility of consumption is higher in the low than high state, meaning the insurance term will be positive, the sign of the cross-subsidisation term will depend on whether the individual's probability of being in the low state is above (in which case it is positive) or below average.[4]

2.2. The Social Planner's Problem

The social welfare impact of a marginal expansion in social insurance depends both on the weights the social planner assigns to each individual and on the effect on the planner's budget.

We focus on the effect of a budget-neutral expansion of social insurance. Hence, the planner faces the constraint $\bar{e}\tau = (1 - \bar{e})b$, meaning an increase in the generosity of benefits must be funded by an increase in taxes. Thus, we write the tax rate as a function of the benefit rate $\tau(b)$. We specify social welfare as the weighted sum of expected utilities:

$$W(b) = \int_i \omega_i V_i(b, \tau(b)) di,$$

where ω_i is the Pareto weight the planner assigns to individual i. We denote the risk-weighted expectation of any variable, x_i in the low and high state, respectively, by $\mathbb{E}^l[x] \equiv \int_i (\frac{1-e_i}{1-\bar{e}}) x_i di$ and $\mathbb{E}^h[x] \equiv \int_i \frac{e_i}{\bar{e}} x_i di$.[5]

We consider the impact of a marginal expansion of social insurance – i.e., a balanced-budget rise in benefits – on social welfare. To do so, we rescale $\frac{dW}{db}$ in money-metric terms as $\frac{d\tilde{W}}{db} = \frac{dW/db}{1-\bar{e}} \times \left(\frac{dW/dy}{\bar{e}}\right)^{-1}$, which represents the impact on social welfare of a balanced-budget \$ expansion in social insurance, relative to the effect on social welfare of a \$ increase in all high state incomes. $\frac{d\tilde{W}}{db}$ satisfies:

$$\frac{d\tilde{W}}{db} = \frac{\mathbb{E}^l[\omega u^{l'}(c^l)] - \mathbb{E}^h[\omega u^{h'}(c^h)]}{\mathbb{E}^h[\omega u^{h'}(c^h)]} - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}, \tag{3}$$

[4] This willingness-to-pay decomposition is similar to that in Finkelstein et al. (2019) and similar in spirit to the decomposition of tax-and-transfer programs into their insurance and redistributive value in Hoynes and Luttmer (2011).

[5] Note, an alternative way to write $\mathbb{E}^h[x]$ is $\mathbb{E}^h[x] = \mathbb{E}[x] + \text{cov}(e_i/\bar{e}, x)$, and analogously for $\mathbb{E}^l[x]$. This formulation emphasises that $\mathbb{E}^h[x]$ depends both on the unweighted expectation of x and its endogenous covariance with risk.

where $\epsilon_{1-\bar{e},b} \equiv \frac{d(1-\bar{e})}{db}\frac{b}{(1-\bar{e})}$ is the elasticity of the aggregate low-state probability with respect to a budget-neutral change in the benefit. Note, this aggregate elasticity is equal to a weighted average of individual-level elasticities; i.e. $\epsilon_{1-\bar{e},b} = \mathbb{E}^l[\epsilon_{1-e,b}]$, where $\epsilon_{1-e,b} = \frac{d(1-e_i)}{db}\frac{b}{1-e_i}$.

The first term on the right-hand side captures the marginal gain of expanding social insurance, and is given by the percentage gap in the risk-weighted expected social marginal utility (i.e., $\omega_i u_i^{s'}(c_i^s)$) in the low and high states. It measures the social value of shifting a marginal dollar from the high to the low state. The second term captures the fiscal cost associated with this transfer, and is equal to the elasticity of the aggregate low state probability with respect to the balanced-budget change in the benefit, scaled by the expected share of people in the high state. This arises as a tax-funded benefit rise lowers the marginal gain to an individual of exerting effort, leading to a higher fraction of individuals arriving in the low state.

The marginal gain from expanding social insurance depends on individuals' demand for insurance, the patterns of cross-subsidisation across individuals and the welfare weights the planner assigns to individuals. In the following proposition we re-write $\frac{d\tilde{W}}{db}$ separating these channels:

PROPOSITION 1. *The impact of a marginal budget-neutral rise in benefits on aggregate welfare satisfies:*

$$\frac{d\tilde{W}}{db} = \underbrace{\bar{\lambda}}_{\substack{\text{aggregate}\\ \text{welfare weight}}} \mathbb{E}^l\left[\underbrace{\lambda}_{\substack{\text{individual}\\ \text{welfare weight}}}\left(\underbrace{\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)}}_{\text{insurance}} + \underbrace{\left(1 - \frac{e/\bar{e}}{(1-e)/(1-\bar{e})}\right)}_{\text{cross-subsidisation}}\right)\right] - \underbrace{\frac{\epsilon_{1-\bar{e},b}}{\bar{e}}}_{\text{moral hazard}} \tag{4}$$

where $\bar{\lambda} \equiv \frac{\mathbb{E}^l[\omega u^{h'}(c^h)]}{\mathbb{E}^h[\omega u^{h'}(c^h)]}$ *is the ratio, across the two states, of risk-weighted high state marginal utility of consumption and* $\lambda_i \equiv \frac{\omega_i u_i^{h'}(c^h)}{\mathbb{E}^l[\omega u^{h'}(c^h)]}$ *is the high-state marginal utility of consumption of individual i relative to the low-state risk-weighted average.*

Equation (4) shows that the marginal gain from social insurance expansion can be written as a weighted average of individuals' willingness-to-pay, along with a set of marginal social welfare weights. $\lambda_i \propto \omega_i u_i^{h'}(c_i^h)$ captures the weight the planner assigns to individual *i*'s willingness-to-pay. Individuals with larger Pareto weights or high-state marginal utility of consumption receive a larger welfare weight. The λ_i's are scaled such that their low-state risk-weighted average is 1 (i.e., $\mathbb{E}^l[\lambda] = 1$). $\bar{\lambda}$ is an aggregate welfare

weight that captures the weight the planner assigns to the population of people that end up in the low state relative to those that end up in the high state.

Aggregating the cross-subsidisation terms in (4) delivers:

$$\frac{d\tilde{W}}{db} = \bar{\lambda}\mathbb{E}^l \left[\lambda \left(\frac{u^{l\prime}(c^l) - u^{h\prime}(c^h)}{u^{h\prime}(c^h)} \right) \right] + (\bar{\lambda} - 1) - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}, \tag{5}$$

where the average of the cross-subsidisation terms (weighted both by low-state risk and welfare weights) is replaced with the aggregate welfare weight, $\bar{\lambda}$, minus 1. This provides a sufficient statistic for the welfare effects of cross-subsidisation. Note, if the planner assigns the same welfare weight to all individuals (i.e., $\lambda_i = 1$) then the effect of cross-subsidisation on social welfare is zero, as, in this case, the planner is indifferent to dollars being redistributed, ex ante, between different individuals. However, for an inequality-averse planner (including pure utilitarianism), this transfer of resources across individuals is relevant for social welfare. If social insurance, on average, transfers resources to individuals with high welfare weights, then $\bar{\lambda} > 1$ and the aggregate effect of cross-subsidisation is to raise the welfare effect of social insurance expansion. Note, this channel remains even if all individuals are able to perfectly hedge against low-state risk (in which case $u_i^{l\prime}(c_i^l) = u_i^{h\prime}(c_i^h)$ for all i, and there is no insurance benefit to program expansion).[6]

Implementation. The insurance term in equation (4) depends on the gap in the marginal utility of consumption between the low and high state. In our baseline implementation we assume that the marginal utility of consumption is state independent and we adopt the "consumption-based" approach to measuring the marginal utility gap, first used by Gruber (1997).

ASSUMPTION 1. *The marginal utility from a given level of consumption is state independent (hence $u_i^{s\prime}(.) = u_i^{\prime}(.)$ for $s \in \{l, h\}$).*

Under assumption (1) a quadratic approximation to the utility function implies:

$$\frac{u_i^{\prime}(c_i^l) - u_i^{\prime}(c_i^h)}{u_i^{\prime}(c_i^h)} \approx \gamma_i \frac{\Delta c_i}{c_i^h}, \tag{6}$$

[6]The impact of cross-subsidisation on social welfare is a tagging effect (Akerlof 1978) of social insurance. Interpreting $e_i/\bar{e}$ and $1-e_i/1-\bar{e}$ in the definition of $\bar{\lambda}$ as densities in the weighted average of social marginal utility, the tagging value exceeds one when $e_i/\bar{e}$ first-order stochastic dominates $1-e_i/1-\bar{e}$ and vice versa.

where $\gamma_i \equiv -\frac{u_i''(c_i^h)c_i^h}{u_i'(c_i^h)}$ is the coefficient of relative risk aversion. This re-expresses the value of insurance to an individual as the product of their coefficient of relative risk aversion and their percentage consumption decline between the high and low state.[7]

Using condition (6), we can re-write the impact of a marginal expansion of social insurance on social welfare as:

$$\frac{d\tilde{W}}{db} \approx \bar{\lambda}\left(\mathbb{E}^l\left[\gamma\frac{\Delta c}{c^h}\right] + \text{cov}^l\left(\lambda, \gamma\frac{\Delta c}{c^h}\right)\right) + (\bar{\lambda} - 1) - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}, \tag{7}$$

where the first two terms on the right-hand side split the welfare gains from providing more insurance into an average effect and a covariance term between risk exposure and individual-level welfare weights.[8]

Equation (7) highlights two new channels through which social insurance reform influences social welfare. The first channel reflects heterogeneity in demand for insurance. When risk exposure and neediness – captured by the individual-level welfare weights (λ_i) – positively co-vary, the social value of the insurance benefit of program expansion increases. The second channel captures redistributive benefits, which operates even in the absence of insurance demand. If individuals in the low state are systematically needier than those in the high state – captured in the aggregate welfare weight ($\bar{\lambda}$) – program expansion yields redistributive benefits. Further, these two channels interact: the insurance benefit of expansion is scaled by $\bar{\lambda}$. Thus, the overall social value of the insurance benefit increases with the relative neediness of the low-state population.

In the standard Baily-Chetty formula (Baily 1978, Chetty 2006), which assumes the existence of a representative agent, both of these new channels are absent. Rather the welfare effect of program expansion is given by $\frac{d\tilde{W}}{db} \approx \gamma\frac{\Delta c}{c^h} - \frac{\epsilon_{1-e,b}}{e}$ and requires measuring the representative agent's insurance value (usually measured as the population average) and comparing it to the moral hazard cost. By contrast, our approach involves measuring two additional sufficient statistics: 1) the aggregate welfare weight, which requires

[7] Alternative approaches to deriving empirical measures of the marginal utility gap, implemented in the UI context, include decomposing unemployment responses into substitution and income effects (Chetty 2008), measuring reservation wage responses (Shimer and Werning 2008) and measuring gaps in the marginal propensity to consume across states (Landais and Spinnewijn 2021). Each of these methods could be used in place of the consumption-based approach to implement equation (4) (or (5)).

[8] If instead of the quadratic expansion in equation (6), we take a cubic expansion, then the right-hand side of equation (7) has two additional additive terms that reflect how the (risk- and welfare-weighted) average curvature of utility varies with consumption, and the covariance of this with consumption declines. See Appendix A.1.

specifying individual-level marginal social welfare weights and taking risk-weighted averages of them, and 2) the covariance between individual welfare weights and risk exposure, which necessitates estimating heterogeneity in consumption drops. As we show below, measuring these additional sufficient statistics is straightforward and can be done using the same datasets employed for implementing the standard formula.

In the following two special cases, we discuss the implications of assuming homogeneous utility functions and demonstrate that the standard formula applies under a specific and unconventional choice of Pareto weights.

Limited heterogeneity. If there is heterogeneity in risk and consumption, but all individuals have the same state-independent utility function and Pareto weights, it is possible to obtain an approximation of the insurance value of program expansion that takes a form resembling the Baily-Chetty formula (e.g., Kolsrud et al. 2018). However, even in the case, the channels captured by equation (7) persist. First, taking a first-order approximation of $u'(c_i^s)$ around $\mathbb{E}[c^s]$ for states $s = \{l, h\}$, yields an approximate form of equation (3) given by $\frac{d\tilde{W}}{db} \approx \frac{u'\left(\mathbb{E}^l[c^l]\right) - u'\left(\mathbb{E}^h[c^h]\right)}{u'\left(\mathbb{E}^h[c^h]\right)} - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}$. In this case the marginal benefit of program expansion is given by the percentage gap in the marginal utility of low- and high-state risk-weighted expected consumption. Taking a further first-order approximation of $u'\left(\mathbb{E}^l[c^l]\right)$ around $\mathbb{E}^h[c^h]$, leads to the approximation $\frac{d\tilde{W}}{db} \approx \gamma \frac{\mathbb{E}^h[c^h] - \mathbb{E}^l[c^l]}{\mathbb{E}^h[c^h]} - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}$.[9] This is similar to the Baily-Chetty formula, but with the key difference that the consumption drop depends on the difference in risk-weighted consumption levels in the high and low state, relative to in the high state. This expression can be decomposed into terms reflecting the insurance and redistributive effects of a marginal social insurance expansion, where, in this case, the welfare weights depend only on risk and consumption and are given by: $\bar{\lambda} = \frac{\mathbb{E}^l[c^h]}{\mathbb{E}^h[c^h]}$ and $\lambda_i = \frac{c_i^h}{\mathbb{E}^l[c^h]}$.[10]

[9] In this formula γ is the coefficient of relative risk aversion evaluated at *average* high-state consumption. Note, an implication of this is that even if all individuals have the same *quadratic* utility function, in which case $\mathbb{E}^s[u'(c^s)] = u'(\mathbb{E}^s[c^s])$, this equation remains an approximation as the coefficient of relative risk aversion necessarily differs between individuals with heterogeneous consumption levels and quadratic utility functions.

[10] Undertaking the decomposition after taking the two first-order approximations described in the text leads to the slightly modified version of equation (7) given by: $\frac{d\tilde{W}}{db} \approx \bar{\lambda}\mathbb{E}^l\left[\gamma \frac{\Delta c}{c^h}\right] + \bar{\lambda}\text{cov}^l\left(\lambda, \gamma \frac{\Delta c}{c^h}\right) + \gamma(\bar{\lambda} - 1) - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}$.

No redistributive motives. If the welfare weights, λ_i, equal 1 for all individuals, equation (7) simplifies to $\frac{d\tilde{W}}{db} \approx \mathbb{E}^l \left[\gamma \frac{\Delta c}{c^h}\right] - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}$ (the case considered by Andrews and Miller 2013). Here, the marginal gain from social insurance expansion equals the low-state risk-weighted average of individuals' insurance value. The planner ignores any additional redistributive effects from the reform. The Pareto weights that lead to $\lambda_i = 1$ are $\omega_i = \frac{1}{u_i'(c_i^h)}$. Relative to the utilitarian weighting of individuals, this implies the planner places relatively more weight on those with *high* high-state consumption. While these weights focus analysis on the average insurance value of program expansion, they encode social preferences that are very different from those typically used to value the social effect of policy reform.

2.3. Extensions

General reforms. So far we have focused on the effects of a rise in a lump-sum benefit funded by a rise in a lump-sum tax. In Appendix A.2.1 we show that our framework can straightforwardly accommodate reforms in which the tax and benefit changes are arbitrary functions of income. This generalisation results in two key changes to equations (3)-(5) and (7).

First, it alters the risk-weighting in the expectation operators. For a flat reform the risk weighting in the low-state expectation is proportional to $(1 - e_i)$. In contrast, for a general reform, it is proportional to $(1 - e_i)\rho^l(y_i)$, where $\rho^l(y_i)$ represents the marginal benefit increase for someone with (high-state) income y_i. Similarly, the risk-weighting in the high-state expectation operator becomes proportional to $e_i\rho^h(y_i)$, where $\rho^h(y_i)$ represents the marginal tax rise for someone with income y_i.[11] This results in two changes to equation (7). First, it alters the planner's aggregation over individual-level insurance values. For instance, all else equal, the more a benefit increase is weighted towards individuals with large consumption declines or high social marginal utility (and hence λ_i), the greater the welfare improvement. Second, it modifies the redistributive effect of the reform, captured by the aggregate welfare weight $\bar{\lambda}$. Specifically, the more a benefit increase is weighted towards individuals with high social marginal utility (or the tax increase towards those with low social marginal utility), the higher $\bar{\lambda}$ will be. For small reforms, the insurance and redistributive effects depend only on the nature of the reform, not the shape of the pre-reform tax and benefit schedules.

[11] For concreteness we consider tax and benefit reforms that depend on high-state income, however, we could equally consider benefit reforms that additionally depend on low-state income. Note, for a flat benefit and tax reform $\rho^l(y_i) = \rho^h(y_i) = 1$.

Second, under more general reforms, the term capturing the incentive cost of the reform changes. For a flat tax and benefit, moral hazard costs are captured by the aggregate elasticity of the low-state probability with respect to the benefit level. However, for other reforms, the relevant elasticity differs. In other words, the moral hazard elasticity is reform-specific (see Appendix A.2.2 for further discussion).

Proportional reforms. A specific alternative reform, which aligns with the design of many social programs, involves raising the replacement rate (i.e., increasing benefits in proportion to high-state income). In this case, $\rho^l(y_i) = y_i$, meaning the low-state risk weighting of expected social marginal utility is proportional to $(1 - e_i)\, y_i$, rather than $(1 - e_i)$ as under a flat reform. Consequently, for an inequality-averse planner, $\bar{\lambda}$ (via its numerator) is reduced relative to the baseline flat reform case. This reflects the fact that an increase in the replacement rate provides larger monetary benefits to higher-income households compared to lower-income households.

Conversely, if the tax is proportional to income – relative to a flat tax – this acts to raise $\bar{\lambda}$. In this case, the risk weighting in the denominator becomes $e_i\, y_i$, rather than e_i, reflecting the greater incidence of a tax rise on high-income individuals.

When both the benefit and tax rates are flat, as discussed above, the incentive costs are captured by $\frac{\mathbb{E}^l[\epsilon_{1-e,b}]}{\bar{e}}$. The numerator represents the low-risk-weighted expectation of the individual-level elasticity of the low-state probability with respect to a marginal benefit increase, while the denominator is the average high-state probability. If instead the benefit and tax rates are proportional to income, the incentive cost takes a similar form but with two key modifications. First, the expectation over the individual-level elasticities is weighted by $(1 - e_i)\, y_i^h$. This adjustment changes the moral hazard cost of reform when individual-level elasticities are correlated with either risk or income. For example, if higher-income individuals are more elastic, an expansion of a proportional system will be more costly, as it provides stronger disincentives to those with higher incomes compared to a flat system. Second, the denominator becomes the *income-weighted* high-state probability, reflecting the incidence of the reform.

Often social insurance programs entail a hybrid benefit that involves replacing a fraction of income up to a cap; that is, for those with incomes below the cap, the benefit is proportional to income, while for those with incomes above the cap, the benefit is lump-sum. In this case a natural program expansion raises the replacement rate for those with incomes below the cap, and raises the lump-sum benefit in proportion to the cap for those with higher incomes. Our framework can easily accommodate this

by combining the flat and proportional marginal program expansion conditions, each applied to the relevant subset of individuals.

Within group budget-balance. If the planner sets an actuarially fair flat tax and benefit that differs across different groups of individuals (based on immutable characteristics) equations (3)-(5) and (7) hold within each group. Now the social insurance expansion only redistributes within, and not across, groups. If the planner can implement such differentiation at the individual level, equations (3) and (7) collapse to an individual specific Baily-Chetty formula.[12] In reality, in most contexts, implementing individual-level tax and benefit rates will be infeasible due to heterogeneity that is not observed by the planner (e.g., in attitudes to work). In addition, the planner's scope to set group-specific policy is likely to be constrained by normative objections to tagging based on fixed characteristics.

State dependence. The approximation to the gap in marginal utilities across states in equation (6) assumes that the marginal utility from a given level of consumption is state independent. However, it may be that an observed level of consumption expenditure yields levels of marginal utilities that vary across states. For instance, leisure time may vary by state and be a complement to or substitute for consumption, there may be state-specific expenditures (e.g., work-related expenditures) that do not directly generate utility (Browning and Crossley 2001), or the extent to which people combine consumption expenditures with home production may vary across states (Aguiar and Hurst 2005). A convenient way to capture state dependence is through a marginal-utility shifter ϕ_i, such that $u_i^{h'}(c) = u_i^{'}(c)$ and $u_i^{l'}(c) = \phi_i u_i^{'}(c)$. We show in Appendix A.2.4, that this leads to a modification in the welfare weights in equation (7) to account for differences in the marginal utility of a given level of consumption across states; they become $\lambda_i = \frac{\phi_i \omega_i u_i^{'}(c_i^h)}{\mathbb{E}^l[\phi \omega u^{'}(c^h)]}$ and $\bar{\lambda} = \frac{\mathbb{E}^l[\phi \omega u^{'}(c^h)]}{\mathbb{E}^h[\omega u^{'}(c^h)]}$.

A related issue is that the time individuals devote to searching out lower prices may vary by state, meaning that the market price of consumption does too. This can also be accommodated with the parameter ϕ_i, which then has the interpretation of

[12] On the other hand, if the planner can differentiate tax rates, but not the benefit level, for instance, because the program is funded using experience rating, the planner must once again aggregate across individual-level insurance and cross-subsidisation values. The difference with equation (7) is, in this case, if tax differentiation leads to differential incidence, individual cross-subsidisation is within-group (who face the same tax liability) and the planner must assign across-group welfare weights to account for differences in tax burdens. We formalise this argument in Appendix A.2.3.

being the price of consumption in the high (relative to the low) state. However, in this case it is necessary to also adjust observed consumption expenditures to account for state-specific price differences (Campos and Reggio 2020).

Richer model of behaviour. The representation of the individual's choice problem in equation (1) is simple; the model is static and the agent makes a single choice over effort, with state-specific consumption being fully determined by their static budget constraint. In Appendix A.2.5 we extend the model to a dynamic environment in which, at each point in the time, the agent makes choices over consumption and an arbitrary number of other actions (which may include elements of search effort, savings behaviour, private insurance purchases, family labour supply decisions etc.) and constraints. We show that, under relatively mild regularity conditions, we obtain a condition very similar to equation (3) (and hence equation (7)). The principal difference is that, while in equation (3) risk-weighted averages are taken across individuals, in the dynamic model they are taken both across time for a given individual and across individuals. The addition of an arbitrary number of additional choices and constraints on the individual problem does not alter the characterisation of the welfare effect of program expansion, as the envelope condition ensures that the marginal utility of consumption of an agent in each state fully captures the private value of shifting a $ between states (Chetty 2006).

3. Unemployment Insurance in the U.S.

Our preceding analysis highlights how individual-level heterogeneity shapes the insurance and redistributive value of social insurance. We now turn to understanding the quantitative importance of these channels for policy design. To do this we apply the theoretical insights that we develop for an abstract social insurance program to the study of a concrete and prominent example of social insurance: unemployment insurance (UI).

In the United States,[13] while the federal government provides guidelines for UI eligibility and benefit generosity, UI is administered at the state level with states determining benefit duration, generosity and eligibility (see Von Wachter 2019 for a discussion). Benefit payments, and entitlement, are computed using workers' (quarterly) employment history and earnings. The modal duration offered by states is

[13] Spinnewijn (2020) provides an overview of various countries around the world.

six months outside of federal recession extensions. Eligible workers receive payments in proportion to their past earnings up to a maximum threshold.

The U.S. UI system is funded through payroll taxes legally incident on employers. The federally mandated level of taxable earnings is $7,000 per year. While most states choose a higher base for taxable earnings, in most cases the tax acts as a head tax on the quantity of workers rather than a proportional earnings tax. A unique feature of the funding of U.S. UI is that tax rates vary across employers as the system is partially experience rated (see Guo and Johnston 2021 for discussion). Recent evidence suggests that firms are unable to pass-through the firm-specific component of taxes to wages (Guo 2024). As a result, in our empirical quantification of UI expansion, we assume that it is the common component of UI taxes that is incident on wages.[14]

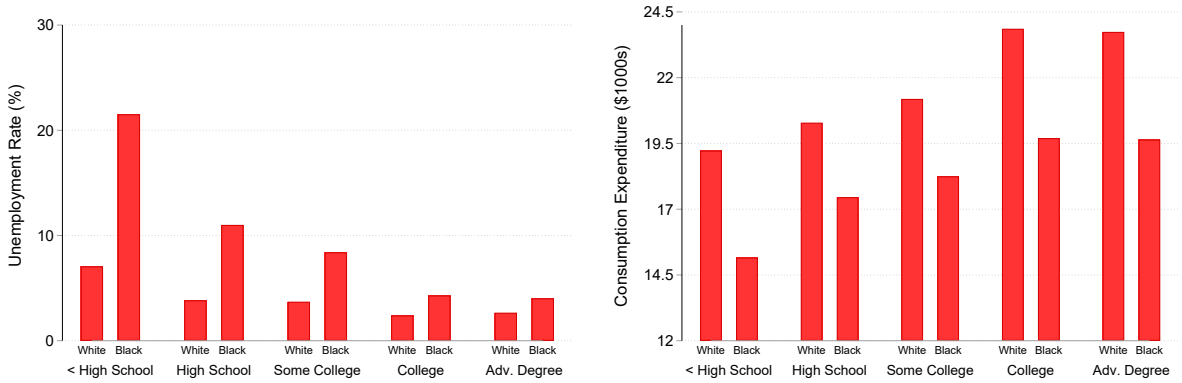
We use data from the Panel Study of Income Dynamics (PSID), a nationally representative survey of the U.S. population. The PSID has surveyed households since 1968 collecting information on demographics as well as labour market outcomes including unemployment, earnings, and wages. We use data from 1999-2019, the biennial so-called "new" PSID, which includes comprehensive information on consumption expenditures.[15]

We drop households from our sample where the reference person does not provide any information on their educational attainment, and we focus on a sample where the reference person is aged between 23 and 60. This drops people from our sample whose labour supply is likely driven by education and retirement decisions. To proxy for UI program eligibility, we exclude households where the reference person never reports being employed, as well as periods immediately following a spell out of the labour force. We use current employment status to construct our measure of unemployment. This yields 55,671 observations (15,342 families observed for an average of 3.6 waves) of which 3,332 observations correspond to the reference person being unemployed (2,211 families for an average of 1.5 waves). For families with changes to the number of adults (such as divorce, marriage, or death) we construct panel identifiers for demographically stable units and omit the year of the change. We focus on the labour market status of the reference person (in some places referring to them as worker). We use a

[14] Guo (2024) show the firm-specific component of the tax instead leads to employment responses. See Spaziani (2024) for an analysis of how experience rating alters the risk borne by employers. Anderson and Meyer (2000) similarly show that firm-specific components are largely incident on firms, but emphasise that employment responses extend to layoffs and subsequent UI claims (and their denials).

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

comprehensive measure of expenditures on non-durable consumption and services excluding expenditures on health care and housing.[16] To account for differences in household composition we equivalise consumption expenditures using the square root of household size. We provide summary statistics for our sample as well as additional details on sample selection and construction of key variables in Appendix B.



(A) Unemployment Rate

(B) Consumption

FIGURE 1. Unemployment Rates and In-Work Consumption Expenditures

Notes: Authors' calculation from the PSID data pooling all households and survey waves. We report observed unemployment rates (left-hand panel) and equivalised consumption expenditures (right-hand panel) by race and education. We omit a third race group due to small sample size. We describe our consumption measure above and report all values in 2019 U.S. dollars. We provide more details in Appendix B. The correlation between observed unemployment status and consumption expenditure on non-durable goods and services in the micro-data is –0.12 and statistically significant at the 0.01 level.

In Figure 1 we present descriptive evidence that employment risk and in-work consumption vary systematically across demographic groups. For each level of education, we report the pooled unemployment rate (left-hand panel) and average in-work consumption expenditures (right-hand panel) separately among white and black households. For white households, the unemployment rate declines monotonically with eduction, from a high of almost 10% for those with less than high school education, to a low of 4% for those with advanced degrees. The pattern for black households is even starker: falling from nearly 30% to 5%. Comparing black and white

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

households with the same education level reveals a well-documented pattern of higher unemployment rates among black households, with the race gap declining in education. For instance, the unemployment rate for black households with less than high school education is three times that of white households. While for those with an advanced degree it is 1.5 times as large.

For average consumption expenditures we see the same patterns in reverse. Both within white and black households, those with more education tend to have larger consumption expenditures. In addition, black households have systematically lower consumption at all levels of education than white households.

The patterns in Figure 1 show that those who belong to groups with relatively high unemployment rates typically have lower consumption expenditures. Our analysis in Section 2 highlights that heterogeneity of this form, along with workers' exposure to unemployment, as reflected in their consumption decline between high and low states, drive the welfare effects of expanding UI.

4. Measuring Employment Risk and Risk Exposure

In this section we outline how we measure heterogeneity in unemployment-induced consumption declines, and in employment risk, which, along with the coefficient of relative risk aversion determine worker risk exposure. In each case we detail our estimation strategy and results, and draw comparison to the existing literature.

4.1. Consumption Declines

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

$$\begin{aligned}\Delta_{i,t}^{FD} &\equiv \log\left(c_{i,t}\right) - \log\left(c_{i,t-2}\right) \\ &= \delta_0 + \delta_1 U_{i,t} + \beta X_{i,t} + \varepsilon_{i,t},\end{aligned} \tag{8}$$

where $U_{i,t}$ is an indicator variable denoting whether the worker is unemployed in period t and $X_{i,t}$ are demographic controls. This is specified in first differences to allow for arbitrary permanent individual-level heterogeneity in consumption levels. The parameter of interest, δ_1, measures the average proportional consumption decline that results from becoming unemployed: $\delta_1 = E\left[\Delta_t^{FD} \mid U_t = 1\right] - E\left[\Delta_t^{FD} \mid U_t = 0\right]$. The

identification assumption underpinning this strategy is that the trend in consumption for the employed acts as a valid counterfactual for the evolution of in-work consumption of the unemployment had they remained employed.

To estimate the distribution of consumption declines we specify a finite mixture approximation to the true distribution (e.g., Heckman and Singer 1984 or Keane and Wolpin 1997). Our approach follows Lewis et al. (2024), who estimate the distribution of marginal propensities to consume out of tax rebates. We assume each worker belongs to one of a finite number of latent types and we augment the specification in equation (8) as follows:

$$\Delta^{FD}_{i,t} = \sum_{g \in G} \mathbb{1}\left[g(i) = g\right] \left(\delta^g_0 + \delta^g_1 U_{i,t}\right) + \beta X_{i,t} + \varepsilon_{i,t}. \tag{9}$$

We include group specific intercepts, δ^g_0, which means that δ^g_1 have the interpretation of the consumption decline for group g at unemployment, and we assume that, conditional on group, controls, and unemployment the error term is mean zero. To implement this we use a Gaussian mixture linear regression (see Quandt 1972), which entails assuming that the errors are normally distributed with group-specific variances.[17]

The key identifying assumption in equation (9) is that, conditional on group, the evolution of consumption for the employed acts as a valid counterfactual for the unobserved trend in consumption had the unemployed remained employed. This is similar to the identification assumption underlying equation (8), though applied within group. Latent types are identified by grouping households with similar distributions of consumption growth in each state. We jointly estimate group membership and the other parameters; see Appendix C for details.

We recover estimates of the group-specific parameters, the common effect of covariates, the group-specific variance of the regression residuals, and the unconditional probability of group assignment (π^g for each group g) by maximum likelihood estimation. Inference on these objects is based on the Fisher Information Matrix. Additionally, we also compute "posterior" probabilities of group assignment, $\hat{\pi}^g_i$, at the individual level, which depend on our parameter estimates and the individuals' sequence of observed regressors and outcome variables. We use these to

[17] Boehm et al. (2025) adopt an alternative non-parametric approach relying on an imputation based estimator, in the spirit of Borusyak et al. (2024), and deconvolution to recover the distribution of marginal propensities to consume. While we could use a similar approach to recover conditional distributions (and therefore covariances), our dataset is an order of magnitude smaller which led us to adopt these parametric restrictions.

simulate the value of unemployment-induced consumption declines within sample with $\Delta c_i/c_i$ as $\sum_{g \in G} \hat{\pi}_i^g \hat{\delta}_1^g$.[18] These results are conditional on the specified number of groups G which we select using the Bayesian Information Criterion.

As discussed in Lewis et al. (2024), this mixture regression can be viewed as a form of clustering regression which jointly (i) groups households together that have similar latent consumption responses to unemployment and employment and (ii) provides estimates of the consumption decline within these groups. Thus, our approach to estimating counterfactual consumption declines for those without observed unemployment spells can be viewed as a form of matching-based identification strategy. Intuitively, the approach can be thought of as entailing the following steps. First, we group together households observed in both states based on their consumption growth when employed and their consumption decline when unemployed. We then 'match' households who are always employed to those who experience similar consumption growth when employed, conditional on observables. Then, we assign them the matched household's observed consumption decline for their counterparts. In practice, we simultaneously estimate probabilistic group assignment and parameters, with the posterior weight, $\hat{\pi}_i^g$, serving as a (convex) imputation weight.[19]

Table 1 reports our estimates from the finite mixture approach alongside estimates from the homogeneous specification (equation (8)). The homogeneous specification yields an estimate of the average consumption decline of 12.8% at the onset of unemployment. Our estimate is slightly larger than those reported by the existing literature (e.g., Gruber 1997 and Hendren 2017) which find declines in food expenditures of between 7 and 10%. There are three differences between our approach and those in the existing literature that can together account for this. First, we measure

[18] An alternative approach would be to specify group membership conditional on the variables that enter our relative unemployment risk model and individual social welfare weights, in which case, we could directly compute moments of interest from these conditional distributions assuming we specified group membership flexibly. In practice, this requires successively finer cuts of the data which is infeasible given our sample size.

[19] Note that if two groups experience similar consumption growth when employed, but different consumption declines (e.g., due to varying abilities to self-insure), our imputation assigns a household with similar consumption growth (who is only observed employed) a consumption decline that is a weighted average of these groups. As we show in Appendix C, these weights depend not only on the average consumption growth, but also its variability, which corresponds to the pass through of shocks to consumption (Blundell et al. 2008). Our use of a mixture model to address the lack of common support in observed unemployment is similar to the use of grouped fixed effects in Bonhomme et al. (2019) to relax the connected set requirement of worker-firm fixed effect estimators (e.g., Abowd et al. 1999). In principal, latent consumption risk may vary over time, and our approach could be extended to allow for a regime-switching model.

consumption declines based on comparing consumption when unemployed with in-work consumption two years before. 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. We show in Appendix D.1 that our results are robust to using a wide range of alternative consumption measures. Third, our sample differs from those used previously, and is drawn from different survey waves. Comparing the response of food expenditures to unemployment in the PSID, East and Kuka (2015) document an increasing trend in the average decline, which they attribute to factors distinct from changes in the survey design.[20]

The implied average consumption decline from our mixture approach, aggregating the group-specific consumption decline δ_1^g, weighted by π^g, is slightly smaller at 11.6%.[21] This comparison of averages, however, ignores the large degree of heterogeneity in the exposure to unemployment-induced consumption declines that we find. The group-specific consumption declines we estimate range from 25.6% to 6.0%. The probability mass of the group with the smallest decline, 6.0%, is almost 40%. This group is relatively well insured against unemployment risk. A second group, with probability mass of 49%, has a consumption drop that equals the homogeneous estimate. The remaining probability mass is accounted for by a group with a much larger consumption decline of 25.6%

[20] Similar to our analysis, East and Kuka (2015) find only small differences by consumption measure. They attribute the rise in declines to a time trend rather than temporal sampling over the business cycle. In addition, our sample selection results in a sample that is less likely to be married and less likely to be white, as evident when comparing summary statistic we report in Table B.1 of Appendix B and Table 1 of Kroft and Notowidigdo (2016), who reproduce the Gruber (1997) estimate.

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

		Mixture Model		
	Homogenous	Group 1	Group 2	Group 3
$U_{i,t}$	−0.128***	−0.256***	−0.128***	−0.060***
	(0.009)	(0.015)	(0.006)	(0.006)
Share ($\hat{\pi}^g$)		0.118***	0.485***	0.397***
		(0.028)	(0.018)	(0.021)
Controls				
Age	*Yes*	*Yes*	*Yes*	*Yes*
Year	*Yes*	*Yes*	*Yes*	*Yes*
N	37, 332	37, 332	37, 332	37, 332

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

* $p < 0.10$, ** $p < 0.5$, *** $p < 0.01$. The first column reports the average consumption decline estimated under the assumption of homogeneous consumption drops (equation (8)). The remaining columns report the consumption declines, δ_1^g, and population shares, π^g, in our Gaussian mixture linear regression framework (equation (9)). In the mixture model standard errors account for uncertainty over both group assignment and parameter estimates conditional on group. In addition to the unemployment indicator, we control for a cubic polynomial in age, a series of year dummies, and the log of the change in family size. Results in Table D.1 provide heteroskedastic robust standard errors for the homogeneous estimate.

Table 1 describes estimates of the degree of unconditional heterogeneity in unemployment-induced consumption declines. However, in the framework we layout in Section 2 the social planner cares about the relationship between workers' consumption drops, their employment risk, and in-work consumption. We therefore compute the posterior predicted consumption decline for each worker, which is computed to minimise the mean squared error of their consumption decline. Figure 2 shows this distribution, which further emphasises the cross-sectional heterogeneity in risk exposure. Using the standard homogeneous estimator, we have a mass point at the estimate (the dashed vertical line). Instead, we find that 51% of households exhibit consumption drops that are smaller than the standard homogeneous estimate, but there is a long right tail of households with larger drops. This variation reflects broad

differences in the ability of households to self-insure unemployment risk.

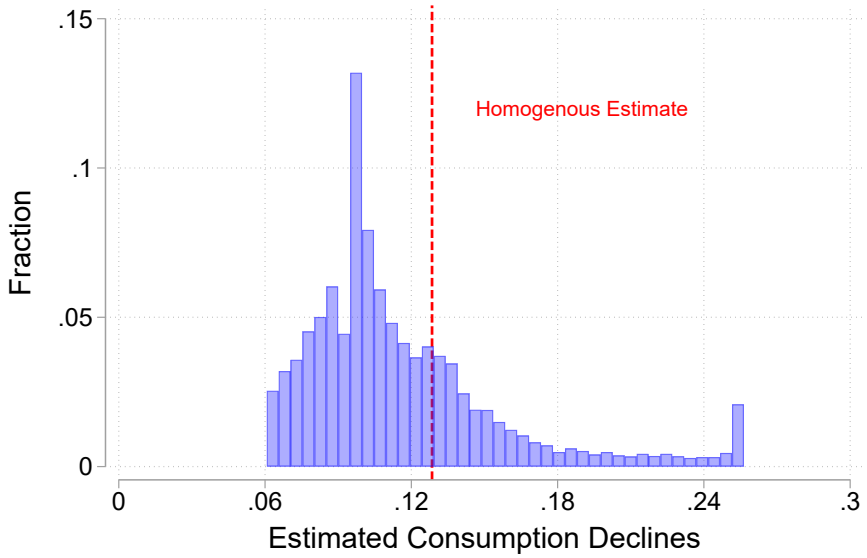


FIGURE 2. Distribution of Estimated Consumption Declines

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

Our finding of substantial heterogeneity in exposure to unemployment-induced consumption declines is consistent with a large body of evidence on how households respond to job-loss. Browning and Crossley (2001) find that the marginal propensity to consume out of UI benefits is largest for those without liquid savings. Using direct information on binding credit constraints, Crossley and Low (2014) find that 5% of job-losers face binding credit constraints and experience an especially large, welfare-reducing, consumption decline at job loss.[22] Ganong and Noel (2019) also show that consumption declines vary systematically with 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 associated with job-

[22]They additionally find widespread excess sensitivity to the income decline at job-loss that is an order of magnitude smaller: 7.7% for unconstrained households compared to an additional 27.1% for those with binding constraints. In other words, those who are able to perfectly smooth do not. Similarly, we find a large number of households experience moderate, but statistically significant consumption declines.

loss primarily by decumulating liquid savings, but also through added worker effects. Finally, both Patterson (2023) and Colarieti et al. (2024) document *marginal* propensities to consume out of unemployment income losses of between 0.5 to 0.6. When different workers face different income declines (for example married and single workers), even constant marginal propensities to consume imply heterogeneous consumption declines. Moreover, Patterson (2023) shows evidence of heterogeneity in these marginal effects by observables.

In Appendix D.2 we report how our estimated unemployment-induced consumption declines correlate with proxies of a household's capacity to self-insure, providing evidence of larger consumption declines for those with low liquid assets and that do not cohabit (and therefore are less likely to benefit from added worker effects).

4.2. Employment Risk

The terms capturing the insurance and redistributive impact of social insurance expansion, in equation (7), depend on risk-weighted expectations in the high and low states, which in the context of UI, correspond to being employed and unemployed. We therefore need to construct individual-level risk weightings. Our approach is to use a statistical model of unemployment to recover rich demographic-specific unemployment rates,[23] which infers how risk varies across individuals based on the observed realisations of the lotteries it induces, and sidesteps the need to specify the full model underlying workers' effort choices.

We specify a logistic regression[24] where the probability of individual i being unemployed at time t depends on the following observed characteristics: a cubic function of their age, indicator variables for their gender and whether they are

[23] An alternative approach is to exploit idiosyncratic and subjective probabilities of job loss (Hendren 2017) and job finding (Spinnewijn 2015; Mueller et al. 2021). Given these measures are not directly available in the same dataset, to use them, we would in practice need impute them as functions of similar demographics.

[24] Consider a dataset that consists of an indicator variable, U_{it}, denoting whether individual i was unemployed during period t, and a vector of attributes, $X_{i,t}$, for individual i that capture the key dimensions over which employment risk varies. With rich enough data we could construct the unemployment risk weight for an individual with attributes X, $1 - e(X)$, from $\widehat{\mathbb{E}[U|X]}$ based on a non-parametric regression for the expected mean function $\mathbb{E}[U|X]$. However, as the PSID data has only a moderate sample size, we instead use a parametric approach.

married, and a set of indicator variables capturing their education status and race.[25]

This regression allows us to recover estimates of their employment probability, $\hat{e}_{i,t}$. To construct relative risk weights we pool all observations into a single cross-section and compute the sample average $\bar{\hat{e}}$ of these predicted probabilities. In the case of a flat reform, our risk weights are then $\frac{\hat{e}}{\bar{\hat{e}}}$ and $\frac{1-\hat{e}}{1-\bar{\hat{e}}}$ for the employed (high) and unemployed (low) states respectively. For other types of reform, which entail differential increases in benefits or taxes across workers, the risk weights will additionally depend on this variation.

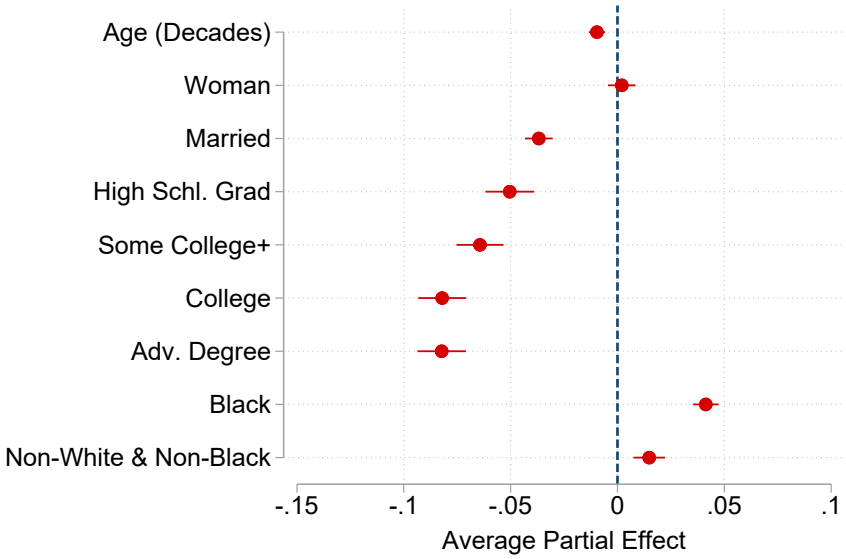


FIGURE 3. Effect of Household Characteristics on Predicted Probability of Unemployment

Notes: Authors' estimates from the PSID data. We estimate the logistic regression described in the text using employment status as the dependent variable. We calculate the marginal effects of each covariate on the unemployment probability using parameter estimates and integrate over the empirical distribution of household characteristics. The centre of each circle corresponds to the point-estimate of the average partial effect (APE) and the horizontal lines span the 95% confidence intervals. The red dashed line indicates the origin. N=55,671 and standard errors are clustered at the household level. We report point-estimates for the APE and the underlying coefficients along with their standard errors in Table D.3 of Appendix D.3.

[25]We show in Appendix D.3 that our estimated employment risks are extremely similar when we explicitly model them as a function of the estimated consumption declines. This suggests that additional factors contributing to latent heterogeneity in consumption declines do not identify persistent employment risk types. In addition, while including time effects improves the overall goodness-of-fit we find no important differences on the log odds ratio of our included regressors; in other words relative employment risks are stable over time.

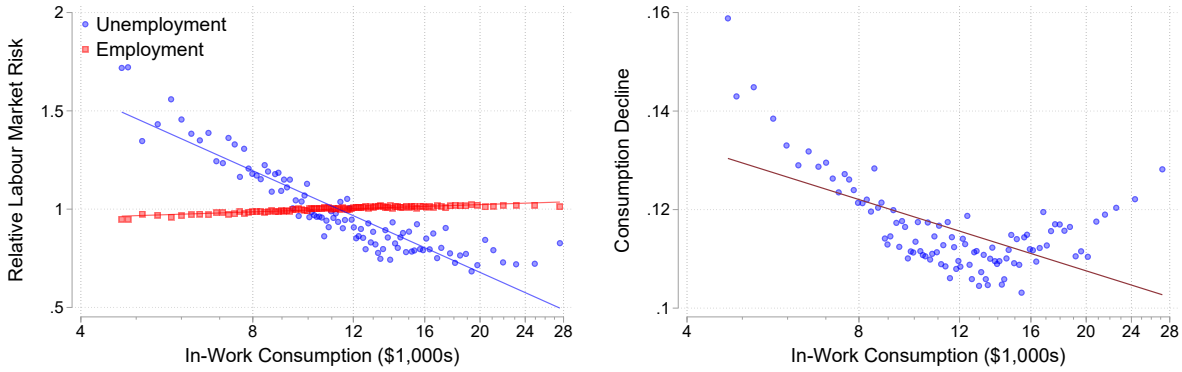
We report the estimated coefficients in Table D.3 of Appendix D.3. Figure 3 reports average partial effects along with their 95% confidence intervals. We find evidence of systematic differences in unemployment rates across demographic groups. Younger households are more likely to be unemployed with a decade decrease in age associated with a 1 percentage point increase in unemployment. We find a strong educational gradient, with those with a college degree 3.2 percentage points less likely to be unemployed than those who graduate high school but do not go to college. Even after controlling for educational achievement we find a large race gap of almost 4.1 percentage points between white and black households. This is more than twice the size of the gap between our third race category and white households. Conditional on these other characteristics we find no effect for gender. While the simple search model we outline in Section 2 rationalises these differences through individuals' effort choices, it is important to stress that as well as reflecting the incentive to find work and enjoy higher income, they will also reflect differences in market thickness, equilibrium congestion externalities, skill-biased labour demand or outright discrimination, which are captured in our model in differences in the cost of exerting effort (i.e. in $\psi_i(.)$).

These patterns are consistent with the descriptive evidence we present in Figure 1 as well as patterns documented in the broader literature. This includes a positive unemployment rate differential between young and old workers (e.g., Choi et al. 2015), unemployment rates that decrease in human capital (e.g., Ashenfelter and Ham 1979, Nickell 1979, and Cairó and Cajner 2018), and evidence of a racial gap in employment outcomes (e.g., Lang and Lehmann 2012) and callbacks (e.g., Bertrand and Mullainathan 2004). While we find no statistically significant employment gap for women, we emphasise that these are women who are the reference person in the PSID and disproportionately single (participation gaps for unmarried men and women are small, e.g., Borella et al. 2023).

4.3. Variation in Risk and Risk Exposure by Ex-Ante Neediness

A worker's willingness-to-pay for UI expansion depends on their consumption decline, which influences their insurance-value, and how their employment risk compares with the population average, which determines the degree to which they are net recipients or contributors to program cross-subsidisation (see equation (2)). The value the planner places on UI expansion depends on how it weights these effects across households. In this section we provide evidence on how employment risk and consumption declines

vary with equivalised in-work consumption, a measure of households' neediness.[26]



(A) Cross-Subsidisation and Consumption

(B) Insurance and Consumption

FIGURE 4. Consumption Exposure, Unemployment Risk, and Employed Consumption

Notes: Authors' calculation from the PSID data. We report measures against equivalised panel average non-durable consumption in the employed state. We plot this on a log-scale. Our measures of risk and consumption declines are estimated as we describe above and correspond to the model results discussed in Figures 3 and 2. We use binned scatter plots with 100 bins in each panel.

Figure 4(A) shows how workers' relative labour market risks, for unemployment and employment, vary by their households' average consumption level when employed. It shows a positive relationship between *relative* employment risk and consumption. Those with the highest (lowest) levels of consumption have the highest (lowest) likelihood of being employed. Mirroring this, there is a negative relationship between relative *unemployment* probabilities and in-work consumption. A doubling of consumption is associated with a decline in the relative unemployment probability of 56 percentage points. As workers tend to be considerably more likely to be in employment than unemployment the slope of the relationship between consumption and relative unemployment risk is steeper than that with relative employment risk. The overall pattern is consistent with the descriptive evidence by education and race we document above in Figure 1.

Figure 4(B) plots how our estimates of the consumption decline co-vary with consumption in the employed state. It shows, over much of the domain of in-work consumption, a positive relationship with risk exposure. Those with the largest consumption declines, and who therefore incur the biggest utility costs from being unemployed, are also systematically more likely to have relatively low consumption

[26] In calculating the marginal welfare effect of UI expansion, we use estimated risk-weights to explicitly aggregate consumption declines. In Appendix D.3, we show how the demographic characteristics that inform risk-weights correlate with consumption declines.

when employed. In other words, the most needy individuals tend also to have the most exposure to unemployment risk, and therefore will benefit the most from the increased insurance provided by UI expansion.

5. The Marginal Welfare Effects of UI Expansion

In this section we combine our estimates of the distribution of consumption declines and unemployment risk with the social insurance design framework we outline in Section 2. This allows us to empirically assess the marginal value of expanding UI benefits (by numerically evaluating equation (7)). To do this, we calibrate the baseline values for the parameters capturing workers' attitudes to risk and the social planner's attitude towards redistribution.

5.1. Risk Aversion and Welfare Weights

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

The social welfare effects of program expansion depend additionally on the marginal social welfare weights, λ_i, which are proportional to the product of the Pareto weight and individual marginal utility of consumption, evaluated at high-state consumption, $\omega_i u_i'(c_i^h)$. Following the optimal tax literature (e.g., Saez 2002) we use a schedule for the social marginal welfare weights, $\omega_i u_i'(c_i^h)$, that is proportional to $\left(c_i^h\right)^{-\psi}$, where $\psi \geq 0$ is a parameter that captures the planner's degree of inequality aversion, and we measure c_i^h based on panel-average equivalised in-work consumption. In our baseline results, we set $\psi = 1$.

By setting $\gamma > \psi$, we encode a set of preferences consistent with evidence that risk aversion over moderate stake risk (e.g., unemployment) is larger than over large stake risk (e.g., redistribution of lifetime wealth), consistent with Rabin (2000). Chetty and Szeidl (2007) microfound this through ex ante consumption commitments, where some components of consumption are costly to adjust, which lead households to respond to moderate shocks by cutting flexible consumption components more deeply than if all consumption components could be freely adjusted. Kaplan and Violante (2014) show that this behaviour can also be microfounded by models where illiquid assets provide

consumption flows (e.g., housing services). We view this calibration as conservative as, relative to setting $\gamma = \psi$, it down weights the redistributive role of social insurance relative to its insurance role. In Section 6 we show how our results vary with different values of (γ, ψ).

5.2. Marginal Gain

Equation (7) separates the effect of marginal UI expansion into i) its impact through redistributing resources across states and individuals, and ii) the moral hazard cost of discouraging work (captured by $\frac{\epsilon_{1-\bar{e},b}}{\bar{e}}$). In Table 2 we focus on the first effect, referring to it as the welfare gain from program expansion, leaving discussion of the moral hazard cost to the following section.[27] We report results for three alternative benefit reforms, assuming in each case the benefit increase is funded by a flat tax increase. This assumption is motivated by the fact that U.S. UI is funded by a payroll tax, with a federally set taxable wage base equal to the first $7000 of earnings. We explore alternative benefit and tax reforms in Section 5.4.

Row (1) in Table 2 shows results for a flat increase in UI benefits. The total welfare gain from this reform is 0.519, indicating that a balanced-budget $1 increase in UI benefits raises welfare by 51.9% as much as a $1 increase in in-work incomes, after adjusting for differences in the size of the pool of employed and unemployed. This welfare gain is significantly higher – both economically and statistically – than the average insurance value of 0.370. The percentage bias resulting from relying solely on the average insurance gain to measure the welfare gain (consistent with the standard Baily-Chetty formula) is -28.8%. In other words, neglecting worker heterogeneity substantially underestimates the welfare gain from benefit expansion.

In Section 2 we highlight two interrelated channels through which worker heterogeneity influences the welfare gains of social insurance expansion. First, there is the redistributive effect of transferring resources from the employed to the unemployed (see Figure 4(A)), captured in the aggregate welfare weight ($\bar{\lambda}$). Second, there is the influence of any correlation in worker-level insurance values and their level of neediness (see Figure 4(B)), captured by the covariance of insurance values and

[27] The table includes 95% confidence intervals based on 400 bootstrap replications, clustered at the household level. The quantification in Table 2 is based on a broad consumption measure of non-durable consumption and services. Previous work on the generosity of UI in the U.S. uses food expenditures, historically available in survey data, to estimate the insurance value of UI. In Appendix D.3 we show we obtain similar results when we instead use changes in food expenditure to estimate unemployment-induced consumption declines.

individual welfare weights ($\text{cov}^l\left(\lambda, \gamma\frac{\Delta c}{c^h}\right)$). Table 2 shows that while both channels contribute to raising the welfare benefit of expansion, the redistributive effect is quantitatively more important.

Row (2) shows the results for a benefit expansion proportional to earnings. The total welfare gain of this reform is 0.303, substantially lower than that of a flat reform. The difference arises because proportional reforms alter the risk-weighting in the expectations in equation (7). For proportional expansions, the weights depend on earnings in addition to unemployment risk, as higher earners benefit more from the expansion is $ terms, conditional on unemployment risk. Table 2 indicates that the most significant quantitative effect is a reduction in the aggregate welfare weight, $\bar{\lambda}$, which falls below 1. The percentage bias from ignoring worker heterogeneity in the case is +16.2%. In other words, neglecting worker heterogeneity overstates the welfare gain from this form of benefit expansion.

The final row of Table 2 considers a reform involving a marginal benefit increase proportional to earnings, capped at 41% of average earnings. For individuals earning above the cap, the benefits increase is proportional to this limit. This reform mirrors a natural expansion of the current U.S. system, which replaces approximately 50% of pre-unemployment earnings up to a cap of 41% of average earnings (Schmieder and Von Wachter 2016). The reform results in a total welfare gain of 0.415, approximately midway between those of the flat and proportional expansions. Like the flat reform, $\bar{\lambda} > 1$, indicating that the redistributive effect of the reform is to raise social welfare. The bias from neglecting worker heterogeneity is -13.3%.

Comparing rows in Table 2 highlights that the welfare effects of a benefit expansion depend critically on the nature of the reform. Differences across reforms may arise from variation in how effectively they channel support to those experiencing the largest consumption declines, or in patterns of program cross-subsidisation across households, as reflected in the aggregate welfare weight, $\bar{\lambda}$. Quantitatively, we find the latter channel to be more influential. While the standard representative agent implementation would assign the same welfare gain to each reform, Table 2 shows we find economically and statistically significant differences in their value.

	(1)	(2)	(3)	(4)	(5)
	Total welfare gain	Unemployed welfare weight	Average insurance gain	Insurance-welfare covariance weight cov.	% bias from neglecting heterogeneity
	(1-(2))+(2)×((3)+(4))	$\bar{\lambda}$	$E^l\left[\gamma\frac{\Delta c}{c}\right]$	$\text{cov}^l\left(\lambda, \gamma\frac{\Delta c}{c^h}\right)$	100*((3)-(1))/(1)
Benefit rise					
Flat	0.519	1.102	0.370	0.008	-28.81
	[0.435, 0.588]	[1.092, 1.112]	[0.295, 0.430]	[0.001, 0.014]	[-32.89, -25.26]
Proportional	0.303	0.959	0.352	0.006	16.20
	[0.237, 0.360]	[0.948, 0.969]	[0.281, 0.411]	[0.001, 0.011]	[10.78, 23.46]
Proportional-capped	0.415	1.036	0.359	0.006	-13.34
	[0.340, 0.476]	[1.028, 1.044]	[0.288, 0.419]	[0.001, 0.012]	[-16.25, -10.30]

TABLE 2. Insurance and Redistributive Gains from UI Expansion

Notes: Authors' calculations from the PSID data. See Proposition 1 and Equation (7) for definition of terms. We set the coefficient of relative risk aversion, γ, equal to 3 and the welfare weight exponent, ψ, equal to 1. Numbers are based on the benefit rise being funded by a lump-sum tax increase. 95% confidence intervals are reported in brackets.

For all forms of program expansion, the covariance between worker-specific welfare weights (λ) and exposure to risk ($\gamma\frac{\Delta c}{c}$) is quantitatively small. Appendix D.3 decompose this covariance effect into the correlation in welfare weights and risk exposure, and the standard deviations in each variable (see Lieber and Lockwood 2019, who employ a similar decomposition in the case of in-kind medical benefits). The standard deviation of the welfare weights reflects inequality in consumption, while the standard deviation of $\gamma\frac{\Delta c}{c}$ reflects heterogeneity unemployed-related risk exposure. The correlation coefficient provides a standardised measure of the extent to which UI targets insurance toward high-need households. Across all three program expansions, we find the correlation coefficient is 0.13-0.17, with the standard deviation in welfare weights approximately 0.38 and the standard deviation in consumption declines around 0.13. Together, these components have a limited overall impact on the social value of UI expansion.

5.3. Moral Hazard Cost

As in the standard Baily-Chetty formula, the moral hazard cost of a flat social insurance expansion in equation (7) is captured by an aggregate elasticity, scaled by the aggregate

employment rate. The relevant elasticity is that of the aggregate unemployment rate with respect to a budget-neutral benefit change. Any worker-level heterogeneity in responses implies that estimates of the elasticity based on local variation (e.g., kinks in benefit schedules) may not coincide with the aggregate elasticity. Nevertheless, an extensive literature exploits both local and non-local variation to estimate this elasticity. The median implied estimate of the moral hazard cost of UI expansion in the U.S. is 0.38 (see Schmieder and Von Wachter 2016). Chetty (2008), who uses cross-state variation in benefit levels, and Card et al. (2015), employing a regression kink design with data from 2003 and 2007, both find moral hazard costs close to this level.[28]

We find that the welfare gain from a flat UI expansion, 0.519, exceeds this central estimate of the moral hazard cost of expansion, and therefore this form of expansion is welfare-enhancing. Note, accounting for worker heterogeneity is crucial to this finding; if we ignored it, measuring the welfare gain based on the average insurance value, we would instead conclude the gains and moral hazard cost of expansion balance.

If all workers have the same elasticity of effort with respect to a $ rise in benefits, then, even if households differ in their consumption levels, the moral hazard cost of a proportional-capped (or proportional) expansion, funded by a flat tax, is equal to that of a flat benefit expansion. In this case, the social gains from a proportional-capped extension slightly outweigh the moral hazard costs, whereas the costs of a proportional expansion exceed its gains.

However, when worker search elasticities are heterogeneous, the correlation between these elasticities and earnings becomes a key determinant of the moral hazard cost of proportional benefit expansions. The degree of this correlation and its potential to meaningfully alter moral hazard costs compared to existing estimates remains an open question. Evidence suggests that responses are stronger at the lower end of the income distribution, implying current estimates of moral hazard costs may serve as an upper bound.[29]

[28] Most studies estimate the elasticity of non-employment or benefit duration with respect to a benefit reform. Schmieder and Von Wachter (2016) convert these estimates into the behavioural cost per $ increase in benefits, which corresponds to the employment rate scaled elasticity in equation (7).

[29] Chetty (2008) provides evidence that the liquidity effects of benefit expansion on worker search are concentrated among liquidity constrained individuals. However, the policy relevant elasticity is the total response rather than the liquidity-induced one. See Schmieder and Von Wachter (2016) for a discussion of this issue. Similarly, Attanasio et al. (2018) find that extensive margin labour supply elasticities in response to permanent wage changes decrease with current wealth. These findings suggest that the aggregated elasticities needed for policy evaluation may be smaller when weighted by employment incomes, as income-weighting lowers the weight placed on the most responsive households. Thus, existing estimates may represent an upper bound. However, Birinci and See (2023) show, using a calibrated search model, that responses to UI generosity can be non-monotonic along these dimensions.

5.4. Alternative Reforms

Proportional tax. The numbers presented in Table 2 are based on a UI expansion funded by a lump-sum tax rise in the employed state. Some U.S. states levy payroll taxes significantly above the federally mandated minimum base (up to a maximum of $50,000 in Washington State). In such cases, it is more natural to consider a budget-balanced UI expansion funded by a proportional tax increase. Whether the tax rise is lump-sum or proportional has no impact on the average insurance or covariance terms (columns (3) and (4) in Table 2). However, a proportional-tax-rise-funded expansion increases the welfare weight the planner assigns to unemployment relative to employment ($\bar{\lambda} = \frac{\mathbb{E}^l[\omega u'(c^h)]}{\mathbb{E}^h[\omega u'(c^h)]}$). The difference arises because, with a lump-sum tax increase, the risk-weighting in the welfare weight denominator is based on employment probabilities $(1 - e_i)$, whereas with a proportional tax increase, it is based on $((1 - e_i)\, y_i)$. The latter reflects the fact that a larger share of the tax increase is incident on better-off individuals, leading to a UI expansion that more strongly redistributes resources to the most needy.

In Table 3 we report the decomposition of the insurance and redistributive gains from UI expansion when funded by a proportional tax. Changing from a flat- to proportional-tax-funded benefit increase, raises the total welfare gains for the flat, proportional and proportional-capped benefit reforms by 0.15. The difference in the marginal gains between the three types of benefit rise is largely unaffected. The difference with a lump-sum-tax-funded expansion is that the UI expansion has a substantially larger positive impact on social welfare (net of moral hazard costs) – leading standard approaches to substantially understate the welfare effects of a proportional-tax-funded program expansion.

	(1)	(2)	(3)	(4)	(5)	(6)
	Total welfare gain $[(1\text{-}(2))+(2)\times((3)+(4))]\times(5)$	Unemployed welfare weight $\bar{\lambda}$	Average insurance gain $E^l\left[\gamma\frac{\Delta c}{c}\right]$	Insurance-welfare covariance weight cov. $\text{cov}^l\left(\hat{\lambda}, \gamma\frac{\Delta c}{c^h}\right)$	Scaling factor	% bias from neglecting heterogeneity $100*((3)\times(5)\text{-}(1))/(1)$
Benefit rise						
Flat	0.664	1.289	0.370	0.008	0.855	-52.39
	[0.584, 0.735]	[1.262, 1.315]	[0.295, 0.430]	[0.001, 0.014]	[0.843, 0.868]	[-57.00, -48.68]
Proportional	0.448	1.122	0.352	0.006	0.855	-32.76
	[0.380, 0.507]	[1.105, 1.138]	[0.281, 0.411]	[0.001, 0.011]	[0.843, 0.868]	[-37.30, -28.70]
Proportional-capped	0.560	1.211	0.359	0.006	0.855	-45.06
	[0.484, 0.624]	[1.188, 1.233]	[0.288, 0.419]	[0.001, 0.012]	[0.843, 0.868]	[-49.85, -41.16]

TABLE 3. Insurance and Redistributive Gains from UI Expansion: Proportional Tax Rise

Notes: Authors' calculations from the PSID data. See Proposition 1 and Equation (7) for definition of terms. We set the coefficient of relative risk aversion, γ, equal to 3 and the welfare weight exponent, ψ, equal to 1. Numbers are based on the benefit rise being funded by a proportional tax increase. Note, we scale effects so they are in the same units as the results in Table 2; see discussion in Appendix A.2.1. 95% confidence intervals are reported in brackets.

Arbitrary reforms. To illustrate how the value of UI expansion varies across a broader set of reforms, we specify the marginal benefit and tax increase as simple non-linear functions of in-work earnings:

$$\rho^h(y_i) = (y_i)^{(1-\nu^h)};\ \nu^h \in (-1, 1) \qquad \rho^l(y_i) = (y_i)^{(1-\nu^l)};\ \nu^l \in (0, 1) \tag{10}$$

$\nu^s = 0$ corresponds to a proportional reform, $\nu^v = 1$ corresponds to a flat reform, and $\nu^s \in (0, 1)$ captures intermediate cases where the rise in the tax or benefit increases less than proportionately with income. We also consider cases where $\nu^h \in (-1, 0)$, to capture progressive tax reforms, in which higher-income people pay proportionately more. These functions are similar in spirit to the parametric forms used to model non-linear tax systems in Feldstein (1969), Benabou (2002) and Heathcote et al. (2017). However, in our case the functions parameterise a marginal reform rather than entire tax and benefit schedules.[30]

Using the schedules in (10), we compute the marginal gain from program expansion for each combination of (ν^h, ν^l), and how the moral hazard cost varies with these parameters. To compute how moral hazard costs vary, we make the following additional

[30]This allows us to consider different forms of reform to a given system. If the reform is a proportional adjustment to the existing schedule, these functions then also correspond to the full schedules.

assumptions. First, we specify the full tax and benefit schedule, assuming a flat tax and proportional-capped benefit with a 41% annual earnings cap, approximating the current U.S. system. Second, we assume homogeneous worker-level search elasticities and set the moral hazard cost of a flat benefit expansion funded by a flat tax increase at 0.38 (the median estimate reported in Schmieder and Von Wachter (2016)). With these assumptions, we can rescale the moral hazard cost to reflect the reform under consideration (see Appendix A.2.2).

In Figure 5, we present a heat map illustrating how the *net* social welfare impact (gain minus moral hazard cost) of UI expansion varies with the type of reform, captured by the parameters (ν^h, ν^l). The dashed lines are iso-welfare curves, indicating parameter combinations that yield the same net social welfare gain. All parameter combinations below the black line correspond to reforms that increase social welfare. The red markers denote the six reforms summarised in Tables 2 and 3.[31] The figure underscores that the social welfare impact of UI reform depends crucially on the nature of the reform. A key strength of our framework is its ability to readily accommodate alternative types of reform using the same datasets commonly employed for the standard analysis.

[31] For the proportional-capped reform, the markers represent parameter values that yield the same marginal gains shown in Tables 2 and 3. While equations (10) do not perfectly nest the proportional-capped system, they provide a close approximation.

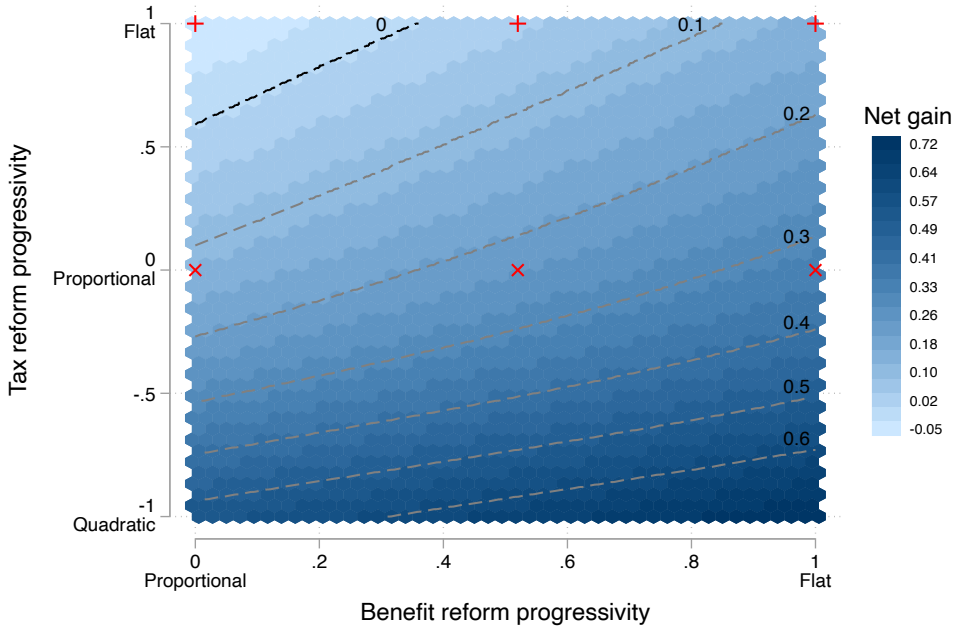


FIGURE 5. Marginal Net Value of Program Expansion by Reform Type

Notes: Authors' calculation from the PSID data. Plot shows how the marginal social value (gain net of moral hazard cost) of expansion varies the progressivity of the benefit rise and tax rise. We set the coefficient of relative risk aversion, γ, equal to 3 and the welfare weight exponent, ψ, equal to 1.

6. Sensitivity to Calibrated Parameters

The results presented in the preceding section are based on fixed values of the coefficient of relative risk aversion ($\gamma = 3$) and the planner's inequality aversion parameter ($\psi = 1$). In this section, we explore the quantitative effects of varying these parameters.[32] We focus on the value of a proportional-capped benefit expansion, funded by a lump-sum tax increase (corresponding to row (3) of Table 2), and how this compares to alternative benefit reforms.

Value of proportional-capped reform. In Figure 6, we present a heat map illustrating how the marginal gain from UI expansion varies with γ and ψ. The dashed lines represent iso-welfare curves, showing parameter combinations that yield the same social welfare gain. The black line corresponds to a welfare gain of 0.38, which equals

[32] Our results extend straightforwardly to the case where the marginal utility from consumption expenditures is state dependent, either due state-specific preferences or prices. In Appendix D.4, we show how our results change if we allow for state-dependence, through a marginal utility shifter $\phi \neq 1$.

the median moral hazard cost estimate. For each value of the coefficient of risk aversion, this line indicates the inequality aversion parameter that rationalises the current U.S. system as optimal.[33] For all parameter combinations to the right of the black line the marginal gain in social welfare exceeds this level of moral hazard cost. The red marker denotes our baseline calibration, which yield a marginal gain of 0.415. If $\gamma = 1$, the marginal value of expansion becomes 39% of this baseline value; if $\gamma = 6$ it rises to 192%. This reflects a well known channel: higher risk aversion increases exposure to risk, thereby raising demand for insurance. A novel feature of our setting is that the marginal value of program expansion also depends on redistributive preferences. For instance, suppose $\gamma = 3$ and $\psi = 0$, which corresponds to the Andrews and Miller (2013) Pareto weights, implying the planner is indifferent to the redistributive consequences of UI expansion. In this case, the marginal value of program expansion is 87% its baseline value ($\psi = 1$), equivalent to the average insurance gain (shown in column (3) and final row of Table 2). Conversely, if $\psi = 3$, the marginal value increases to 124% of its baseline value. This rise is primarily driven by stronger inequality-averse social preferences acting to raise the aggregate welfare weight assigned to the unemployed, thereby amplifying the benefits of channeling resources to this group. Figure 6 underscores the quantitative significance of incorporating redistributive implications when evaluating UI expansion.[34]

Comparison with alternative reform. Table 2 highlights economically and statistically significant differences in the marginal welfare gains of flat, proportional and proportional-capped benefit expansions, driven primarily by differences in the aggregate welfare weight, $\bar{\lambda}$. This reflects that proportional-capped, and to a greater degree proportional, expansions provide relatively larger cash benefits to better-off households compared to a flat expansion.

[33] More specifically, it rationalises the existing level of flat tax and benefit, provided in proportion to $\min\{y, \kappa\}$, where κ is the earnings cap, as the optimal conditional on this tax and benefit structure.

[34] We report these results for combinations $\psi \leq \gamma$ which is consistent with consumption commitments delivering increased *local* risk-aversion. Absent consumption commitments, $\psi = \gamma$ corresponds to a standard utilitarian social welfare function evaluated behind the veil of ignorance.

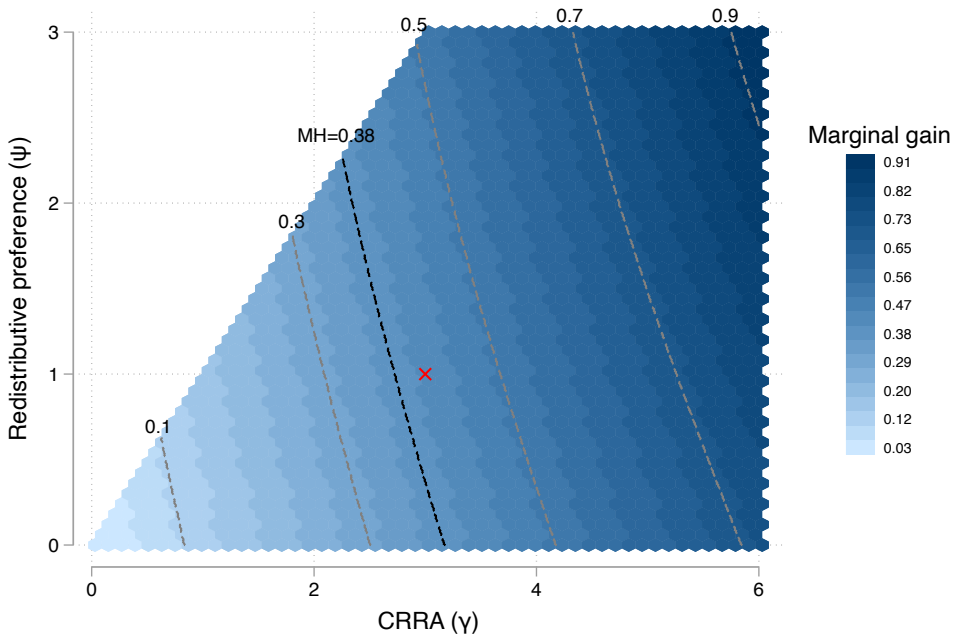


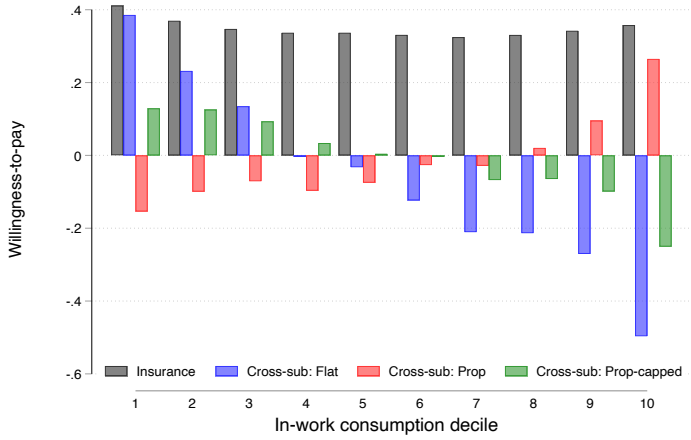
FIGURE 6. Marginal Value of Program Expansion for Different Values of Risk and Inequality Aversion

Notes: Authors' calculation from the PSID data. Plot shows marginal gain of expansion of the proportional-capped system, funded by a lump-sum tax, for different values of the coefficient of relative risk aversion, γ, and the inequality aversion parameter, ψ. We report values for parameter values $\gamma \in [\psi, 6]$ and $\psi \in [0, 3]$ and use a red cross to indicate our baseline calibration.

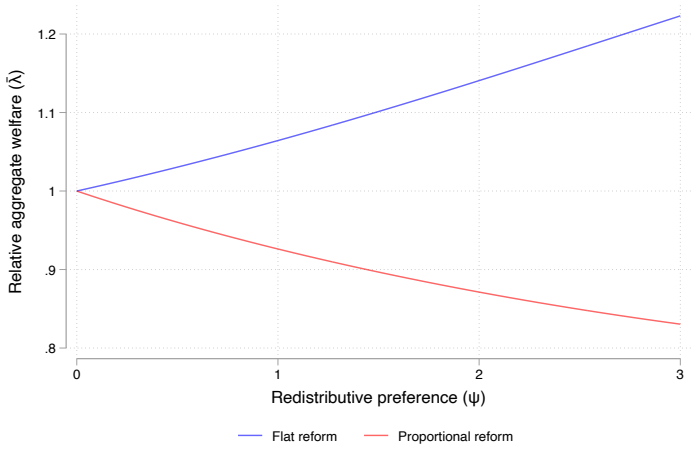
Figure 7(A) illustrates this by showing how the insurance and cross-subsidisation components of households' willingness-to-pay for UI expansion (as defined in equation (2)), vary across deciles on the equivalised in-work consumption distribution. While the insurance component is relatively flat across the consumption distribution, the cross-subsidisation component exhibits a decreasing (in the case of a flat reform) or increasing (in the case of proportional and proportional-capped reforms) relationship. These differential patterns of cross-subsidisation across the consumption distribution drive the variation in $\bar{\lambda}$ and are the primary mechanism underlying the differences in welfare gains across reforms.

The impact of this heterogeneity on the social value of the reforms depends on the planner's degree of inequality aversion (ψ). Figure 7(B) illustrates this by showing how the aggregate welfare weights for flat and proportional expansions, relative to the proportional-capped expansion, vary with ψ. When $\psi = 0$, the planner is indifferent to redistributing resources across households and, consequently, to patterns of UI cross-subsidisation. However, when $\psi > 0$, indicating social preferences entail at least some

degree of aversion to inequality, differences in resource allocation across households become a key determinant of the social value of the reforms.



(A) Willingness-to-pay, by Employed Consumption



(B) Social Value of Cross-Subsidisation, Relative to Proportional-Capped Reform

FIGURE 7. Social and Private Values of Cross-Subsidisation

Notes: Authors' calculation from the PSID data. Panel (A) shows how insurance and cross-subsidisation components of households' willingness-to-pay for UI vary by deciles of equivalised panel average non-durable consumption in the employed state (at $\psi = 1$). Flat, Prop and Prop capped refer to flat, proportional and proportional-capped benefit expansions respectively. See Appendix D.4 for a similar graph with earnings on the horizontal axis. Panel (B) shows how the aggregate welfare weight $\bar{\lambda}$ of each reform, relative to that of a proportional-capped reform, varies the inequality aversion parameter, ψ. We set the coefficient of relative risk aversion, γ, equal to 3. Numbers are based on the benefit rise being funded by a lump-sum tax increase.

7. Conclusion

In this paper, we develop and empirically implement a tractable framework to measure the impact of changes in the generosity of a social insurance program on social welfare. Our approach accounts for heterogeneity in individual-level insurance values and program cross-subsidisation, allowing an inequality-averse social planner to assign different welfare weights based on individuals' levels of need. In our application to U.S. unemployment insurance, we demonstrate that accounting for the redistributive effects of program expansion is essential for accurately assessing its value. For example, under inequality-averse social preferences, the marginal value of flat benefit expansions exceeds their value in the standard representative agent framework because to the positive social value attributed to the transfer of resources to relatively needy households. Our results also highlight that the patterns of cross-subsidisation from program expansion – and, consequently, their social value – depends critically the structure of the reform.

We focus on how incorporating individual-level heterogeneity within a social welfare framework affects the evaluation of benefit levels. Natural extensions to our model include quantifying the social gains from varying benefit levels over the economic cycle (Kroft and Notowidigdo 2016), benefit duration (Schmieder et al. 2012), and their time profile (Kolsrud et al. 2018), while accounting for program cross-subsidisation and the social value of redistribution. In addition, our framework can readily accommodate heterogeneity in take-up (Anderson and Meyer 1997) and the effective progressivity this induces (Lachowska et al. 2022), and could be extended to accommodate heterogeneity in labour market-tightness externalities (Landais et al. 2018).

Our model shows that, with ex-ante heterogeneous agents, the incentive costs of a lump-sum expansion of social insurance continue to be captured by an aggregate labour supply elasticity. However, when benefit levels are income-dependent and individuals' behavioural responses are heterogeneous, moral hazard costs become more complex than the standard duration elasticities estimated in the existing literature. An important avenue for future research is to determine the extent to which accounting for this mitigates or amplifies the incentive costs arising from moral hazard.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica 67*(2), 251–333.

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

Aguiar, M., E. Hurst, and L. Karabarbounis (2013). Time use during the great recession. *American Economic Review 103*(5), 1664–1696.

Akerlof, G. A. (1978). The economics of" tagging" as applied to the optimal income tax, welfare programs, and manpower planning. *American Economic Review 68*(1), 8–19.

Allcott, H., B. B. Lockwood, and D. Taubinsky (2019). Regressive Sin Taxes, with an Application to the Optimal Soda Tax. *Quarterly Journal of Economics 134*(3), 1557–1626.

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.

Andrews, I. and C. Miller (2013). Optimal Social Insurance with Heterogeneity. *mimeo*, 26.

Arapakis, K., E. French, J. B. Jones, and J. McCauley (2023). Medical spending risk among retired households by race. *Ann Arbor, MI. University of Michigan Retirement and Disability Research Center (MRDRC) Working Paper*.

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.

Benabou, R. (2002). Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency? *Econometrica 70*(2), 481–517.

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.

Birinci, S. and K. See (2023). Labor market responses to unemployment insurance: The role of heterogeneity. *American Economic Journal: Macroeconomics 15*(3), 388–430.

Blomqvist, Å. and H. Horn (1984, August). 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.

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

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.

Bonhomme, S., T. Lamadon, and E. Manresa (2019). A distributional framework for matched employer employee data. *Econometrica 87*(3), 699–739.

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.

Borusyak, K., X. Jaravel, and J. Spiess (2024). Revisiting event-study designs: robust and efficient estimation. *Review of Economic Studies*, rdae007.

Braun, R. A., K. A. Kopecky, and T. Koreshkova (2017). Old, sick, alone, and poor: A welfare analysis of old-age social insurance programmes. *Review of Economic Studies 84*(2), 580–612.

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.

Card, D., A. Johnston, P. Leung, A. Mas, and Z. Pei (2015). The effect of unemployment benefits on the duration of unemployment insurance receipt: New evidence from a regression kink design in missouri, 2003–2013. *American Economic Review 105*(5), 126–130.

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

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. Elsevier.

Chetty, R. and A. Szeidl (2007). Consumption commitments and risk preferences. *Quarterly Journal of Economics 122*(2).

Choi, S., A. Janiak, and B. Villena-Roldán (2015). Unemployment, participation and worker flows over the life-cycle. *Economic Journal 125*(589), 1705–1733.

Cohen, J. P. and P. Ganong (2024). Disemployment effects of unemployment insurance: A meta-analysis. Technical report, National Bureau of Economic Research.

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.

Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: series B (methodological) 39*(1),

1–22.

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.

Feldstein, M. S. (1969). The effects of taxation on risk taking. *Journal of Political Economy 77*(5), 755–764.

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.

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.

Havránek, T. (2015). Measuring intertemporal substitution: The importance of method choices and selective reporting. *Journal of the European Economic Association 13*(6), 1180–1204.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal tax progressivity: An analytical framework. *Quarterly Journal of Economics 132*(4), 1693–1754.

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.

Hoynes, H. W. and E. F. Luttmer (2011). The insurance value of state tax-and-transfer programs. *Journal of Public Economics 95*(11-12), 1466–1484.

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

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Keane, M. P. and K. I. Wolpin (1997). The career decisions of young men. *Journal of Political Economy 105*(3), 473–522.

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.

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.

Krueger, A. B. and B. D. Meyer (2002). Labor supply effects of social insurance. *Handbook of Public Economics 4*, 2327–2392.

Lachowska, M., I. Sorkin, and S. A. Woodbury (2022). Firms and unemployment insurance take-up. Technical report, National Bureau of Economic Research.

Landais, C., P. Michaillat, and E. Saez (2018). A macroeconomic approach to optimal unemployment insurance: Theory. *American Economic Journal: Economic Policy 10*(2), 152–181.

Landais, C. and J. Spinnewijn (2021). The value of unemployment insurance. *The 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.

Lewis, D., D. Melcangi, and L. Pilossoph (2024). Latent heterogeneity in the marginal propensity to consume. Technical report, National Bureau of Economic Research.

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

Lockwood, L. M. (2024). Health insurance and consumption risk. Technical report, Working Paper.

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

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.

Piketty, T. and E. Saez (2013). Optimal labor income taxation. In *Handbook of Public Economics*, Volume 5, pp. 391–474. Elsevier.

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.

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

Rabin, M. (2000). Risk Aversion and Expected-Utility Theory: A Calibration Theorem. *Econometrica 68*(5), 1281–1292.

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

Saez, E. (2002). Optimal income transfer programs: intensive versus extensive labor supply responses. *Quarterly Journal of Economics 117*(3), 1039–1073.

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.

Online Appendix

Appendix A. Technical Appendix

A.1. Section 2.2 Derivations

The planner's problem is

$$\max_b W(b) = \int_i \omega_i V_i(b, \tau(b)) di$$

where

$$\begin{aligned} \tau(b) &= \frac{1-\bar{e}}{\bar{e}} b \\ V_i(b,\tau) &= \max_e e u_i^h(c_i^h) + (1-e) u_i^l(c_i^l) - \psi_i(e) \\ c_i^h &= y_i - \tau(b) \\ c_i^l &= z_i + b \end{aligned}$$

The first-order condition is:

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

Note:

$$\frac{d\tau}{db} = \frac{(1-\bar{e})}{\bar{e}} \left(1 + \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}\right)$$

where $\epsilon_{1-\bar{e},b} \equiv \frac{d(1-\bar{e})}{db} \frac{b}{(1-\bar{e})}$. Therefore:

$$\frac{dW}{db} = (1-\bar{e}) \left[\int_i \frac{(1-e_i)}{1-\bar{e}} \omega_i u_i^{l\prime}(c_i^l) di - \int_i \frac{e_i}{\bar{e}} \omega_i u_i^{h\prime}(c_i^h) di \left(1 + \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}\right) \right]$$

Define:

$$\mathbb{E}^h[\omega u_i^{h\prime}(c^h)] = \int_i \frac{e_i}{\bar{e}} \omega_i u_i^{h\prime}(c_i^h) di$$

$$\mathbb{E}^l[\omega u_i^{l'}(c^l)] = \int_i \frac{(1-e_i)}{1-\bar{e}} \omega_i u_i^{l'}(c_i^l) di.$$

Consider the welfare effect of a marginal rise in all high-state incomes (or equivalently a marginal tax reduction):

$$\Lambda = \int_i e_i \omega_i u_i^{h'}(c_i^h) di = \bar{e}\mathbb{E}^h[\omega U_i^{h'}(c^h)]$$

We use this to rescale $\frac{dW}{db}$ according to:

$$\frac{d\tilde{W}}{db} = \frac{dW/db}{1-\bar{e}} \Big/ \frac{\Lambda}{\bar{e}}.$$

$\frac{d\tilde{W}}{db}$ has the interpretation of the impact on social welfare of a balanced-budget \$ expansion in social insurance relative to the effect on social welfare of a \$ flat increase in all high state incomes.

Hence, the marginal value of social insurance expansion (equation (3)) is:

$$\frac{d\tilde{W}}{db} = \frac{\mathbb{E}^l[\omega u^{l'}(c^l)] - \mathbb{E}^h[\omega u^{h'}(c^h)]}{\mathbb{E}^h[\omega u^{h'}(c^h)]} - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}$$

Household *i*'s willingness-to-pay for a marginal social insurance expansion is defined in equation (2) and repeated here:

$$\theta_i = \left(\frac{u_i^{l'}(c_i^l) - u_i^{h'}(c_i^h)}{u_i^{h'}(c_i^h)} \right) + \left(1 - \frac{e_i/\bar{e}}{1-e_i/1-\bar{e}} \right).$$

Note that:

$$\begin{aligned}
\mathbb{E}^l[\omega u^{l'}(c^l)] - \mathbb{E}^h[\omega u^{h'}(c^h)] =& \mathbb{E}^l\left[\omega\left(u^{l'}(c^l) - \frac{e_i/\bar{e}}{1-e_i/1-\bar{e}} u^{h'}(c^h)\right)\right] \\
=& \mathbb{E}^l\left[\omega u^{h'}(c^h)\left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} + \left(1 - \frac{e_i/\bar{e}}{1-e_i/1-\bar{e}}\right)\right)\right] \\
=& \mathbb{E}^l[\omega u^{h'}(c^h)]\mathbb{E}^l\left[\lambda\left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} + \left(1 - \frac{e_i/\bar{e}}{1-e_i/1-\bar{e}}\right)\right)\right],
\end{aligned}$$

where $\lambda_i \equiv \frac{\omega_i u_i^{h'}(c^h)}{\mathbb{E}^l[\omega u^{h'}(c^h)]}$. Hence we obtain equation (4),

$$\frac{d\tilde{W}}{db} = \bar{\lambda}\mathbb{E}^l\left[\lambda\left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} + \left(1 - \frac{e_i/\bar{e}}{1-e_i/1-\bar{e}}\right)\right)\right] - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}},$$

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

$$\mathbb{E}^l\left[\lambda\left(1 - \frac{e_i/\bar{e}}{1-e_i/1-\bar{e}}\right)\right] = 1 - \mathbb{E}^h[\lambda] = 1 - \frac{1}{\bar{\lambda}}$$

Combining with equation (4) gives equation (5):

$$\frac{d\tilde{W}}{db} = \bar{\lambda}\mathbb{E}^l\left[\lambda\left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)}\right)\right] + (\bar{\lambda} - 1) - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}.$$

Assuming the utility function is state independent and taking a quadratic approximation to the utility function: $\frac{u_i'(c_i^l) - u_i'(c_i^h)}{u_i'(c_i^h)} \approx \gamma_i \frac{\Delta c_i}{c_i^h}$ where $\Delta c_i \equiv c_i^h - c_i^l$ and $\gamma_i \equiv -\frac{u_i''(c_i^h)c_i^h}{u_i'(c_i^h)}$, we can re-write equation (5) as in equation (7)

$$\frac{d\tilde{W}}{db} \approx \bar{\lambda}\left(\mathbb{E}^l\left[\gamma\frac{\Delta c}{c^h}\right] + \mathrm{cov}^l\left(\lambda, \gamma\frac{\Delta c}{c^h}\right)\right) + (\bar{\lambda} - 1) - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}.$$

Alternative, taking the cubic approximation $\frac{u_i'(c_i^l) - u_i'(c_i^h)}{u_i'(c_i^h)} \approx \gamma_i \frac{\Delta c_i}{c_i^h} + \frac{1}{2}\gamma_i\zeta_i\left(\frac{\Delta c_i}{c_i^h}\right)^2$, where $\zeta_i \equiv -\frac{u_i'''(c_i^h)c_i^h}{u_i''(c_i^h)}$, we obtain:

$$\frac{d\tilde{W}}{db} \approx \bar{\lambda}\left(\mathbb{E}^l\left[\gamma\frac{\Delta c}{c^h}\right] + \mathrm{cov}^l\left(\lambda, \gamma\frac{\Delta c}{c^h}\right) + \frac{1}{2}\mathbb{E}^l\left[\gamma\zeta\left(\frac{\Delta c}{c^h}\right)^2\right] + \frac{1}{2}\mathrm{cov}^l\left(\lambda, \gamma\zeta\left(\frac{\Delta c}{c^h}\right)^2\right)\right) + (\bar{\lambda} - 1) - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}.$$

A.2. Section 2.3 Derivations

A.2.1. Arbitrary Tax and Benefit System

Suppose the tax and benefit schedules are both arbitrary functions of high-state income, so state-specific consumption is:

$$c_i^h = y_i - f^h(y_i)$$

$$c_i^l = z_i + f^l(y_i)$$

and the planner's budget-balance constraint is:

$$\int_i e_i f^h(y_i)di = \int_i (1-e_i) f^l(y_i)di.$$

Consider an arbitrary marginal budget-balanced reform that leads to the following changes in state-specific consumption:

$$dc_i^h = -d\tau_i = -\rho^h(y_i)d\mathbb{T}$$
$$dc_i^l = db_i = \rho^l(y_i)d\mathbb{b}.$$

$d\tau_i$ is the tax change for agent i, which is proportional to a function of their income given by $\rho^h(y_i)$. Similarly, db_i is the agent i benefit change, proportional to $\rho^l(y_i)$. As the reform is budget-balanced, we have

$$\left(\int_i e_i \rho^h(y_i)di\right) d\mathbb{T} = \left(\int_i (1-e_i)\rho^l(y_i)di\right) d\mathbb{b} + \left(\int_i d(1-e_i)(f^h(y_i) + f^l(y_i))di\right),$$

or

$$\frac{d\mathbb{T}}{d\mathbb{b}} = \frac{\int_i (1-e_i)\rho^l(y_i)di}{\int_i e_i \rho^h(y_i)di}\left(1 + \frac{1}{\int_i (1-e_i)\rho^l(y_i)di}\int_i \frac{d(1-e_i)}{d\mathbb{b}}(f^h(y_i) + f^l(y_i))di\right).$$

Denote:

$$\tilde{y}^h = \int_i e_i \rho^h(y_i)di$$
$$\tilde{y}^l = \int_i (1-e_i)\rho^l(y_i)di$$
$$\tilde{y}_i^l = \int_i \frac{(1-e_i)\rho^l(y_i)}{\int_{i'}(1-e_{i'})\rho^l(y_{i'})di'}di$$

Define the individual-level elasticity of the low-state probability with respect to a marginal (budget-balanced) benefit rise:

$$\epsilon_{1-e_i, b_i} \equiv \frac{d(1-e_i)}{db_i}\frac{f^l(y_i)}{1-e_i}.$$

Finally, noting that $db_i = \rho^l(y_i)d\mathbb{b}$, we can define the following moral hazard term, which captures the impact of the reform on the planner's budget that arises due to behaviour change;

$$\begin{aligned} MH =& \frac{1}{\int_i (1-e_i)\rho^l(y_i)di} \int_i \frac{d(1-e_i)}{d\mathbb{b}} (f^h(y_i) + f^l(y_i))di \\ =& \frac{1}{\int_i (1-e_i)\rho^l(y_i)di} \int_i (1-e_i)\rho^l(y_i) \frac{d(1-e_i)}{db_i} \frac{f^l(y_i)}{1-e_i} \frac{f^h(y_i)+f^l(y_i)}{f^l(y_i)} di \\ =& \int_i \tilde{y}_i^l \epsilon_{1-e_i,b_i} \frac{f^h(y_i)+f^l(y_i)}{f^l(y_i)} di \end{aligned} \tag{A.1}$$

This allows us to write:

$$\frac{d\mathbb{T}}{d\mathbb{b}} = \frac{\tilde{y}^l}{\tilde{y}^h}(1+MH).$$

The planner's first-order condition is:

$$\frac{dW}{db} = \tilde{y}^l \left[\int_i \frac{(1-e_i)\rho^l(y_i)}{\tilde{y}^l} \omega_i u_i^{l\prime}(c_i^l) di - \int_i \frac{e_i \rho^h(y_i)}{\tilde{y}^h} \omega_i u_i^{h\prime}(c_i^h) di \, (1+MH) \right]$$

Define the state-specific weighted average social marginal social utilities:

$$\widetilde{\mathbb{E}}^h[\omega u^{h\prime}(c^h)] = \int_i \frac{e_i \rho^h(y_i)}{\tilde{y}^h} \omega_i u_i^{h\prime}(c_i^h) di$$

$$\widetilde{\mathbb{E}}^l[\omega u^{l\prime}(c^l)] = \int_i \frac{(1-e_i)\rho^l(y_i)}{\tilde{y}^l} \omega_i u_i^{l\prime}(c_i^l) di.$$

Note, the individual-level weights in the high-state expectation are the expected marginal tax liability, $e_i\rho^h(y_i)$ and the weights in the low-state expectation are the expected marginal benefit rise, $(1 - e_i)\rho^l(y_i)$ (where in each case the weights are normalised to sum to 1).

Consider the impact on welfare of a marginal tax reduction:

$$\widetilde{\Lambda} = \int_i e_i \rho^h(y_i) \omega_i u_i^{h\prime}(c_i^h) di = \tilde{y}^h \widetilde{\mathbb{E}}^h[\omega U^{h\prime}(c^h)]$$

We use this to rescale $\frac{dW}{db}$ according to:

$$\frac{d\tilde{W}}{db} = \frac{dW/db}{\tilde{y}^l} \Big/ \frac{\widetilde{\Lambda}}{\tilde{y}^h}.$$

Hence the marginal value of social insurance expansion is:

$$\frac{d\tilde{W}}{db} = \frac{\widetilde{\mathbb{E}}^l[\omega u^{l'}(c^l)] - \widetilde{\mathbb{E}}^h[\omega u^{h'}(c^h)]}{\widetilde{\mathbb{E}}^h[\omega u^{h'}(c^h)]} - MH \tag{A.2}$$

Following the steps outlined in Appendix A.1, we have:

$$\frac{d\tilde{W}}{db} = \bar{\mathbb{\lambda}} \left(\widetilde{\mathbb{E}}^l \left[\gamma \frac{\Delta c}{c^h} \right] + \widetilde{\text{cov}}^l \left(\mathbb{\lambda}, \gamma \frac{\Delta c}{c^h} \right) \right) + (\bar{\mathbb{\lambda}} - 1) - MH \tag{A.3}$$

where in case of arbitrary reform, $\bar{\mathbb{\lambda}} \equiv \frac{\widetilde{\mathbb{E}}^l[\omega u^{h'}(c^h)]}{\widetilde{\mathbb{E}}^h[\omega u^{h'}(c^h)]}$ and $\mathbb{\lambda}_i \equiv \frac{\omega u^{h'}(c^h)}{\widetilde{\mathbb{E}}^l[\omega u^{h'}(c^h)]}$.

Comparing across different tax reforms. Note in equations (A.2) and (A.3) we scale the effect of a social insurance reform under consideration by the effect of an increase in high-state incomes that are proportional to the tax change that funds the reform (e.g., a flat income rise in the case of a flat tax rise, or proportional income rises in the case of a proportional tax rise). To quantitatively compare the impact of reforms that are financed by different types of tax reform, we need to express welfare effects in the same units. In all our results we do this by multiplying $\frac{d\tilde{W}}{db}$ by $\frac{\widetilde{\mathbb{E}}^h[\omega u^{h'}(c^h)]}{\mathbb{E}^h[\omega u^{h'}(c^h)]}$, which means all our quantitative results have the interpretation of the effect on social welfare of the specific reform under consideration, relative to the effect on social welfare of a flat rise in all incomes.

A.2.2. Moral Hazard Cost

Equation (A.1) provides a general formula for the moral hazard costs of an arbitrary budget-balanced marginal reform to an arbitrary tax and benefit schedule. Here we consider some special cases, and then outline assumptions under which we can rescale a single moral hazard estimate to capture differences across different reforms.

Special cases

Below we write down the moral hazard costs in four special cases, corresponding to flat (proportional) changes to flat (proportional) tax and benefit system. Note, in each case $\rho^s(y_i) = f^s(y_i)$ for $s = \{l, h\}$.

1: Flat system. Consider a system that entails a flat tax and benefit, so $f^h(y_i) = \tau^f$ and $f^l(y_i) = b^f$. And consider a flat rise in the benefit, so $db_i = db^f$. In this case:

$$\begin{aligned}
\tilde{y}^h &= \int_i e_i di = \bar{e} \\
\tilde{y}^l &= \int_i (1 - e_i) di = 1 - \bar{e} \\
\epsilon_{1-e_i,b_i} &= \frac{d(1-e_i)}{db^f} \frac{b^f}{1-e_i} \\
MH^{ff} &= \left(\int_i \left(\frac{1-e_i}{1-\bar{e}} \right) \epsilon_{1-e_i,b} di \right) \frac{\tau^f + b^f}{b^f} = \frac{1}{\bar{e}} \mathbb{E}^l [\epsilon_{1-e,b}]
\end{aligned}$$

Combining with equation (A.2) gives equation (3).

2: Flat tax and proportional benefit. Consider a system that entails a flat tax and proportional benefit, so $f^h(y_i) = \tau^f$ and $f^l(y_i) = b^p y_i$.[35] And consider a rise in the benefit replacement rate, so $db_i = db^p y_i$. In this case:

$$\begin{aligned}
\tilde{y}^h &= \int_i e_i di = \bar{e} \\
\tilde{y}^l &= \int_i (1 - e_i) y_i di \\
\epsilon_{1-e_i,b_i} &= \frac{d(1-e_i)}{d(b^p y_i)} \frac{b^p y_i}{1-e_i} = \frac{d(1-e_i)}{db^p} \frac{b^p}{1-e_i} \\
MH^{fp} &= \left(\int_i \left(\frac{(1-e_i) y_i}{\int_{i'} (1-e_{i'}) y_{i'} di'} \right) \epsilon_{1-e_i,b} \frac{\tau^f + b^p y_i}{b^p y_i} di \right) \\
&= \frac{1}{\bar{e}} \tilde{\mathbb{E}}^l \left[\epsilon_{1-e,b} \left(\frac{\tilde{y}^l + \bar{e} y}{y} \right) \right]
\end{aligned}$$

3: Proportional system. Consider a system that entails a proportional tax and benefit, so $f^h(y_i) = \tau^p y_i$ and $f^l(y_i) = b^p y_i$. And consider a rise in the benefit replacement rate,

[35] We write the benefit as proportional to high state income. For a proportional-capped system, y_i is replaced with $y_i^\kappa = \min\{y_i, \kappa\}$.

$db_i = db^p\, y_i$. In this case:

$$\begin{aligned}
\tilde{y}^h &= \int_i e_i\, y_i di \\
\tilde{y}^l &= \int_i (1-e_i)\, y_i di \\
\epsilon_{1-e_i,b_i} &= \frac{d(1-e_i)}{d(b^p\, y_i)} \frac{b^p\, y_i}{1-e_i} = \frac{d(1-e_i)}{db^p} \frac{b^p}{1-e_i} \\
MH^{pp} &= \left(\int_i \frac{(1-e_i)\, y_i}{\int_{i'} (1-e_{i'})\, y_{i'} di} \epsilon_{1-e_i,b} di \right) \frac{\tau^p + b^p}{b^p} = \frac{1}{\widetilde{e}} \tilde{\mathbb{E}}^l [\epsilon_{1-e,b}]
\end{aligned}$$

where $\widetilde{e} = \int_i \frac{y_i}{\int_i' y_i' di'} e_i di$.

4: Proportional tax and flat benefit. Consider a system that entails a proportional tax and flat benefit, so $f^h(y_i) = \tau^p\, y_i$ and $f^l(y_i) = b^f$. And consider a flat rise in the benefit, so $db_i = db^f$. In this case:

$$\begin{aligned}
\tilde{y}^h &= \int_i e_i\, y_i di \\
\tilde{y}^l &= \int_i (1-e_i) di = 1 - \bar{e} \\
\epsilon_{1-e_i,b_i} &= \frac{d(1-e_i)}{db^f} \frac{b^f}{1-e_i} \\
MH^{pf} &= \left(\int_i \left(\frac{1-e_i}{1-\bar{e}} \right) \epsilon_{1-e_i,b} \frac{\tau^p\, y_i + b^f}{b^f} di \right) \\
&= \mathbb{E}^l \left[\epsilon_{1-e,b} \left(\frac{(1-\bar{e})\, y + \tilde{y}^h}{\tilde{y}^h} \right) \right]
\end{aligned}$$

Note, if $\epsilon_{1-e_i,b} = \epsilon_{1-e,b}$ for all i, meaning there is no heterogeneity in the elasticity of he low-state probability with respect to the benefit reform, then:

$$\begin{aligned}
MH^{ff} &= MH^{fp} = \frac{\epsilon_{1-e,b}}{\bar{e}} \\
MH^{pf} &= MH^{pp} = \frac{\epsilon_{1-e,b}}{\widetilde{e}}
\end{aligned}$$

Conditions under which a single moral hazard estimate is sufficient for arbitrary reforms

Here we re-write equation (A.1) in the case when the following two conditions apply:

(a) The current tax and benefit schedules take the form

$$f^h(y_i) = \tau$$
$$f^l(y_i) = y_i^{\kappa} b,$$

where $y_i^{\kappa} = \min\{y_i, \kappa\}$. Note, budget-balance requires:

$$\left(\int_i e_i di\right)\tau = \left(\int_i (1-e_i)\, y_i^{\kappa} di\right) b$$

(b) The elasticity of the low-state probability with respect to a marginal benefit rise is common across agents – i.e., $\frac{d(1-e_i)}{db_i}\frac{b_i}{1-e_i} = \epsilon_{1-e,b}$ for all i – and the moral hazard cost from expansion of a flat system, $\frac{\epsilon_{1-e,b}}{\bar{e}}$, equals 0.38.

Condition (a) sets the tax and benefit schedules to approximate the current U.S. UI system (where κ equals 41% of average earnings). Condition (b) restricts heterogeneity in labour supply responses to benefit changes and sets the standard flat system UI moral hazard cost equal to the median estimate from Schmieder and Von Wachter (2016).

Under these conditions we can rewrite the moral hazard cost from an arbitrary marginal reform as:

$$\begin{aligned}
MH =& \frac{1}{\int_i (1-e_i)\rho^l(y_i) di}\epsilon_{1-e,b}\int_i (1-e_i)\rho^l(y_i)\frac{\tau + y_i^{\kappa} b}{y_i^{\kappa} b} di, \\
=& \epsilon_{1-e,b}\int_i \frac{(1-e_i)\rho^l(y_i)}{\int_{i'}(1-e_{i'})\rho^l(y_{i'})di'}\left(\frac{1}{y_i^{\kappa}}\frac{\int_{i'}(1-e'_i)\, y_{i'}^{\kappa} di'}{\int_{i'} e'_i di'} + 1\right) di \\
=& \frac{\epsilon_{1-e,b}}{\int_i e_i di}\left(\left(\int_i \frac{(1-e_i)\rho^l(y_i)}{\int_{i'}(1-e_{i'})\rho^l(y_{i'})di'}\frac{1}{y_i^{\kappa}} di\right)\left(\int_i (1-e_i)\, y_i^{\kappa} di\right) + \int_i e_i di\right) \\
=& 0.38 \times \left(\tilde{E}^l\left[\frac{1}{y_i^{\kappa}}\right] \times \left(\int_i (1-e_i)\, y_i^{\kappa} di\right) + \bar{e}\right)
\end{aligned}$$

A.2.3. Differentiated Program Funding

Suppose there is a single lump-sum benefit rate b, but group-specific lump-sum tax rates that must clear group-specific balanced-budget conditions. Then for all g:

$$\begin{aligned}(1-\bar{e}_g)b &= \bar{e}_g\tau_g \\ \frac{d\tau_g}{db} &= \frac{(1-\bar{e}_g)}{\bar{e}_g}\left(1+\frac{\epsilon_{1-\bar{e}_g}}{\bar{e}_g}\right),\end{aligned}$$

where $\bar{e}_g = \int_{i\in g} e_i di$. Therefore:

$$\frac{dW}{db} = \sum_g (1-\bar{e}_g)\left[\int_{i\in g}\frac{(1-e_i)}{1-\bar{e}_g}\omega_i u_i'(c_i^l)di - \int_{i\in g}\frac{e_i}{\bar{e}_g}\omega_i u_i'(c_i^h)di\left(1+\frac{\epsilon_{1-\bar{e}_g,b}}{\bar{e}_g}\right)\right].$$

Define:

$$\begin{aligned}\mathbb{E}_g^h[\omega u_i^{h'}(c^h)] &= \int_{i\in g}\frac{e_i}{\bar{e}_g}\omega_i u_i^{h'}(c_i^h)di \\ \mathbb{E}_g^l[\omega u_i^{l'}(c^l)] &= \int_{i\in g}\frac{(1-e_i)}{1-\bar{e}_g}\omega_i u_i^{l'}(c_i^l)di,\end{aligned}$$

Consider the welfare effect of a marginal rise in all high-state incomes:

$$\Lambda = \int_i e_i\omega_i u_i^{h'}(c_i^h)di = \bar{e}\mathbb{E}^h[\omega U_i^{h'}(c^h)]$$

We use this to rescale $\frac{dW}{db}$ according to:

$$\frac{d\tilde{W}}{db} = \frac{dW/db}{1-\bar{e}}\Big/\frac{\Lambda}{\bar{e}}.$$

Hence the marginal value of social insurance expansion is:

$$\begin{aligned}\frac{d\tilde{W}}{db} = &\sum_g \frac{1-\bar{e}_g}{1-\bar{e}}\frac{1}{\mathbb{E}^h[\omega u^{h'}(c^h)]}\left(\mathbb{E}_g^l[\omega u^{l'}(c^l)] - \mathbb{E}_g^h[\omega u^{h'}(c^h)]\right) \\ &- \sum_g \frac{1-\bar{e}_g}{1-\bar{e}}\frac{\mathbb{E}_g^h[\omega u^{h'}(c^h)]}{\mathbb{E}^h[\omega u^{h'}(c^h)]}\frac{\epsilon_{1-\bar{e}_g,b}}{\bar{e}_g}\end{aligned}$$

Household i's willingness-to-pay for a marginal social insurance expansion (based on

the perturbation $db = 1$, $d\tau = \frac{1-\bar{e}_g}{\bar{e}_g}$) is:

$$\theta_i = \left(\frac{u_i^{l'}(c_i^l) - u_i^{h'}(c_i^h)}{u_i^{h'}(c_i^h)} \right) + \left(1 - \frac{e_i/\bar{e}_g}{(1-e_i)/(1-\bar{e}_g)} \right).$$

Note that:

$$\begin{aligned}
\mathbb{E}_g^l[\omega u^{l'}(c^l)] - \mathbb{E}_g^h[\omega u^{h'}(c^h)] =& \mathbb{E}_g^l \left[\omega \left(u^{l'}(c^l) - \frac{e_i/\bar{e}_g}{(1-e_i)/(1-\bar{e}_g)} u^{h'}(c^h) \right) \right] \\
=& \mathbb{E}_g^l \left[\omega u^{h'}(c^h) \left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} + \left(1 - \frac{e_i/\bar{e}_g}{(1-e_i)/(1-\bar{e}_g)} \right) \right) \right] \\
=& \mathbb{E}^l[\omega u^{h'}(c^h)] \mathbb{E}_g^l \left[\lambda \left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} + \left(1 - \frac{e_i/\bar{e}_g}{(1-e_i)/(1-\bar{e}_g)} \right) \right) \right],
\end{aligned}$$

where $\lambda_i \equiv \frac{\omega_i u_i^{h'}(c^h)}{\mathbb{E}^l[\omega u^{h'}(c^h)]}$. Splitting $\frac{d\tilde{W}}{db}$ = Gain – Cost, (and using the definition $\bar{\lambda} = \frac{\mathbb{E}^l[\omega u^{h'}(c^h)]}{\mathbb{E}^h[\omega u^{h'}(c^h)]}$) we have

$$\begin{aligned}
\text{Gain} =& \sum_g \frac{1-\bar{e}_g}{1-\bar{e}} \frac{\mathbb{E}^l[\omega u^{h'}(c^h)]}{\mathbb{E}^h[\omega u^{h'}(c^h)]} \mathbb{E}_g^l \left[\lambda \left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} + \left(1 - \frac{e_i/\bar{e}_g}{(1-e_i)/(1-\bar{e}_g)} \right) \right) \right] \\
=& \bar{\lambda} \mathbb{E}^l \left[\lambda \left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} + \left(1 - \frac{e_i/\bar{e}_g}{(1-e_i)/(1-\bar{e}_g)} \right) \right) \right] \\
=& \bar{\lambda} \mathbb{E}^l \left[\lambda \left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} \right) \right] + \bar{\lambda} - \bar{\lambda} \int_i \frac{1-e_i}{1-\bar{e}} \lambda_i \frac{e_i}{\bar{e}_g} \frac{1-\bar{e}_g}{1-e_i} di \\
=& \bar{\lambda} \mathbb{E}^l \left[\lambda \left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} \right) \right] + \bar{\lambda} - \bar{\lambda} \sum_g \frac{1-\bar{e}_g}{1-\bar{e}} \int_{i \in g} \frac{e_i}{\bar{e}_g} \lambda_i di \\
=& \bar{\lambda} \mathbb{E}^l \left[\lambda \left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} \right) \right] + \bar{\lambda} \left(1 - \sum_g \frac{1-\bar{e}_g}{1-\bar{e}} \frac{\mathbb{E}_g^h[\omega u^{h'}(c^h)]}{\mathbb{E}^l[\omega u^{h'}(c^h)]} \right) \\
=& \bar{\lambda} \mathbb{E}^l \left[\lambda \left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} \right) \right] + \bar{\lambda} \left(1 - \sum_g \frac{1-\bar{e}_g}{1-\bar{e}} \frac{1}{\bar{\lambda}_g} \right)
\end{aligned}$$

where $\bar{\lambda}_g = \frac{\mathbb{E}_l[\omega u^{h'}(c^h)]}{\mathbb{E}_g^h[\omega u^{h'}(c^h)]}$ is the welfare weight assigned to the unemployed relative to the

employed in group g. We also have

$$\text{Cost} = \bar{\lambda} \sum_g \frac{1-\bar{e}_g}{1-\bar{e}} \frac{1}{\bar{\lambda}_g} \frac{\epsilon_{1-\bar{e}_g,b}}{\bar{e}_g}$$

A.2.4. State Dependence

Using the definitions of λ_i and $\bar{\lambda}$, we can re-write (4) as:

$$\frac{d\tilde{W}}{db} = \frac{1}{\mathbb{E}^h[\omega u^{h'}(c^h)]} \mathbb{E}^l \left[\omega u^{h'}(c^h) \left(\frac{u^{l'}(c^l) - u^{h'}(c^h)}{u^{h'}(c^h)} + \left(1 - \frac{e_i/\bar{e}}{(1-e_i)/(1-\bar{e})} \right) \right) \right] - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}},$$

We model state dependence through ϕ_i, where $u_i^{h'}(c) = u_i'(c)$ and $u_i^{l'}(c) = \phi_i u_i'(c)$. In this case a quadratic approximation to the utility function is:

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

Hence,

$$\begin{aligned}
\frac{d\tilde{W}}{db} =& \frac{1}{\mathbb{E}^h[\omega u'(c^h)]} \mathbb{E}^l \left[\phi \omega u'(c^h) \left(\gamma \frac{\Delta c}{c^h} + \left(1 - \frac{1}{\phi} \frac{e_i/\bar{e}}{(1-e_i)/(1-\bar{e})} \right) \right) \right] - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}} \\
=& \bar{\lambda}^\phi \mathbb{E}^l \left[\lambda^\phi \left(\gamma \frac{\Delta c}{c^h} + \left(1 - \frac{1}{\phi} \frac{e_i/\bar{e}}{(1-e_i)/(1-\bar{e})} \right) \right) \right] - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}
\end{aligned}$$

where $\bar{\lambda}^\phi = \frac{\mathbb{E}^l[\omega \phi u'(c^h)]}{\mathbb{E}^h[\omega u'(c^h)]}$ and $\lambda^\phi = \frac{\omega_i \phi_i u_i'(c_i^h)}{\mathbb{E}^l[\omega \phi u'(c^h)]}$. Note:

$$\mathbb{E}^l \left[\lambda^\phi \left(1 - \frac{1}{\phi} \frac{e_i/\bar{e}}{(1-e_i)/(1-\bar{e})} \right) \right] = 1 - \mathbb{E}^h \left[\frac{\lambda^\phi}{\phi} \right] = 1 - \frac{1}{\bar{\lambda}^\phi},$$

meaning that:

$$\frac{d\tilde{W}}{db} \approx \bar{\lambda}^\phi \left(\mathbb{E}^l \left[\gamma \frac{\Delta c}{c^h} \right] + \text{cov}^l \left(\lambda^\phi, \gamma \frac{\Delta c}{c^h} \right) \right) + (\bar{\lambda}^\phi - 1) - \frac{\epsilon_{1-\bar{e},b}}{\bar{e}}.$$

This is similar to equation (7), but with welfare weights adjusted to take account of state dependence.

State-Specific Prices. Suppose the utility function is state independent (so $u^{h'}(c) = u^{l'}(c) = u'(c)$), but the price of consumption is state specific. In particular, normalise the price of low-state consumption to 1 and let the price of high-state consumption be ϕ, so that the state-specific budget constraints are:

$$\begin{aligned}\phi c_i^h &= y_i - \tau \\ c_i^l &= z_i + b\end{aligned}$$

The impact of the tax and benefit change $d\tau = \frac{1-\bar{e}}{\bar{e}}$, $db = 1$ on individual i's expected utility is:

$$dV_i = (1 - e_i)u_i'(c_i^l) - e_i u_i'(c_i^h)\frac{1}{\phi}\frac{1-\bar{e}}{\bar{e}},$$

and their willingness-to-pay (given by $\theta_i = \frac{dV_i/(1-e_i)}{\frac{dV_i}{dy_i}/e_i}$) is:

$$\theta_i = \left(\frac{\phi u_i'(c_i^l) - u_i'(c_i^h)}{u_i'(c_i^h)}\right) + \left(1 - \frac{e_i/\bar{e}}{(1-e_i)/(1-\bar{e})}\right).$$

Note, that the state-specific prices transforms the marginal utility gap in the same way as the state-specific marginal-utility shifter. However, in this case, when using the consumption-based measure of the marginal utility gap it is necessary to correct observed expenditure changes for the state-specific price difference. Specifically, let x_i^s denote state-specific expenditures (which are recorded in the data). These are related to consumption by

$$\begin{aligned}\frac{\Delta c_i}{c_i^h} &= \frac{x_i^h/\phi - x_i^l}{x_i^h/\phi} \\ &= 1 - \phi\left(\frac{x_i^l}{x_i^h}\right) \\ &= \phi\left(\frac{\Delta x_i}{x_i^h}\right) - (\phi - 1)\end{aligned}$$

Hence, we have:

$$\frac{d\tilde{W}}{db} \approx \bar{\lambda}^{\phi}\left(\mathbb{E}^{l}\left[\gamma\left(\phi\left(\frac{\Delta x_i}{x_i^h}\right)-(\phi-1)\right)\right]+\text{cov}^{l}\left(\lambda^{\phi},\gamma\phi\left(\frac{\Delta x_i}{x_i^h}\right)\right)\right)+(\bar{\lambda}^{\phi}-1)-\frac{\epsilon_{1-\bar{e},b}}{\bar{e}}$$

A.2.5. General Model

Here we extend the static model we lay out in Section 2 to a much richer dynamic environment similar to that in Chetty (2006).

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

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

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

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

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

$$\dot{A}_{i,t}(\varphi_{i,t}) = \xi_{i,t}(\varphi_{i,t})\left(w_{i,t}-\tau\right)+(1-\xi_{i,t}(\varphi_{i,t}))b+f_{i,t}(x_{i,t}(\varphi_{i,t}))-c_{i,t}(\varphi_{i,t})$$

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

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

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

$$\begin{aligned}\max \int_t \int_{\varphi_{i,t}} \left(\xi_{i,t}(\varphi_{i,t}) u_i^h(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t})) + (1-\xi_{i,t}(\varphi_{i,t})) u_i^l(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))\right) dF_{i,t}(\varphi_{i,t}) dt \\ + \int_t \int_{\varphi_{i,t}} \lambda_{it}^A(\varphi_{i,t}) \left(\xi_{i,t}(\varphi_{i,t}) \left(w_{i,t} - \tau\right) + (1-\xi_{i,t}(\varphi_{i,t}))b + f_{i,t}(x_{i,t}(\varphi_{i,t})) - c_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t}) dt \\ + \int_{\varphi_{i,1}} \lambda_{i1}^A(\varphi_{i,1})(A_{i1}(\varphi_{i1}) - \bar{A}_i) dF_{i,1}(\varphi_{i,1}) \\ + \sum_{n=1}^{N} \int_t \int_{\varphi_{i,t}} \lambda_{it}^n(\varphi_{i,t}) g_{it}^n(c_{it}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t}) dF_{i,t}(\varphi_{i,t}) dt\end{aligned}$$

Denote the maximum function of this problem $V_i(b, \tau)$, written as a function of the policy parameters (b, τ). We assume the following regularity conditions:

ASSUMPTION A1 (Regularity Conditions). *Assume*

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

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

iii. $V_i(b, \tau)$ *is differentiable*

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

The government's budget constraint must hold in expectation, which means:

$$\tau \int_i \int_t \int_{\varphi_{i,t}} \xi_{i,t}(\varphi_{i,t}) dF_{i,t}(\varphi_{i,t}) dt di = b \int_i \int_t \int_{\varphi_{i,t}} \left(1 - \xi_{i,t}(\varphi_{i,t})\right) dF_{i,t}(\varphi_{i,t}) dt di$$

Let $\Xi_i = \int_t \int_{\varphi_{i,t}} \xi_{i,t}(\varphi_{i,t}) dF_{i,t}(\varphi_{i,t}) dt$ denote the expected fraction of their life agent i spends in the high state and let $\bar{\Xi} = \int_i \Xi_i di$ denote the average time spent in the high state across all agents. We can re-write the budget constraint:

$$\tau(b) = \frac{1 - \bar{\Xi}}{\bar{\Xi}} b$$

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

$$\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial b} = -(1 - \xi_{i,t}(\varphi_{i,t}))\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}$$
$$\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial \tau} = \xi_{i,t}(\varphi_{i,t})\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}$$
$$\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,s}(\varphi_{i,s})} = 0 \quad \text{if } t \neq s$$

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 government's problem is $\max_b W(b) = \int_i \omega_i V_i(b, \tau(b))di$. The effect of a marginal program expansion is:

$$\frac{dW}{db} = -\int_i \omega_i \left[\int_t \int_{\varphi_{i,t}} \left(\xi_{i,t}(\varphi_{i,t})\lambda^A_{i,t}(\varphi_{i,t}) - \sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t})\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial \tau} \right) dF_{i,t}(\varphi_{i,t})dt\frac{d\tau}{db} \right.$$
$$\left. + \int_t \int_{\varphi_{i,t}} \left((1 - \xi_{i,t}(\varphi_{i,t}))\lambda^A_{i,t}(\varphi_{i,t}) + \sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t})\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial b} \right) dF_{i,t}(\varphi_{i,t})dt \right] di$$

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

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

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

$$\sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t})\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial \tau} = \sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t})\xi_{i,t}(\varphi_{i,t})\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}$$
$$\sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t})\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial b} = -\sum_{n=1}^N \lambda^n_{i,t}(\varphi_{i,t})(1 - \xi_{i,t}(\varphi_{i,t}))\frac{\partial g^n_{i,t}(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})}$$

Hence, we can re-write dW/db:

$$\frac{dW}{db} = -\int_i \omega_i \left[\int_t \int_{\varphi_{i,t}} \xi_{i,t}(\varphi_{i,t})\frac{\partial u^h_i(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t})dt \right] di\frac{d\tau}{db}$$

$$+\int_i \omega_i \left[\int_t \int_{\varphi_{i,t}} (1-\xi_{i,t}(\varphi_{i,t})) \frac{\partial u_i^l(c_{i,t}(\varphi_{i,t})), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t}) dt \right] di$$

Define the risk-weighted expected marginal utility (taking the expectation over agents, states and time):

$$\mathbb{E}^h \left[\omega \frac{\partial u^h(c,x)}{\partial c} \right] = \int_i \left[\omega_i \int_t \int_{\varphi_{i,t}} \frac{\xi_{i,t}(\varphi_{i,t})}{\bar{\Xi}} \frac{\partial u_i^h(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t}) dt \right] di$$

$$\mathbb{E}^l \left[\omega \frac{\partial u^l(c,x)}{\partial c} \right] = \int_i \left[\omega_i \int_t \int_{\varphi_{i,t}} \frac{(1-\xi_{i,t}(\varphi_{i,t}))}{1-\bar{\Xi}} \frac{\partial u_i^l(c_{i,t}(\varphi_{i,t})), x_{i,t}(\varphi_{i,t}))}{\partial c_{i,t}(\varphi_{i,t})} dF_{i,t}(\varphi_{i,t}) dt \right] di$$

and note: $\frac{d\tau}{db} = \frac{1-\bar{\Xi}}{\bar{\Xi}} \left(1 + \frac{\epsilon_{1-\Xi,b}}{\bar{\Xi}} \right)$. Hence:

$$\begin{aligned} \frac{dW}{db} =& \bar{\Xi} \left[\mathbb{E}^l \left[\omega \frac{\partial u^h(c,x)}{\partial c} \right] - \mathbb{E}^h \left[\omega \frac{\partial u^l(c,x)}{\partial c} \right] \right] \\ & - \bar{\Xi} \mathbb{E}^h \left[\omega \frac{\partial u^h(c,x)}{\partial c} \right] \frac{\epsilon_{1-\Xi,b}}{\bar{\Xi}} \end{aligned}$$

The welfare effect of a marginal unfunded tax cut is:

$$\begin{aligned} \Lambda &= \int_i \left[\omega_i \int_t \int_{\varphi_{i,t}} \left(\xi_{i,t}(\varphi_{i,t}) \lambda_{i,t}^A(\varphi_{i,t}) - \sum_{n=1}^N \lambda_{i,t}^n(\varphi_{i,t}) \frac{\partial g_{i,t}^n(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial \tau} \right) dF_{i,t}(\varphi_{i,t}) dt \right] di \\ &= \int_i \left[\omega_i \int_t \int_{\varphi_{i,t}} \left(\xi_{i,t}(\varphi_{i,t}) \lambda_{i,t}^A(\varphi_{i,t}) - \sum_{n=1}^N \lambda_{i,t}^n(\varphi_{i,t}) \xi_{i,t}(\varphi_{i,t}) \frac{\partial g_{i,t}^n(c_{i,t}(\varphi_{i,t}), x(\varphi_{i,t}); \varphi_{i,t})}{\partial c_{i,t}(\varphi_{i,t})} \right) dF_{i,t}(\varphi_{i,t}) dt \right] di \\ &= \int_i \left[\omega_i \int_t \xi_{i,t}(\varphi_{i,t}) \int_{\varphi_{i,t}} \frac{\partial u_i^h(c_{i,t}(\varphi_{i,t}), x_{i,t}(\varphi_{i,t}))}{\partial c} dF_{i,t}(\varphi_{i,t}) dt \right] di \\ &= \bar{\Xi} \mathbb{E}^h \left[\omega \frac{\partial u^h(c,x)}{\partial c} \right] \end{aligned}$$

Recaling $\frac{dW}{db}$ according to:

$$\frac{d\tilde{W}}{db} = \frac{dW/db}{1-\bar{\Xi}} \Big/ \frac{\Lambda}{\bar{\Xi}},$$

we have:

$$\frac{d\tilde{W}}{db} = \frac{\mathbb{E}^l\left[\omega \frac{\partial u^l(c,x)}{\partial c}\right] - \mathbb{E}^h\left[\omega \frac{\partial u^h(c,x)}{\partial c}\right]}{\mathbb{E}^h\left[\omega \frac{\partial u^h(c,x)}{\partial c}\right]} - \frac{\epsilon_{1-\Xi,b}}{\bar{\Xi}}$$

Appendix B. Description of PSID Data

We use the Panel Study of Income Dynamics to estimate the sufficient statistics that allow us to quantify the welfare effects of unemployment insurance expansion. We use data from 1999-2019, the biennial so-called "new" PSID, which includes high quality information on consumption expenditures and asset holdings. For our baseline sample we focus on non-immigrant households.[36] We include both single and married households and do not select the sample on the basis of the reference person's gender.[37] Our sample has non-missing information on key demographics (age, education, and state of residence). We focus only on a sample of those where the reference person is aged between 23 and 60 and is participating in the labour force. Thus, we exclude observations where the reference person is retired, permanently disabled and neither working or looking for work, on sick (or maternity) leave or temporarily laid off, as well as those in education, who are a homemaker or in prison. In addition, to proxy for UI program eligibility we exclude those who never report employment, as well as periods immediately following a period out of the labour force. We include in our sample periods when the reference person is either in employment or unemployment, and we do not condition our sample on the employment status of their spouse. This is because we are explicitly interested in the unemployment risk facing households, including any self-insurance that may occur through spousal labour supply (Blundell et al. 2016).

To reduce the influence of measurement error, we also drop observations with extremely high asset values, as well as observations that exhibit large fluctuations in key outcomes of interest. Following, Blundell et al. (2016) we exclude households with net worth greater than $20 million to limit the impact of outliers and drop observations reporting extreme jumps (the bottom and top 0.25^{th} percentile) in wages, income and

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

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

consumption to limit the role of measurement error. We do not use data displaying extreme "jumps" from one period to the next as we view this as most likely due to measurement error. A "jump" is defined as an extremely positive (negative) change from $t-2$ to t, followed by an extreme negative (positive) change from t to $t+2$. Formally, for each variable x, we construct the biennial log difference $\Delta_2 \log(x_t)$, and drop the relevant variables for observation in the bottom 0.25 percent of the product $\Delta_2 \log(x_t)\Delta_2 \log(x_{t-2})$. Furthermore, in our analysis of permanent income we do not use earnings and wage data when the implied hourly wage is below one-half of the state minimum wage. Table B.1 provides summary statistics for our sample of interest.

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

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

Baseline: Non-Durable Expenditure including Services. Our baseline measure includes a broad range of non-durable expenditures and services, as long as those

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

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

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

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

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

Total (Non-Housing) Expenditure. We combine our baseline measure with the additional services described in the preceding paragraph (Service Expenditure) to produce a household-level series for total non-housing expenditures. As the PSID consumption categories are not designed to have full coverage, the name total expenditure is a misnomer. Furthermore, we continue to exclude the purchase of

durables, such as vehicles, and memory goods (Hai et al. 2020), such as vacations, that have durable-like properties.

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

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

Appendix C. Maximum Likelihood Estimator for Latent Heterogeneity

In the main text we specify group-specific heterogeneity as (equation (9) above):

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

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

Let D collect the vector of individual-specific indicators defined by the partition $g(i)$ and $\mathbb{X}$ collect the covariates and unemployment indicators. Then, were the assignment

	All	Employed	Unemployed
Ref. Person's Age (years)	39.57	39.69	37.57
	(10.31)	(10.29)	(10.32)
Married	0.53	0.55	0.30
	(0.50)	(0.50)	(0.46)
Number of Children	1.03	1.02	1.15
	(1.21)	(1.20)	(1.37)
Share White	0.55	0.56	0.32
	(0.50)	(0.50)	(0.47)
Share Black	0.13	0.13	0.13
	(0.33)	(0.33)	(0.33)
Ref. Person's Schooling (years)	13.78	13.84	12.76
	(2.45)	(2.44)	(2.41)
Unemployed	0.06	0.00	1.00
	(0.24)	(0.00)	(0.00)
Non-Durable Expenditure	20494.14	20852.40	14866.69
	(10760.22)	(10769.71)	(8,883.31)
Food Expenditure	9,597.02	9,741.22	7,331.98
	(5,750.13)	(5,771.88)	(4,866.04)
Services Expenditure	16225.48	16583.59	10600.17
	(13271.94)	(13343.16)	(10618.02)
Total Expenditure (Non-Housing)	25718.44	26210.59	17987.66
	(16315.74)	(16355.79)	(13495.65)
Total Expenditure (Incl. Housing)	38474.06	39287.00	25704.44
	(24913.67)	(24979.28)	(19898.34)
N	55671	52339	3332
Waves	3.65	3.45	1.51
Unique Households	15270	15188	2211

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

to groups known with certainty, the complete-data likelihood is given by

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

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

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

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

where

$$\pi_i^g = Pr(d_{i,g} = 1|\Delta_i^{FD}, \mathbb{X}_i) = \frac{\pi^g \prod_t f\left(\Delta_{i,t}^{FD}|\mu^g(U_{i,t}, X_{i,t}), (\sigma^g)^2\right)}{\sum_{h=1}^{G} \pi^h \prod_t f\left(\Delta_{i,t}^{FD}|\mu^h(U_{i,t}, X_{i,t}), (\sigma^h)^2\right)} \tag{C.4}$$

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

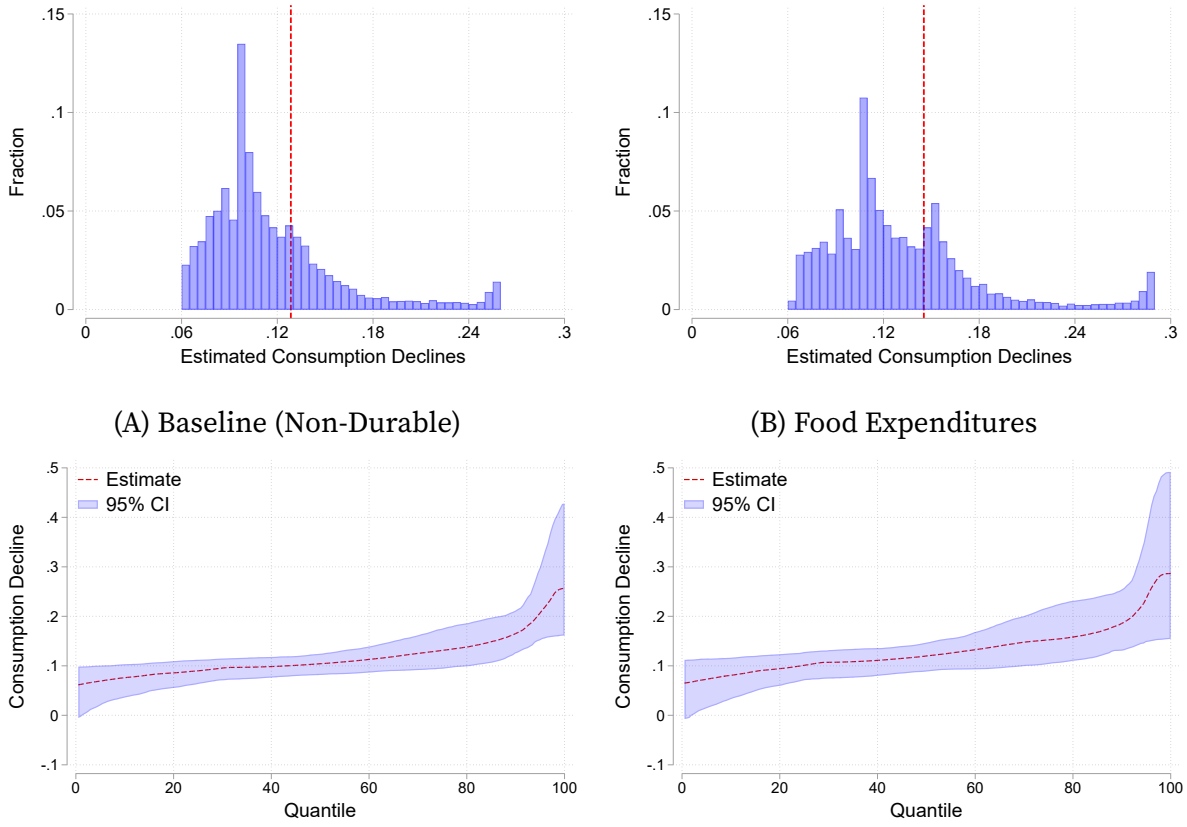
As Lewis et al. (2024) highlight, the likelihood in (C.3) admits a sequential estimation procedure using the Expectation-Maximization (E-M) algorithm Dempster et al. (1977). We implement this estimation procedure using the R package `flexmix`. Due to the local convergence properties of the E-M algorithm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

Appendix D. Additional Empirical Results

D.1. Robustness of Consumption Decline Estimates to Alternative Consumption Measures

In addition to measuring consumption using our baseline expenditure measures, we also consider robustness to alternative consumption series. Table D.1 shows that

estimates of the average consumption decline are similar across these alternative measure and Figure D.1 shows how the distribution of consumption declines differs between our baseline measure and a measure of food expenditure as well as reporting the confidence interval for our estimate of the distribution.



(A) Baseline (Non-Durable)

(B) Food Expenditures

(C) Baseline (Non-Durable)

(D) Food Expenditures

FIGURE D.1. Distribution of Estimated Consumption Declines for Alternative Consumption Series

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

	Baseline	Alternative Measures			
	Non-durable	Food Expenditure	Services	Total (Non-Housing)	Total (Incl. Housing)
$U_{i,t}$	-0.128^{***}	-0.145^{***}	-0.142^{***}	-0.128^{***}	-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, 332	37, 067	37, 319	37, 362	37, 368

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

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

D.2. Observable Heterogeneity in Consumption Declines

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

	Consumption Decline	log(Variance)
Second Quartile Liquidity	−0.009***	−0.254***
	(0.001)	(0.044)
Third Quartile Liquidity	−0.013***	−0.392***
	(0.001)	(0.046)
Top Quartile Liquidity	−0.011***	−0.219***
	(0.001)	(0.051)
Second Quartile Earnings	−0.007***	−0.174***
	(0.001)	(0.045)
Third Quartile Earnings	−0.010***	−0.234***
	(0.001)	(0.049)
Top Quartile Earnings	−0.008***	−0.165***
	(0.001)	(0.056)
Home Owner	−0.013***	−0.274***
	(0.001)	(0.042)
Married	−0.011***	−0.102**
	(0.001)	(0.046)
N	47, 841	47, 841

TABLE D.2. Consumption Declines ($\Delta c/c$) and Proxies of Access to Self-Insurance

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

We focus on proxies of a household's capacity to self-insure measured in periods of employment. This enables us to assess evidence for whether our estimated consumption declines are higher for those less able to self-insure.[38] To proxy for the availability of private savings, by which households may smooth consumption, we include indicators for each quartile of the distribution of liquid wealth and an indicator for whether or not the household owns their home. To construct liquid wealth, we follow Carroll and Samwick (1997) and sum the value of cash (including checking accounts) and savings, stocks owned outside of a retirement account, and bonds and government treasuries. These are the most liquid assets and can be used to smooth shocks in the short run. We additionally include indicators for quartiles of earned income, which captures both earning capacity and the size of the income loss a household will experience in unemployment. Finally, we include an indicator for marriage and cohabitation, which proxies for the additional insurance provided by the added work effect or spousal insurance.

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

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

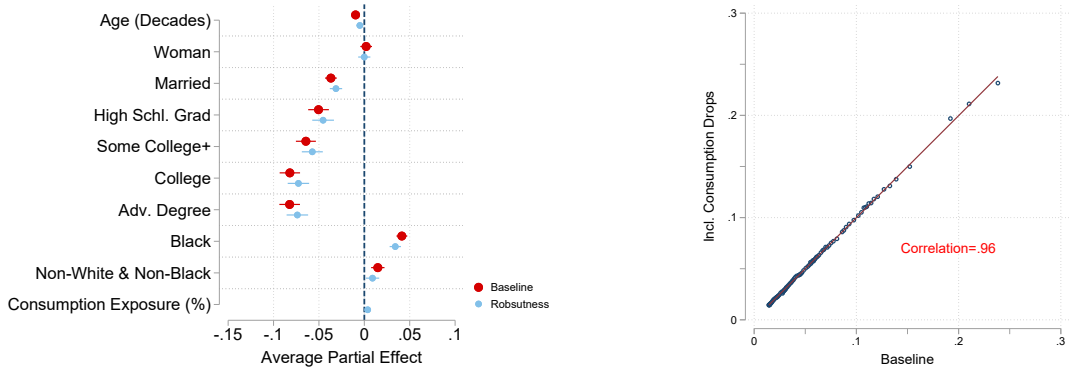
[38] As Lewis et al. (2024) note in the context of marginal propensities to consume, a key advantage of using latent types to capture heterogeneity is that it implicitly allows for various observables to be included jointly, without loss of statistical power from the interaction of successively smaller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

consumption declines.

D.3. Modeling the Correlation between Employment Risk and Consumption Declines

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

In our main empirical implementation we model employment risk as a function of observables. In this subsection we show that directly modeling the correlation between employment risk and our estimate of individual-specific consumption drops yields similar estimates of employment risk. To do so, we add to our model of employment risk (described in Section 4.2) a cubic function of the idiosyncratic consumption exposure.



(A) Average Partial Effects (B) Individual Predictions

FIGURE D.2. Correlation Between Consumption Exposure and Unemployment Risk

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

Figure D.2 reports the results of this exercise. Panel (A) shows that the implied average partial effects in the extended model are very similar for all regressors included in the baseline model. These effects are slightly attenuated towards zero as the additional measure of consumption exposure absorbs some of their explanatory power. However, the changes in effect sizes are economically insignificant. The effect of consumption exposure is statistically significant, but small. On average, an increase in the consumption exposure of 10 percentage points is associated with an increase in employment risk of 3.7 percentage points.

Panel (B) displays the average employment risk in our extended model as a function of our baseline estimates. There are no clear or systematic deviations in the predictions

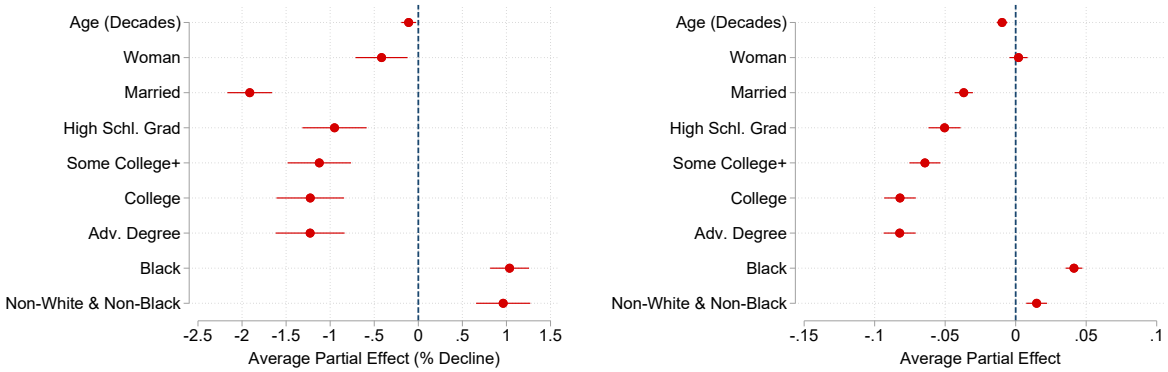
	Coefficient Estimates	Average Partial Effect
Age (Decades)	–1.368	–0.010***
	(0.985)	(0.002)
Age^2	0.263	
	(0.254)	
Age^3	–0.017	
	(0.021)	
Woman	0.037	0.002
	(0.060)	(0.003)
High Schl. Grad	–0.636***	–0.050***
	(0.068)	(0.006)
Some College+	–0.889***	–0.064***
	(0.069)	(0.006)
College	–1.318***	–0.082***
	(0.090)	(0.006)
Adv. Degree	–1.325***	–0.082***
	(0.093)	(0.006)
Black	0.755***	0.041***
	(0.054)	(0.003)
Non-White & Non-Black	0.329***	0.015***
	(0.077)	(0.004)
Constant	0.236	
	(1.224)	
N	55, 671	55, 671

TABLE D.3. Logit Estimation of e_i

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

across models: values lie on or close to the 45-degree line and the predictions have a positive correlation of almost 1. Given we estimate the consumption declines in an alternative estimation procedure which complicates the problem of inference, we choose to use the simpler model of employment risk in our baseline analysis.

Figure 4 shows how both the estimated consumption declines and labour market risks vary with in-work consumption. Here we provide additional evidence on the link between unemployment risk and consumption declines. We regress our individual-level estimates of consumption declines on the same demographic and education information that we use to model unemployment risk. This directly illustrates the covariance between weights and consumption declines in our risk-weighted expectation.



(A) Consumption Decline

(B) Unemployment Risk

FIGURE D.3. Consumption Exposure, Unemployment Risk and Demographics

Notes: Authors' calculation from the PSID data. Our baseline measures of risk and consumption declines are estimated as we describe above and correspond to the model results summarised in Figures 2 and 3. We predict consumption declines as a function of the same variables we included in our estimated unemployment risk and report the average partial effects for both.

Figure D.3 reports the average partial effects from this exercise (panel A), alongside our baseline estimates of average partial effects for unemployment risk (panel B, repeating Figure 3 in the main text). We find that factors that are systematically associated with smaller consumption declines are also associated with smaller probabilities of becoming unemployed and vice-versa. Consistent with the results shown above in Appendix D.2, we find a larger effect of marriage on consumption declines relative to employment risk, and a smaller education gradient.

D.4. Marginal Welfare Gain of UI Expansion: Additional Results

In Table D.4 we decompose the covariance effects from Table 2 into the product of standard deviations and the correlation coefficient. Table D.5 reports the insurance-redistribution decomposition (corresponding to Table 2) in the case of food-expenditure based consumption drops. Table D.6 show the impact of using a higher-order (cubic) approximation for the utility function on our baseline decomposition results. Figure D.4 shows how the insurance and cross-subsidisation components of households' willingness-to-pay for UI expansion vary across deciles of the in-work earnings distribution.

	(1)=(2)×(2)×(3)	**(2)**	**(3)**	**(4)**
	Covariance	Correlation	Standard deviation:	
	$\text{cov}^l\left(\hat{\lambda}, \gamma\frac{\Delta c}{c^h}\right)$	$\text{corr}^l\left(\hat{\lambda}, \gamma\frac{\Delta c}{c^h}\right)$	$\sigma^l\left(\hat{\lambda}\right)$	$\sigma^l\left(\gamma\frac{\Delta c}{c^h}\right)$
Benefit rise				
Flat	0.008	0.173	0.375	0.131
	[0.001, 0.014]	[0.012, 0.191]	[0.365, 0.385]	[0.077, 0.245]
Proportional	0.006	0.130	0.388	0.123
	[0.001, 0.011]	[0.025, 0.155]	[0.378, 0.398]	[0.075, 0.226]
Proportional-capped	0.006	0.140	0.364	0.126
	[0.001, 0.012]	[0.016, 0.161]	[0.355, 0.373]	[0.076, 0.234]

TABLE D.4. Decomposition of Covariance Effect

Notes: Authors' calculations from the PSID data. See Proposition 1 and Equation (7) for definition of terms. We set the coefficient of relative risk aversion, γ, equal to 3 and the welfare weight exponent, ψ, equal to 1. Numbers are based on the benefit rise being funded by a flat tax increase. 95% confidence intervals are reported in brackets.

	(1)	(2)	(3)	(4)
	Total welfare gain (1-(2))+(2)×((3)+(4))	Unemployed welfare weight $\bar{\lambda}$	Average insurance gain $E^l\left[\gamma\frac{\Delta c}{c}\right]$	Insurance-welf. wei. covariance $\text{cov}^l\left(\hat{\lambda}, \gamma\frac{\Delta c}{c^h}\right)$
Benefit rise				
Flat	0.569	1.100	0.419	0.007
	[0.456, 0.667]	[1.090, 1.110]	[0.324, 0.510]	[-0.001, 0.014]
Proportional	0.348	0.959	0.400	0.007
	[0.260, 0.427]	[0.947, 0.968]	[0.310, 0.483]	[0.001, 0.012]
Proportional-capped	0.464	1.035	0.409	0.006
	[0.366, 0.552]	[1.027, 1.043]	[0.319, 0.496]	[-0.000, 0.011]

TABLE D.5. Insurance and Redistributive Gains from UI Expansion: Food-Based Consumption Drops

Notes: Authors' calculations from the PSID data. See Proposition 1 and Equation (7) for definition of terms. We set the coefficient of relative risk aversion, γ, equal to 3 and the welfare weight exponent, ψ, equal to 1. Numbers are based on the benefit rise being funded by a flat tax increase. 95% confidence intervals are reported in brackets.

	(1)	(2)	(3)	(4)	(5)
	Baseline:		Prudence		
	Total welfare gain	Unemployed welfare weight $\bar{\lambda}$	Total adjustment (2)×((4)+(5))	Average effect $\frac{1}{2}E^l\left[\gamma\zeta\left(\frac{\Delta c}{c}\right)^2\right]$	Covariance effect $\frac{1}{2}\text{Cov}^l\left(\lambda, \gamma\zeta\left(\frac{\Delta c}{c}\right)^2\right)$
Benefit rise					
Flat	0.519	1.102	0.119	0.103	0.005
	[0.435, 0.588]	[1.092, 1.112]	[0.077, 0.176]	[0.067, 0.151]	[0.000, 0.012]
Proportional	0.303	0.959	0.092	0.093	0.004
	[0.237, 0.360]	[0.948, 0.969]	[0.062, 0.135]	[0.063, 0.135]	[0.000, 0.009]
Proportional-capped	0.415	1.036	0.104	0.097	0.004
	[0.340, 0.476]	[1.028, 1.044]	[0.068, 0.153]	[0.064, 0.142]	[0.000, 0.010]

TABLE D.6. Insurance and Redistributive Gains from UI Expansion: Higher-order Approximation

Notes: Authors' calculations from the PSID data. See Proposition 1 and Equation (7) and Appendix A.1 definition of terms. We set the coefficient of relative risk aversion, γ, equal to 3, the coefficient of relative prudence, ζ, equal to 4, and the welfare weight exponent, ψ, equal to 1. Numbers are based on the benefit rise being funded by a flat tax increase. 95% confidence intervals are reported in brackets.

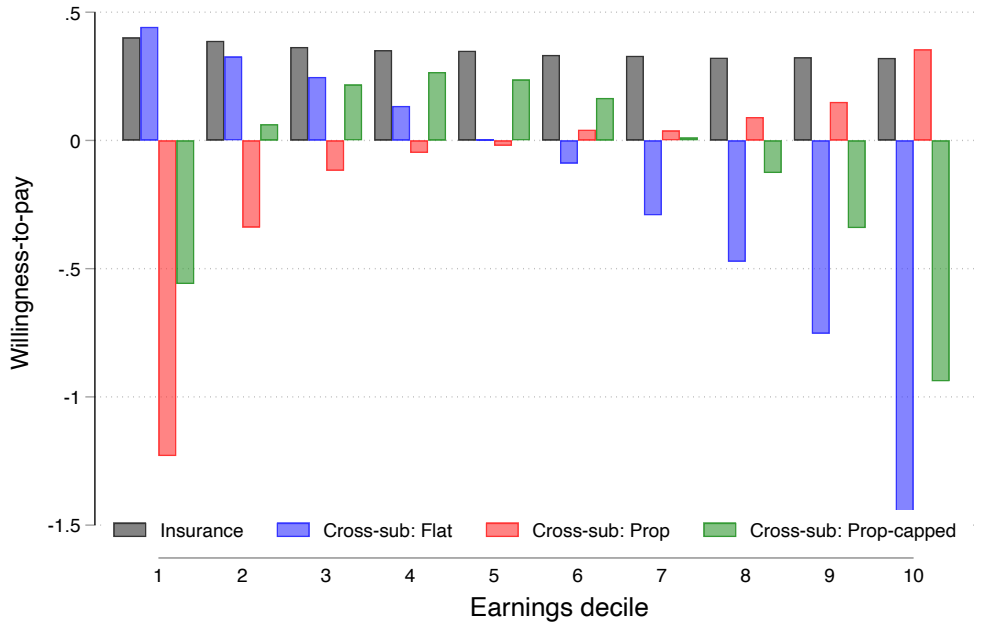


FIGURE D.4. Willingness-to-pay, by Earnings

Notes: Authors' calculation from the PSID data. Bars shows how insurance and cross-subsidisation components of households' willingness-to-pay for UI vary by deciles of the earnings distribution (at $\psi = 1$). Flat, Prop and Prop capped refer tp flat, proportional and proportional-capped benefit expansions respectively. We set the coefficient of relative risk aversion, γ, equal to 3. Numbers are based on the benefit rise being funded by a lump-sum tax increase.

State dependence. In our baseline implementation we assume utility is state independent. However, we show our implementation can easily incorporate state dependence through a marginal-utility shifter, such that $u^{h'}(c) = u'(c)$ and $u^{l'}(c) = \phi u'(c)$. Suppose ϕ is constant across workers. In this case it scales the aggregate welfare weight the planner assigns to the unemployed, $\bar{\lambda}$. If the marginal utility of consumption is higher in the unemployed state (for instance, due to a complementarity between consumption and leisure time), all else equal, this will raise $\bar{\lambda}$ and hence the marginal value of program expansion.

Aguiar et al. (2013) document that the unemployed spend more time shopping, while a number of papers provide evidence that the unemployed pay lower prices than the employed. For instance, Kaplan and Menzio (2015) provide evidence with U.S. data that, on average, the unemployed pay between 1.5% and 4.6% less, depending on whether one or both household-heads are unemployed, and Campos and Reggio (2020)

provide evidence using Spanish data that the unemployed pay 1.5% less. This can also be captured by the state-dependence parameter ϕ, which then has the interpretation as the ratio of employed- to unemployed-state prices. However, in this case it is necessary to also adjust the measured "consumption drop", so that it reflects the change in consumption (rather than expenditures) on becoming unemployed. If the unemployed pay lower prices, the unadjusted estimate of the consumption drop will overstate the true fall in consumption. See Appendix A.2.4.

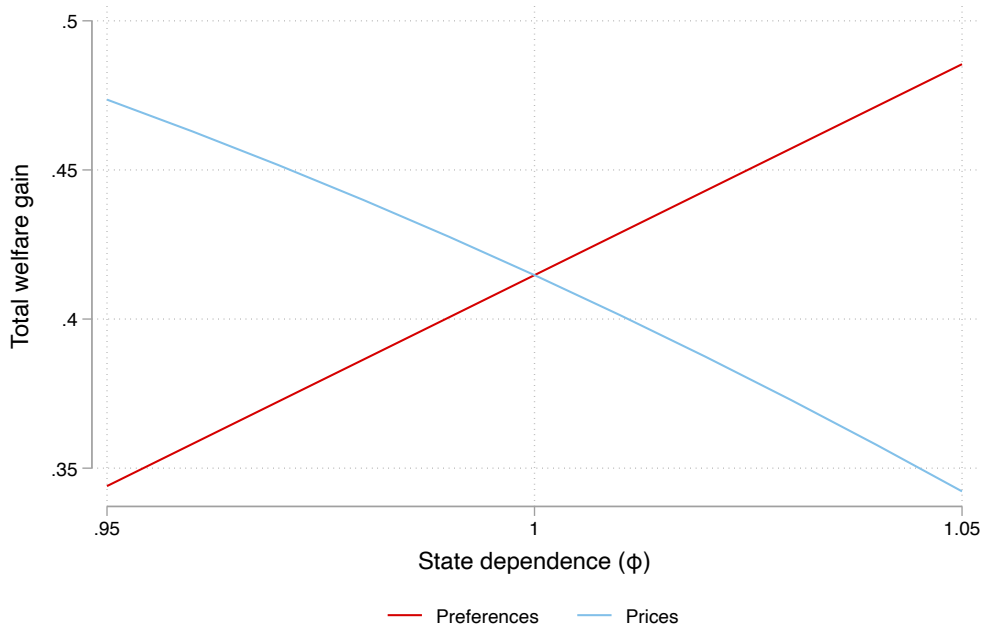


FIGURE D.5. Marginal Value of Program Expansion under State Dependence

Notes: Authors' calculation from the PSID data. Plot show marginal value of expansion of the proportional-capped system, funded by a lump-sum tax, for different values of state dependence, ϕ. "Preferences" refer to the case in which state dependence is due to state-specific preferences. "Prices" refer to the case in which state dependence is due to state-specific prices. We set the coefficient of relative risk aversion, γ, equal to 3 and the welfare weight exponent, ψ, equal to 1.

In Figure D.5 we show how the marginal value of UI program expansion varies with the state dependence parameter, both when it represents pure preference-based state dependence and differences in prices paid. In the former case, the higher is ϕ, the larger is the value of expanding UI. When ϕ represents differences in prices paid there are two off-setting effects. The higher is ϕ the cheaper is consumption in the unemployed relative to the employed state and therefore the greater is the value of UI expansion. The magnitude of this channel is the same as the pure preference channel. However,

this is off-set by the adjustment in measured consumption drops that accounts for the wedge between observed expenditure and consumption due to lower prices paid by the unemployed. When the coefficient of relative risk aversion is 3, as in our baseline, the latter effect is larger in magnitude than the former, explaining why the "Prices" line is downward sloping. For all values of ϕ displayed in Figure D.5 and over both preference- and price-based state dependence, the social value of program expansion is higher (by 0.05-0.06) than that based on the average insurance value alone (i.e., ignoring the redistributive effects of UI).