The Value of Unemployment Insurance

Camille Landais — LSE
Johannes Spinnewijn* — LSE

July 14, 2020

Abstract

Due to the absence of unemployment insurance (UI) choices, the traditional approach to estimating the value of UI is to infer it from the observed consumption response to job loss under some assumption on risk preferences. Exploiting the rich data and unique policy context in Sweden, we propose two alternative approaches that relax this assumption and we implement all three methods on the same sample of workers. The first approach considers the difference in marginal propensity to consume (MPC) when unemployed vs. employed, which allows to identify the difference in prices to smooth consumption in the respective states. The second approach exploits UI choices embedded in the Swedish UI system in a Revealed Preference approach. While the drop in consumption expenditures is relatively small ($\sim$ 13 percent), we find that the MPC is around 25 percent higher when unemployed than employed, translating into a marginal value of transfers that is at least 60 percent higher when unemployed than employed. This high value of UI is confirmed by our RP estimates and indicates substantial risk aversion given the relatively small drop in consumption expenditures.

Keywords: Unemployment Insurance, Consumption Smoothing, Revealed Preference, MPC
JEL codes: H20, J64

*We would like to thank Richard Blundell, Amy Finkelstein, Peter Ganong, Francois Gerard, Jon Gruber, Hamish Low, Ben Moll, Luigi Pistaferri, Emmanuel Saez, Frans Spinnewyn, Jon Steinsson, Dmitry Taubinsky, Silvana Tenreyro, Danny Yagan and seminar participants at Berkeley, Bristol, Bonn, Cambridge, Columbia, IFS, IIPF, LSE, MIT, Oxford, Princeton, Stanford, Tinbergen, TSE, and Uppsala for helpful discussions and suggestions. We also thank Henriette Druba, Arnaud Dyevre, Miguel Fajardo-Steinhauser, Jack Fisher, Alice Lapeyre, Mathilde Munoz, Will Parker and Quirin Von Blomberg for excellent research assistance. We acknowledge financial support from ERC grants 679704 and 716485.

1 Introduction

Social insurance programs that protect workers against adverse shocks take up a substantial share of government expenditures. As a consequence, the potential negative impact of these programs on workers' employment has been put under scrutiny and is the topic of a large and ever-growing literature [see reviews in Krueger and Meyer [2002], Chetty and Finkelstein [2013] and Schmieder and Von Wachter [2016]]. As the distortionary costs of social insurance programs are found to be high, one would expect the insurance value of these programs to be high too in order to be able to justify their generosity [see Baily [1978]; Chetty [2006]]. However, the evidence on the value of social insurance is lagging behind the evidence on its costs. Conceptually, it is easy to understand the value of providing more insurance against an adverse shock, like, for instance, unemployment: the value of insurance is simply captured by the marginal rate of substitution (MRS) between employment and unemployment consumption. Yet, it is remarkably difficult to estimate in practice. The main reason is that social insurance programs, like unemployment insurance (UI), are often mandated, leaving little or no choice for its beneficiaries. This reduces the ability of researchers to identify the value of these programs by applying direct revealed preference methods.

The traditional approach in the literature, famously implemented by Gruber [1997] in the context of unemployment, is to focus on consumption smoothing. The estimated drop in consumption in response to an adverse event can be scaled by workers' risk aversion to get an estimate of the value of providing additional insurance against that event. As the consumption drops at job loss are consistently found to be small, the implied value of insurance would be low for conventional levels of risk aversion. Given the large moral hazard responses to UI estimated in the literature, the consumption-based (CB) approach suggests that UI policies are too generous. This conclusion relies on assumptions on preferences, which are difficult to relax in practice, and has been contested in more recent work [e.g., Chetty and Looney [2006]; Chetty [2008]; Landais [2015]; Hendren [2017]].

This paper proposes and implements two novel methods to estimate the value of unemployment insurance, and then compares their results to the standard CB approach. A major advantage of our setting is that we can deliver all three implementations, not just in the same context, but for the very same workers. Instead of considering the change in consumption levels, our first method considers the change in marginal propensities to consume when becoming unemployed. We show that this identifies the relative price of increasing consumption when unemployed vs. employed and bounds the MRS between employment and unemployment consumption. This approach is robust to some important challenges for the CB approach, but requires comparable sources of income variation both when unemployed and employed. Our second method studies the value of insurance as revealed by a worker's insurance choices. Using a revealed preference argument, the price paid for expected coverage, taking a worker's unemployment risk into account, identifies her MRS. This method, however, requires data on unemployment risks and UI choices.

To implement and compare all three methods, we take advantage of the uniqueness of the Swedish setting, which combines granular data on consumption expenditures with the availability of coverage choice in the UI policy. We start our empirical investigation by revisiting the analysis of consumption dynamics around job loss. We use the registry-based measure of consumption expenditures from Kolsrud et al. [2017], constructed as a residual from the household's budget constraint, thanks to the availability of comprehensive and detailed information on income and assets. In line with prior work, the estimated consumption drops are relatively small, and translate into moderate values of the MRS. The mark-ups that workers are willing to pay for transferring a marginal krona from employment to unemployment, controlling for unemployment risk, are between 10 to 50 percent for a range of commonly used risk aversion values. Our results, however, show that almost all consumption protection is offered by the UI transfers. Liquid assets play a limited role, and the take-up of debt in fact decreases when unemployed. Also, the earnings of other household members do not significantly increase in response to a job loss. Taken together, this evidence suggests that the observed lack of private consumption smoothing may not be driven by its low value to workers, but simply by its high cost.

To get at the value of insurance more directly, we propose a new approach based on the difference in marginal propensities to consume when unemployed and employed. The approach leverages the fact that the marginal propensity to consume (MPC) out of extra income is directly related to the shadow price of increasing consumption. This shadow price can for example be driven by the interest rate faced by the individual seeking to borrow and/or the wage of spousal employment. The higher the shadow price of consumption, the more a worker will increase her consumption when given an extra krona of income. This relation between the shadow price and the MPC is mitigated by the curvature in preferences over consumption, but we control for this by considering the relative MPC when unemployed compared to when employed. The relative price of consumption smoothing in the two states then allows us to bound the value of consumption smoothing between the two states at the margin. We derive the assumptions under which this MPC approach provides a lower bound on the value of UI. We show that the use of MPCs rather than consumption levels makes the method robust to the important challenge for the CB approach to convert wedges in consumption expenditures into wedges in marginal utilities, which requires information on the curvature of consumption preferences and thus about the role of for example work-related expenditures, committed expenditures, non-durable goods, home production, etc.

To estimate the state-specific MPCs, we exploit the large variation in welfare transfers that municipalities provide differentially across household types (e.g., household size, age, and income) and over time. Using a first-difference model and the same sample of job losers used in the CB approach, we estimate large MPCs, both when employed and unemployed, but find that the MPC is significantly larger when unemployed. In our baseline specification, the MPC when unemployed equals .551 and is about 25 percent higher than when employed. We show that this difference is robust across a wide range of specifications and confirm the high MPC when unemployed exploiting variation in UI benefits in a regression-kink design. The MPC estimates translate into a price of increasing consumption that is about 60% higher during unemployment compared to employment. These values provide a lower bound on the mark-up that workers are willing to pay to transfer a krona from employment to unemployment, which is substantially higher than the range of values we find from the CB approach.

We finally take advantage of the presence of consumer choice in the UI system in Sweden to estimate the insurance value using a Revealed Preference (RP) approach. All Swedish workers are given the choice between a basic flat benefit level and income-related unemployment benefits against a uniform premium. The specific challenge is to retrieve a worker's revealed value of insurance coming from her MRS rather than from her unemployment risk. We predict workers' unemployment risk using a rich set of observables, including arguably exogenous risk shifters [Landais et al., 2017], and exploit the risk variation to estimate workers' MRS. In this RP approach, we try to account for risk misperceptions and other potential choice frictions, that may confound the estimates of the MRS. Overall, we find the revealed MRS to be substantially higher than with the CB implementation, corroborating the high value of UI we get from the MPC approach. Interestingly, our RP approach also reveals large dispersion in MRS. We provide evidence suggesting that both preference heterogeneity and heterogeneity in choice frictions seem to play a role.

The high value of UI, implied by both the MPC and RP approach, would justify setting UI benefits at a generous level, even when the corresponding moral hazard cost is high. These policy recommendations are opposite to what one would conclude based on the CB approach, at least for the values of relative risk aversion conventionally used in the literature ($\sim \gamma = 1 - 4$). To reconcile the high average value of UI with the modest consumption drops, we would need higher, but arguably plausible risk aversion levels ($\sim \gamma = 4 - 8$). In principle, state-dependence in workers' preferences or the measurement of consumption expenditures could also help bridging the wedge between the results of the different approaches. However, a comparison of the observed consumption drop at job loss and the imputed consumption drop based on our MPC estimates suggests that the role for state-dependence is limited.

The paper will proceed as follows. After setting up the model in Section 2, we present the MPC approach and compare it to the traditional CB approach and alternative RP approaches in Section 3. We then discuss the data and context in Section 4 before implementing the respective approaches in Sections 5-7. The last sections put the results together and conclude.

Related literature The gap between the literature on the value and cost of social insurance was the motivation of Gruber [1997]'s original study of consumption smoothing by unemployed workers now two decades ago. Even today the gap is still wide, but there are notable exceptions. A number of papers have studied the value of UI, either in the spirit of the CB approach [e.g, Browning and Crossley [2001], Stephens [2001], Ganong and Noel [2017], Kolsrud et al. [2018]] or using so-called 'optimization approaches' [e.g., Shimer and Werning [2007], Chetty [2008], Landais [2015], Hendren [2017]] developed to overcome challenges of the CB approach. The latter work considers other margins that workers adjust to protect against unemployment (e.g., search effort, reservation wage, precautionary savings, household labor supply) and use behavioral responses (to UI benefit or to unemployment risk changes) to infer the value of UI. Our MPC approach is closely related to this, but centered on consumption, which is the directly relevant margin of adjustment encompassing all other margins, and comparing responses when unemployed and employed. The former avoids having to take a stance on which margin of adjustment is binding.

The literature studying the value of social insurance and transfers extends beyond UI, with studies using CB approaches [e.g., Autor et al. [2017]], optimization approaches [e.g., Finkelstein et al. [2015], RP approaches [Cabral and Cullen [2016], Finkelstein et al. [2017]], Fadlon and Nielsen [2018]] and more structural approaches [e.g., Low and Pistaferri [2015], Finkelstein et al. [2015], Low et al. [2018]]. Our work is to the best of our knowledge unique by implementing these approaches in the same setting and on the same sample.

Our analysis also contributes to the large literature studying consumption insurance and consumption responses to income shocks more generally [see Jappelli and Pistaferri [2010]], and two rapidly growing strands within that literature using registry-based measures of consumption [e.g., Koijen et al. [2014], Kolsrud et al. [2017], Eika et al. [2017]] and estimating MPCs [e.g., Kreiner et al. [2016], Kekre [2017], Di Maggio et al. [2018]]. While several papers use MPC estimates to learn about plausible models of consumption behavior [see Nakamura and Steinsson [2018]], our focus is on the value of insurance that the MPCs reveal. In the context of UI, two notable examples are Ganong and Noel [2017] and Gerard and Naritomi [2019], who document and explain the lack of anticipation of UI benefit exhaustion and the excess sensivity in consumption to liquidity. Finally, while RP approaches are commonplace in the insurance literature, our combination of methods allows us to shed new light on the role of preferences vs. behavioral frictions, which is a central topic in the health insurance literature [e.g., Abaluck and Gruber [2011], Handel [2013], Handel and Kolstad [2015], Spinnewijn [2017]].

2 Conceptual Framework

We set up a stylized model of unemployment to define our object of interest and to present the different approaches to estimating the value of unemployment insurance.

2.1 Setup

Our baseline model is static and considers an agent who is either employed or unemployed. The respective states are denoted by $s \in \{e, u\}$. When employed the agent has disposable income y_e, which depends on her earnings and the taxes she pays. When unemployed the agent has disposable income y_u, which depends on the unemployment benefits she receives. The agent's expected utility equals

$$V = \pi(z)\, v_u(c_u, x_u, z) + (1 - \pi(z))\, v_e(c_e, x_e, z), \tag{1}$$

where c_u and c_e denote consumption when unemployed and employed respectively. The variables z and x represent two types of actions taken by the agent:

The first type refers to the actions the agent undertakes to reduce her unemployment risk, for example effort to avoid job loss or to find a job when unemployed. We assume that the probability of unemployment equals $\pi(z)$ where z is costly and affects utility in both states.

The second type of actions refers to the various means an agent can use to smooth consumption between employment and unemployment. This includes a worker's precautionary savings, access to credit, formal and informal insurance arrangements, household labor supply, etc. We refer to x_s as the resources used to increase or decrease consumption relative to the income y_s in state s. We allow the price of increasing resources p_s to be state-dependent. That is,

$$c_s = y_s + \frac{1}{p_s} x_s \text{ for } s = e, u. \tag{2}$$

The agent maximizes her expected utility given the state-specific budget constraints. In an interior optimum, she equalizes the utility of an extra krona of consumption to the utility cost of raising that extra krona of revenue in any given state:

$$\frac{\partial v_s(c_s, x_s, z)}{\partial c} = -p_s \frac{\partial v_s(c_s, x_s, z)}{\partial x}. \tag{3}$$

For tractability, we will assume that preferences are separable, but allow them to be state-dependent:

Assumption 0. $\frac{\partial^2 v_s}{\partial a_k \partial a_l} = 0$ *for* $a_k, a_l \in \{c, x, z\}$ *and* $a_k \neq a_l$.

Our stylized framework allows for tractable characterizations that ease the comparison of the different methodologies to estimate the value of UI. The general representation is meant to capture different models of resources used to smooth consumption. The static representation naturally fits a model with household labor supply x_s, where the resource cost of increasing consumption is to increase the household hours of work (e.g., Fadlon and Nielsen [2018]; Hendren [2017]). In this case, p_s is the inverse of the household's marginal wage, which may change with a member's employment status. An alternative resource to smooth consumption are financial assets and credit. To illustrate how our insights generalize, we briefly introduce a dynamic extension of our model with intertemporal consumption smoothing.[1]

Dynamic Application In a dynamic setting, an agent's expected utility can be written as:

$$V_{s,t}(A_t) = \max_{A_{t+1}, c_{s,t}, z_{s,t}} \left\{ \tilde{v}_s(c_{s,t}, z_{s,t}) + \beta \left[\tilde{\pi}_s(z_{s,t}) V_{u,t+1}(A_{s,t+1}) + (1 - \tilde{\pi}_s(z_{s,t})) V_{e,t+1}(A_{s,t+1}) \right] \right\},$$
$$\text{s.t. } c_{s,t} = A_t + y_{s,t} - \frac{A_{s,t+1}}{R_{s,t}} \text{ and } A_{s,t+1} \geq \bar{A},$$

for $s \in \{e, u\}$ and t. The asset (or debt) holdings left for the next period $A_{s,t+1}$ are the endogenous resource used to change consumption today relative to the available cash-on-hand, $A_t + y_{s,t}$. The marginal resource cost of increasing consumption today thus equals the marginal present value of

[1] Our framework can also be extended to incorporate private insurance choices, where the resource cost of increasing consumption when unemployed is to lower consumption when employed, as discussed in Section 3. This corresponds to introducing an *ex ante* choice to buy Arrow-Debreu securities x_s that pay out $1/p_s$ in state s per krona spent, occurring with probability π_s.

next period's asset holdings. Optimal consumption smoothing is governed by an Euler equation,

$$\frac{\partial \tilde{v}_s(c_{s,t}, z_{s,t})}{\partial c} = R_{s,t} \times \beta\left[\tilde{\pi}_s(z_{s,t})V'_{u,t+1}(A_{s,t+1}) + (1-\tilde{\pi}_s(z_{s,t}))V'_{e,t+1}(A_{s,t+1})\right] \tag{4}$$

$$\equiv p_{s,t} \times \beta E_s(V'_{\tilde{s},t+1}(A_{s,t+1})). \tag{5}$$

Since $V'_{\tilde{s},t+1}(A_{s,t+1}) = \frac{\partial}{\partial c}\tilde{v}_s(c_{\tilde{s},t+1}, z_{\tilde{s},t+1})$ by the envelope condition, the Euler equation highlights the trade-off in utility between more consumption today and more consumption in the future. Comparing this to condition (3), the price of consumption today $p_{s,t}$ is simply the gross interest rate $R_{s,t}$. The interest rate is plausibly larger when unemployed as liquidity or borrowing constraints are more likely to bind.

2.2 Consumption Smoothing and the MRS

The marginal rate of substitution (MRS) describes how much consumption workers are willing to give up when employed to increase their consumption when unemployed,

$$MRS = \frac{\frac{\partial v_u(c_u, x_u, z)}{\partial c}}{\frac{\partial v_e(c_e, x_e, z)}{\partial c}}.$$

The value of extra unemployment benefits is fully determined by the MRS.[2] As a result of the envelope theorem, a small change in x (or in z) in response to a change in UI has only a second order impact on the agent's own welfare.[3] This implies that the welfare impact of a small increase in state-specific income y_s depends only on its direct effect, captured by the state-specific marginal utility of consumption. Hence, conditional on knowing the MRS, there is no need to know how much an increase in UI crowds out private consumption smoothing to estimate the value of UI. Neither would we need to know the means used to smooth consumption or the income shock underlying the job loss.

When setting the unemployment benefit levels, their value is traded off against the fiscal externality due to the reduced incentives to avoid unemployment. In our stylized model of unemployment, optimal UI is characterized by a Baily-Chetty formula, $MRS = 1 + \varepsilon_{\frac{\pi}{1-\pi}}$, where the fiscal externality is captured by $\varepsilon_{\frac{\pi}{1-\pi}}$, which equals the elasticity of the unemployment risk $\pi/(1-\pi)$ wrt to a tax-funded increase in UI (Baily [1978]; Chetty [2006]).

Consumption-Based Approach The standard approach to estimating the MRS is to link it to the difference in consumption between employment and unemployment. The basic idea is that, everything else equal, a worker values UI more the larger the drop in consumption she would be

[2] In a dynamic context, the expression extends to the average marginal utility of consumption over the beneficiaries of the unemployment benefits when evaluating either the average generosity of UI [see Chetty [2006]] or the dynamic profile of UI [see Kolsrud et al. [2018]].

[3] The envelope theorem requires concavity and differentiability of both v and π. See Chetty [2006] for further discussion.

exposed to when becoming unemployed. We refer to this standard approach as the consumption-based (CB) approach. Based on a Taylor expansion of the marginal utility of consumption, the MRS tends to be approximated by

$$MRS \cong 1 + \sigma^c \times [c_e - c_u]\,, \tag{6}$$

showing clearly how the MRS depends on the drop in consumption, scaled by the curvature of the consumption preferences $\sigma^c = -\frac{\partial^2 v/\partial c^2}{\partial v/\partial c}$. With the right information on this curvature, the CB approach is remarkably easy to implement. This information, however, is essential, but hard to come by. The relevant curvature depends on how consumption expenditures are measured and which consumption categories respond to unemployment, as we discuss in detail in Section 3.2.[4] This challenge is further complicated by the fact that the observed drop in consumption depends both on a worker's preferences and the price she faces [Chetty and Looney [2006, 2007]]. She may not smooth consumption much, either because it is expensive or she doesn't value it. The difference is of course essential for inferring the value of UI and at the root of the challenge for the CB approach when information on preferences is not readily available.

3 The MPC Approach

This section presents a novel approach to identify the price of consumption smoothing at the margin and bound the value of extra consumption smoothing through UI. This MPC approach allows to relax the assumptions on preferences that are required for the CB approach. We first present the MPC approach in our baseline model and then evaluate its robustness relative to the CB approach in extensions of the baseline model. We also compare the MPC approach to a revealed preference (RP) approach, which we implement as well, and other related optimization approaches recently proposed in the literature.

3.1 Characterization

We first show that the price of smoothing consumption in a given state p_c can be inferred from the marginal propensity to consume (MPC) out of income received in that state, dc_s/dy_s. Intuitively, the higher the shadow price of income in a given state, ceteris paribus, the higher the marginal propensity to consume out of income in that state. More formally, by implicit differentiation of the optimality condition in (3), we find

$$\left[\frac{\partial^2 v_s}{\partial c^2} + p_s \frac{\partial^2 v_s}{\partial x \partial c}\right] \times dc_s + \left[\frac{\partial^2 v_s}{\partial x \partial c} + p_s \frac{\partial^2 v_s}{\partial x^2}\right] p_s \times [dc_s - dy_s] + \left[\frac{\partial^2 v_s}{\partial z \partial c} + p_s \frac{\partial^2 v_s}{\partial z \partial x}\right] \times dz = 0,$$

[4] We present the full Taylor expansion in Appendix A.1. The approximation relies on the higher-order derivatives of the utility functions being small and requires the marginal utility of consumption to be state-independent, conditional on consumption and the resources used, and to be separable in consumption and resources [see Chetty [2006]]. We come back to this in Section 3.2.

where we have dropped the arguments of the utility function. Assuming separable preferences and using the optimality condition (3) again, we can re-express the MPC as

$$\frac{dc_s}{dy_s} = \frac{p_s \frac{\partial^2 v_s}{\partial x^2} / \frac{\partial v_s}{\partial x}}{-\frac{\partial^2 v_s}{\partial c^2} / \frac{\partial v_s}{\partial c} + p_s \frac{\partial^2 v_s}{\partial x^2} / \frac{\partial v_s}{\partial x}}. \tag{7}$$

Equation (7) indicates that the MPC is simply a function of the price of consumption on the one hand and the curvature of the utility function wrt consumption and resources on the other hand. In a given state, a larger share of extra income is consumed the higher the cost of generating extra income in that state, as captured by p_s. This effect does, however, get mitigated when the marginal return to consumption decreases more rapidly than the marginal cost of generating income increases. Expressed as an odds ratio and rescaling the MPC when unemployed relative to the MPC when employed, we can state the following result:

Lemma 1. *Under Assumption 0, we have*

$$\frac{O_{mpc_u}}{O_{mpc_e}} = \frac{p_u}{p_e} \times \frac{\sigma_u^x / \sigma_u^c}{\sigma_e^x / \sigma_e^c},$$

where $O_{mpc_s} \equiv \frac{dc_s/dy_s}{1-dc_s/dy_s}, \sigma_s^c \equiv -\frac{\partial^2 v_s}{\partial c^2} / \frac{\partial v_s}{\partial c}$ *and* $\sigma_s^x \equiv \frac{\partial^2 v_s}{\partial x^2} / \frac{\partial v_s}{\partial x}$, *all evaluated at the respective* (c_s, x_s, z) *.*

The Lemma demonstrates that by rescaling the MPCs we can infer the relative price p_u/p_e without requiring information on preferences per se. What we need instead is an assumption on how the relative curvature changes between unemployment and employment. A first natural assumption is that the relative curvature remains constant across states. This is the case when both preferences over consumption and resources are represented by exponential functions, like for CARA preferences. That is, $v_s(c,x) = \frac{\exp(-\tilde{\sigma}^c c)}{-\tilde{\sigma}^c} - \frac{\exp(\tilde{\sigma}^x x)}{\tilde{\sigma}^x}$. This assumes state-independence in the curvature parameters $\tilde{\sigma}^j$, so that $\sigma_s^c/\sigma_s^x = \tilde{\sigma}^c/\tilde{\sigma}^x$ for $s = e, u$. Note that any state-specific scalars in the utility function would not change this result. A second step is to note that as long as the relative curvature is higher when unemployed, the ratio of MPCs provides a lower bound on the relative price. The corresponding assumption on the preference curvatures is:

Assumption 1. $\frac{\sigma_u^c}{\sigma_u^x} / \frac{\sigma_e^c}{\sigma_e^x} \geq 1$, *where the curvatures are evaluated at* (c_u, x_u, z) *and* (c_e, x_e, z) *respectively.*

Simply because income is lower when unemployed, our premise is that more state-specific resources are used when unemployed to smooth consumption (i.e., $x_u \geq x_e$), while consumption remains lower (i.e., $c_e \geq c_u$). In this case, it becomes sufficient for utility over consumption and resources to satisfy DARA for Assumption 1 to be satisfied. This includes CRRA preferences: $v_s(c,x) = \frac{c^{1-\tilde{\gamma}^c}}{1-\tilde{\gamma}^c} - \frac{x^{1+\tilde{\gamma}^x}}{1+\tilde{\gamma}^x}$, so that $\sigma_s^c/\sigma_s^x = (\tilde{\gamma}^c/\tilde{\gamma}^x)/(x_s/c_s)$ for $s = e, u$.

The second step is to link the relative prices of smoothing consumption to the MRS. Indeed, optimizing workers equalize the marginal utility of consumption and the marginal cost of raising revenue, both when unemployed and employed, as stated in condition (3). Putting the optimality conditions in both states together, we get:

Lemma 2. *At an interior optimum, we have*

$$MRS = \frac{p_u}{p_e} \times \frac{\partial v_u / \partial x}{\partial v_e / \partial x}, \tag{8}$$

where the marginal resource costs are evaluated at (c_u, x_u, z) *and* (c_e, x_e, z) *respectively.*

The MRS crucially depends on the difference in prices across states, but it also depends on the difference in the marginal resource cost $\partial v_s / \partial x$. Again, as income drops when losing one's job, more state-specific resources are used when unemployed (i.e., $x_u > x_e$) so that we expect the marginal resource cost to be higher when unemployed. Stated differently, workers will smooth consumption less, and thus the MRS will be higher, when either the relative price of increasing consumption is higher or the relative disutility of increasing resources is higher when unemployed. Hence, the price ratio is expected to provide a lower bound on the MRS. The corresponding assumption equals:

Assumption 2. $\frac{\partial v_u}{\partial x} / \frac{\partial v_e}{\partial x} \geq 1$*, where the marginal resource costs are evaluated at* (c_u, x_u, z) *and* (c_e, x_e, z) *respectively.*

The assumption on the sign of the wedge seems naturally satisfied when only the resources are different between unemployment and employment. However, state-specific differences in preferences could lead to a violation of Assumption 2. For example, in a model where the partner of the unemployed worker increases his or her earnings to smooth the loss in household earnings, the assumption can be violated when the worker's unemployment actually decreases the disutility of work for the partner, for example when leisure time is substitutable between partners.[5]

Putting the two steps together we can use the MPCs to derive a lower bound on the MRS under the stated assumptions:

Proposition 1. *Under Assumptions 0, 1 and 2, we have*

$$MRS \geq \frac{O_u^{mpc}}{O_e^{mpc}}. \tag{9}$$

Proposition 1 shows that beyond the stated assumptions we do not need any information on the curvature of consumption preferences. This is a major advantage relative to the CB approach. A limitation of the MPC approach is that it only provides a lower bound, picking up the difference in prices, but not the potential wedge in resource costs. Moreover, the implementation of the MPC approach is more demanding than the CB approach. To estimate the differential response in consumption, we need exogenous variation in state-contingent income both when employed and unemployed.

[5] We further discuss state-specific preferences in Section 3.2.

Dynamic Application We briefly illustrate how the lower-bound argument and assumptions extend to the dynamic application where MPCs helps to uncover the difference in interest rates that apply when unemployed and employed. For tractability, we continue to assume that consumption smoothing behavior does not interact with the efforts underlying the unemployment probability (see Assumption 0). Lemma 1 then continues to apply for the marginal propensity to consumption with respect to a transitory, unanticipated shock in income, $dc_{s,t}/dy_{s,t}$. That is,

$$\frac{\frac{dc_{s,t}}{dy_{s,t}}}{1-\frac{dc_{s,t}}{dy_{s,t}}} = R_{s,t}\left(\frac{E_s\left[V''_{\tilde{s},t+1}(A_{s,t+1})\right]}{E_s\left[V'_{\tilde{s},t+1}(A_{s,t+1})\right]} \Big/ \frac{v''_s(c_{s,t})}{v'_s(c_{s,t})}\right), \tag{10}$$

as we derive in Appendix A.1.[6] Hence, the relative odds ratios of the MPCs is equal to the relative interest rates R_u/R_s multiplied by the ratio of the relative curvatures of the instantaneous utility and continuation utility. With CARA consumption preferences, the instantaneous and continuation utility have constant curvature [see Spinnewijn [2015]], implying that Assumption 1 would be satisfied with equality and the relative MPCs exactly identify the relative interest rates.

Using the Euler conditions in (4) when unemployed and employed, we can then relate the relative interest rates to the MRS,

$$\frac{\frac{\partial v_u(c_{u,t})}{\partial c}}{\frac{\partial v_e(c_{e,t})}{\partial c}} = \frac{R_{u,t}}{R_{e,t}}\frac{E_u(V'_{\tilde{s},t+1}(A_{u,t+1}))}{E_e(V'_{\tilde{s},t+1}(A_{e,t+1}))}. \tag{11}$$

The cost of using future assets at the margin to increase consumption today is generally expected to be higher when unemployed than when employed, $E_u(V'_{\tilde{s},t+1}(A_{u,t+1})) > E_e(V'_{\tilde{s},t+1}(A_{e,t+1}))$, which would imply that Assumption 2 is satisfied as well. First, given her lower current income when unemployed, a worker will draw down her assets more to increase her current consumption when she is unemployed than when employed and thus leave fewer assets for the next period $A_{u,t+1} < A_{e,t+1}$. Second, future income is expected to be lower too when unemployed (and may be so permanently). With concave utility, both forces push the difference in marginal resource costs when unemployed and employed in the same direction. While the MPC approach only picks up the difference in interest rates and not this difference in resource costs underlying the MRS, the lower-bound argument continues to hold.

3.2 Comparison to CB Approach

We already highlighted the important advantage of the MPC approach relative to the CB approach allowing us to relax the information requirements on preferences. We further discuss the robustness of the MPC approach relative to the CB approach in the presence of state-dependence preferences or consumption expenditures and show how the two approaches can be combined to gauge its importance.

[6] All the appendix material from this paper is available online in Landais and Spinnewijn [2020].

Consumption Expenditures An important challenge to implement the CB approach in practice is the measurement of the drop in consumption expenditures and the relevant preference parameters the measured drop needs to be scaled by.

First, work- and job search-specific expenditures may affect the observed drop in consumption expenditures $c_e - c_u$, but not the drop in consumption relevant for utility [Browning and Crossley, 2001]. In particular, assuming $v_s(c_s) = v(c_s - \phi_s)$ with state-specific expenditure ϕ_s, the MRS becomes

$$MRS \cong 1 + \sigma^c \times [c_e - c_u - (\phi_e - \phi_u)]. \tag{12}$$

While state-specific expenditures affect the observed consumption drop, they do not affect the MPCs and thus the validity of the MPC approach, as we show in Appendix A.2. A related issue is the possibility that unemployed workers complement their consumption expenditures with home production or time to shop for lower prices [Aguiar and Hurst, 2005]. We can model this assuming $v_s(c_s) = v(\eta_s c_s)$. While this state-dependence again affects the observed drop in consumption expenditures, it provides another source of state-specific variation in the price of consumption, which is exactly what is picked up by the MPC approach as we show in Appendix A.2.[7]

Second, the type of consumption expenditures that are reduced determines the relevant curvature σ^c to scale the consumption drop to obtain an estimate of the MRS. This scalar will be higher when a large share of expenditures is committed and thus the drop in consumption expenditures is highly concentrated [Chetty and Szeidl, 2007]. However, it can also be lower when expenditures on durable goods can be temporarily reduced [Browning and Crossley, 2009]. Through their impact on the curvature in consumption preferences the importance of specific consumption categories will also affect the MPC. However, following Lemma 1, the implementation of the MPC approach would only be affected if the ratio of MPCs across states changes. We illustrate this in Appendix A.2 by characterizing the MPC approach with two types of consumption categories, i.e., $v_s(c_s) = g(c_s^1) + h(c_s^2)$.

In general, a reason why the MPC approach is arguably more robust compared to the CB approach is that some key confounders for the CB approach do not affect the relationship between the MPC and the state-specific price, and if they do, they may well do so in a similar way when employed and unemployed and have limited impact on the ratio of state-specific MPCs.

State-Specific Preferences A common challenge when assessing the value of UI for both the MPC and CB approach is the presence of state-specific preferences. It is important to make a distinction between state-specific preferences over consumption and resources.

State-specific resource preferences (e.g., $\frac{\partial v_s}{\partial x} = \theta_s^x \frac{\partial v}{\partial x}$) have no impact on the CB approach, but affect the wedge between the MRS and the price-ratio, as shown in Lemma 2. Take again the example of spousal labor supply. When spousal labor supply is increased during unemployment, we

[7] As we discuss in the Appendix, the MRS should be scaled by the ratio $\frac{\eta_u}{\eta_e}$ to estimate the value of UI transfers in order to capture the relative efficiency with which a krona is spent when unemployed vs. employed. In the case of home production or lower shopping prices, we expect $\eta_u > \eta_e$, so that the lower-bound argument on the value of UI continues to hold.

expect this marginal resource cost to be higher. However, state-specific preferences can decrease the resource costs when unemployed vs. employed when for example partners' leisure time are substitutes (i.e., $\theta_u^x < \theta_e^x$). In this case, state-specific preferences reduce the wedge in resource costs and can in principle reverse it when they are strong enough, so that Assumption 2 is violated. The second part of the lower-bound argument then no longer holds.

State-specific consumption preferences (e.g., $\frac{\partial v_s}{\partial c} = \theta_s^c \frac{\partial v}{\partial c}$) directly affect the MRS itself. Optimizing individuals continue to set it equal to its cost, accounting for the price and the disutility of the resources used. So state-specific consumption preferences do not overturn the lower-bound argument, but they can still weaken the lower bound itself (when $\theta_u^c > \theta_e^c$).[8] State-specific consumption preferences also affect the relationship between the MRS and the observed consumption drop, which is central for the CB implementation. Applying the same Taylor approximation with state-specific consumption preferences, we get

$$MRS \cong \frac{\theta_u^c}{\theta_e^c} \left\{1 + \sigma^c \times [c_e - c_u]\right\}. \tag{13}$$

Hence, information on state-dependence becomes essential to translate the observed drop in consumption expenditures into an estimate of the MRS.

Empirical Test We can gauge the importance of state-specific consumption expenditures or preferences by combining the empirical moments from the CB and MPC approaches. Assuming $\frac{\partial v_s}{\partial c} = \theta_s^c v'(c_s - \phi_s)$, we can approximate the drop in consumption upon job loss by differentiating equation (3) with respect to income and accounting for the change in state-specific expenditures/preferences:

$$\Delta c \cong \frac{dc_u}{dy_u}\Delta y + \left[1 - \frac{dc_u}{dy_u}\right]\left[\frac{1}{\sigma^c}\frac{\theta_e^c - \theta_u^c}{\theta_u^c} + [\phi_e - \phi_u]\right]. \tag{14}$$

Hence, we can test for the presence of state-dependence by comparing the observed drop in consumption expenditures upon job loss $\Delta c = c_e - c_u$ to the imputed drop in consumption, based on the drop in disposable income $\Delta y = y_e - y_u$ and the estimated MPC $\frac{dc_u}{dy_u}$.[9] If the observed drop is larger than the imputed drop, either the employment-specific expenditures are higher than the unemployment-specific expenditures, $\phi_e > \phi_u$, or a worker values unemployment consumption less than employment consumption, $\theta_e^c > \theta_u^c$. In both cases, a naive implementation of the CB approach using equation (6) would overstate the MRS. The MPC approach, however, continues to provide a lower bound on the MRS.

[8] Note that the lower-bound argument still relies on our premise that workers have a preference for smoothing consumption relative to the income loss when unemployed. In theory, with state-specific consumption preferences, the marginal utility of consumption when unemployed could be so low that the opposite is true. However, this would be an extreme case, in which more resources are used to further widen the wedge in consumption relative to the wedge in earnings between employment and unemployment, which we deem implausible.

[9] Note that the original CB implementation in Gruber [1997] did not consider the drop in consumption upon job loss as such, but in fact estimated the consumption response to UI benefits to then impute the consumption drop due to the net-income lost when unemployed.

3.3 Comparison to RP Approach

The central idea of the MPC approach is that the MPCs identify the shadow price of income when unemployed vs. employed to then bound the value of extra consumption when unemployed vs. employed. In principle, when the price of transferring consumption between unemployment and employment is known, we can bound workers' valuations *directly* using an RP approach. The Swedish unemployment policy offers a UI choice that allows us to do exactly that and provide an alternative RP estimate of the MRS.

Revealed Preference Approach Our stylized model can easily be extended to incorporate insurance choice. The resource cost of increasing consumption when unemployed is then simply to lower consumption when employed. In particular, consider the possibility to get extra UI coverage at rate p_u/p_e, which can be interpreted as the state-specific prices of Arrow-Debreu securities. A worker will buy the extra UI coverage only if

$$\pi \frac{\partial v_u}{\partial c} \frac{1}{p_u} + (1-\pi) \frac{\partial v_e}{\partial c} \frac{1}{p_e} \geq 0. \tag{15}$$

Importantly, the willingness to take up the extra insurance does not only depend on the relative price p_u/p_e, but also on the worker's unemployment risk π. Indeed, the expected price per unit of coverage is lower the higher one's unemployment risk, $\tilde{p} \equiv \frac{p_u}{p_e} \times \frac{1-\pi}{\pi}$. Note that workers may well reduce their efforts and other consumption smoothing behavior when taking up extra insurance, but these changes are of second-order importance for the decision to take up the extra insurance. Hence, a worker will only take up the extra UI coverage if the MRS exceeds the expected price,

$$MRS \geq \tilde{p} \equiv \frac{p_u}{p_e} \times \frac{1-\pi}{\pi}, \tag{16}$$

and vice versa.[10]

The RP approach is the most direct approach to estimate the MRS and in principle allows us to relax any parametric assumptions on preferences. Moreover, when exogenous variation in expected prices is available, we can go beyond bounding the MRS and uncover the MRS distribution in the population. However, the RP approach adds two very strong data requirements relative to the MPC approach. First, it requires the availability of insurance choices, which are often absent in the context of social insurance and originally motivated the CB approach.[11] Second, it requires information on individuals' risk to estimate the expected price of coverage determining their choice. The RP approach also extends to discrete insurance choices observed in practice, but this entails two extra caveats. First, for discrete insurance choices moral hazard responses may no longer

[10] In an Arrow-Debreu market, individuals would buy coverage up to the point that the MRS equals the expected price. This corresponds to condition (8) with $\frac{\partial v_u/\partial x}{\partial v_e/\partial x} = \frac{\pi_e}{\pi_u}$. With actuarially fair pricing ($p_s = \pi_s$), a risk-averse individual (with otherwise state-independent preferences) would fully insure the income risk she faces (Arrow [1963], Mossin [1968]).

[11] Exceptions are for example Cabral and Cullen [2016] in the context of long-term disability insurance and Finkelstein et al. [2017] in the context of Medicaid.

be of second-order importance. As a result, we would underestimate the MRS when predicting unemployment risk conditional on buying the extra coverage and overestimate the MRS when predicting it conditional on not buying it. Second, we are no longer identifying the MRS at the margin, but an average MRS for the supplemental coverage workers can buy. Since for risk-averse workers the MRS is decreasing in the coverage level, the average MRS would overestimate the MRS evaluated at the margin for workers who buy the insurance and underestimate it for workers who don't.[12]

Optimization Approaches We briefly characterize the recently proposed approaches based on behavioral responses that *indirectly* reveal the MRS in the context of our stylized model. The underlying derivations are provided in Appendix A.3.

Closely related to the MPC approach are the approaches proposed by Chetty [2008] and Landais [2015], who study responses in unemployment risk (instead of consumption) and consider the differential response to changes in unemployment benefits relative to other income changes. Like the MPC approach, these differential unemployment responses allow inferring the relative marginal utility of consumption when unemployed and employed without requiring information on the curvature of consumption preferences (as needed for the CB approach). In our stylized model, from implicit differentiation of the optimality condition for job search effort z, we find

$$\frac{\frac{d\pi(z)}{dy_u}}{\frac{d\pi(z)}{dy_e}} = MRS \times \frac{2\frac{dc_u}{dy_u} - 1}{2\frac{dc_e}{dy_e} - 1}. \tag{17}$$

When assuming that UI does not affect other consumption smoothing means (i.e., $\frac{dc_s}{dy_s} = 1$) like in Chetty [2008] and Landais [2015], the differential unemployment response exactly identifies the MRS. This continues to be true when the MPCs are equal when unemployed and employed. However, when the consumption response is larger when unemployed than when employed (i.e., $\frac{dc_u}{dy_u} > \frac{dc_e}{dy_e}$), the differential unemployment response provides an upperbound on the MRS.

Closer to the CB approach, Fadlon and Nielsen [2018] and Hendren [2017] show how instead of looking at wedges in consumption, one can also look at wedges in resources used when employed and unemployed, like for example changes in spousal labor supply. If the particular resource is optimized at the margin, the wedge allows to estimate the marginal value of transfers when unemployed and employed. In our stylized model, assuming $\frac{\partial v_s(c,x)}{\partial x} = \theta_s^x \frac{\partial v(c,x)}{\partial x}$ and using the optimality condition 3, we can approximate

$$MRS \cong \frac{p_u}{p_e} \frac{\theta_u^x}{\theta_e^x} \left\{1 + \sigma^x \times [x_u - x_e]\right\} \tag{18}$$

Importantly, this approximation still requires information on the curvature and/or state-dependence in preferences, but now not over consumption but over the resources used. Moreover, simply scaling the wedge in resources with the preference parameters would under-estimate the MRS when the

[12]We formally develop these arguments in Appendix A.4 and come back to the implementation issues in Section 7.

price of increasing consumption is higher when unemployed than employed (i.e., $p_u > p_e$).[13]

As shown in the context our stylized model, the MPC approach does not rely on whether actions z are undertaken to reduce the unemployment risk or what type of resource x is used to smooth consumption. An important advantage of the MPC approach relative to the other optimization approaches is therefore that - as long as the assumptions in Proposition 1 are satisfied - it does not require to take a stance on the actions taken to reduce the unemployment risk or the resources used to smooth consumption. By construction the MPC reveals the shadow price of increasing consumption given the resource used at the margin. Still, any information on differential unemployment responses or wedges in resources could in principle be used in combination with the MPCs to obtain more precise estimates of the MRS.[14]

Behavioral Frictions Importantly, all revealed preference approaches rely on the optimization by workers and the absence of behavioral frictions. A growing literature, however, documents the importance of behavioral frictions in insurance choices [e.g., Abaluck and Gruber [2011], Handel [2013], Handel and Kolstad [2015]], but also specifically in consumption and job search choices in the context of unemployment [e.g., Ganong and Noel [2017], Gerard and Naritomi [2019]]. These frictions would be wrongly attributed to workers' valuation of insurance and thus bias our estimates of the MRS. For the implementation of our RP approach, rather than knowing a worker's unemployment risk, we need to know the worker's perception of the unemployment risk at the time he or she decides to buy the insurance. Additional data (or assumptions) are required to deal with this challenge, as we discuss further in Section 7. For the MPC approach, we can invoke a similar robustness argument as before, which applies when behavioral frictions affect the MPC in a similar way when employed and unemployed.[15] The CB approach is arguably more robust to behavioral frictions, although their presence may further complicate the estimation of the relevant preference parameters.

It is important to note that beyond confounding the estimation of the MRS, frictions may cause the MRS to be no longer sufficient for determining the welfare gain from more generous unemployment benefits, as shown in Spinnewijn [2015] in the context of biased beliefs. An evaluation of the unemployment policy may then require the estimation of these behavioral frictions and call for different interventions that target these frictions directly.

[13] Note that Hendren [2017] proposes to use responses (in consumption or other resources) to changes in the anticipated unemployment risk (rather than comparing the wedge between employment and unemployment), which also requires information on the curvature of preferences, but is robust to state-specific preferences (or expenditures).

[14] We elaborate further on this in Appendix A.2, where we consider an extension of our model with multiple resources and show how the MPC approach indeed depends on the price of the resource used at the margin, but also on potential changes in this price when using more resources.

[15] For example, when workers under-estimate the resource cost to increase their consumption, but do so in the same way when employed vs. unemployed.

4 Empirical Implementation: Data and Context

We exploit the unique institutional setting and data in Sweden to estimate the value of UI using the MPC approach, and comparing it to the CB approach and the RP approach in the exact same context on the same sample of individuals.

The three approaches offer complementary ways of estimating the value of insurance. They differ not only in their theoretical assumptions, but also in their data and empirical implementation requirements. While the CB approach only requires precise information on consumption and employment status, the MPC approach requires in addition exogenous variation in income, when employed and unemployed, to study how consumption responds to changes in income in both states. The RP approach requires information on UI choices and unemployment risk. We note that the most fruitful approach will therefore depend on the data available and the implementation assumptions required given the specific context.

We present here the institutional background and data allowing us, by merging data from various registers in Sweden, to implement all three approaches on the same sample of workers.

4.1 Unemployment Insurance Policy

Sweden is with Iceland, Denmark and Finland, one of the only four countries in the world to have a voluntary UI scheme derived from the "Ghent system". The existence of choice in the Swedish UI system means we have the rare ability to implement the RP approach to identify the value of insurance.

The Swedish unemployment insurance system offers two levels of coverages in case of unemployment: a basic and a comprehensive coverage. To be eligible to receive any benefit, displaced workers must have worked for at least six months prior to being laid-off. The basic plan works like a minimum mandate. It provides a low coverage, i.e. a floor of $320SEK$ a day ($\approx 35USD$), regardless of the level of pre-unemployment earnings. Benefits for the basic coverage are funded through payroll taxes paid by all workers. Workers can opt in for a comprehensive UI benefit coverage. Under this comprehensive plan, a worker gets 80% of their earnings replaced, up to a cap. In practice, the level of the cap applies to about 50% of unemployed workers and the average unemployment benefit received under the comprehensive plan is twice the benefit level provided by the basic plan. Workers are free to opt in or out of the comprehensive UI plan at any time, but need to have been contributing for 12 consecutive months at the start of their unemployment spell to be eligible to receive the additional coverage.[16] During our period of study, the UI premium for the additional coverage was set uniformly at $100SEK$ a month, and around 85% of all workers were

[16] Overall, the potential for frictions associated with opting in and out seems limited. Enrolling in the supplemental UI coverage is done by filling out a short form, which can be obtained online or in direct contact with the UI funds. The premium is paid monthly and enrolled members can select between receiving monthly invoices or paying via direct debit. To opt out of the plan, individuals can fill out a form, analogous to the procedure of opting in, or stop paying their premia for three months. There are no waiting periods associated with opting in or out, and the processing time for such requests are typically limited to a few days.

contributing to an unemployment fund to get the comprehensive coverage.[17] Further institutional details underlying the UI choice are provided in Landais et al. [2017].

4.2 Data & Sample

We use administrative data from various Swedish registers that can be linked and cover the universe of Swedish individuals.

Registry-based Consumption Measure To measure consumption, we use the registry-based measure of annual household consumption expenditures for the universe of Swedish households created for all years 2000 to 2007 by Kolsrud et al. [2017]. The construction of this measure is based on the identity coming from the household's budget constraint between consumption expenditures and income net of changes in assets. The quality of this measure relies on the comprehensiveness of income and asset data in Sweden.[18] The longitudinal dataset LISA contains exhaustive disaggregated information on all earnings, all taxes and transfers and capital income on an annual basis. Data on wealth comes from the wealth tax register (Förmögenhetsregistret), which covers the asset portfolios for the universe of Swedish individuals with detailed information on the stock of all financial assets (including debt) and real assets as of December of each year. Data on asset balances is complemented with data on financial asset transactions (KURU) and real estate transactions from the housing registries. We refer the reader to Kolsrud et al. [2017] for further details on the construction of this measure and assessment of its robustness and consistency compared to survey measures of expenditures.[19]

Data on Unemployment (Risk) Unemployment data comes from the HÄNDEL register of the Public Employment Service (PES, Arbetsförmedlingen), with records for the universe of unemployment spells from 1990 to 2015, and were merged with the ASTAT register from the UI administration (IAF, Inspektionen för Arbetslöshetsförsäkringen) in Sweden. The data contain information on the date the unemployed registered with the PES (which is a pre-requisite to start receiving UI benefits), eligibility to receive UI benefits (for both the basic and comprehensive coverage), earnings used to determine UI benefits, weekly information on benefits received,

[17]The administration of the comprehensive UI coverage is done by 27 UI funds (*Kassa's*) but the government, through the Swedish Unemployment Insurance Board (IAF), supervises and coordinates the entire UI system. Both the premia and benefit levels of the basic and comprehensive coverage are fully determined by the government.

[18]See also Browning and Leth-Petersen [2003]; Koijen et al. [2014]; Eika et al. [2017], for similar applications of the registry-based measure of consumption in Sweden.

[19]We note again that our imputed measure of consumption is capturing, like most survey measures, expenditures rather than consumption, which poses a particular challenge for the CB approach as discussed before. Our registry-based measure of consumption may feature other measurement issues, discussed in Kolsrud et al. [2017]. In particular we do not observe informal transfers. Using data from the ULF survey, which records informal transfers between family and friends, we found that these transfers are small, and do not happen more frequently upon job loss (see Figure B-7). All details on the data and programs used to create this measure of consumption can also be found at: `http://sticerd.lse.ac.uk/_new/research/pep/consumption/default.asp`.

unemployment status and participation in labor market programs.[20]

We use three types of data on workers' unemployment risk. First, we have data on realized unemployment risk from the unemployment history data. Second, we have rich data on the determinants of unemployment risk. This data comes from information on demographics (age, gender, marital status and family composition, education, place of work, industry, occupation) from the LISA register. We use two additional labor market registers that reveal important information about unemployment risk. The matched employer-employee register (RAMS), from 1985 to 2015, reports monthly earnings for the universe of individuals employed in establishments of firms operating in Sweden. We use this register to compute tenure and tenure ranking for each employee as well as firm level unemployment risk. We also use the layoff-notification register (VARSEL) which records, for years 2002 to 2012, all layoff notifications emitted by firms. Following Landais et al. [2017], we flexibly combine all observable sources of risk variation together in a comprehensive prediction model of individuals' unemployment risk. We model the total number of days unemployed in $t+1$ based on observable risk determinants in year t, using a zero-inflated Poisson model. The model includes all the demographic characteristics plus the average firm layoff risk, the full history of the firm layoff notifications, and the relative tenure ranking of the individual. We start by including a rich set of interactions between all the variables, and optimize our model using a forward stepwise selection algorithm. From this model we get a yearly individual measure of predicted unemployment risk for each worker, which we can then use to back out the MRS in the RP approach, using equation 16. Finally, we have data on elicited beliefs coming from the Household Market and Nonmarket Activities (HUS) panel survey. The HUS survey asks individuals questions that provide information on perceived unemployment risk. We exploit responses to the question "How likely is it that you will keep your current job next year?", with the answer ranging form 0% to 100%, and compare the elicited beliefs to the realized outcomes.

Data on UI choices We finally use UI fund membership information for the universe of workers in Sweden aged 18 and above, from 2002 to 2009, based on two distinct sources. The first source is tax data for the period 2002 to 2006, during which workers paying UI premia received a 40% tax credit. This source records the total amount of UI premia paid for each year. From this source, we define a dummy variable for buying the comprehensive coverage in year t as reporting any positive amount of premia paid in year t. We combine this data with a second source of information, coming from UI fund data that *Kassa's* sent to the IAF. This data contains a dummy variable indicating whether an individual aged 18 and above in Sweden is contributing premia for the comprehensive coverage as of December of each year from 2005 until 2009.

Main Sample of Analysis We create a sample of individuals for which all three approaches (CB, MPC and RP) can be implemented. This enables us to compare the valuations of UI implied

[20]We define unemployment as a spell of non-employment, following an involuntary job loss, and during which an individual has zero earnings, receives unemployment benefits and reports searching for a full time job (see Kolsrud et al. [2018]).

by these approaches not only in the same context, but on the very same individuals. Our baseline sample is composed of individuals who experience a first unemployment spell between 2002 and 2007. To create the sample, we start from the universe of layoffs in the PES data for years 1990-2015. We only consider the first layoff observed in this period per individual. We restrict the sample to individuals who are aged 25-55 at the time of layoff and who are eligible for any UI coverage (basic or comprehensive) according to the 6 months work requirement prior to being laid-off. We further restrict the sample to individuals who are unemployed in December in the year of being laid off, as this is the month when all other demographics, income, tax and wealth information are observed and reported in the registry data. Consumption is at the household level, where we fix composition of the household as of event time -1, the year prior to being laid off. We exclude households where more than one member experiences an unemployment spell between 2002 and 2007. This leaves us with a baseline sample containing 164,248 individuals experiencing their first unemployment shock between 2002 and 2007, matched with all other members of their households.

In Table 1, we provide summary statistics on demographics, income and wealth, and unemployment details for individuals in this baseline sample. All statistics are computed in the year prior to the start of their unemployment spell. The table shows that most individuals in our sample have relatively few means of smoothing consumption. Most of them enter unemployment with close to zero net wealth, high levels of debt, and little liquid assets as a fraction of their annual household consumption.[21] But most individuals (91%) in our sample are contributing to the comprehensive coverage prior to job loss.

5 Consumption-Based Approach

Before implementing our new MPC approach, we start by revisiting the consumption-based implementation in the Swedish context. This will provide us with a benchmark estimate to compare our alternative implementations to.

The CB implementation relies solely on the estimation of the relative consumption drop at unemployment $\frac{\Delta c}{c}$. In practice, we identify this consumption drop using an event study strategy, based on the following model:

$$C_{it} = \alpha_i + \nu_t + \sum_{j=-N_0}^{N_1} \beta_j \cdot \mathbb{1}[J_{it} = j] + \varepsilon_{it} \tag{19}$$

where households are indexed by i and $t = 1, .., T$ denote the calendar year of observation. α_i is a household fixed-effect and ν_t is a time effect. $J_{it} = t - E_{it}$ denotes event time, that is the time in year relative to the occurrence of job loss. $[-N_0; N_1]$ is a window of dynamic effects around the unemployment treatment event. The treatment group is composed of all the households from our baseline sample described in section 4.2 above. We follow the recent literature on event studies

[21] Liquid assets are total household bank holdings in liquid accounts. Debt is total household debt including student loans.

[Borusyak and Jaravel, 2016; Kolsrud et al., 2017; Freyaldenhoven et al., 2018], and introduce a control group that never experiences treatment. This control group, created using nearest-neighbor matching based on pre-event characteristics, allows for the identification of time effects ν_t independently of the dynamic treatment effect of the event $\{\beta_j\}_{j=-N_0}^{N_1}$.[22]

We estimate specification (30) using fixed-effect regressions and taking event time $t = -1$ as the reference category. Figure 1 plots our estimates for the event time dummies $\hat{\beta}_j$, scaled by the average predicted consumption in $t = -1$, $\widehat{C}_{-1}$, from specification (30). In line with existing evidence, the graph shows that consumption expenditures experience a significant drop at job loss of around 10%. This drop is persistent over time: 5 years after layoff, consumption expenditures are still about 10% lower than their pre-unemployment level.

The figure displays how *annual* consumption evolves around the unemployment event. The implementation of the CB approach requires that we translate these estimates into measures of the *flow* drop in consumption when unemployed $c_u - c_e$ (see Kolsrud et al. [2018]). For this purpose, we adopt a simple parametric approach, and use the fact that in event year 0, individuals are all observed unemployed in December, but differ in the time in months M_i they have spent unemployed in that year. An individual having spent M_i month unemployed in year 0 will have an annual consumption in year 0 equal to $(12 - M_i) \cdot c_e + M_i \cdot c_u$, and an annual consumption in year -1 equal to $12 \cdot c_e$. The change in annual consumption between year -1 and year 0 is equal to $M_i \cdot (c_u - c_e)$, and therefore a linear function of the number of months spent unemployed in year 0.

We start by illustrating non-parametrically how time spent unemployed in event year 0 relates to annual consumption drops in year 0. We split the sample in 6 bins of M_i, and estimate specification (30) for each group. Appendix Figure B-2 reports the estimates $\hat{\beta}_0/\hat{C}_{-1}$ of the percentage drop in annual consumption in year 0 for each bin of M_i. The graph reveals that the relationship between time spent unemployed in year 0 and the annual drop in consumption in year 0 is well approximated by a linear line with intercept at zero, which corresponds to our simple parametric model. Based on this evidence, we can estimate the following specification:[23]

$$C_{it} = \alpha_i + \nu_t + \sum_{\substack{j=-N_0 \\ j\neq 0}}^{N_1} \beta_j \cdot \mathbb{1}[J_{it} = j] + \beta_0 \cdot M_i \cdot \mathbb{1}[J_{it} = 0] + \varepsilon_{it} \tag{20}$$

In the above regression, the flow drop in consumption at unemployment $\frac{c_u - c_e}{c_e}$ is identified by

[22] We adopt the following matching strategy. For each calendar year t, we take all individuals who receive the event in that particular year ($E_{it} = t$), and find a nearest neighbor from the sample of all individuals who never receive treatment. Individuals are matched exactly on age, gender, region of residence in $t-1$ (21 cells), level of education in $t-1$ (10 cells) and family structure in $t-1$ (12 cells), and by propensity score on their number of dependent children in $t-1$, 12 industry dummies in $t-1$ and their earnings in $t-1$, $t-2$ and $t-3$. Appendix Figure B-1 displays the evolution of household consumption around event time for both control and treatment groups.

[23] We note that the consumption profile in Kolsrud et al. [2018] decreases less and is more concave, which corresponds to a smaller drop in consumption, especially early in the unemployment spell. The sample in Kolsrud et al. [2018] only includes workers with comprehensive coverage and pre-unemployment earnings around the comprehensive benefit kink. Worker who only have basic coverage tend to have shorter unemployment spells and are thus over-represented early in the spell.

$\frac{12 \cdot \hat{\beta}_0}{\widehat{C}_{-1}}$. This approach provides us with a baseline estimate of $-.129$ (.028), as reported in Figure 1. Available estimates in the literature find average consumption drops at unemployment of 5 to 12% (e.g. Gruber [1997], Browning and Crossley [2001], Ganong and Noel [2017]). Our estimate is at the upper end of the spectrum but otherwise very comparable to the existing evidence. The finding of a relatively moderate drop of around 10% in household consumption expenditure at unemployment therefore seems very robust across contexts and sources of expenditure measurement.

MRS estimates We can scale the estimated consumption drop by the curvature in preferences over consumption to get an estimate of $MRS \cong 1 + \gamma \times \frac{\Delta c}{c}$ with $\gamma = \sigma c$ equal to the relative risk aversion. This follows the standard CB implementation, assuming state-independence in marginal utilities. We report estimates of the MRS for γ ranging from 1 ($MRS = 1.13$) to 4 ($MRS = 1.51$) in Figure 1, which are values commonly used in the literature. The bottom-line of the CB implementation is that given the relatively small drop in consumption, the value of a marginal krona of insurance is small as well, even for presumably high levels of risk aversion.

As noted by Gruber [1997] and Hendren [2017], the validity of the standard CB implementation depends on the extent to which job losses are anticipated. The consumption drop at job loss understates the insurance value when the job loss has been anticipated and workers have taken precautionary actions as a result. To gauge the severity of the issue, we start by studying how much individual unemployment risk gets revealed through changes in observables in the years prior to job loss. We report in Appendix Figure B-3 estimates from specification (30) where we use as the outcome our measure of predicted unemployment risk, based on our rich model of observable determinants of unemployment risk in Sweden. The graph shows a significant yet quite modest increase in the predicted unemployment risk measure in the two years prior to layoff. Following Hendren [2017] we can then relate this change in risk to the change in consumption in the two years prior to job loss, and obtain an alternative measure of the MRS from anticipatory behaviors alone. For a risk aversion parameter $\gamma = 1$, our implementation gives a large MRS, of about 2.1, but very imprecisely estimated, with a 95% confidence interval spanning MRS values from 0 to 5. This lack of precision is due to the small magnitude of anticipation of job loss in our context, both in terms of underlying risk, and in terms of anticipatory consumption changes.

Mechanisms We can leverage the granularity of our data to explore the means used to smooth consumption at job loss and reveal further information about the potential value of unemployment insurance. Appendix Figure B-4 decomposes the consumption expenditures of households into various components (transfers and taxes, spousal earnings, consumption out of asset and consumption out of debt). We find that most of the consumption smoothing is done by transfers. On net, assets play only a limited role, reducing the drop in consumption by only 2%. While household do draw down their assets when unemployed, consumption out of debt decreases and thus contributes negatively to consumption smoothing, indicative of tighter borrowing constraints during unemployment. Also spousal earnings do not significantly contribute to consumption smoothing. Appendix Figure

B-5 documents heterogeneity in consumption responses to unemployment. We find that consumption drops are particularly sensitive to the level of liquid wealth and the generosity of transfers received. We provide more detail on the regression specifications and results in Appendix B.3.

While not conclusive, the combined evidence suggests that liquidity and borrowing constraints bind for a significant fraction of unemployed, indicating a high shadow price of smoothing consumption after job loss. In the next section, we analyze the marginal propensities to consume exactly with the aim of identifying this price.

6 MPC Approach

The MPC approach relies on identifying and comparing the marginal propensity to consume out of income when employed and when unemployed. This requires a source of exogenous variation in income, that applies similarly both to employed and unemployed people. We leverage the existence of significant variation in local welfare transfers in our context.

6.1 Local Welfare Transfers in Sweden

In Sweden, municipalities are responsible for an important fraction of total welfare transfers called social assistance ("*social bidrag*"). Social assistance is regulated by federal law, but the interpretation and enactment is delegated to each municipality. By federal law, the function of *social bidrag* is to offer an income guarantee, potentially depending on household income and assets and a restricted set of other household characteristics. The National Board of Health and Welfare (*Socialstyrelsen*) provides recommendations on the level of the income guarantee for different household types (defined by the number of adults and the number and age of children), that operate as effective minima.[24] However, it is up to the municipalities to set the exact level of this guarantee, and the precise conditionality and means testing attached to it.

The transfers received by household i with income y_{it}, liquid assets a_{it} and characteristics $\mathbf{X}_{it}$ living in municipality m at time t can be described as $B_{imt} = G_{mt}(\mathbf{X}_{it}) - \tau^y_{mt}(y_{it}) - \tau^a_{mt}(a_{it})$, where municipalities set how the level of the income guarantee G varies with observable characteristics $\mathbf{X}$ and the phase-out rates τ^y and τ^a.[25] The characteristics $\mathbf{X}$, restricted by federal law, are the composition of the household (the number of dependents and the age of the dependent children), and whether children are at school and get free lunch at school. Hence, in a given municipality and year t, the level of welfare transfers received by a household is fully determined by its vector

[24] In 2019, the minimum level of the income guarantee for a single individual without children is 3090 SEK per month, and 5570 SEK per month for a married couple without children. This minimum is increased further depending on the number and age of the children in the household. For example, each child under age 1 in the household increases the minimum guarantee by 2130 SEK/month, while each child aged 11-14 years increases the minimum guarantee by 3440 SEK/month. The details of the National Board of Health and Welfare recommendations with the tables for the minima by family types can be found here: `https://www.socialstyrelsen.se/hittarattmyndighet/ekonomisktbistand/riksnormen`.

[25] Note that in practice the phase-out rates τ^y and τ^a can be non-linear functions of income and assets. We account for this by entering both income and wealth non-parametrically (using deciles) instead of linearly. Note also that only liquid assets (not real estate wealth) are taken into account in the benefit formula.

of characteristics $\mathbf{V}_{it} = \{\mathbf{X}_{it}, y_{it}, a_{it}\}$:

$$B_{imt} = \sum_k \tau^k_{mt} \cdot V^k_{it}, \tag{21}$$

where V^k_{it} is the k-th component of vector $\mathbf{V}_{it}$ and τ^k_{mt} represents how the schedule of welfare transfers depends on this characteristic in municipality m at time t. In practice, we include in $\mathbf{V}_{it}$ marital and cohabitation status of the household head, dummies for the number of adults in the household, dummies for the number of children in the household and their age, and dummies for the decile of disposable income (excluding local transfers) and for the decile of net liquid assets of the household.

In our sample, 13% of households receive positive *social bidrag*, and the average annual transfer per household conditional on receipt is 24,981 SEK2003. Because of the discretion given to municipalities, there is a significant amount of variation in the generosity of local welfare transfers across municipalities. This is illustrated in Appendix Figure C-1. To control for compositional differences across municipalities, we residualize transfers B_{imt} received by household i in year t on the vector of observable characteristics $\mathbf{V}_{it}$,

$$B_{imt} = \sum_k \bar{\tau}^k \cdot V^k_{it} + \tilde{B}_{imt}.$$

The figure then plots the average residualized transfer $\tilde{B}_{imt}$ in each municipality over the period 2000-2007. The map shows a large amount of variation in the average residual generosity of welfare transfers between municipalities. For example, the urban municipalities in Stockholm, Gothenburg or Malmö in the South, but also some less populated municipalities in the North are significantly more generous. Of course, this variation in average generosity may reflect some endogenous policy choice in the municipalities in relation to differences in the cost of living or differences in unobserved characteristics of its inhabitants.

Importantly for our purpose, there is also a significant amount of residual variation in transfers within households within municipalities. This variation stems from two sources. First, municipalities set the τ^k_m from formula (21) differently (i.e., the functions $G_m(\mathbf{X})$, $\tau^y_m(y)$ and $\tau^a_m(a)$ are different across municipalites). Therefore when the characteristics $\mathbf{V}$ of a household change, for instance because a child in the family gets older, or income changes, this will trigger different adjustments in B across different municipalities. This within-household variation in transfers provides an opportunity to identify the MPC. The intuition for identification is the following. Take two families with identical characteristics $\mathbf{V}$, one is living in municipality m and the other in municipality m'. Say for instance they are married and have one child of age 10 in year t. In year $t+1$, the child turns 11. This will trigger different variation in B between t and $t+1$ in m and m' because of differences in τ^{age} (i.e. the way G depends on the age of children) between m and m'. The identifying assumption is that differences in τ^k (i.e. τ^{age} in this case) are not correlated with other unobserved heterogeneity across municipalities that differentially affects consumption depending on the child's age. The second source of within household variation stems from variations in τ_m^k over time within municipalities due to local reforms. Here the identifying assumption is that the reform of social transfers is not implemented in response to changes at the municipality level that are correlated with household consumption, nor jointly with other local reforms directly affecting household consumption.[26]

Figure 2 illustrates these rich sources of identifying variation, showing how τ^k, for specific household characteristics V^k of vector $\mathbf{V}$, differs across municipalities. To visualize this, we focus on year 2006, and residualize welfare transfers of individual i in municipality m using the following specification (where we dropped the t subscript):

$$B_{im} = \sum_{j \neq k} \tau_m^j \cdot V_i^j + \bar{\tau}^k V_i^k + \tilde{B}_{im}.$$

We then plot in a map the statistics $\mathbf{E}[\tilde{B}_{im}|V^k = v^k] - \mathbf{E}[\tilde{B}_{im}|V^k = v'^k]$ for each municipality m. From formula (21) which defines welfare transfers B_{im}, we have that $\tilde{B}_{im} = (\tau_m^k - \bar{\tau}^k)V_i^k$. If municipalities set the same τ^k, then the statistics should be equal to zero in all municipality. Differences in these statistics across municipalities reflect the fact that τ^k are set differently across municipalities.

In panel A of Figure 2, we start by showing differences in the way municipalities set $\tau^{children}$, the generosity of B as a function of the number of dependent children. The map shows, for all municipalities, the difference in average residual benefits $\tilde{B}_{im}$ in thousands of SEK for a household with 2 children vs 3 or more children. There is significant variation in $\tau^{children}$: some municipalities give significantly more (up to SEK20k per year) in welfare benefits for the arrival of a third child, everything else equal. Panel B shows variation in τ^{age}, the generosity of B as a function of the age of dependent children. The map shows, for all municipalities, the difference in average residual benefits $\tilde{B}_{im}$ in thousands of SEK for household with similar structure and number of dependents, whose youngest child is between 4 to 6 years old versus 11 to 15 years old. There is again significant variation in τ^{age}: some municipalities give significantly more (up to SEK20k per year) in welfare benefits for older children compared to younger children, everything else equal. Panel C illustrates the significant variation in the income phase-out rate τ^y of welfare transfers for households of similar structure. It plots for all municipalities the difference in average residual benefits $\tilde{B}_{im}$ in thousands of SEK for similar household with income in the bottom quintile vs the second quintile of the household income distribution. In some municipalities, this increase in income gets taxed at a high marginal rate, while in other municipalities, there is almost no change in transfers. Interestingly, the maps of the residual variation in $\tau^{children}$, τ^{age} and τ^y exhibit significant differences. For example, municipalities that are more generous for larger families are not necessarily the ones with the lower income phase-out rates.

[26] Note that part of this variation seems to be driven by electoral changes in local political majorities in municipalities. There is indeed ample anecdotal evidence that social-democrats favor increasing the generosity of social bidrag transfers when controlling a municipality, while the center-right parties encourage reductions in local welfare transfer generosity.

In panel D, we also provide evidence of the significant geographical variation in the evolution of the welfare benefits schedule over time. The panel plots the growth rate of residualized transfers $\tilde{B}_{imt}$ between 2000 and 2007 across municipalities. For this purpose, we residualized transfers according to the following specification $B_{imt} = \sum_k \tau_m^k \cdot V_{it}^k + \nu_t^0 + \tilde{B}_{imt}$, where ν_t^0 are year fixed-effects. We then plot $\mathbf{E}[\tilde{B}_{imt}|t = 2007] - \mathbf{E}[\tilde{B}_{imt}|t = 2000]$ for each municipality m. Note that we also reweight each observation so as to keep the distribution of characteristics V_{it}^k fixed in each municipality. The map suggests significant variation within municipality over time in the generosity of welfare transfers.

6.2 MPC: Implementation

Our strategy to identify marginal propensities to consume is to use within-municipality within-household variation in local welfare transfers stemming from variation, documented above, in the ways municipalities set the schedule of their transfers as a function of characteristics $\mathbf{V}$ across household types and over time. We keep for this analysis the same sample as the one used for the CB approach in Section 5 above. The sample contains only individuals who become unemployed at some point, and are observed either employed (prior to their unemployment spell) or unemployed (during their unemployment spell). The goal is to compare their MPC out of local transfers when they are employed versus when they are unemployed. The strength of our approach is to estimate MPC in both states on the same individuals in the same sample using a unique source of variation in transfers in both states.[27]

We estimate the following specification in first-differences to control for household fixed-effects α_i:

$$C_{imt} = \alpha_i + \nu_t + \eta_m + \mathbf{V}_{it}'\beta + \mu_e \cdot \tilde{B}_{imt} + \mu_u \cdot \tilde{B}_{imt} \cdot \frac{M_i}{12} \cdot \mathbb{1}[J_{it} = 0] + \eta \cdot \frac{M_i}{12} \cdot \mathbb{1}[J_{it} = 0] + \upsilon_{imt}. \quad (22)$$

ν_t and η_m are time and municipality fixed effects. $\mathbf{V}_{it}$ is the vector described above of households characteristics that determine welfare transfers. We are interested in the impact of the residual transfer $\tilde{B}_{imt}$ on consumption when unemployed vs. unemployed. $\mathbb{1}[J_{it} = 0]$ is again a dummy for unemployment event time being equal to zero (i.e., an indicator variable for one member of the household being observed experiencing an unemployment spell in December of year t). As in the consumption-based approach, because we observe consumption at annual frequency, we control for the time spent unemployed by interacting $\mathbb{1}[J_{it} = 0]$ with the fraction of the year the individual has spent unemployed $\frac{M_i}{12}$. The variable $\tilde{B}_{imt}$ is the residual from the following regression of household local welfare transfers B_{imt} on the vector of households characteristics $\mathbf{V}_{it}$:

$$B_{imt} = \nu_t + \eta_m + \sum_{l \in \mathbf{K_0}} \tau_{mt}^l V_{it}^l + \sum_{k \in \mathbf{K_1}} \tau_m^k V_{it}^k + \sum_{j \in \mathbf{K}^c} \bar{\tau}^j V_{it}^j + \tilde{B}_{imt}. \quad (23)$$

[27] This contrasts with previous implementations of "optimization approaches", which rely on comparing statistics that are estimated on different samples due to data limitations and identification constraints (e.g., the liquidity vs. moral hazard effect in Chetty [2008]).

The municipality fixed-effects η_m do absorb the average difference in generosity across municipalities. The above specification then interacts characteristics V^l, for $l \in \mathbf{K_0}$, with municipality times year fixed-effects. This means that the variation within municipality over time in benefits according to characteristics $l \in \mathbf{K_0}$ is fully absorbed, and these characteristics do not participate to identification. Characteristics V^k, for $k \in \mathbf{K_1}$, are interacted with municipality fixed-effects only. This means that these characteristics participate to identification only to the extent that there is variation within municipality over time in the generosity of the schedule for these characteristics. Finally, for the remaining characteristics $j \in \mathbf{K^c}$, identification exploits both benefit variation across municipalities and over time in the schedule. By changing the characteristics included in $\mathbf{K_0}$, $\mathbf{K_1}$ and $\mathbf{K^c}$, we can therefore shut down particular sources of variation, and focus identification on specific dimensions of the welfare schedule. Our identifying assumption in all cases is that the residual variation in transfers $\tilde{B}_{imt}$ is orthogonal to the dynamics of household consumption. We probe into the credibility of this assumption below.

As a baseline, we exploit variations stemming from differences across municipalities and over time in the schedule for all characteristics $\mathbf{V}$. That is, we do not interact any characteristics V^k with municipality fixed effects, nor with municipality times year fixed effects in (6.2), such that $\mathbf{K_0} = \emptyset$ and $\mathbf{K_1} = \emptyset$, and all characteristics are included in $\mathbf{K^c}$. Figure 3 shows, for this baseline specification, the relationship between the first-difference in residualized transfers $\tilde{B}_{imt}$ and the first-difference in annual household consumption in a bin-scatter plot. The sample is split between households prior to the unemployment shock and households who experience unemployment in year t.[28] We find a positive and rather linear relationship between consumption and transfers, where the steep slope is indicative of a large marginal propensity to consume out of transfers for both groups. Importantly, the graph clearly displays a significantly steeper slope for the households in the unemployed group than for the households in the employed group, suggesting a significantly higher MPC for the former group compared to the latter.

Table 2 reports our results for the MPC out of local transfers when employed $\hat{\mu}_e$ and when unemployed $\hat{\mu}_e + \hat{\mu}_u$, from specification (6.2) estimated in first-differences. Standard errors are clustered at the municipality times year level. Column (1) corresponds to our baseline specification, when we residualize local welfare transfers on the vector of characteristics $\mathbf{V}$ and control for year and municipality fixed effects but without any interaction. The Table confirms that the estimated MPC out of local transfers is large and significantly larger when unemployed (.55) than prior to unemployment (.44).

The remainder of Table 2 explores the sensitivity of our results to exploiting different sources of underlying variation in $\tilde{B}_{imt}$. This can first be done by adding different characteristics in the set $\mathbf{K_1}$ in residualization (6.2). In column (2), we interact income and asset deciles with municipality fixed effects, so as to primarily exploit the variation in τ^{age} and $\tau^{children}$, arising from how welfare transfers differently account for the family structure of the household across municipalities. In

[28]Note that in our analysis, we systematically trim the first-difference in household consumption, omitting the top and bottom 5% of ΔC_{imt}.

column (3), we instead add family structure dummies interacted with municipality fixed effects: this specification exploits variation in the phase-out rates τ^y and τ^a conditional on the family structure. In column (4), we add both income and family structure dummies interacted with municipality fixed effects. The identifying variation now stems from changes in the average generosity of transfers within municipality over time due to local reforms. In column (5), we then add municipality times year fixed effects in the residualization of transfers:

$$B_{imt} = \nu^0_{mt} + \sum_j \bar{\tau}^j V^j_{it} + \tilde{B}_{imt}.$$

By controlling for the year-by-year change in average generosity of transfers within municipality, we shut down the variation in average transfers stemming from local reforms over time, in case we are concerned that these reforms are endogenous. Finally, in columns (6) and (7), we fully shut down the variation stemming from income and wealth profile of the schedule of benefits. We do so by incorporating income and wealth to the set $\mathbf{K_0}$. In other words, we introduce income times year times municipality fixed effects (column (6)), and wealth times year times municipality fixed effects (column (7)) in specification (6.2).

In all specifications, we find that the MPC out of welfare transfers is large in both states, when employed and unemployed. Furthermore, in all specifications, we find a larger MPC when unemployed than when employed. The difference is significant and stable across specifications with the MPC estimates when unemployed being around 25% higher than when employed. The only specification where the difference is not as stark ($\sim$ 10%) and no longer statistically significant at the 5% level is that of column (4) where we restrict attention to variation from local reforms over time. We also note that our estimates of the MPC in both states exploit the same source of variation in income, but remain two distinct local average treatment effects, where the weights depend on the specific workers contributing to estimation. It is reassuring nevertheless to find consistent results across all seven columns of Table 2, given that the workers bringing identification are likely to be very different across specifications.

MRS estimates Under the assumptions stated in Proposition 1, the relative odds ratios of state-specific MPCs provide a lower bound on the MRS. Table 2 also reports our estimates of this lower bound, following formula (9), and the corresponding standard errors, using two alternative approaches: the Delta-method, and a block-bootstrap computation where the clusters that are sampled with replacement are all observations at the municipality-level. Note that in the block-bootstrap procedure, we re-estimate equation () and () jointly at each iteration, and therefore account for the fact that residual benefits are estimated. In our baseline specification (column (1) of Table 2), the lower bound on the MRS is 1.59 with standard errors of .26 and .22 using the Delta-method and bootstrapping respectively. In the alternative specifications, the estimate for the lower bound is above 1.5 for most and slightly below for some. The outlier is when we only use variation stemming from changes in average transfers due to local reforms over time, which provides an estimate of 1.19 (column (4)). The underlying MPC estimate when unemployed for this specification is .454, which is at the low side, also compared to the estimates we provide in the further robustness checks.

Overall, the MPC approach thus suggests that the value of unemployment insurance is larger than what can be inferred from the traditional CB approach. The estimates from the MPC approach come with large standard errors, that include the estimated MRS from the CB approach, but they arguably provide a lower bound. The finding of a large MRS is in line with prior work implementing optimization approaches (e.g., Chetty [2008], Landais [2015]).[29] Importantly, we can confirm this in a setting where the consumption drop upon job loss, observed for the same sample of workers is relatively small.[30]

Mechanisms The high MPC estimates resonate with the prior evidence in Section 5, suggestive of lacking liquidity and binding credit constraints for the unemployed. The higher MPC during unemployment suggests that the price of increasing consumption at the margin is indeed higher when unemployed than when employed. We can investigate this further by studying how the various components of total household expenditures respond to the increase in welfare benefits, both when unemployed and employed. We decompose consumption into 5 components:

$$C = Y + B + C_{\text{Assets}} + C_{\text{Debt}} + C_{\text{Residual}},$$

where Y is total household labor income net of taxes, B are local welfare transfers, C_{Assets} is consumption out of assets, C_{Debt} is consumption out of debt, and C_{Residual} is the residual (including for example other transfers or taxes). Appendix Figure C-6 reports the estimated change of each of these consumption components in response to a change in welfare benefits, using our baseline specification (see column (1) of Table 2). We see that, in response to a 1 krona increase in their welfare benefits, employed individuals reduce their net household labor income by about .42 kroner, reduce their consumption out of debt by about .08 kroner, and increase their savings (reduced consumption out of assets) by about .05 kroner. Consumption from the residual consumption part is unaffected. This suggests that when employed, the main margin for raising an additional krona of consumption is household labor supply. For unemployed individuals, we find that in response to a 1 krona increase in their welfare benefits, they reduce their net household earnings by about .21 kroner. They increase their savings (reduced consumption out of assets) by about .15 kroner, which is somewhat higher than when employed, but their consumption out of debt is almost unaffected, again suggesting that unemployed individuals face credit constraints that make it difficult for them

[29] In particular, we can infer from footnote 34 in Chetty [2008] that his estimates correspond to an average MRS of 2.5. Landais [2015] also finds a large MRS of 1.88 (.02).

[30] One concern in comparing the MRS obtained from the MPC and the CB approach above is that the marginals out of which the MPCs are identified here are not the same as the average unemployed in Figure 1. To alleviate this concern, we re-estimate the consumption drop at job loss, using specification (20) on the subsample of individuals who ever receive local welfare transfers, (and therefore constitute our population of marginals in the MPC approach). We find an average consumption drop of 13.1%, extremely similar to the 12.9% drop estimated for our baseline sample.

to raise an additional krona of consumption using debt.[31]

6.3 Robustness

To establish the robustness of our results, we first probe into the validity of our identifying assumption that the residual variation in welfare benefits $\tilde{B}_{imt}$ is orthogonal to the dynamics of household consumption. For this, we build a covariate index based on observable characteristics available in the registry data, that correlate with consumption, but do not enter the benefit formula of welfare transfers. We use a linear combination of the education level of the household members, the age of the household head and his or her industry, and the lagged total debt and real estate wealth of the household, with the coefficients obtained from regressing consumption on these covariates. We then test for the presence of a significant correlation between $\tilde{B}_{imt}$ and this covariate index. Appendix Figure C-2 shows a binscatter of the relationship between the residual $\tilde{B}_{imt}$ in our baseline specification and the covariate index. Panel A shows this relationship in our whole sample. The graph displays no significant correlation between observable heterogeneity and $\tilde{B}_{imt}$. Panel B splits the sample by employment status, and shows that this absence of correlation holds equally well in the employed and the unemployed state.

To further investigate the orthogonality between past household consumption dynamics and our identifying variation in benefits $\tilde{B}_{imt}$, we follow in Appendix Figure C-4 an event-study design, where we isolate large changes in individuals' residualized benefits. We defined an event as experiencing a sudden increase in residual transfer of more than 12,500SEK between year t and year $t+1$, which corresponds roughly to the top 10% of the distribution of first differences in residualized benefits.[32] Appendix Figure C-4 panel A shows the evolution of residualized benefits around the event. It shows that the variation captured by these shocks is large, sudden, and does not correlate with the previous dynamics of transfers. In panel B, we show the evolution of consumption around the event. The first thing to notice is that there does not seem to be evidence of pre-trends: the shocks in residualized benefits do not correlate with past consumption dynamics. The second thing to notice is that the shock in welfare benefits is associated with a sudden increase in consumption These clear dynamic patterns bring credibility to the MPC we identify from variation in residualized transfers. Note that from this event-study design, we can compute an implied MPC corresponding to the estimated change in consumption in year 0, divided by the estimated change in benefits in year 0. We find an MPC of .456 (.093), which is very similar to our baseline estimated MPCs in Table 2.

In a similar spirit, we have also investigated how consumption and benefits evolve around the event of a large municipal reform of welfare transfers. To do this, we conducted event study designs where events are defined at the municipality level. We defined an event as a year in which the average residual transfer in municipality m experiences a sudden increase of more than 12,500SEK.

[31] Note that for the unemployed, we find a small negative effect for the residual part of consumption of -.08 kroner which may capture some crowding out of other transfers.

[32] To implement this event-study strategy, as well as that of Appendix Figure C-5, we follow the definition of residualization corresponding to column (1) of Table 2.

These events correspond to comprehensive reforms at the municipal level of the schedules of welfare benefits. We found 8 municipalities experiencing such events over our sample period. Appendix Figure C-5 shows the evolution of average benefits (panel A) and of average consumption (panel B) at the municipality level around the time of the event, following an event study specification with year and municipality fixed effects. The absence of pre-trends in both panel A and B confirms that the identifying variation brought about by these reforms is not endogenous to the past dynamics of benefits nor to the past dynamics of consumption.

A second line of concern is that $\tilde{B}_{imt}$ may be correlated with employment status. While we directly control for employment status in specification (6.2), the presence of non-linear effects between consumption and transfers could still be problematic. If the underlying distribution of $\tilde{B}_{imt}$ differs when unemployed and employed, such non-linearities could result in different estimates of the MPC in our linear specification. In panel A of appendix Figure C-3, we plot the distribution of our baseline residual variation $\tilde{B}_{imt}$ by employment status. The figure shows that the distribution of our identifying variation in welfare transfers is very similar across employment status. This alleviates the concern that the difference in our MPC estimates while employed and unemployed are simply driven by different distributions of underlying variation in transfer.

A last concern we address is the presence of selective migration: if individuals move to more generous municipalities in response to a negative consumption shock, this may introduce bias in our MPC estimates. Panel B of appendix Figure C-3 displays the distribution of $\tilde{B}_{imt}$, splitting the sample between movers (households who moved municipality in year t) and stayers. We find no significant correlation between $\tilde{B}_{imt}$ and the probability of moving, which indicates that our identifying variation in transfers is immune to the bias of selective migration.

6.4 External Validity

We also evaluate the external validity of our MPC estimates. We start by using an alternative identification strategy to estimate the MPC for the same sample of workers, but only applicable when they are unemployed. For this purpose, we take advantage of the variation in unemployment benefits due to the presence of a kink in the Swedish UI benefit schedule: individuals receive a replacement rate of 80% of their previous daily wage up to a cap. Over the period 2002 to 2007, the cap in daily UI benefits was fixed at 680SEK, meaning that the relationship between UI benefits and daily wage w exhibited a kink at $w = 850SEK$. This offers a credible source of exogenous variation in income that can be exploited in a regression kink design, as discussed in Kolsrud et al. [2018]. The identifying variation is displayed in Figure 4 panel A, which plots, in our main sample over the period 2002 to 2007, a binscatter of the relationship between the daily wage and the average replacement rate. The latter is computed as the average benefit received during unemployment from the IAF data divided by the daily wage. The graph shows first that the replacement rate is close to exactly 80% on the left hand side.[33] The graph also displays a clear

[33] Note that the reason why the replacement rate is a bit below 80% is that some workers have their UI benefits reduced due to sanctions.

kink at $w = 850SEK$, with the replacement rate declining sharply, as benefits are capped. We use this kinked relationship and treat it as a fuzzy RKD around the 850SEK threshold. Figure 4 panel B plots the average change in consumption ΔC_i between the year the individual is unemployed and the year prior to the start of the spell by bins of daily wages. The graph shows evidence of a large non-linearity in the relationship between daily wage and the consumption drop at unemployment. This translates into an estimate of the MPC while unemployed of .63 (.16), which is very robust across specifications. This estimate is very similar to our MPC estimate while unemployed using the local welfare transfer variation of .55 (.02). This evidence provides additional credibility to our estimates of the MPC for the Swedish unemployed. For the sake of brevity, we report all the details of the estimation, as well as robustness and sensitivity analyses in Appendix D.

A second way to probe into the external validity of our results is to compare them to estimates available in the MPC literature. The average MPC in our sample is large (around .45), but comparable to estimates in other countries or settings. The empirical literature on the consumption responses to unanticipated income changes can be broadly divided into two groups. The first group exploits abrupt policy changes as quasi-natural experiments, such as income tax rebates (e.g., Johnson et al. [2006], Parker et al. [2013]) or credit card limit increases (e.g., Gross and Souleles [2002], Gross et al. [2016]) and finds large average MPCs, between .4 and .6, that are similar in magnitude to the ones found in our setting.[34] Another branch of research focuses on survey-based responses to hypothetical increases in household resources (e.g., Jappelli and Pistaferri [2014]) and also finds large MPCs, around .5.

Finally, a nascent literature has documented the presence of significant heterogeneity in MPCs. While most of this literature focuses on heterogeneity by cash on hand, a small number of papers also report heterogeneity by employment status. Jappelli and Pistaferri [2014] find significantly higher MPC for unemployed compared to employed.[35] Bunn et al. [2018] also find larger MPC out of negative income shocks for unemployed compared to employed people. Although the variation in unemployment status used in these papers is cross-sectional and not within individual as in our setting, the results corroborate our findings that MPCs are larger in the unemployed state compared to the employed state.

7 Revealed Preference Approach

This section now turns to the insurance choices embedded in the Swedish UI setting, which can be used in combination with workers' predicted unemployment risk to implement the Revealed Preference approach outlined in Section 3.3. While the requirement to observe insurance choices makes the RP approach not as generally applicable as the MPC approach, the choice setting in

[34] Johnson et al. [2006] find a MPC between 40% and 60%, while the total response of spending estimated to the 2008 tax rebate in Parker et al. [2013] is 50 to 90 percent of the total payments. The estimated MPC out of liquidity in Gross et al. [2016] is 0.37.

[35] The average estimated MPC in Jappelli and Pistaferri [2014] is .48, very close to our average MPC. Controlling for observables, the MPC is 7 percentage points higher for unemployed in their baseline estimation.

Sweden provides a unique opportunity to verify whether the higher MRS estimates from the MPC approach can be confirmed.

7.1 RP: Implementation

The implementation of the RP approach requires, in addition to data on workers' insurance choices, knowledge of the prices ($\frac{p_u}{p_e}$) and the unemployment risk ($\frac{1-\pi_i}{\pi_i}$) at the time of the insurance choice.

The relative price $\frac{p_u}{p_e}$ is equal to the premium to coverage ratio $\frac{\tau_1-\tau_0}{b_1-b_0}$ for the extra insurance the comprehensive plan (b_1, τ_1) provides relative to the basic plan (b_0, τ_0). While the basic level b_0 is flat and identical for all workers, the comprehensive level b_1 replaces 80% of pre-unemployment earnings, but is capped and cannot drop under the flat benefit level b_0 either. We therefore predict the extra coverage $b_1 - b_0$ a worker would receive based on her earnings level.[36] On average the comprehensive plan increases the net daily UI benefit by 188 SEK. As explained in section 4.1, to be eligible for the comprehensive coverage, workers need to pay an additional premium $\tau_1 - \tau_0$, which was equal to a net annual amount of 720 SEK and identical across workers during our sample period.[37]

To predict the unemployment risk, we use the risk model introduced in Section 4.2 which predicts the number of days spent unemployed in year $t + 1$ based on a very rich set of observable characteristics Z in year t. We estimate the model separately on the workers buying comprehensive and basic coverage to account for the presence of moral hazard, as discussed in 3.3. We convert our prediction of the number of days spent unemployed in year $t + 1$ into a binary risk and obtain a measure of the expected price per unit of coverage of individual i, $\tilde{p}_i = \frac{p_u}{p_e}\frac{1-\pi(Z_i)}{\pi(Z_i)}$.[38] In other words, this expected price is the ratio of the predicted days spent employed times the extra premium to the predicted days spent unemployed times the extra coverage. We compute the expected price per unit of coverage $\tilde{p}_i$ for the entire population of Swedish workers aged 25 - 55 between 2002 and 2007.

The differences in unemployment risk lead to important variation in the expected price per unit of coverage across workers, as illustrated in Panel A of Figure 5. The figure divides workers in different cells based on observable characteristics and plots the average expected prices, based on the predicted risk under comprehensive coverage, against the share of people buying comprehensive coverage by cell. On average, workers facing a higher expected price are less likely to buy the

[36] To be precise, we start by estimating the relationship between annual income and daily benefits received when unemployed using the data on income and benefits of people who lost their jobs. The need to estimate this relationship arises from the imperfect compatibility of the daily wage data (available in the PES registries, only for unemployed individuals) and the annual income data (reported in the LISA registries, for all workers). We thus define benefits as a kinked function of annual income, with constant replacement rate up to the 850 SEK daily income threshold (210,000 SEK annually). This provides us with an individual-level potential benefit level under comprehensive coverage, $\widehat{b}_{1,i}$. We then subtract the basic (daily) benefit level b_0 (and the daily premium $\tau_1 - \tau_0$, which the unemployed continue to pay to remain eligible).

[37] In computing the price of insurance p_u/p_e for a given worker, we account for the fact that the premia were tax-deductible at 40%, while the unemployment benefits are taxed.

[38] To be precise, we first estimate individuals' risks by predicting their number of unemployed days in the forthcoming year, $\widehat{d}_{i,t+1}$, using our zero-inflated Poisson model of unemployment risk and convert this into a risk odds ratio, using 260 working days a year.

comprehensive coverage, but there is substantial variation in take-up. For some cells, the expected price to transfer income from employment to unemployment is high, with implied mark-ups of more than 100%. The share of workers buying comprehensive coverage is lower at those high prices, but still sizeable. Of course, some determinants of the unemployment risk may also directly affect the willingness to buy insurance.[39]

Bounds As discussed in Section 3.3, we can use the expected price to bound an MRS for the supplemental coverage workers can buy. The expected price using the predicted unemployment risk under comprehensive coverage provides a (conservative) lower bound on the MRS of the workers buying the comprehensive coverage. While for the workers who stick to the basic coverage, the expected price using the predicted unemployment risk under basic coverage provides a (conservative) upper bound on the MRS. Hence, regardless of the determinants of the variation in unemployment risk, we can simply average these individual bounds to provide non-parametric bounds on the average MRS for the workers choosing comprehensive and basic coverage respectively. This gives a lower bound of $\mathbb{E}[\tilde{p}_1 | b_1] = .69$ for the MRS of the workers on comprehensive coverage and an upper bound of $\mathbb{E}[\tilde{p}_0 | b_0] = 2.15$ for the MRS of the workers on basic coverage. The workers on basic coverage face a price that entails a substantial mark-up above the actuarially fair price. In contrast, the workers on comprehensive coverage, who face higher unemployment risk, but pay the same premium, pay a price that is substantially below the actuarially fair price. The bounds are loose and only exclude rather extreme risk-loving preferences (or state-dependent preferences that favour employment consumption) for the average worker on comprehensive coverage. The opposite is true for the average worker on basic coverage. In any case, we cannot reject homogeneity in the MRS, neither can we exclude a wide dispersion in the MRS.

MRS Estimates We now use the rich variation in expected prices across workers to estimate the distribution of MRS using a linear choice model: an individual i buys comprehensive insurance at time t if and only if

$$X_{it}\beta - \gamma\tilde{p}\left(Z_{it}\right) + \varepsilon_{it} \geq 0. \tag{24}$$

The corresponding marginal rate of substitution equals

$$MRS\left(X_{it}\right) = \frac{X_{it}\beta}{\gamma}.$$

We allow for a rich set of controls X to affect the MRS, which includes demographics (age, gender, family type, presence of children), deciles of household disposable income, and education in our baseline specification. The error term ε_{it} is assumed to follow a logistic distribution. Identification relies on the presence of risk shifters in the set of variables Z in our predicted risk model that do not affect MRS directly. We take advantage of the presence of two risk shifters that stem from the

[39] Landais et al. [2017] show, for example, that age is a driver of advantageous selection. Older workers are less likely to be unemployed, but more likely to buy comprehensive coverage.

specificities of the Swedish labor market institutions, and are arguably exogenous to individuals' preferences. First, we use the fact that Sweden applies a strict version of the last-in-first-out principle, which creates exogenous variation in the probability of layoff by tenure ranking at the establishment times occupation level. Second, we use the fact that when a firm wants to layoff more than 5 workers in a 6-month period, it needs to emit a layoff notification to the public employment service. These notifications therefore capture idiosyncratic variation in firm's business conditions that are plausibly exogenous to individuals' risk preferences. We therefore include in our risk model both tenure ranking and the full history of firm layoff notifications, as well as their interaction, as shifters of workers' risk.[40]

Table 3 summarizes the regression output, using the whole population to estimate the risk and choice model and predicting workers' unemployment risk under comprehensive coverage.[41] In the baseline specification, the average MRS is estimated to be 3.13, implying that workers are on average willing to pay more than a 200% mark-up to get comprehensive coverage. This is substantially higher than the CB estimates, but also above the estimates we found in the MPC approach. The RP estimates also indicate substantial heterogeneity in MRS. While for 10% of workers the estimated MRS is lower than .91, at the other end of the distribution 10% of workers have MRS above 4.72.

7.2 Robustness

Our RP estimates rely on the adequacy of the choice model and the absence of information or choice frictions that bias our estimation of the MRS. A particular concern, given the growing evidence on biases in perceived risks [e.g., Sydnor [2010], Spinnewijn [2015]], is that workers perceive their risk differently from what is predicted by our risk model. Appendix Figure E-1 gauges the potential severity in our context by comparing true and perceived risks using ex-ante (elicited) beliefs and ex-post (reported) employment outcomes in the HUS survey in Sweden. We find that workers who report a 1 percent higher probability to keep their job are only .26 (.05) percent more likely to keep their job. We therefore explore the sensitivity of the estimation results to the particular risk model we use in two alternative specifications. In the first specification, we predict risk using variation that we believe is most salient. The specific risk shifters we use are the unemployment history of a worker and the layoff rate of its current employer in recent years.[42] In the second specification, we convert the predicted risk using our baseline model into an estimate of the perceived risk, assuming

[40]For more institutional details on these sources of risk variation and a thorough discussion of the credibility of their exogeneity to individuals' preferences, see Landais et al. [2017].

[41]As discussed in Section 3.3, in the presence of moral hazard, the RP estimates provide still a lower bound on workers' MRS when using the prediced risk model under comprehensive coverage. In Appendix E we also re-estimate all specifications using the predicted risk model under the basic coverage, providing an upper bound on workers' MRS.

[42]We note that a trade-off arises between using salient shifters and satisfying the exclusion restriction that the used risk shifters do not affect the MRS directly. We also note that we explored specifications where we use risk shifters to instrument for $\tilde{p}$ (rather than using specific risk shifters to compute alternative measures of π_i and of $\tilde{p}$). We did not find that instrumentation made a significant difference to the estimates nor to the standard errors.

an imperfect correlation of .26 based on the HUT estimate.[43]

The change in estimates is comparable for both specifications that try to account for the wedge between predicted and perceived risk. First, for both alternative specifications the price elasticity of insurance choice has increased, as reported in columns (2) and (3) of Table 3. Second, the estimated MRS is lower under both specifications. When correcting for risk mis-perceptions, the mean MRS is 2.13, which is closer to our estimate from the MPC approach. Also the 75th percentile for example shifts down from 4.21 to 2.73. The results when using salient shifters of risk are in between the baseline risk model and the perceived risk model, with a mean MRS of 2.43 and the 75th percentile at 3.43. Panel B of Figure 5 compares the distributions of the estimated MRS for the different risk models, using only the individuals in the baseline sample used for the CB and MPC implementation. The figure confirms how the MRS distribution shifts down and becomes less dispersed under the alternative risk models. The lower mean and variance in the MRS when correcting for the risk misperceptions are reminiscent of the findings in Handel and Kolstad [2015].

We then further explore the robustness of our results in Table 3 using the perceived risk model. Compared to column (3), columns (4) and (5) allow for extra controls affecting the MRS in the choice model. The estimates are similar when controlling for asset holdings (including net worth, liquid assets and debt) and for industry and region fixed effects. Column (6) allows the price effect to depend on income. Overall, workers with lower income are more price-elastic and allowing for this interaction increases the estimated mean MRS. Column (7) controls for persistence in choices, either capturing switching costs or persistent preferences, by including a dummy for whether the insurance choice is the same as in the prior year. The estimated effect is neither significant, nor does it affect the estimated price elasticity. Columns (8) and (9) further control for switching costs, by restricting the sample to workers who have switched jobs between 2002 and 2006 as their cost from switching insurance is arguably smaller. Again, both the price elasticity and the implied MRS remain similar, even when restricting to the years in which workers actually switched jobs.

Importantly, the dispersion in the estimated MRS distribution remains substantial across all specifications, as reported in the bottom panel of Table 3. Clearly, the difficulty remains to understand whether heterogeneous valuations capture deep structural heterogeneity in risk preferences, or some form of heterogeneous frictions, which is also evidenced by the sensitivity of our estimates to controlling for perceived risks. Appendix Figure E-3 correlates the estimated MRS with various observable characteristics (age, gender, wealth,...). We use the regression specification shown in column (10) in Table 3, which also includes individuals' cognitive ability as measured by army

[43]The low coefficient estimate when regressing true job loss on perceived job loss can not only be driven by a low correlation between true and perceived risks, but also by a high variance in perceived risks relative to the variance in true risks. Mueller et al. [2018] separate the two in the context of job finding probabilities rather than job loss probabilities, using the Survey of Consumer Expectations, and find a correlation coefficient that is approximately the OLS estimate, which underlies our choice to assume an imperfect correlation equal to the OLS estimate in our context. Interestingly, in a regression of job loss on elicited beliefs about job loss, using the Survey of Consumer Expectations in the US, we find an OLS estimate of .27 (.08), similar in magnitude to the OLS estimate in the HUS survey in Sweden. Note that the perceived unemployment risk relevant for the insurance choice combines the perceived job loss probability and job finding probability when unemployed, but we do not have any survey information on the latter in the Swedish context.

enlistments tests. While the correlations shown in Figure E-3 panels A to E can be consistent with substantial heterogeneneity in preferences, panel F strongly indicates that part of the variance in the estimated MRS is due to heterogeneity in frictions. It documents a strong negative correlation between MRS and cognitive ability, with workers who are estimated to have high MRS having lower cognitive ability on average. We provide more details on this analysis in Appendix E.

8 Comparison of Results

Having implemented three alternative approaches to estimate the MRS on the same sample of individuals, we now briefly compare the results and try to reconcile the different findings.

Figure 6 puts the MRS estimates together for the three approaches. For the CB approach, we highlight again the range of values of the MRS between 1.13 (.03) and 1.51 (.11), which correspond to the CB estimates using a relative risk aversion of $\gamma = 1$ and $\gamma = 4$ respectively (see Figure 1). For the MPC approach, we report the ratio of the MPCs expressed as odds ratios for our preferred specification using the within municipality variation in local welfare transfers, both across household types and over time (column (1) of Table 2). This ratio provides a lower bound on the MRS of 1.59 (.22). For the RP approach, we report a mean of 2.13 (.02), together with the estimated distribution of MRS. These estimates are obtained using the parametric RP model, where we use the perception-adjusted predicted risks under comprehensive coverage (column (3) of Table 3). According to the CB approach, for commonly used values of risk aversion, the mark-up workers are willing to pay to transfer a krona of consumption from employment to unemployment would not be more than 50 percent. This mark-up is at least 60% using the MPC approach and even larger than 100% using the RP approach. As a result, the estimated value of insurance is substantially larger using the MPC and RP approaches than using the CB approach.

We contrast our estimates for the value of UI with plausible values for the moral hazard (MH) costs of UI. The MH cost is equal to 1 plus the unemployment elasticity with respect to the UI generosity, $\varepsilon_{\frac{\pi}{1-\pi}}$ (see subsection 2.2). In the same Swedish context, Kolsrud et al. [2018] find an elasticity of 1.5. Schmieder and Von Wachter [2016] summarize estimates of $\varepsilon_{\frac{\pi}{1-\pi}}$ from 18 studies from 5 different countries, and find a median of estimate of 0.53. We therefore show $1.5 - 2.5$ as a range of credible estimates of the moral hazard cost of UI in our context. Following the Baily-Chetty formula, an increase in the generosity of UI is desirable only if the insurance value is higher than the MH cost. The CB estimates of the UI value are small in comparison to available estimates of the costs of UI. Even in the higher range, the estimated values do not exceed the more conservative estimates of the MH costs. Hence, according to the CB approach, the Swedish UI system is too generous. In contrast, both the MPC and RP approaches suggest that the average valuation of UI is actually comparable, if not higher than the moral hazard cost, so that the generous UI benefits in Sweden may well be optimal.

An important remaining question is what can explain this discrepancy between the CB approach and the two alternative approaches? A first explanation is that the relevant risk aversion to scale the consumption drop to obtain MRS estimates is higher than previously thought. A second explanation is that unemployed workers suffer more from the observed consumption drop due to state-dependent expenditures or preferences. While the RP approach is immune to state-dependence (and also the lower-bound argument for the MPC approach can continue to hold), the implied MRS using the CB approach crucially depends on it. *A priori* state-dependence could play either way for the CB approach, as discussed in subsection 3.2, and it has been hard to get any empirical traction on the sign and magnitude of the actual wedge it introduces.

To gauge the importance of state-dependence, we first follow prior work in studying how specific types of expenditures evolve around job loss. Using the Swedish survey of consumption expenditures (HUT), we run event studies for different expenditure categories. Appendix Figure B-6 shows that some expenditures that are complementary to spending time away from home (e.g., restaurants and hotels) decrease significantly during unemployment, while other expenditures that are likely complements to job search and unemployment (e.g., telecom, education and health expenses) increase. Some other expenditures seem not to change (e.g., housing). It is clear that translating these differential dynamic patterns of expenditure categories into an estimate of state-dependence relevant for welfare is difficult. An alternative approach was proposed in section 3.2, simply comparing the empirical moments from the CB and the MPC approach. Following equation (14), state-dependence drives a wedge between the drop in consumption at job loss Δc and the imputed drop in consumption $\frac{dc_u}{dy_u}\Delta y$ based on the observed drop in disposable income and the estimated MPC when unemployed. Computing the drop in net earnings at job loss Δy from our estimates from Appendix Figure B-4 and using our baseline estimate of the MPC in unemployment $\frac{dc_u}{dy_u}$ (column (1) in Table 2), we find an imputed drop in consumption $\frac{dc_u}{dy_u}\Delta y = .11$, which is very close to the estimated drop in consumption expenditures at job loss (12.9%). If we use the MPC in unemployment estimated from our RKD approach (see Figure 4), we find an imputed drop of .126, almost indistinguishable from the estimated drop at job loss. This evidence suggests that state-dependence plays a relatively limited role and is unlikely to account for the significant difference in MRS estimates between the CB approach and the RP and MPC approaches.[44]

In the absence of significant state-dependence, the gap in estimated MRS suggests the presence of high levels of risk aversion. The relative risk aversion parameter γ that would close the gap between our CB estimate and our MPC and RP estimates equals 4.5 and 8.7 respectively.[45] We note that there is little consensus about γ in the literature, which in part motivated the use of alternatives to the CB approach (see Chetty and Finkelstein [2013]). While the required γ's to reconcile our approaches may seem large, specific mechanisms like the importance of committed

[44] We note that equation (14) implicitly assumes that the nature of the income variation due to job loss and the income variation underlying the MPC is similar in nature. For example, as evidenced by Figure 1, the drop in income and in consumption at job loss exhibits some persistence. In the face of a more persistent income shock upon job loss, the estimated drop at job loss will therefore tend to overestimate the drop Δc that is conceptually relevant for the reconciliation of equation (14).

[45] We note again that the RP approach provides an estimate of the average MRS for the supplemental coverage workers can buy. This may overstate the MRS at the margin for workers under comprehensive coverage, implying that the required γ for reconciliation may be lower.

expenditures (e.g., Chetty and Szeidl [2007]) could help rationalize such high levels of risk aversion at job loss.

9 Conclusion

The challenges to the standard consumption-based approach in estimating the value of insurance have inspired us to develop an alternative, arguably more robust approach, based on the estimation of marginal propensities to consume when employed and unemployed. This MPC approach uncovers the price of smoothing consumption, which in turn enables to bound the value of insurance at the margin. Importantly our method can be generalized to any other social insurance program. Our focus has been on the insurance value of UI, following the tradition in the literature. We have been mostly ignoring the corresponding redistributive value of social insurance, which may be equally important to determine its social desirability. In order to make progress in identifying the value of redistribution, we note that differences in MPCs can also provide useful information.

Beyond the CB and MPC approaches, we have also exploited the uniqueness of the Swedish context to implement a third approach. This revealed preference approach requires the observation of insurance choices and information on individuals' (perceived) unemployment risk. We have implemented all three approaches for the same sample of workers and in the same setting. Our analysis shows that the value of UI is high despite small observed drops in consumption at job loss.

We have aimed in this paper to improve our understanding of the average value of UI, but much more remains to be done to understand heterogeneity in the valuation of UI. The CB approach, unfortunately, offers limited hope to identify such heterogeneity. Like for estimating the average value, the important constraint is that the drop in consumption at job loss is endogenous to risk preferences (Chetty and Looney [2007], Andrews and Miller [2013]): individuals with high MRS will take action to reduce their drop in consumption. By examining how the MRS from the RP approach correlates with realized drops in consumption at job loss, we can provide some direct evidence of this mechanism. In Appendix Figure E-4, we split our baseline sample in cells of observable characteristics and report the estimated average drop in consumption at job loss for households in that cell against the average MRS in the cell estimated from the RP approach in the year prior to job loss. The graph shows that the relationship is decreasing indeed: individuals with the highest MRS experience lower drops in consumption on average. This asks for caution when using cross-sectional heterogeneity in consumption smoothing as a guide for differentiating social transfers, while recommendations based on within-individual changes in consumption, for example, over the unemployment spell (Kolsrud et al. [2018]), are arguably more robust.

Going forward, the RP approaches seem to offer more promise to uncover heterogeneity in MRS, but the challenge remains to separate preferences from information and choice frictions. Any embedded choice in social insurance insurance programs, in principle, will allow us to further improve their evaluation, but whether the introduction of (unregulated) choice is desirable will crucially depend on the presence of these other frictions.

References

Abaluck, Jason and Jonathan Gruber, "Choice Inconsistencies Among the Elderly: Evidence from Plan Choice in the Medicare Part D Program," *American Economic Review*, 2011, *101* (4), 1180–1210. 5, 16

Aguiar, Mark and Erik Hurst, "Consumption versus Expenditure," *Journal of Political Economy*, 2005, *113* (5), pp. 919–948. 12

Andrews, Isaiah and Conrad Miller, "Optimal social insurance with heterogeneity," *Working Paper*, 2013. 39, 89

Arrow, Kenneth, "Uncertainty and the welfare economics of medical care," *The American Economic Review*, 1963, *53* (5), 941–973. 14

Autor, David, Andreas Ravndal Kostol, Magne Mogstad, and Bradley Setzler, "Disability Benefits, Consumption Insurance, and Household Labor Supply," Working Paper 23466, National Bureau of Economic Research June 2017. 5

Baily, Martin Neil, "Some aspects of optimal unemployment insurance," *Journal of Public Economics*, 1978, *10* (3), 379–402. 2, 7

Borusyak, Kirill and Xavier Jaravel, "Revisiting event study designs," *Working Paper*, 2016. 21

Browning, Martin and Søren Leth-Petersen, "Imputing consumption from income and wealth information," *The Economic Journal*, 2003, *113* (488), 282–301. 18

_ and Thomas F Crossley, "Unemployment insurance benefit levels and consumption changes," *Journal of Public Economics*, 2001, *80* (1), 1–23. 4, 12, 22

_ and Thomas F. Crossley, "Shocks, Stocks and Socks: Smoothing Consumption Over a Temporary Income Loss," *Journal of the European Economic Association*, 2009, *7* (6), pp. 1169–1192. 12

Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico, "The consumption response to positive and negative income shocks," *Journal of Monetary Economics*, 2018, *96*, 1–15. 32

Cabral, Marika and Mark Cullen, "Estimating the Value of Public Insurance Using Complementary Private Insurance," *Forthcoming American Economic Journal: Economic Policy*, 2016. 5, 14

Chetty, Raj, "A general formula for the optimal level of social insurance," *Journal of Public Economics*, 2006, *90* (10-11), 1879–1901. 2, 7, 8

_ , "Moral hazard versus liquidity and optimal unemployment insurance," *Journal of Political Economy*, 2008, *116* (2), 173–234. 2, 4, 15, 26, 29, 61

_ and Adam Looney, "Consumption smoothing and the welfare consequences of social insurance in developing economies," *Journal of Public Economics*, 2006, *90* (12), 2351 – 2356. 2, 8

_ and _ , "Income Risk and the Benefits of Social Insurance: Evidence from Indonesia and the United States," in T. Ito and A. Rose, eds., *Fiscal Policy and Management in East Asia, NBER-EASE, Volume 16*, University of Chicago Press, October 2007, pp. 99–121. 8, 39, 89

— and Adam Szeidl, "Consumption Commitments and Risk Preferences," *The Quarterly Journal of Economics*, 2007, *122* (2), pp. 831–877. 12, 39

— and Amy Finkelstein, "Social insurance: Connecting theory to data," in "Handbook of Public Economics," Vol. 5, Elsevier, 2013, pp. 111–193. 2, 38

Eika, Lasse, Magne Mogstad, and Ola Vestad, "What can we learn about household consumption from information on income and wealth," 2017. 5, 18

Fadlon, Itzik and Torben Heien Nielsen, "Household labor supply and the gains from social insurance," *Forthcoming Journal of Public Economics*, 2018. 5, 6, 15, 62

Finkelstein, Amy, Nathaniel Hendren, and Erzo F.P. Luttmer, "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment," *Forthcoming Journal of Political Economy*, June 2015. 5

— , — , and Mark Shepard, "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts," *Forthcoming American Economic Review*, 2017. 5, 14

Freyaldenhoven, Simon, Christian Hansen, and Jesse M Shapiro, "Pre-event trends in the panel event-study design," *Working paper*, 2018. 21

Ganong, Peter and Pascal Noel, "Consumer spending during unemployment: Positive and normative implications," *Forthcoming American Economic Review*, 2017. 4, 5, 16, 22

— and Simon Jaeger, "A Permutation Test for the Regression Kink Design," *Journal of the American Statistical Association*, 2018, *113* (522), 494–504. 81, 82

Gerard, Francois and Joana Naritomi, "Job Displacement Insurance and Consumption: Evidence from Brazil," *working paper*, 2019. 5, 16

Grönqvist, Hans, J Peter Nilsson, Per-Olof Robling et al., "Early lead exposure and outcomes in adulthood," *Working Paper, IIES*, 2017. 90

Gross, David B and Nicholas S Souleles, "Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data," *The Quarterly journal of economics*, 2002, *117* (1), 149–185. 32

Gross, Tal, Matthew J Notowidigdo, and Jialan Wang, "The marginal propensity to consume over the business cycle," Technical Report, National Bureau of Economic Research 2016. 32

Gruber, Jonathan, "The consumption smoothing benefits of unemployment insurance," *The American Economic Review*, 1997, *87* (1), 192–205. 2, 4, 13, 22

Handel, B., "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 2013, *103* (7), 2643–2682. 5, 16

— and J. Kolstad, "Health Insurance For Humans: Information Frictions, Plan Choice, and Consumer Welfare," *American Economic Review*, 2015, *105* (8), 2449–2500. 5, 16, 36

Hendren, Nathaniel, "Knowledge of future job loss and implications for unemployment insurance," *American Economic Review*, 2017, *107* (7), 1778–1823. 2, 4, 6, 15, 16, 22, 62, 67

Jappelli, Tullio and Luigi Pistaferri, "The Consumption Response to Income Changes," *Annual Review of Economics*, September 2010, *2* (1), 479–506. 5

— and —, "Fiscal policy and MPC heterogeneity," *American Economic Journal: Macroeconomics*, 2014, *6* (4), 107–36. 32

Johnson, David S, Jonathan A Parker, and Nicholas S Souleles, "Household expenditure and the income tax rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610. 32

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239. 69

Kekre, Rohan, "Unemployment Insurance in Macroeconomic Stabilization," *Working Paper*, 2017. 5

Koijen, Ralph, Stijn Van Nieuwerburgh, and Roine Vestman, "Judging the quality of survey data by comparison with truth as measured by administrative records: Evidence from Sweden," in "Improving the Measurement of Consumer Expenditures," University of Chicago Press, 2014, pp. 308–346. 5, 18

Kolsrud, Jonas, Camille Landais, and Johannes Spinnewijn, "Studying consumption patterns using registry data: lessons from Swedish administrative data," *Working Paper*, 2017. 3, 5, 18, 21, 50

—, —, Peter Nilsson, and Johannes Spinnewijn, "The optimal timing of unemployment benefits: Theory and evidence from Sweden," *American Economic Review*, 2018, *108* (4-5), 985–1033. 4, 7, 19, 21, 31, 37, 39, 64, 80, 81

Kreiner, Claus Thustrup, Søren Leth-Petersen, and Peer Ebbesen Skov, "Tax Reforms and Intertemporal Shifting of Wage Income: Evidence from Danish Monthly Payroll Records," *American Economic Journal: Economic Policy*, 2016, *8* (3), 233–257. 5

Krueger, Alan B. and Bruce D. Meyer, "Labor Supply Effects of Social Insurance," Handbook of Public Economics, Elsevier 2002. 2

Landais, Camille, "Assessing the welfare effects of unemployment benefits using the regression kink design," *American Economic Journal: Economic Policy*, 2015, *7* (4), 243–278. 2, 4, 15, 29, 61

— and Johannes Spinnewijn, "Online Appendix to The Value of Unemployment Insurance," http://econ.lse.ac.uk/staff/clandais/cgi-bin/Articles/valueUI_onlineappendix.pdf 2020. 11

—, Arash Nekoei, Peter Nilsson, David Seim, and Johannes Spinnewijn, "Risk-based selection in unemployment insurance: Evidence and implications," *CEPR Discussion Paper*, 2017, (12364). 4, 18, 19, 34, 35

Low, Hamish and Luigi Pistaferri, "Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off," *American Economic Review*, October 2015, *105* (10), 2986–3029. 5

—, Costas Meghir, Luigi Pistaferri, and Alessandra Voena, "Marriage, Labor Supply and the Dynamics of the Social Safety Net," Working Paper 24356, National Bureau of Economic Research February 2018. 5

Maggio, Marco Di, Amir Kermani, and Kaveh Majlesi, "Stock Market Returns and Consumption," Working Paper 24262, National Bureau of Economic Research January 2018. 5

Mossin, Jan, "Aspects of rational insurance purchasing," *Journal of Political Economy*, 1968, *76* (4), 553–568. 14

Mueller, Andreas I, Johannes Spinnewijn, and Giorgio Topa, "Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence and Bias," Working Paper 25294, National Bureau of Economic Research November 2018. 36

Nakamura, Emi and Jón Steinsson, "Identification in Macroeconomics," *The Journal of Economic Perspectives*, 2018, *32* (3), 59–86. 5

Parker, Jonathan A, Nicholas S Souleles, David S Johnson, and Robert McClelland, "Consumer spending and the economic stimulus payments of 2008," *American Economic Review*, 2013, *103* (6), 2530–53. 32

Schmieder, Johannes F and Till Von Wachter, "The effects of unemployment insurance benefits: New evidence and interpretation," *Annual Review of Economics*, 2016, *8*, 547–581. 2, 37

Shimer, Robert and Iván Werning, "Reservation Wages and Unemployment Insurance," *The Quarterly Journal of Economics*, 2007, *122* (3), 1145–1185. 4

Spinnewijn, Johannes, "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs," *Journal of the European Economic Association*, 2015, *13* (1), 130–167. 11, 16, 35

_ , "Heterogeneity, Demand for Insurance and Adverse Selection," *American Economic Journal: Economic Policy*, 2017, *9* (1), 308–343. 5

Stephens, Melvin, "The Long-Run Consumption Effects Of Earnings Shocks," *The Review of Economics and Statistics*, February 2001, *83* (1), 28–36. 4

Sydnor, Justin, "(Over)insuring Modest Risks," *American Economic Journal: Applied Economics*, October 2010, *2* (4), 177–99. 35

Figures

Figure 1: Estimated Consumption Dynamics around Start of Unemployment Spell

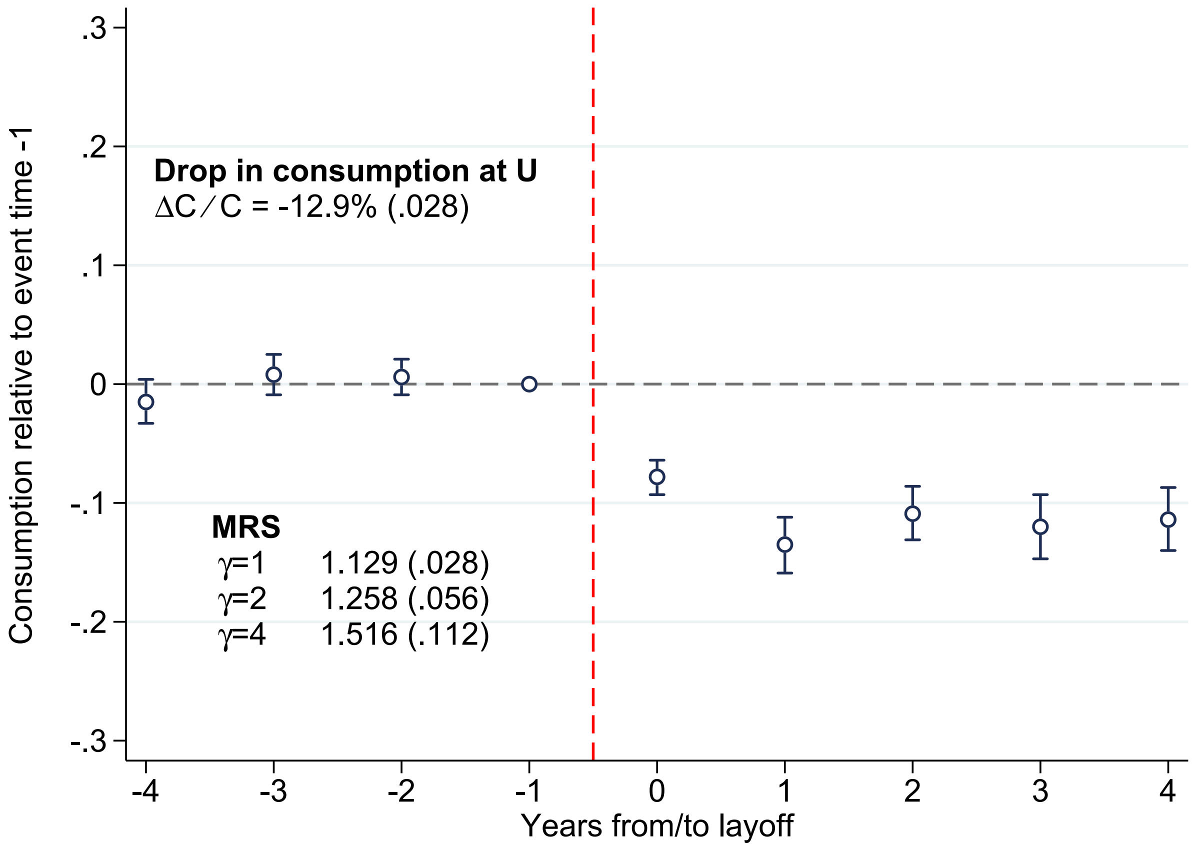

Notes: The figure reports event study estimates of household annual consumption around the time when a household member loses her job. Coefficients and confidence intervals come from specification (30) run on the sample of treated individuals and a control group of individuals obtained from nearest-neighbor matching on pre-event characteristics. All point estimates are expressed as a fraction of average total household consumption as of event year -1. We restrict the sample to individuals aged 25 to 55, who are eligible for any form of UI at the time of the event and who are unemployed in December of the year in which they lose their job for the first time. We also report on the graph an estimate of the drop in flow consumption at unemployment $\Delta C/C$ estimated using the parametric approach of specification (20) We convert this estimate of $\Delta C/C$ into a measure of the MRS, following the standard version of the consumption-based implementation, which is to assume that third and higher order terms of the utility function are negligible and that there is no state dependent utility. We report the corresponding MRS for three different values of risk-aversion γ. See text for details.

Figure 2: Identifying Variation in Local Transfers

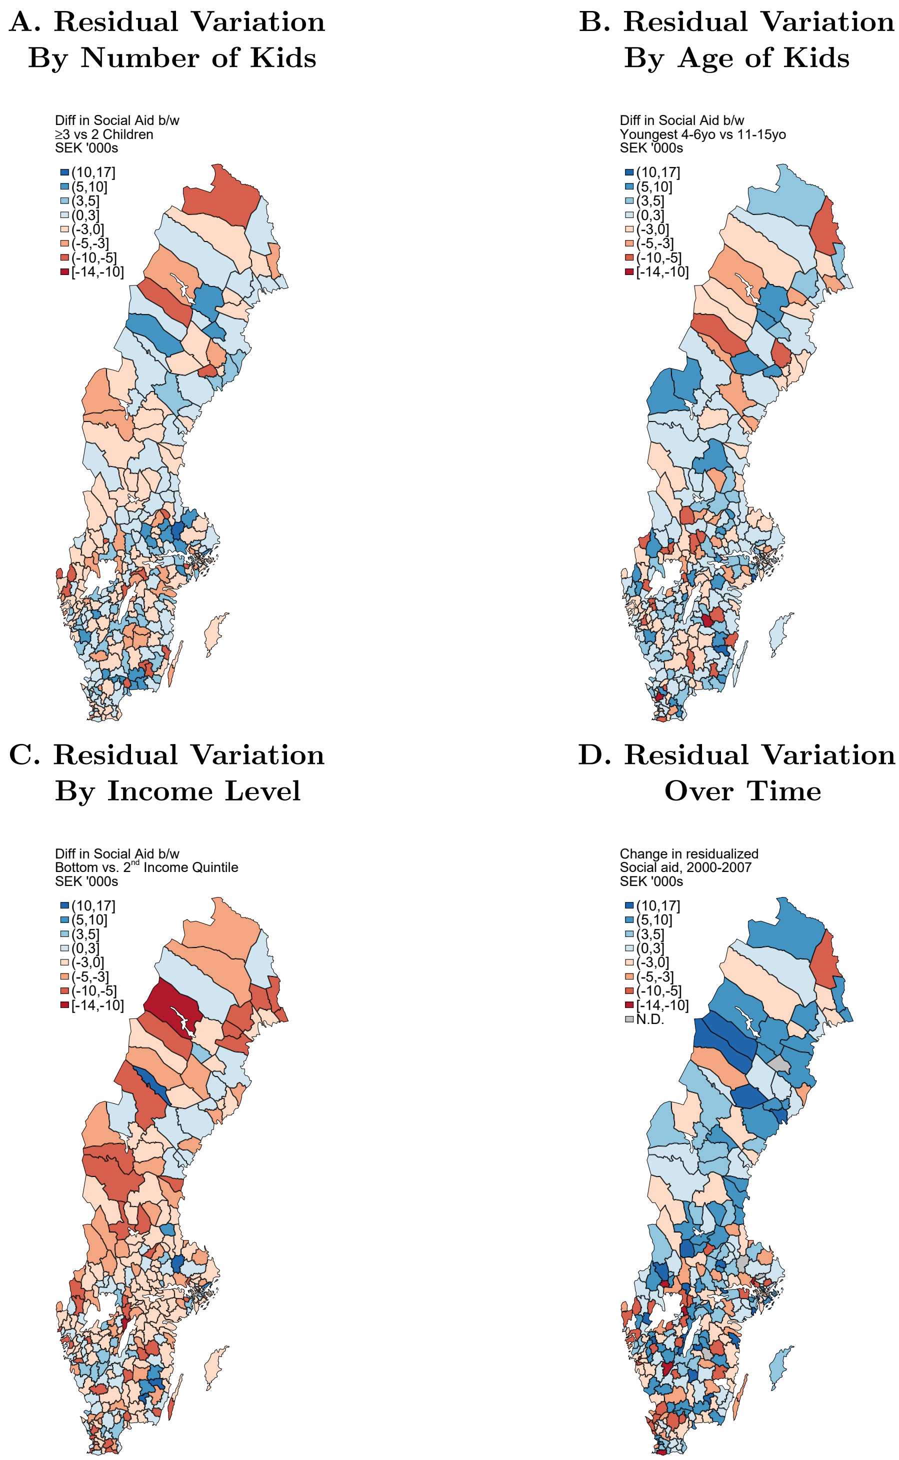

Notes: The Figure provides evidence of the variation in the way Swedish municipalities set local welfare transfers ("*social bidrag*"). By law, transfers are functions of characteristics $\mathbf{V}$, which include the number of dependents, the age of the dependent children, the liquid assets and income of the household: $B_{imt} = \sum_k \tau^k_{mt} \cdot V^k_{it}$. Municipalities are free to set τ^k_{mt} as they see fit. To visualize how τ^k differs across municipalities, we fix time in year 2006, and residualize welfare transfers using specification:

$$B_{im} = \sum_{j \neq k} \tau^j_m \cdot V^j_i + \bar{\tau}^k V^k_i + \tilde{B}_{im}$$

We then plot $\mathbf{E}[\tilde{B}_{im}|V^k = v^k] - \mathbf{E}[\tilde{B}_{im}|V^k = v'^k]$ for each municipality m. Differences in this statistics across municipalities reflect the fact that τ^k is set differently across municipalities. Panel A maps differences in $\tau^{children}$, the generosity of B as a function of the number of dependent children. The map shows the difference in average residual benefits $\tilde{B}_{imt}$ in thousands of SEK for a household with 2 children vs 3 or more children. Panel B shows variation in τ^{age}, i.e. how B varies with the age of dependent children. It shows the difference in average residual benefits $\tilde{B}_{im}$ for household with similar structure and number of dependents, whose youngest child is between 4 to 6 years old versus 11 to 15 years old. Panel C illustrates variation in the income phase-out rate τ^y. It plots the difference in average residual benefits $\tilde{B}_{im}$ for household with income in the bottom quintile vs the second quintile of the household income distribution. Panel D plots the growth rate of residualized transfers $\tilde{B}_{imt}$ between 2000 and 2007 across municipalities. See text for details.

Figure 3: Relationship Between First-Difference in Residual Local Transfers and First-Difference in Consumption by Employment Status

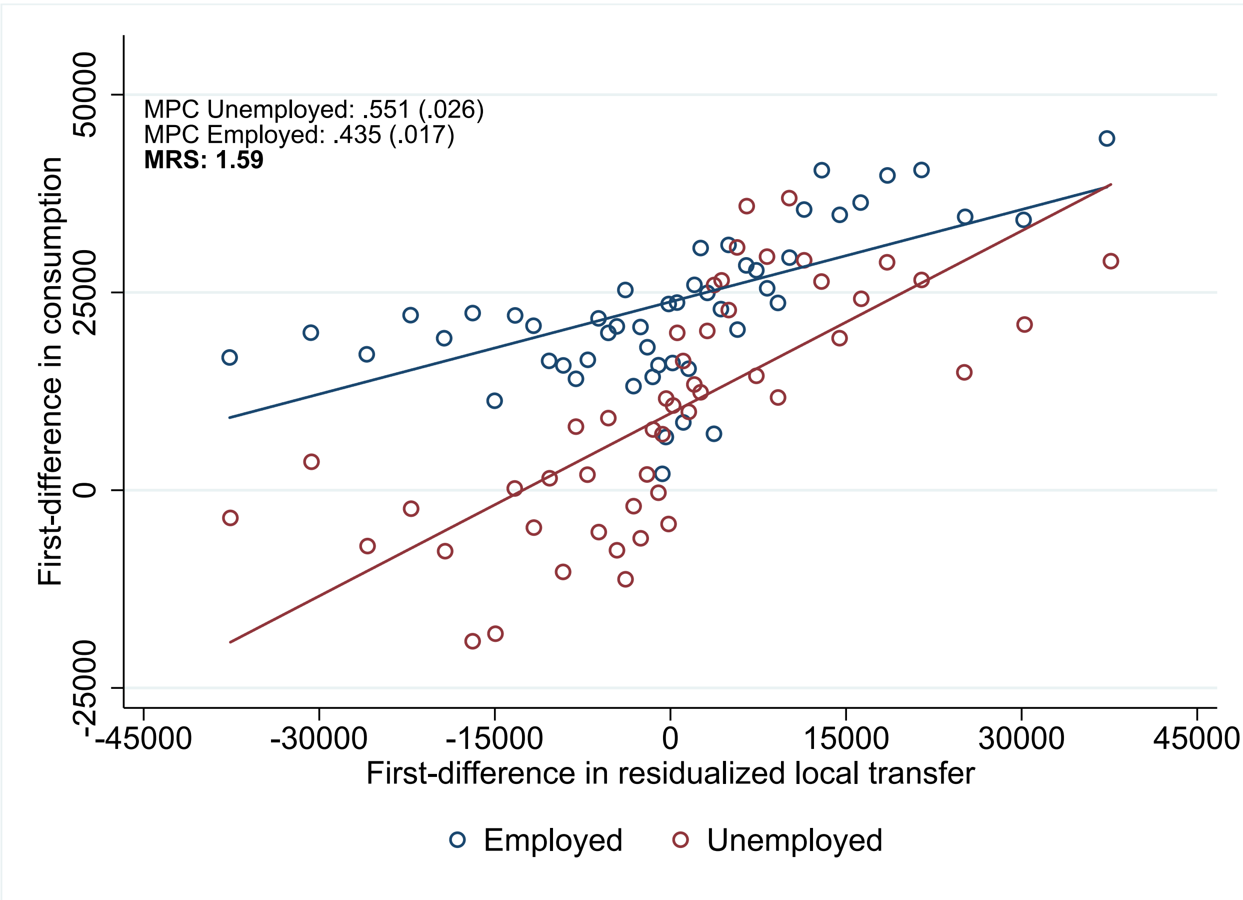

Notes: The graph is a bin-scatter plot of the relationship between the first-difference in residualized transfers $\tilde{B}_{imt}$ and the first-difference in annual household consumption, splitting the sample between households prior to the unemployment shock and households who experience unemployment in year t. The sample is the same as the one used for the CB approach in section 5 above (i.e. the sample contains only individuals who are becoming unemployed at some point, and are observed either employed or unemployed). The variable $\tilde{B}_{imt}$ is the residual from a regression of a household local welfare transfers B_{imt} on the vector of households characteristics $\mathbf{V}_{it}$, plus time and municipality fixed effects. We winsorize the first-difference in household consumption, omitting the top and bottom 5% of ΔC_{imt}. The graph shows a positive and quite linear relationship between consumption and transfers, indicative of a relatively large marginal propensity to consume out of transfers for both groups. The graph also displays a significantly steeper slope for the households in the unemployed group than for the households in the employed group, suggesting a significantly higher MPC for the former group compared to the latter.

Figure 4: Regression kink design: Effect of UI benefits variation on consumption at unemployment

A. UI Benefit Replacement Rate

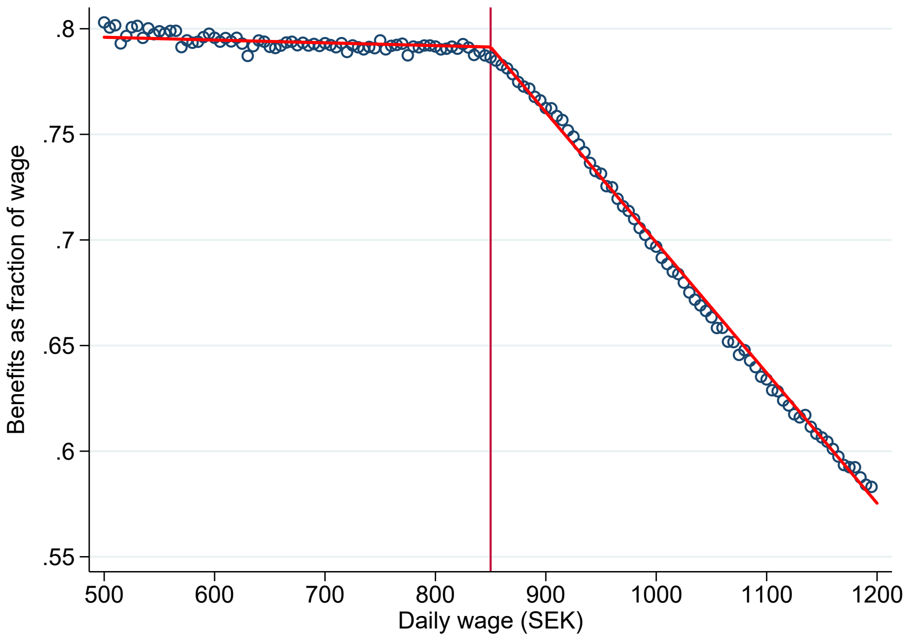

B. Consumption Drop at Unemployment

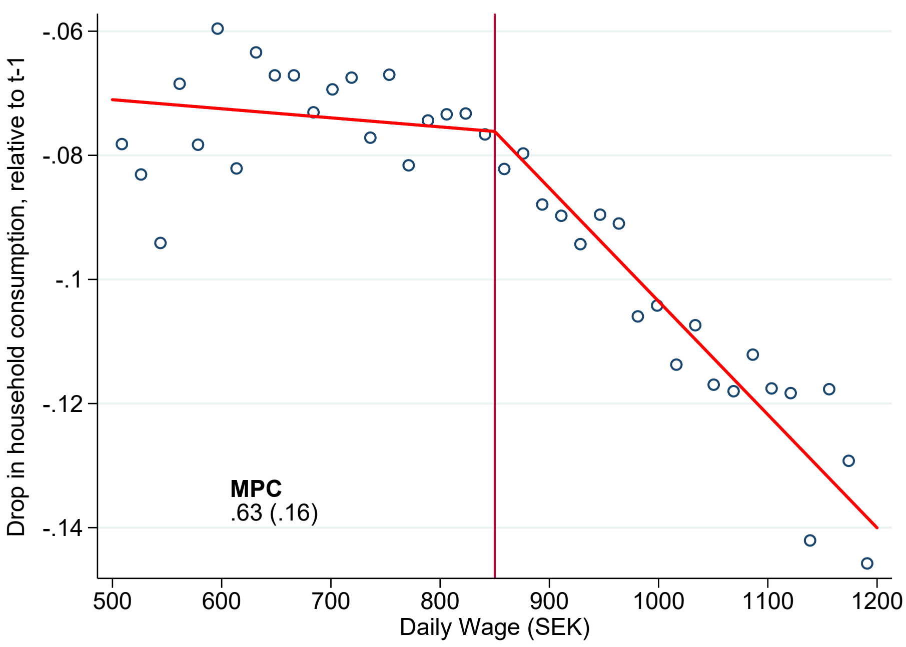

Notes: The Figure presents our RK design and main result. The design relies on the presence of a cap in UI benefits: the replacement rate is 80% of previous daily wage, up to a cap, when daily wages = 850SEK. Panel A plots, in our main sample over the period 2002 to 2007, a binscatter of the relationship between the daily wage and the average replacement rate. The latter is computed as the average benefit received during unemployment from the IAF data divided by the daily wage. The graph displays a clear kink at $w = 850SEK$, with the replacement rate declining sharply, as benefits are capped. We use this kinked relationship and treat it as a fuzzy RKD around the 850SEK threshold. Panel B plots the average drop in consumption ΔC_i at unemployment residualized on a set of dummies for the number of months spent unemployed M_i and the vector of characteristics $\mathbf{X}$ which includes year, age gender, education, region, family structure, and industry fixed effects. We scale consumption change by the average consumption in the year prior to unemployment in each bin. The graph shows evidence of a significant non-linearity in the relationship between daily wage and the consumption drop at unemployment. We also report on the panel our baseline RK estimates, using a bandwidth of 300, of the $MPC = \frac{\hat{\beta}_1}{\hat{\eta}_1} = .63$ (.16).

Figure 5: Non-parametric and Parametric RP Estimation

A. Expected Price vs. Insurance Coverage

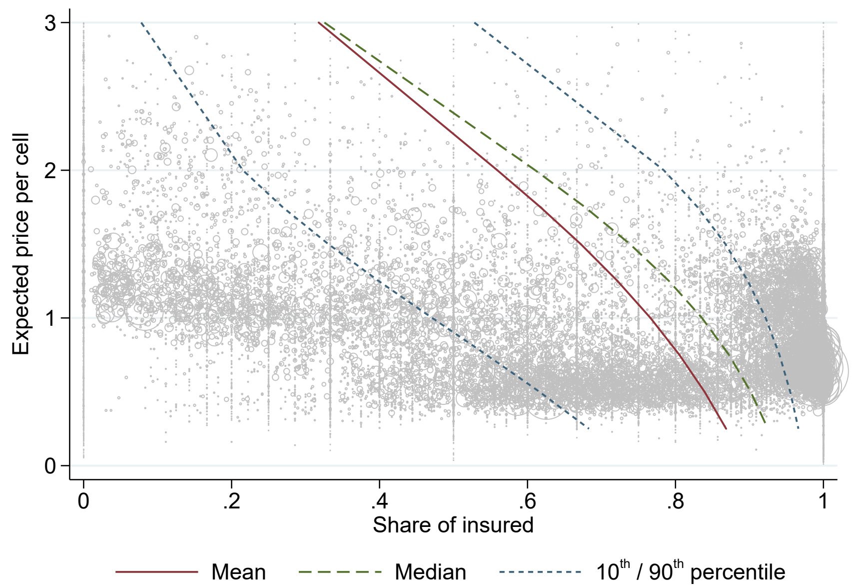

B. Parametric Estimation of MRS Distribution

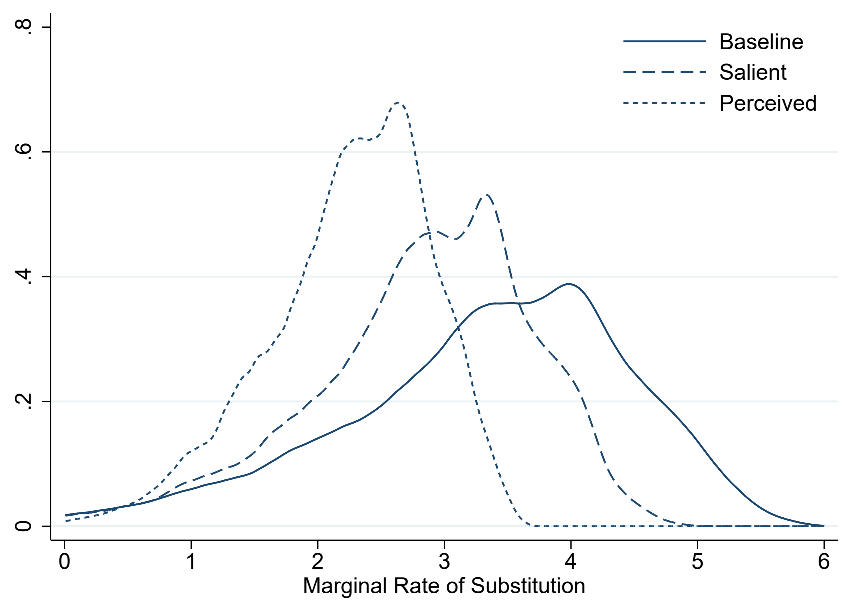

Notes: Panel A shows a scatter plot of the average expected price and share buying comprehensive insurance coverage for workers grouped by cells based on a rich set of observables. The expected price is calculated given the predicted risk under comprehensive coverage. Panel A of Appendix Figure E-2 shows the same plot using the predicted risk under basic coverage. The cells are defined by the intersections of 3 income groups, 3 age groups, 5 marital statuses, 20 regions, 9 education levels, 10 industries, 2 genders, 2 union membership statuses, 2 halves of firm level risk, 2 types of layoff histories (ever unemployed and never unemployed), and 2 halves of firm tenure ranks. Cell sizes on the graph are proportional to the number of individuals within them. Panel B shows the distribution of MRS estimated under different risk models. The solid line is based on the baseline risk model (see column (1) of Table 3). The long dashed line is based on the salient risk shifters (see column (2) of Table 3). The short dashed line used the perceived risk model (see column (3) of Table 3). In contrast with Table 3, the distributions are shown for the baseline sample of workers experiencing their first recorded unemployment spell between 2002 and 2007, as also used for the CB and MPC implementation. The full line in Panel A superimposes the implied demand from the parametric MRS estimation based on the perceived risk model (see column (3) in Table 3), showing the average share of individuals predicted to buy comprehensive coverage at different prices.

Figure 6: Comparison of MRS Estimates Across Different Approaches for the Baseline Sample

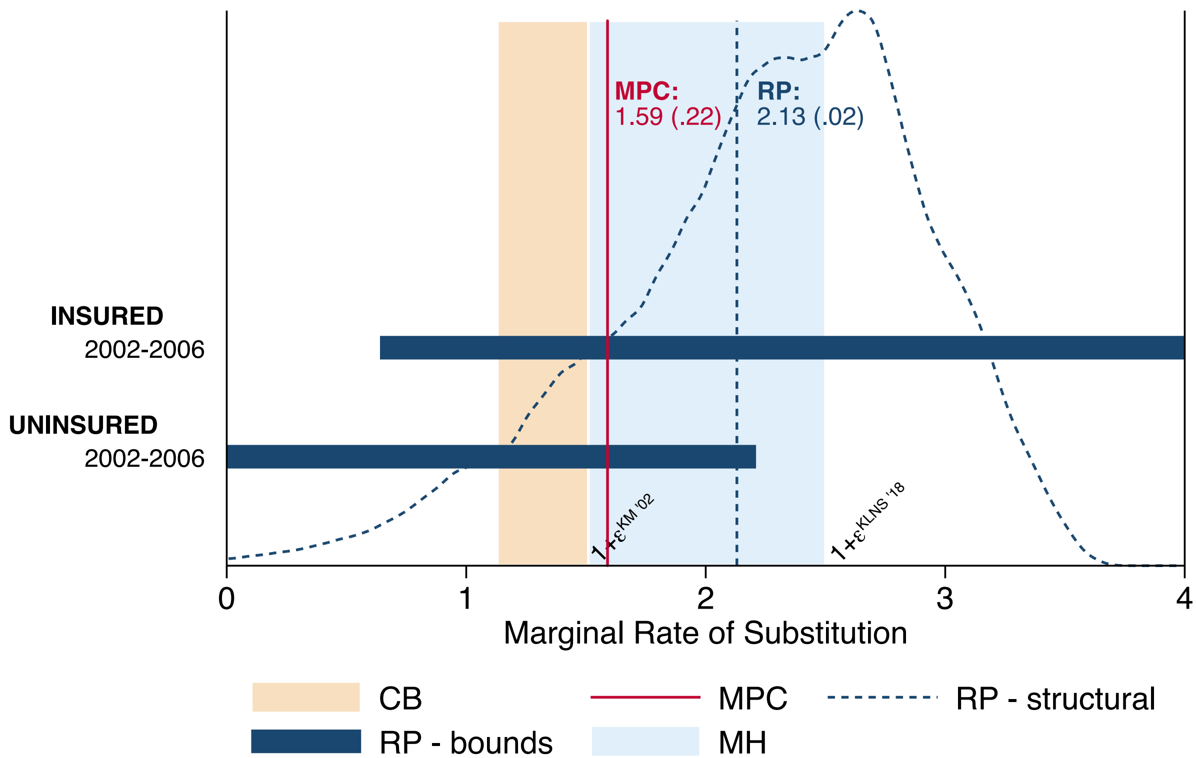

Notes: The graph compares the estimates of the MRS from the three different approaches, all implemented for the same baseline sample of workers. The region shaded in orange represents the range of MRS estimates from the CB approach, based on a consumption drop of 12.9% and relative risk aversion $\gamma \in [1; 4]$. The red line represents the lower-bound estimate on the MRS from the MPC approach, estimating the state-specific MPCs using the variation in local transfers. The dashed line shows the distribution of MRS estimated using the RP approach. The MRS estimation is using the perceived risk model, with the risks estimated under comprehensive UI coverage (see column (3) of Table 3. The mean MRS is represented by the vertical dashed line with standard error obtained using the delta method. The blue bars indicate the non-parametric upper and lower bound on MRS, as discussed in Section 7.1, using the predicted unemployment risk under basic and comprehensive coverage respectively. For comparison, the area shaded in grey represents a plausible range of moral hazard cost estimates, as discussed in Section 8.

Tables

Table 1: Baseline Sample: Summary statistics

	Mean	P10	P50	P90
	I. Demographics			
Age	37.3	26	37	51
Fraction female	0.42			
Fraction with kids	0.51			
Fraction with higher education	0.22			
	II. Income and Wealth, 2003 SEK (K)			
Individual Earnings	210.5	74.8	200.6	340.7
Household Disposable Income	323	107.5	265.5	586.7
Household Net wealth	377.4	-257.2	6.4	1348.6
Household Total Debt	420	0	210.2	1063.8
Household Liquid Assets	53.1	0	0	135
Household Consumption	317.6	104.9	247.8	599.7
	III. Unemployment			
Duration, days	381.8	89	287	819
Predicted risk of unemployment, days	8.5	3.9	6.4	13
Fraction with comprehensive coverage	0.91			
N	164,248			

Notes: The Table reports summary statistics for all individuals in our baseline sample in the year prior to their unemployment spell. Our baseline sample consists of 164,248 individuals experiencing their first recorded unemployment spells. Individuals must be between age 25-55 at the time of their first unemployment spell, and their first spell happens between 2002 and 2007. We exclude households where more than one member experiences an unemployment spell between 2002 and 2007. We compute layoffs from the PES data and exclude quits. We restrict the sample to individuals who are eligible for any UI coverage (basic or comprehensive) according to the 6 months work requirement prior to being laid-off. We further restrict the sample to individuals who are unemployed in December in the year of being laid off, as this is the month when all other demographics, income, tax and wealth information are observed and reported in the registry data. Earnings, income, and wealth are all measured in constant 2003 SEK, in the year prior to the unemployment spell. Household disposable income includes all earnings and income plus all transfers net of taxes. Liquid assets are total household bank holdings in liquid accounts. Debt is total household debt including student loans. Consumption is annual total expenditures at the household level from our registry-based measure (see text and Kolsrud et al. [2017] for details), where we fix composition of the household as of the year prior to being laid off. Duration of unemployment is the duration of the actual spell. Predicted risk is the measure obtained from our zero-inflated Poisson model of predicted total number of days spent unemployed in year t based on observables in year $t-1$. See text for details.

Table 2: Marginal Propensity to Consume Out of Local Welfare Transfers By Unemployment Status

	First Difference in Household Consumption						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
MPC employment	0.435	0.439	0.414	0.413	0.431	.421	.421
95% CI	[.402,.468]	[.405,.472]	[.380,.447]	[.377,.448]	[.398,.464]	[.385,.456]	[.387,.455]
MPC unemployment	0.551	0.544	0.519	0.454	0.547	.500	.509
95% CI	[.500,.601]	[.494,.594]	[.468,.570]	[.404,.505]	[.497,.598]	[.450,.550]	[457,.562]
p-value Test $MPC_u = MPC_E$	0.001	0.002	0.002	0.224	0.001	.023	.011
Lower Bound on MRS	1.591	1.528	1.530	1.185	1.596	1.377	1.430
SE (delta method)	(0.262)	(0.248)	(0.254)	(0.198)	(0.262)	(.234)	(.238)
SE (bootstrap)	(0.218)	(0.204)	(0.211)	(0.156)	(0.218)	(.181)	(.200)
Residualization:							
$\mathbf{V}_{it}$	×	×	×	×	×	×	×
Municipality FE	×	×	×	×	×	×	×
Income × Municipality FE		×		×			
Wealth × Municipality FE							
Family type × Municipality FE			×	×			
# & age child. × Municipality FE			×	×			
Year × Municipality FE					×	×	×
Income ×Year × Municipality FE						×	
Wealth × Year × Municipality FE							×
Observations	89673	89673	89673	89673	89673	89673	89673
Adjusted R-squared	0.026	0.026	0.024	0.021	0.025	0.023	0.023

Notes: The Table reports estimates of the marginal propensity to consume out of variation in local welfare transfers when employed and when unemployed. The sample is the same in the CB approach from section 5 above (i.e. the sample contains only individuals who are becoming unemployed at some point, and are observed either employed or unemployed). Identification exploits variation in transfers within municipality within household that comes from differences in the schedules τ_m^k. The table reports our results for the MPC out of local transfers when employed $\hat{\mu}_e$ and when unemployed $\hat{\mu}_e + \hat{\mu}_u$, from specification (6.2) estimated in first-differences. Column (1) corresponds to our baseline specification, when we residualize local welfare transfers on the vector of characteristics $\mathbf{V}$ plus year and municipality fixed effects. This residualization exploits variation stemming from τ_m^k, for all k as well as local reforms in the schedules over time. In the remainder of the table, we explore the sensitivity of our results to exploiting sources of variation stemming from specific τ^k. This is done by adding additional controls in the residualization to shut down particular dimensions of variations in local transfers, as indicated at the bottom of the table. Details for the different specifications are provided in the main text. For all specifications, we also report our estimate for the lower bound on the MRS, which, following formula (9), is equal to $\frac{\hat{\mu}_e+\hat{\mu}_u}{1-(\hat{\mu}_e+\hat{\mu}_u)}/\frac{\hat{\mu}_e}{1-\hat{\mu}_e}$, and its corresponding standard error, using two alternative approaches: the Delta-method, and a block-bootstrap computation.

Table 3: Insurance Choice Model Estimation & Implied MRS Distribution

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Baseline	Salient				Perceived Risk				
Coefficient on price (γ)	-0.704	-0.909	-1.2	-1.175	-1.085		-1.197	-1.114	-0.947	-1.147
	(0.004)	(0.004)	(0.006)	(0.006)	(0.006)		(0.006)	(0.013)	(0.024)	(0.007)
Bottom						-1.474				
Income Quartile						(0.0081)				
2^{nd}						-1.14				
						(0.007)				
3^{rd}						-0.572				
						(0.011)				
Top						-1.003				
						(0.012)				
Financial Variables				×						×
Region and Industry FE					×					
UI Choice Persistence							×			
Cognitive ability										×
Observations	1,052,294	1,205,844	1,052,294	1,052,294	1,034,364	1,052,294	1,052,294	310,316	97,381	862,100
MRS Distribution										
Mean	3.13	2.43	2.13	2.2	2.39	2.91	2.14	2.17	2.38	2.37
10^{th}	0.91	0.4	0.89	0.87	0.82	0.91	0.89	0.97	0.88	1.11
25^{th}	2.55	1.73	1.8	1.73	1.82	1.95	1.8	1.86	2.05	1.88
50^{th}	3.51	2.81	2.33	2.4	2.59	2.69	2.34	2.38	2.66	2.55
75^{th}	4.21	3.43	2.73	2.91	3.2	3.72	2.74	2.76	3.09	3.06
90^{th}	4.72	3.87	3.03	3.25	3.65	5.44	3.04	3.03	3.39	3.41

Notes: The table reports estimates from a logit regression using the linear choice model, $X_{it}\beta - \gamma\tilde{p}\left(Z_{it}\right) + \varepsilon_{it} \geq 0$. The risk is predicted under comprehensive coverage. Appendix Table 4 reports on the same estimations using risk predicted under basic coverage. The choice model is estimated on the entire population, except in columns (8) and (9). The panels also report moments of the corresponding distributions of MRS. The MRS for each individual is then calculated as $MRS\left(X_{it}\right) = \frac{X_{it}\beta}{\gamma}$. Column (1) to (3) show the estimates for different risk models. Column (1) uses the baseline risk model, as discussed in Section 4. Column (2) uses only salient risk shifters to predict unemployment risk. Columns (3) uses the perceived risk model. That is, the risk is predicted under the baseline model, but adjusted to account for plausible risk misperceptions, $\hat{\pi}_i = \bar{\pi} + .26(\pi_i - \bar{\pi})$, before calculating the expected price $\tilde{p}$ (see Appendix Figure E-1). Columns (4) to (10) perform further sensitivity checks of the estimation using the perceived risk model. Column (4) adds deciles of financial variables (net wealth, bank holdings and total debt) in the choice model. Region and Industry Fixed effects are added in column (5). In column (6) the expected price of UI is interacted with quartiles of income in $t-1$. Columns (7) to (9) present three alternatives to account for inertia in UI choices. Column (7) includes a dummy for persistence in UI choices between years $t-1$ and t. Column (8) restricts the sample to individuals that experienced a change in job at some point in the 2002-2007 period. In column (9) the sample is further restricted to years in which a change in job took place. Finally, a measure of cognitive ability administered by the Swedish Armed Forces on all enlisted individuals is included in column (10). This specification is used for the heterogeneity analysis reported in Figure E-3.

Proof of Lemma 1: Assuming an interior optimum for the consumption choice in each state, we have

$$\frac{\partial v_s\left(c_s, x_s\right)}{\partial c}=-p_s \frac{\partial v_s\left(c_s, x_s\right)}{\partial x}.$$

Now, by substituting $x_s=p_s\left[c_s-y_s\right]$ in the optimality condition, we obtain

$$\frac{\partial v_s\left(c_s, p_s\left[c_s-y_s\right]\right)}{\partial c}+p_s \frac{\partial v_s\left(c_s, p_s\left[c_s-y_s\right]\right)}{\partial x}=0,$$

and by implicit differentiation,

$$\left[\frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial c^2}+2 p_s \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x \partial c}+\left(p_s\right)^2 \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x^2}\right] d c_s- \\ \left[p_s \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x \partial c}+\left(p_s\right)^2 \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x^2}\right] d y_s=0.$$

Assuming separability in preferences, and using the optimality condition to substitute for p_s, we obtain

$$\left[\frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial c^2}-p_s \frac{\frac{\partial v_s\left(c_s, x_s\right)}{\partial c}}{\frac{\partial v_s\left(c_s, x_s\right)}{\partial x}} \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x^2}\right] d c_s-p_s \frac{\frac{\partial v_s\left(c_s, x_s\right)}{\partial c}}{\frac{\partial v_s\left(c_s, x_s\right)}{\partial x}} \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x^2} d y_s=0$$

and thus

$$\frac{d c_s}{d y_s}=\frac{p_s \sigma_s^x}{\sigma_s^c+p_s \sigma_s^x},$$

where $\sigma_s^c=-\frac{\partial^2 v_s\left(c_s, x_s\right) / \partial c^2}{\partial v_s\left(c_s, x_s\right) / \partial c}$ and $\sigma_s^x=\frac{\partial^2 v_s\left(c_s, x_s\right) / \partial x^2}{\partial v_s\left(c_s, x_s\right) / \partial x}$. Hence,

$$p_s=\frac{d c_s / d y_s}{1-d c_s / d y_s} \frac{\sigma_s^c}{\sigma_s^x}. \tag{25}$$

Putting this together for both states, we get the expression in Lemma 1.

Proof of Lemma 2: Assuming an interior optimum for the consumption choice in each state, we have

$$\frac{\partial v_s\left(c_s, x_s\right)}{\partial c}=-p_s \frac{\partial v_s\left(c_s, x_s\right)}{\partial x}.$$

Putting this together for both state, this immediately implieas

$$\frac{\frac{\partial v_u\left(c_u, x_u\right)}{\partial c}}{\frac{\partial v_e\left(c_e, x_e\right)}{\partial c}}=\frac{p_u}{p_e} \times \frac{\frac{\partial v_u\left(c_u, x_u\right)}{\partial x}}{\frac{\partial v_e\left(c_e, x_e\right)}{\partial x}}. \tag{26}$$

Appendix A Technical Appendix

This technical appendix provides the Taylor expansion underlying the CB approach and proofs of the Lemmas underlying the MPC approach in the static and dynamic model. We also set up the model extensions referred to in Section 3 and demonstrate the robustness of the MPC and RP approach in more detail.

A.1 Proofs and Derivations

Taylor expansion: A Taylor expansion of the marginal consumption utility when unemployed $\partial v_u / \partial c$ around (c_e, x_e) gives

$$\frac{\partial v_u(c_u, x_u)}{\partial c} = \frac{\partial v_u(c_e, x_e)}{\partial c} - \frac{\partial^2 v_u(c_e, x_e)}{\partial c^2}[c_e - c_u] - \frac{\partial^2 v_u(c_e, x_e)}{\partial c \partial x}[x_e - x_u] + ...$$

where the omitted higher-order terms depend on third- and higher-order derivatives of the utility function. Assuming that preferences are separable in consumption and resources ($\partial^2 v_s(c,x) / (\partial c \partial x) = 0$), we can approximate

$$\begin{aligned} \frac{\partial v_u(c_u, x_u)}{\partial c} &\cong \frac{\partial v_u(c_e, x_e)}{\partial c} - \frac{\partial^2 v_u(c_e, x_e)}{\partial c^2}[c_e - c_u] \\ &= \frac{\partial v_e(c_e, x_e)}{\partial c}\left[\frac{\frac{\partial v_u(c_e,x_e)}{\partial c}}{\frac{\partial v_e(c_e,x_e)}{\partial c}} - \frac{\frac{\partial^2 v_u(c_e,x_e)}{\partial c^2}}{\frac{\partial^2 v_e(c_e,x_e)}{\partial c^2}} \frac{\frac{\partial^2 v_e(c_e,x_e)}{\partial c^2}}{\frac{\partial v_e(c_e,x_e)}{\partial c}}[c_e - c_u]\right]. \end{aligned}$$

Using $\frac{\partial v_s(c,x)}{\partial c} = \theta_s^c \frac{\partial v(c,x)}{\partial c}$ and $\sigma^c = -\frac{\partial^2 v(c,x)/\partial c^2}{\partial v(c,x)/\partial c}$, we obtain

$$\frac{\frac{\partial v_u(c_u,x_u)}{\partial c}}{\frac{\partial v_e(c_e,x_e)}{\partial c}} \cong \frac{\theta_u^c}{\theta_e^c}[1 + \sigma^c [c_e - c_u].$$

This proves the approximation in equation (13) and simplifies to equation (6) for $\theta_u^c = \theta_e^c$.

Proof of Lemma 1: Assuming an interior optimum for the consumption choice in each state, we have

$$\frac{\partial v_s(c_s, x_s)}{\partial c} = -p_s \frac{\partial v_s(c_s, x_s)}{\partial x}.$$

Now, by substituting $x_s = p_s[c_s - y_s]$ in the optimality condition, we obtain

$$\frac{\partial v_s(c_s, p_s[c_s - y_s])}{\partial c} + p_s \frac{\partial v_s(c_s, p_s[c_s - y_s])}{\partial x} = 0,$$

and by implicit differentiation,

$$\left[\frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial c^2}+2 p_s \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x \partial c}+\left(p_s\right)^2 \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x^2}\right] d c_s- \left[p_s \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x \partial c}+\left(p_s\right)^2 \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x^2}\right] d y_s=0.$$

Assuming separability in preferences, and using the optimality condition to substitute for p_s, we obtain

$$\left[\frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial c^2}-p_s \frac{\frac{\partial v_s\left(c_s, x_s\right)}{\partial c}}{\frac{\partial v_s\left(c_s, x_s\right)}{\partial x}} \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x^2}\right] d c_s-p_s \frac{\frac{\partial v_s\left(c_s, x_s\right)}{\partial c}}{\frac{\partial v_s\left(c_s, x_s\right)}{\partial x}} \frac{\partial^2 v_s\left(c_s, x_s\right)}{\partial x^2} d y_s=0$$

and thus

$$\frac{d c_s}{d y_s}=\frac{p_s \sigma_s^x}{\sigma_s^c+p_s \sigma_s^x},$$

where $\sigma_s^c=-\frac{\partial^2 v_s\left(c_s, x_s\right) / \partial c^2}{\partial v_s\left(c_s, x_s\right) / \partial c}$ and $\sigma_s^x=\frac{\partial^2 v_s\left(c_s, x_s\right) / \partial x^2}{\partial v_s\left(c_s, x_s\right) / \partial x}$. Hence,

$$p_s=\frac{d c_s / d y_s}{1-d c_s / d y_s} \frac{\sigma_s^c}{\sigma_s^x}. \tag{27}$$

Putting this together for both states, we get the expression in Lemma 1.

Proof of Lemma 2: Assuming an interior optimum for the consumption choice in each state, we have

$$\frac{\partial v_s\left(c_s, x_s\right)}{\partial c}=-p_s \frac{\partial v_s\left(c_s, x_s\right)}{\partial x}.$$

Putting this together for both state, this immediately implieas

$$\frac{\frac{\partial v_u\left(c_u, x_u\right)}{\partial c}}{\frac{\partial v_e\left(c_e, x_e\right)}{\partial c}}=\frac{p_u}{p_e} \times \frac{\frac{\partial v_u\left(c_u, x_u\right)}{\partial x}}{\frac{\partial v_e\left(c_e, x_e\right)}{\partial x}}. \tag{28}$$

Derivation of MRS and MPC for Dynamic Application: At an interior optimum, the first-order condition with respect to consumption equals,

$$\frac{\partial v_s(c_{s,t})}{\partial c}=\beta R_{s,t}\left[\tilde{\pi}_s V_{u,t+1}^{\prime}(A_{s,t+1})+(1-\tilde{\pi}_s) V_{e,t+1}^{\prime}(A_{s,t+1})\right].$$

As we assume separability between the consumption decision and job search behavior, we can fully ignore the latter. Using $V_{s,t}^{\prime}(A_t)=\partial v_s\left(c_{s,t}\right) / \partial c$, the optimal consumption decision can be re-written as an Euler equation,

$$\frac{\partial v_s(c_{s,t})}{\partial c} = \beta R_{s,t}\left[\tilde{\pi}_s \frac{\partial v_u(c_{u,t+1})}{\partial c} + (1-\tilde{\pi}_s)\frac{\partial v_e(c_{e,t+1})}{\partial c}\right],$$

highlighting the trade-off between consumption today and consumption in the future, where $R_{s,t}$ determines the rate at which consumption in state s at time t can be substituted for consumption at time $t+1$. The MRS can be written as

$$\frac{\frac{\partial v_u(c_{u,t})}{\partial c}}{\frac{\partial v_e(c_{e,t})}{\partial c}} = \frac{R_{u,t}}{R_{e,t}}\frac{E_u(V'_{\tilde{s},t+1}(A_{u,t+1}))}{E_e(V'_{\tilde{s},t+1}(A_{e,t+1}))}.$$

As we discuss in the paper, conditional on the same history, the marginal cost of using future assets to increase consumption today is generally expected to be higher when unemployed than when employed, $E_u(V'_{\tilde{s},t+1}(A_{u,t+1})) > E_e(V'_{\tilde{s},t+1}(A_{e,t+1}))$. First, given lower current income, a worker will draw down their assets more when she is unemployed than when employed, $A_{u,t+1} < A_{e,t+1}$. Second, future income is expected to be lower too when unemployed and may be so permanently. Both forces increase the marginal continuation value at $t+1$ when unemployed rather than employed at t. To clarify this further, consider a worker becoming unemployed at time t in two extreme cases. In the first case, job loss implies a transitory shock in income from y_e to y_u that only lasts for one period, but the expected path of future earnings is the same. In the second case, job loss implies a permanent shock in income from y_e to y_u. In the first case, the marginal value of assets at $t+1$ is higher, because the worker will have used some of his or her assets in t. In the second case, the worker will not use any assets at t but reduce his or her consumption in accordance with the permanent drop in income. However, the marginal value of assets at $t+1$ is higher because of the decrease in permanent income. In both cases, the MPC approach will only pick up the difference in interest rates faced when unemployed vs. employed. In both cases, this price may will be higher when unemployed, because the unemployed will have depleted more of her assets in the former case or because she faces tighter borrowing constraints given the drop in future earnings in the latter case.

We verify that we get the same formula for the MPC in this model. From implicit differentiation of the Euler condition, we get

$$\frac{dA_{s,t+1}}{dy_{s,t}} = \frac{v''_s(c_{s,t})}{\beta R_{s,t}E_s\left[V''_{\tilde{s},t+1}(A_{s,t+1})\right] + v''_s(c_{s,t})/R_{s,t}}.$$

And from the budget constraint we know,

$$\frac{dc_{s,t}}{dy_{s,t}} = 1 - \frac{1}{R_{s,t}}\frac{dA_{s,t+1}}{dy_{s,t}}.$$

Hence,

$$\frac{\frac{dc_{s,t}}{dy_{s,t}}}{1-\frac{dc_{s,t}}{dy_{s,t}}} = R_{s,t}\left(\frac{E_s\left[V''_{\tilde{s},t+1}(A_{s,t+1})\right]}{E_s\left[V'_{\tilde{s},t+1}(A_{s,t+1})\right]} \Big/ \frac{v''_s(c_{s,t})}{v'_s(c_{s,t})}\right).$$

Derivation of Empirical Test: We provide the derivation underlying the empirical test. We allow for state-dependence in consumption preferences and expenditures and resource preferences:

$$v_s\left(c_s, x_s\right) = \theta_s^c v\left(c_s - \phi_s\right) - \theta_s^x h\left(x_s\right).$$

By implicit differentiation of the first-order condition,

$$\theta_s^c v'(c_s - \phi_s) - p_s \theta_s^x h'\left(p_s\left[c_s - y_s\right]\right) = 0,$$

we get

$$\begin{aligned}\left[\theta_s^c v''(c_s-\phi_s) - p_s^2\theta_s^x h''\left(x_s\right)\right]\Delta c + p_s^2\theta_s^x h''\left(x_s\right)\Delta y& \\ + v'(c_s-\phi_s)\Delta\theta^c - \theta_s^c v''(c_s-\phi_s)\Delta\phi - p_s h'\left(x_s\right)\Delta\theta^x &\cong 0.\end{aligned}$$

The change in consumption when becoming unemployed depends on the difference in income, but also on the difference in state-specific preferences and expenditures. Note that we ignore the change in prices, which naturally has a negligible effect on consumption if the substitution and income effect of a change in prices is similar. Re-arranging this condition and using the characterization of the MPC,

$$\frac{dc_s}{dy_s} = \frac{-p_s^2\theta_s^x h''\left(x_s\right)}{\theta_s^c v''(c_s-\phi_s) - p_s^2\theta_s^x h''\left(x_s\right)} = \frac{1}{1+\frac{\sigma^c}{p_s\sigma^x}},$$

we obtain

$$\Delta c = \frac{dc_s}{dy_s}\Delta y + \left[1-\frac{dc_s}{dy_s}\right]\left[\frac{\Delta\theta^c}{\theta_s^c}\frac{v'(c_s-\phi_s)}{-v''(c_s-\phi_s)} + \Delta\phi\right] - \frac{dc_s}{dy_s}\frac{1}{p_s}\frac{h'\left(x_s\right)}{h''\left(x_s\right)}\frac{\Delta\theta^x}{\theta_s^x}.$$

Ignoring state-dependence in resource preferences, we get expression (14) in the main text.

A.2 MPC Approach: Robustness

We consider extensions of our stylized model and study the robustness of the MPC approach.

State-specific Expenditures We can introduce state-specific expenditures, like work or job search-related expenditures, by requiring individuals to purchase an exogenous amount of consumption ϕ_s in any given state s. The setup is as follows,

$$\max_{c_s,x_s}\ u_s(c_s-\phi_s) - v_s(x_s) \text{ s.t. } c_s+\phi_s = y_s + \frac{1}{p_s}x_s \text{ for } s\in\{e,u\}.$$

Optimality is characterized by,

$$u_s'(c_s - \phi_s) = p_s v_s'(x_s),$$

which by implicit differentiation leads to the same expression for the MPC,

$$\frac{dc_s}{dy_s} = \frac{1}{1 - \frac{1}{p_s^2}\frac{u_s''(c_s-\phi_s)}{v_s''(x_s)}} = \frac{1}{1 + \frac{1}{p_s}\frac{\sigma_u^c}{\sigma_u^x}}.$$

In this model, state-specific expenditures will affect the observed consumption drop Δc_s between employment and unemployment, while only the drop in net consumption $\Delta\left[c_s - \phi_s\right]$ is relevant for the MRS. The ratio of MPC odds ratios, however, still identifies the relative prices and bounds the MRS as in our stylized model,

$$MRS = \frac{O_u^{mpc}}{O_e^{mpc}} \times \frac{\frac{\sigma_u^c}{\sigma_u^x}}{\frac{\sigma_e^c}{\sigma_e^x}} \times \frac{\frac{\partial v_u(c_u,x_u)}{\partial x_u}}{\frac{\partial v_e(c_e,x_e)}{\partial x_e}}.$$

Home Production Another reason for the measurement or utility of expenditures to be state-specific is that the way expenditures convert into utility-relevant consumption depends on the state. Examples are the substitution towards home production and lower shopping prices when more time is available. In both cases, a given level of expenditures provides more utility. We can model this as follows,

$$\max_{c_s,x_s} \; u_s(\eta_s c_s) - v_s(x_s) \;\; \text{s.t.} \;\; c_s = y_s + \frac{1}{p_s}x_s \text{ for } s \in \{e,u\},$$

where c_s are the observed expenditures and η_s scales the expenditures into utility-relevant consumption. Optimality is now characterized by

$$u_s'(\eta_s c_s) = \frac{p_s}{\eta_s} v_s'(x_s),$$

from which we can derive,

$$\frac{dc_s}{dy_s} = \frac{1}{1 - \frac{\eta_s^2}{p_s^2}\frac{u_s''(c_s)}{v_s''(x_s)}} = \frac{1}{1 + \frac{\eta_s}{p_s}\frac{\sigma_u^c}{\sigma_u^x}}.$$

The MPC depends on the state-specific consumption scalar η. The marginal propensity to *spend* out of income is smaller when consumption is cheaper, either because of the low price of increasing resources or because of the low price of the consumption goods. The state-specific prices will affect the observed drop in expenditures Δc_s between employment and unemployment, while it is the drop in consumption $\Delta\left[\eta_s c_s\right]$ that is relevant for the MRS. The ratio of the MPC odds ratios, however, still identifies the relative prices and bounds the MRS as defined in our stylized model,

$$MRS = \frac{O_u^{mpc}}{O_e^{mpc}} \times \frac{\frac{\sigma_u^c}{\sigma_u^x}}{\frac{\sigma_e^c}{\sigma_e^x}} \times \frac{\frac{\partial v_u(c_u,x_u)}{\partial x_u}}{\frac{\partial v_e(c_e,x_e)}{\partial x_e}}.$$

However, we should account for the fact that the return to a krona is different when unemployed vs. employed and scale the marginal rate of substitution to determine the value of a krona when unemployed vs. employed. That is, $\frac{\eta_u u_u'(\eta_u c_u)}{\eta_e u_e'(\eta_e c_e)}$. The MPC ratio is still a lower bound on this object if $\eta_u > \eta_e$ (and assumptions 0-2 hold), but it may no longer be a lower bound when $\eta_e > \eta_u$.

Multiple consumption categories To study the robustness of our approach when allowing for multiple consumption categories, we introduce a second good into our setup:

$$\max_{c_{s,1},c_{s,2},x_s} u_s(c_{s,1}) + g_s(c_{s,2}) - v_s(x_s) \text{ s.t. } c_{s,1} + q_s c_{s,2} = y_s + \frac{1}{p_s} x_s \text{ for } s \in \{e, u\}.$$

We allow the utility function and the relative prices of the consumption goods to be different, but assume separability to keep expressions tractable. This is also a reduced-form way to think about expenditures on durable goods, for which the curvature of preferences is smaller as their impact does not only depend on the current investments, but also on past and future investments, or about committed expenditures, which affect the prefence curvature over the non-committed expenditures that can be changed in response to shocks.

Optimality is characterized by

$$\begin{aligned} F(c_{s,1}, c_{s,2}; y) &= u_s'(c_{s,1}) - p_s v_s'(p_s c_{s,1} + p_s q_s c_{s,2} - p_s y_s) = 0, \\ G(c_{s,1}, c_{s,2}; y) &= g_s'(c_{s,2}) - p_s q_s v_s'(p_s c_{s,1} + p_s q_s c_{s,2} - p_s y_s) = 0. \end{aligned}$$

Using implicit differentiation, we find

$$\begin{aligned} \frac{dc_{s,1}}{dy_s} &= -\frac{\begin{vmatrix} F_{c_{s,2}} & F_{y_s} \\ G_{c_{s,2}} & G_{y_s} \end{vmatrix}}{\begin{vmatrix} F_{c_{s,2}} & F_{c_{s,1}} \\ G_{c_{s,2}} & G_{c_{s,1}} \end{vmatrix}} = -\frac{\begin{vmatrix} -p_s^2 q_s v_s''(x_s) & p_s^2 v_s''(x_s) \\ g_s''(c_{s,2}) - p_s^2 q_s^2 v_s''(x_s) & p_s^2 q_s v_s''(x_s) \end{vmatrix}}{\begin{vmatrix} -p_s^2 q_s v_s''(x_s) & u_s''(c_{s,1}) - p_s^2 v_s''(x_s) \\ g_s''(c_{s,2}) - p_s^2 q_s^2 v_s''(x_s) & -p_s^2 q_s v_s''(x_s) \end{vmatrix}} \\ &= -\frac{-p_s^4 q_s^2 v_s''(x_s)^2 - p_s^2 v_s''(x_s) g_s''(c_{s,2}) + p_s^4 q_s^2 v_s''(x_s)^2}{p_s^4 q_s^2 v_s''(x_s)^2 - u_s''(c_{s,1}) g_s''(c_{s,2}) + u_s''(c_{s,1}) p_s^2 q_s^2 v_s''(x_s) + g_s''(c_{s,2}) p_s^2 v_s''(x_s) - p_s^4 q_s^2 v_s''(x_s)^2} \\ &= \frac{1}{1 - \frac{u_s''(c_{s,1})}{p_s^2 v_s''(x_s)} + \frac{u_s''(c_{s,1}) q_s^2}{g_s''(c_{s,2})}} \\ &= \frac{p_s \frac{\sigma_s^x}{\sigma_s^{c1}}}{1 + p_s[\frac{\sigma_s^x}{\sigma_s^{c1}} + \frac{\sigma_s^x}{\sigma_s^{c2}} MRS_s^{c_1,c_2}]}, \end{aligned}$$

where $MRS_s^{c_1,c_2} = \frac{u_s'(c_{s,1})}{g_s'(c_{s,2})} (= \frac{1}{q_s})$. By symmetry,

$$\frac{dc_{s,2}}{dy_s} = \frac{p_s \frac{\sigma_s^x}{\sigma_s^{c2}} MRS_s^{c_1,c_2}}{1 + p_s[\frac{\sigma_s^x}{\sigma_s^{c1}} + \frac{\sigma_s^x}{\sigma_s^{c2}} MRS_s^{c_1,c_2}]}.$$

Therefore,

$$\frac{d[c_{s,1}+c_{s,2}]}{dy_s}=\frac{p_s[\frac{\sigma_s^x}{\sigma_s^{c1}}+\frac{\sigma_s^x}{\sigma_s^{c2}}MRS_s^{c_1,c_2}]}{1+p_s[\frac{\sigma_s^x}{\sigma_s^{c1}}+\frac{\sigma_s^x}{\sigma_s^{c2}}MRS_s^{c_1,c_2}]}.$$

Hence, while for the CB approach knowing the role of the response for different consumption categories $\Delta c_j/c_j$ is relevant to know which preference parameter to use (e.g., committed expenditures, durable goods), the MPC continues to depend on the state-specific price. The ratio of MPC odds ratios identifies the price ratio if the curvature of preferences over the consumption goods relative to the resources used remains constant across states. Translating this into the MRS, we obtain

$$MRS=\frac{O_u^{mpc}}{O_e^{mpc}}\times\frac{\frac{\sigma_e^x}{\sigma_e^{c1}}+\frac{\sigma_e^x}{\sigma_e^{c2}}MRS_e^{c_1,c_2}}{\frac{\sigma_u^x}{\sigma_u^{c1}}+\frac{\sigma_u^x}{\sigma_u^{c2}}MRS_u^{c_1,c_2}}\times\frac{\frac{\partial v_u(x_u)}{\partial x_u}}{\frac{\partial v_e(x_e)}{\partial x_e}}.$$

The second factor cancels out again for CARA preferences $v_s(c_s)=\frac{\exp(-\tilde{\sigma}^{c_1}c_{s,1})}{-\tilde{\sigma}^{c_1}}+\frac{\exp(-\tilde{\sigma}^{c_2}c_{s,2})}{-\tilde{\sigma}^{c_2}}$, assuming $q_u=q_e$.

It is useful to note that the above implementation uses total consumption expenditures. If we only observe a partial measure of consumption expenditures, we would have

$$\frac{\frac{dc_{s,1}}{dy_s}}{1-\frac{dc_{s,1}}{dy_s}}=\frac{\frac{p_s\sigma_s^x}{\sigma_s^{c1}}}{1+p_sMRS_s^{c_1,c_2}\frac{\sigma_s^x}{\sigma_s^{c2}}}=\frac{\frac{p_s\sigma_s^x}{\sigma_s^{c1}}}{1+p_sq_s\frac{\sigma_s^x}{\sigma_s^{c2}}}.$$

Now the impact of a higher price p_s on the measured MPC is smaller, since the opposing effect through the denominator (unless of course when $q_s=1/p_s$). As a consequence, the ratio of partial MPC odds ratios would underestimate the price ratio p_u/p_e and thus provides a weaker lower bound on the MRS. This indicates that a more comprehensive measure of expenditures is preferable.

Multiple Resources We now extend our model to allow for different means to smooth consumption. We model this in a parsimonious way by making the price of using extra resources endogenous to the level of resources used, $p(x)$. Workers first use the resources that are available at the lowest price (relative to its utility costs). That is, $p'(x)\geq 0$. This also endogenously introduces $p(x_u)\geq p(x_e)$ as $x_u>x_e$. The setup is as follows,

$$\max_{c_s,x_s}\ u_s(c_s)-v_s(x_s)\ \ s.t.\ \ c_s=y_s+\int_0^{x_s}\frac{1}{p(z)}dz \text{ for } s\in\{e,u\},$$

where $p_s'(z)\geq 0$. Optimality is characterized by

$$u_s'(c_s)=p(x_s)v_s'(x_s).$$

From this and the budget constraint, we derive

$$dc_s = dy_s + \frac{1}{p(x_s)} dx_s,$$
$$v_s''(x_s) dx_s = u_s''(c_s) \frac{1}{p(x_s)} dc_s - u_s'(c_s) \frac{p'(x_s)}{p(x_s)^2} dx_s.$$

Combining these, we find

$$\frac{dc_s}{dy_s} = \frac{1 + \frac{p'(x_s)}{p(x_s)^2} \frac{u_s'(c_s)}{v_s''(x_s)}}{1 - \frac{1}{p(x_s)^2} \frac{u_s''(c_s)}{v_s''(x_s)} + \frac{u_s'(c_s) p'(x_s)}{v_s''(x_s) p(x_s)^2}} = \frac{1}{1 + \frac{1}{p(x_s)} \frac{\sigma_s^c}{\sigma_s^x + \varepsilon_s^{p,x}} / x_s},$$

where $\varepsilon_s^{p,x} = \frac{p'(x_s)}{p(x_s)} x_s$. This shows that the MPC reveals the price of the resource used at the margin, but also depends on the curvature in preferences and any potential changes in prices when changing the resources used. Translating this into the MRS, we get

$$MRS = \frac{O_u^{mpc}}{O_e^{mpc}} \times \frac{\frac{\sigma_e^x + \varepsilon_e^{p,x}/x_e}{\sigma_e^c}}{\frac{\sigma_u^x + \varepsilon_u^{p,x}/x_u}{\sigma_u^c}} \times \frac{\frac{\partial v_u(x_u)}{\partial x_u}}{\frac{\partial v_e(x_e)}{\partial x_e}}.$$

When the local change in price is small or $\varepsilon_e^{p,x}/x_e > \varepsilon_u^{p,x}/x_u$, the same assumptions as in Proposition 1 imply that the the ratio of MPC odds ratios provides a lower bound on the MRS. This clearly holds when the price elasticity is constant.

A further extension could be to consider multiple resources in the dynamic extension, where to smooth consumption the resources used at the margin can be devoted *ex ante*, *contemporaneously* or from the *future*. While in principle this does not need to be distinguished, the means that are used to smooth consumption in response to job loss can still guide the implementation of the MPC. For example, the MPC approach can account for the use of precautionary means, but this would require using anticipated variation in state-contingent income so that workers can adjust precautionary means in response. Similarly, if the relevant margin is intertemporal, we would need transitory variation in income. While, if the relevant margin is contemporaneous, it is sufficient to have variation in state-contingent income, regardless of whether it is anticipated or transitory. In practice, there may be a trade-off between finding anticipated and exogenous variation in income, but this issue becomes less binding if the relevant margins of adjustment are *ex post* means of consumption smoothing (e.g., household labour supply, credit).

A.3 Alternative Optimization Approaches

We first present the alternative approach proposed by Chetty [2008] and Landais [2015] in the context of our model. We assume the following separable preferences:

$$V = \pi(z) v_u(c_u, x_u) + (1 - \pi(z)) v_e(c_e, x_e) - z.$$

The first-order condition with respect to effort z equals

$$\pi'(z)\left[v_u(c_u,x_u)-v_e(c_e,x_e)\right]=1.$$

Consider now a change in y_u. By implicit differentiation of this FOC, we find

$$\pi''(z)\left[v_u(c_u,x_u)-v_e(c_e,x_e)\right]dz+\pi'(z)\left[\frac{\partial v_u}{\partial c_u}dc_u-\frac{\partial v_u}{\partial x_u}dx_u\right]=0.$$

Rewriting this, we get

$$\frac{d\pi(z)}{dy_u}=\frac{-\left(\pi'(z)\right)^2\left[\frac{\partial v_u(c_u,x_u)}{\partial c_u}\frac{dc_u}{dy_u}-\frac{\partial v_u(c_u,x_u)}{\partial x_u}\frac{dx_u}{dy_u}\right]}{\pi''(z)\left[v_u(c_u,x_u)-v_e(c_e,x_e)\right]}.$$

From the budget constraint, we have

$$\frac{dx_u}{dy_u}=\left[\frac{dc_u}{dy_u}-1\right]p_u.$$

From the optimality condition, we also have

$$\frac{\partial v_u(c_u,x_u)}{\partial c_u}=-p_u\frac{\partial v_u(c_u,x_u)}{\partial x_u}.$$

Using these expressions, we can rewrite

$$\begin{aligned}\frac{d\pi(z)}{dy_u} &= \frac{-\left(\pi'(z)\right)^2\left[\frac{\partial v_u(c_u,x_u)}{\partial c_u}\frac{dc_u}{dy_u}-\frac{\partial v_u(c_u,x_u)}{\partial x_u}\left[\frac{dc_u}{dy_u}-1\right]p_u\right]}{\pi''(z)\left[v_u(c_u,x_u)-v_e(c_e,x_e)\right]}\\ &= \frac{-\left(\pi'(z)\right)^2\frac{\partial v_u(c_u,x_u)}{\partial c_u}\left[2\frac{dc_u}{dy_u}-1\right]}{\pi''(z)\left[v_u(c_u,x_u)-v_e(c_e,x_e)\right]}.\end{aligned}$$

We can do the analogue derivation for y_e and we immediately get the expression in the main text:

$$\frac{\frac{d\pi(z)}{dy_u}}{\frac{d\pi(z)}{dy_e}}=\frac{\frac{\partial v_u(c_u,x_u)}{\partial c_u}\left[2\frac{dc_u}{dy_u}-1\right]}{\frac{\partial v_e(c_e,x_e)}{\partial c_e}\left[2\frac{dc_e}{dy_e}-1\right]}.$$

We can also characterize the alternative approach proposed by Fadlon and Nielsen [2018] and Hendren [2017] in the context of our model. From the optimality conditions, we have

$$\frac{\frac{\partial v_u(c_u,x_u)}{\partial c_u}}{\frac{\partial v_e(c_e,x_e)}{\partial c_e}}=\frac{p_u}{p_e}\frac{\frac{\partial v_u(c_u,x_u)}{\partial x_u}}{\frac{\partial v_e(c_e,x_e)}{\partial c_e}}.$$

Using the analogue Taylor expansion for $\frac{\partial v_u(c_u,x_u)}{\partial x_u}$ as for $\frac{\partial v_u(c_u,x_u)}{\partial c_u}$, as shown in Appendix A.1, assuming again $\frac{\partial v_s(c,x)}{\partial x} = \theta_s^x \frac{\partial v(c,x)}{\partial x}$ and using notation $\sigma^x = \frac{\partial^2 v(c,x)/\partial x^2}{\partial v(c,x)/\partial x}$, we immediately obtain

$$\frac{\frac{\partial v_u(c_u,x_u)}{\partial c_u}}{\frac{\partial v_e(c_e,x_e)}{\partial c_e}} \cong \frac{p_u}{p_e}\frac{\theta_u^x}{\theta_e^x}\left[1+\sigma^x\left[x_u-x_e\right]\right].$$

A.4 RP Approach: Robustness

We show how the RP approach generalizes to a model with a discrete insurance choice.

As discussed in the main text, the marginal impact of extra insurance on the individual's expected utility equals

$$dV = \pi(z)\frac{\partial v_u(c_u,x_u)}{\partial c}\frac{1}{p_u} - (1-\pi(z))\frac{\partial v_e(c_e,x_e)}{\partial c}\frac{1}{p_e}.$$

While the individual may change her search behavior z and consumption smoothing behavior x_s in response, the impact on the optimizing individual's welfare is of second-order importance by the envelope theorem. Hence, when offered insurance at the margin, an individual will buy this if

$$dV \geq 0 \Leftrightarrow \frac{\frac{\partial v_u(c_u,x_u)}{\partial c}}{\frac{\partial v_e(c_e,x_e)}{\partial c}} \geq \frac{p_u}{p_e}\times\frac{1-\pi(z)}{\pi(z)}.$$

Consider now the discrete choice between the securities $\left(x_u^1, x_e^1\right)$ and $\left(x_u^0, x_e^0\right)$ where $\Delta x_u = x_u^1 - x_u^0 = -\Delta x_e = x_e^0 - x_u^1 > 0$. So plan 1 provides more coverage than plan 0. The state-specific prices are still p_u and p_e respectively.

We can write the welfare gain as the integral over the marginal gains when moving from plan x^0 to plan x^1 at rate $dx_u = -dx_e$. For each marginal gain, we can invoke the envelope theorem to conclude that only the direct impact on the worker's expected utility will be of first-order. Each marginal gain, however, will be evaluated at the counterfactual effort and consumption levels the worker would choose given the intermediate plan $(\tilde{x}_u, \tilde{x}_e)$. Using short-hand notation to denote these choices, we can write

$$\Delta V = \int_{x_u^0}^{x_u^1}\left[\pi(\tilde{x})\frac{\partial v_u(c_u(\tilde{x}))}{\partial c}\frac{1}{p_u} - (1-\pi(\tilde{x}))\frac{\partial v_e(c_e(\tilde{x}))}{\partial c}\frac{1}{p_e}\right]d\tilde{x}_u.$$

We know the welfare gain is positive for workers who choose plan 1. We can find intermediate consumption levels $\bar{c}_u \in \left[c_u\left(\tilde{x}^0\right), c_u\left(\tilde{x}^1\right)\right]$ and $\bar{c}_e \in \left[c_e\left(\tilde{x}^1\right), c_e\left(\tilde{x}^0\right)\right]$, such that

$$\Delta V = \left[\int_{x_u^0}^{x_u^1}\pi(\tilde{x})\,d\tilde{x}_u\right]\frac{\partial v_u(\bar{c}_u)}{\partial c}\frac{1}{p_u} - \left[\int_{x_u^0}^{x_u^1}(1-\pi(\tilde{x}))d\tilde{x}_u\right]\frac{\partial v_e(\bar{c}_e)}{\partial c}\frac{1}{p_e}$$

and thus

$$\Delta V \geq 0 \Leftrightarrow \frac{\frac{\partial v_u(\bar{c}_u)}{\partial c}}{\frac{\partial v_e(\bar{c}_e)}{\partial c}} \geq \frac{p_u}{p_e}\frac{\int_{x_u^0}^{x_u^1}(1-\pi(\tilde{x}))d\tilde{x}_u}{\int_{x_u^0}^{x_u^1}\pi(\tilde{x})\,d\tilde{x}_u} \Rightarrow \frac{\frac{\partial v_u(\bar{c}_u)}{\partial c}}{\frac{\partial v_e(\bar{c}_e)}{\partial c}} \geq \frac{p_u}{p_e}\frac{1-\pi\left(x^1\right)}{\pi\left(x^1\right)}.$$

Due to moral hazard, the unemployment risk is increasing in coverage and thus highest under plan x^1. Hence,

$$\Delta V \geq 0 \Rightarrow \pi\left(x^1\right)\left[\int_{x_u^0}^{x_u^1} \frac{\partial v_u\left(c_u\left(\tilde{x}\right)\right)}{\partial c} d\tilde{x}_u\right]\frac{1}{p_u} - \left(1-\pi\left(x^1\right)\right)\left[\int_{x_u^0}^{x_u^1} \frac{\partial v_e\left(c_e\left(\tilde{x}\right)\right)}{\partial c} d\tilde{x}_u\right]\frac{1}{p_e} \geq 0.$$

So for workers who buy plan 1, we have

$$\frac{\int_{x_u^0}^{x_u^1} \frac{\partial v_u(c_u(\tilde{x}))}{\partial c} d\tilde{x}_u}{\int_{x_u^0}^{x_u^1} \frac{\partial v_e(c_e(\tilde{x}))}{\partial c} d\tilde{x}_u} \geq \frac{p_u}{p_e} \times \frac{1-\pi\left(x^1\right)}{\pi\left(x^1\right)}.$$

The expected price using the predicted risk under x^1 provides a lower bound on the 'average' MRS for workers opting for x^1. This average MRS captures the ratio of the average marginal utility gain from increasing consumption when unemployed in moving from low-coverage x^0 to high-coverage x^1 to the average marginal utility losses from the corresponding decrease in consumption when employed. The same argument makes that the expected price using the predicted risk under x^0 provides an upper bound on the MRS for workers opting for x^0.

We also note that we simplified the unemployment risk to be binary. In practice, unemployment risk is more complex with people differing in their probability of job loss and the time spent unemployed conditional on job loss. Moreover, the benefits typically depend on the length of the ongoing unemployment spell. All of this affect the value and thus willingness to buy UI. See also Kolsrud et al. [2018].

Appendix B Consumption-Based Approach: Additional Results & Figures

B.1 Baseline Implementation: Details

Figure B-1: Consumption dynamics around start of unemployment spell: Treated & matched control households

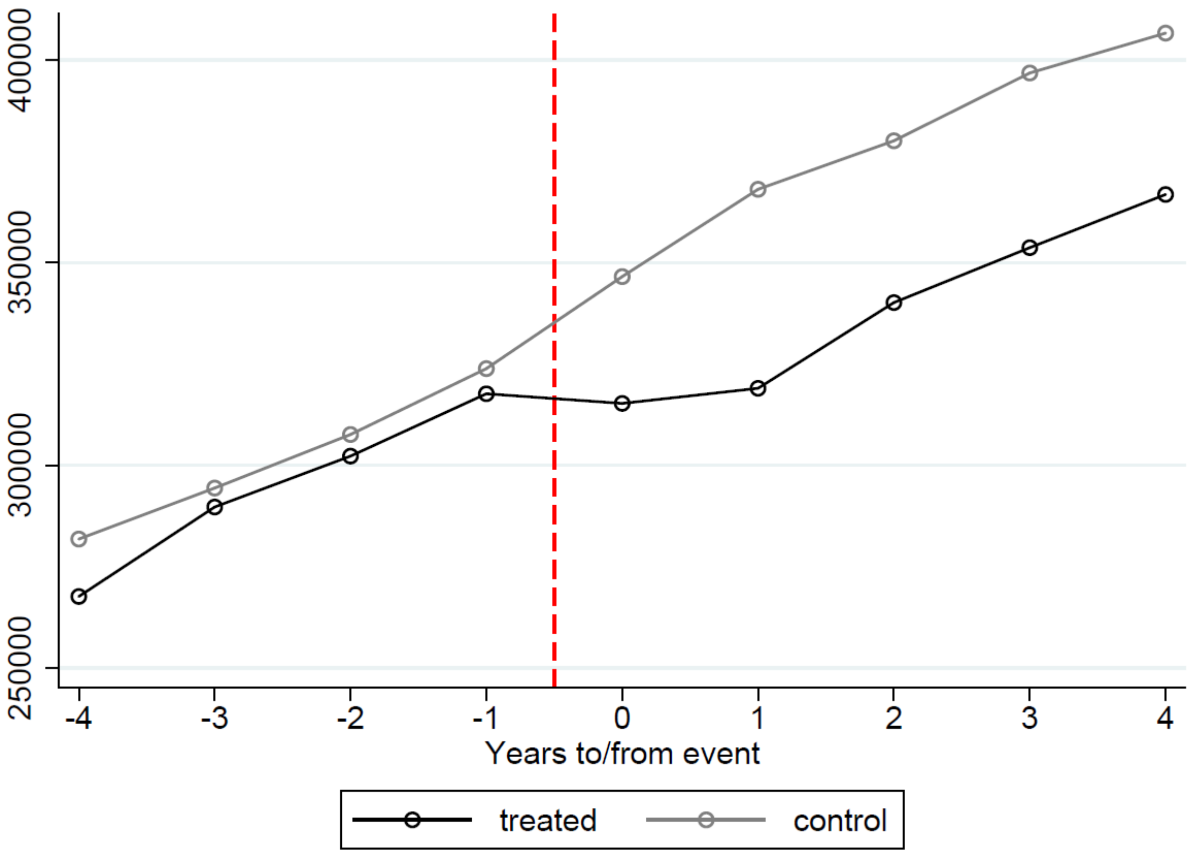

Notes: The figure reports the evolution of average household annual consumption in constant SEK2003 around the time when a household member loses her job. Event time is the time in years relative to the occurrence of the first job loss. The treatment group is composed of all the households from our baseline sample described in section 4.2 above. Individuals are aged between 25 to 55 at the time of job loss, and eligible for any form of UI at the time of the event. We introduce a control group that never experiences treatment. This control group is created using nearest-neighbor matching based on pre-event characteristics. We adopt the following matching strategy. For each calendar year t, we take all individuals who receive the event in that particular year ($E_{it} = t$), and find a nearest neighbor from the sample of all individuals who never receive treatment. Individuals are matched exactly on age, gender, region of residence in $t-1$ (21 cells), level of education in $t-1$ (10 cells) and family structure in $t-1$ (12 cells), and by propensity score on their number of dependent children in $t-1$, 12 industry dummies in $t-1$ and their earnings in $t-1$, $t-2$ and $t-3$. Consumption is annual total household expenditures from our registry-based measure. The structure of the household is determined as of event year -1 and kept constant throughout event times. See text for details.

Figure B-2: Estimated drop in annual consumption in year of job loss as a function of time spent unemployed

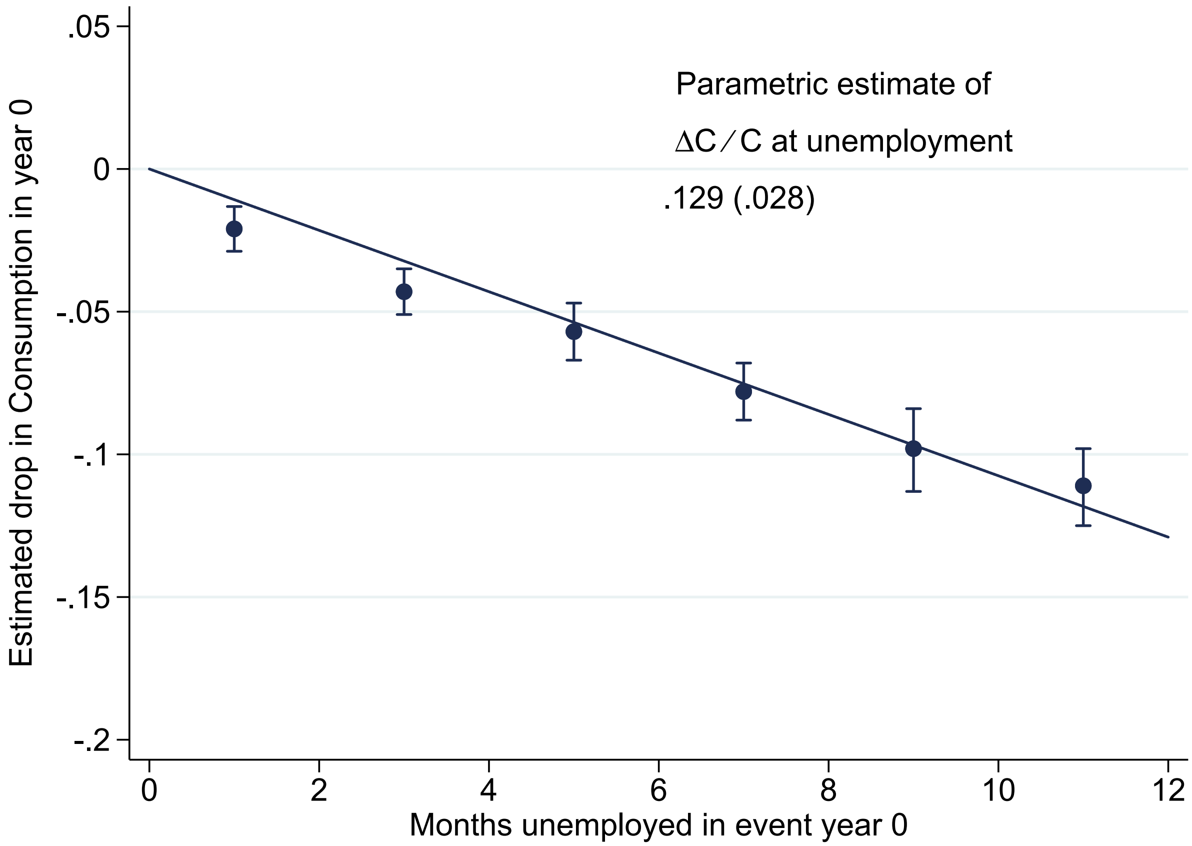

Notes: The graph displays the relationship between the drop in annual household consumption in event year 0 and the number of months spent unemployed in event year 0. In our sample, in event year 0, individuals are all observed unemployed in December, but differ in the time in months M_i they have spent unemployed in that year. We split the sample in 6 bins of M_i, and estimate specification (30) for each group. The figure reports the estimates $\hat{\beta}_0/\hat{C}_{-1}$ of the percentage drop in annual consumption in year 0 for each bin of M_i. The graph reveals that the relationship between time spent unemployed in year 0 and the annual drop in consumption in year 0 is indeed very close to being linear with an intercept equal to zero. This evidence motivates our use of the parametric model of equation (20) to identify the *flow* drop in consumption when unemployed $c_u - c_e$ from our annual consumption measure. See text for details.

B.2 Anticipation

Figure B-3: Estimated Change in Predicted Unemployment Risk around Start of Unemployment Spell

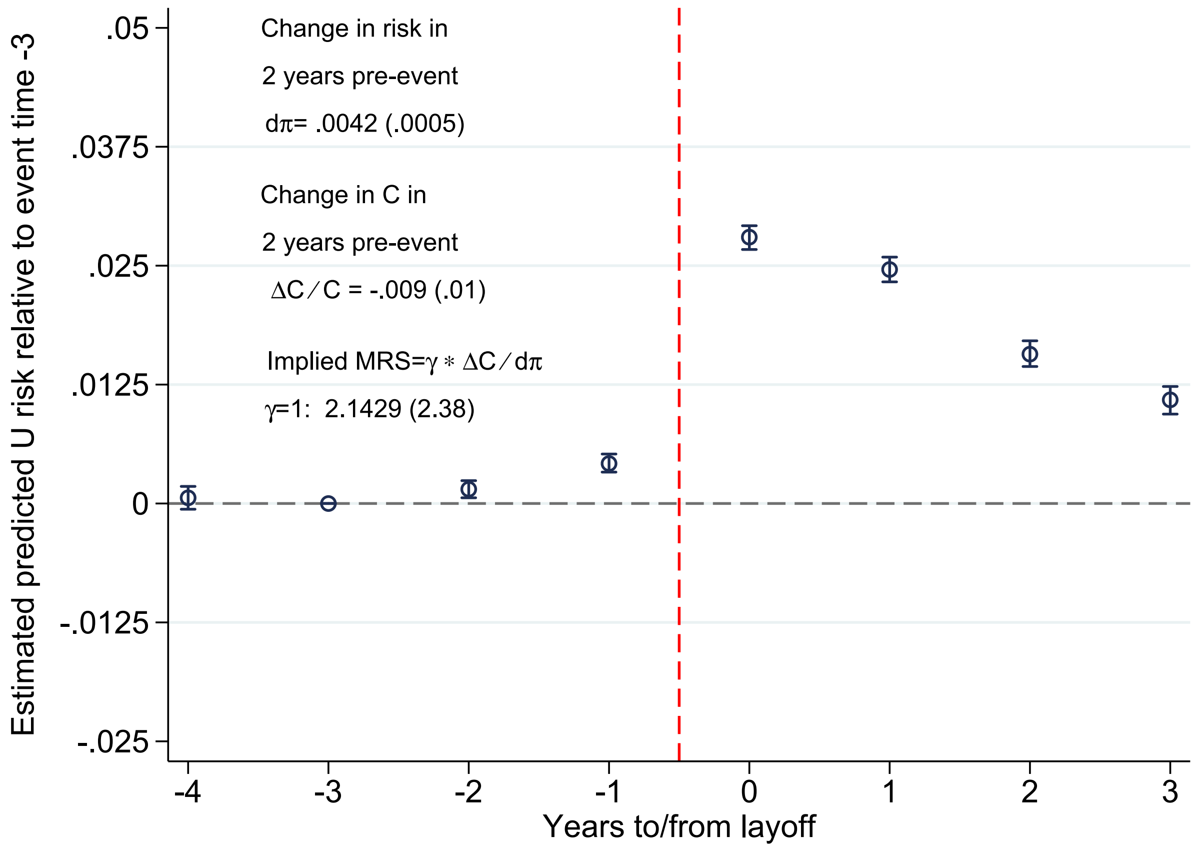

Notes: The graph studies how much individual unemployment risk gets revealed through changes in observables in the years around job loss. It reports estimates from specification (30) where we use as the outcome our measure of predicted unemployment risk, based on our rich model of observable determinants of unemployment risk in Sweden. See section 4.2 for details on our predicted risk model. The sample is the same as our sample used for Figure 1. Following Hendren [2017], we can then relate the estimated change in risk prior to job loss to the change in consumption in the two years prior to job loss estimated from Figure 1, and obtain an alternative measure of the MRS from anticipatory behaviors alone. We report on the graph our results from this implementation, which gives a large but very imprecisely estimated MRS, of 2.14.

B.3 Consumption Smoothing Mechanisms

The granularity of our data enables us to go beyond the standard implementation of the CB approach, and explore dimensions of consumption expenditure dynamics at unemployment that reveal further useful information about the value of insurance.

Decomposition of Consumption Smoothing We start by decomposing consumption expenditures of household i in year t into the following five components: earnings of the individual subject to the unemployment shock (E^u_{it}), spousal earnings (E^{-u}_{it}), all transfers net of taxes paid (T_{it}), consumption out of assets ($-\Delta A_{it}$), consumption out of debt (ΔD_{it}),

$$C_{it} = E^u_{it} + E^{-u}_{it} + T_{it} - \Delta A_{it} + \Delta D_{it}.$$

We then use this decomposition to document the respective role of each margin in smoothing consumption at job loss. For this purpose, we estimate specification (20) replacing consumption by each component of total household expenditures. In Figure B-4, we report for each component X the estimate of the change in that component at job loss, scaled by the consumption level prior to unemployment ($\frac{12 \cdot \hat{\beta}^X_0}{\widehat{C}_{-1}}$). Upon unemployment, individuals experience a loss of earnings amounting to more than 50% of their pre-unemployment household expenditures. Total transfers (including UI) net of all taxes paid, however, increase massively, an increase equivalent to more than 35% of pre-unemployment household consumption. The government transfers thus explain most of the difference between the drop in earnings relative to the drop in consumption.

The response of consumption out of assets and debt is suggestive of the importance of liquidity and borrowing constraints in explaining consumption dynamics at unemployment. Consumption out of assets increases at job loss, by about 7% of pre-unemployment consumption on average, and thus represents a significant source of consumption smoothing. Consumption out of debt, however, decreases significantly at job loss, by about 5% of pre-unemployment consumption. This implies that rather than taking out more debt to smooth consumption, on average workers reduce their consumption from debt when becoming unemployed. On net, the use of assets reduces the drop in consumption by only 2%. Figure B-4 also reveals the very limited role played by the added worker effect in our context: the contribution of changes in spousal earnings to consumption smoothing is almost negligible.

Overall, most of the consumption smoothing is done by transfers, leaving a much more limited role to the other adjustment margins. This lack of significant additional consumption smoothing through self-insurance mechanisms could be interpreted as revealing the low value that workers place on average to getting extra insurance. Yet, one could also interpret this evidence as suggesting that the price of increasing consumption through self insurance such as spousal labor supply or debt is particularly large when unemployed.

Heterogeneity Prior work has mostly focused on average drops in consumption at unemployment, due to small sample size in consumption surveys. Our registry-based measure provides the statistical power for a rich analysis of heterogeneity in consumption drops at unemployment. To analyse how heterogeneity along some dimension H affects the drop in expenditures at job loss, we discretize H into bins, and fully interact regressors in specification (20) with bin dummies. We run the following specification:

$$C_{it} = \alpha_i + \nu_t + \sum_{h} \sum_{\substack{j=-N_0 \\ j \neq 0}}^{N_1} \beta_{jh} \cdot \mathbb{1}[J_{it} = j] \cdot \mathbb{1}[H_i = h] + \sum_{h} \beta_{0h} \cdot M_i \cdot \mathbb{1}[J_{it} = 0] \cdot \mathbb{1}[H_i = h] + \varepsilon_{it} \quad (29)$$

In Figure B-5 we report the estimates of the effect $\frac{12 \cdot \hat{\beta}_{0h}}{\widehat{C}_{-1,h}}$ of variation in dimension H on the drop in consumption at unemployment, when all dimensions H are entered simultaneously into specification (29). We focus on demographic characteristics (age, marital status), as well as characteristics affecting the ability to smooth consumption over the spell (wealth, and portfolio composition at the start of the spell, UI replacement rate, etc.). All estimates are relative to the baseline category for each dimension H.[46] Results confirm the presence of a substantial amount of heterogeneity in consumption drops along these observable characteristics, and the importance of liquidity and borrowing constraints in particular. While the overall level of net wealth itself does not seem to have much of an effect, the allocation of wealth matters a lot. Indeed, having more liquid assets is associated with significantly less severe drops, in line with the notion of wealthy hand-to-mouth consumers in Kaplan and Violante [2014]. Moreover, having more debt at the start of a spell is associated with larger drops in consumption, with the most severe drop suffered by the workers who are most indebted. The heterogeneity analysis also confirms the important role played by transfers in smoothing consumption at unemployment: having a replacement rate below the maximum of 80% is associated with a significantly larger drop in consumption at job loss.

[46]Individuals in the baseline group are less than 35 years old at the start of the spell, married and in the bottom quartile of the wealth, income and debt distribution pre-unemployment. They have no liquid assets at the start of the spell, but receive a UI replacement rate of 80%, which is the maximum replacement rate under the comprehensive coverage.

Figure B-4: Decomposition of the Estimated Drop in Consumption at Unemployment

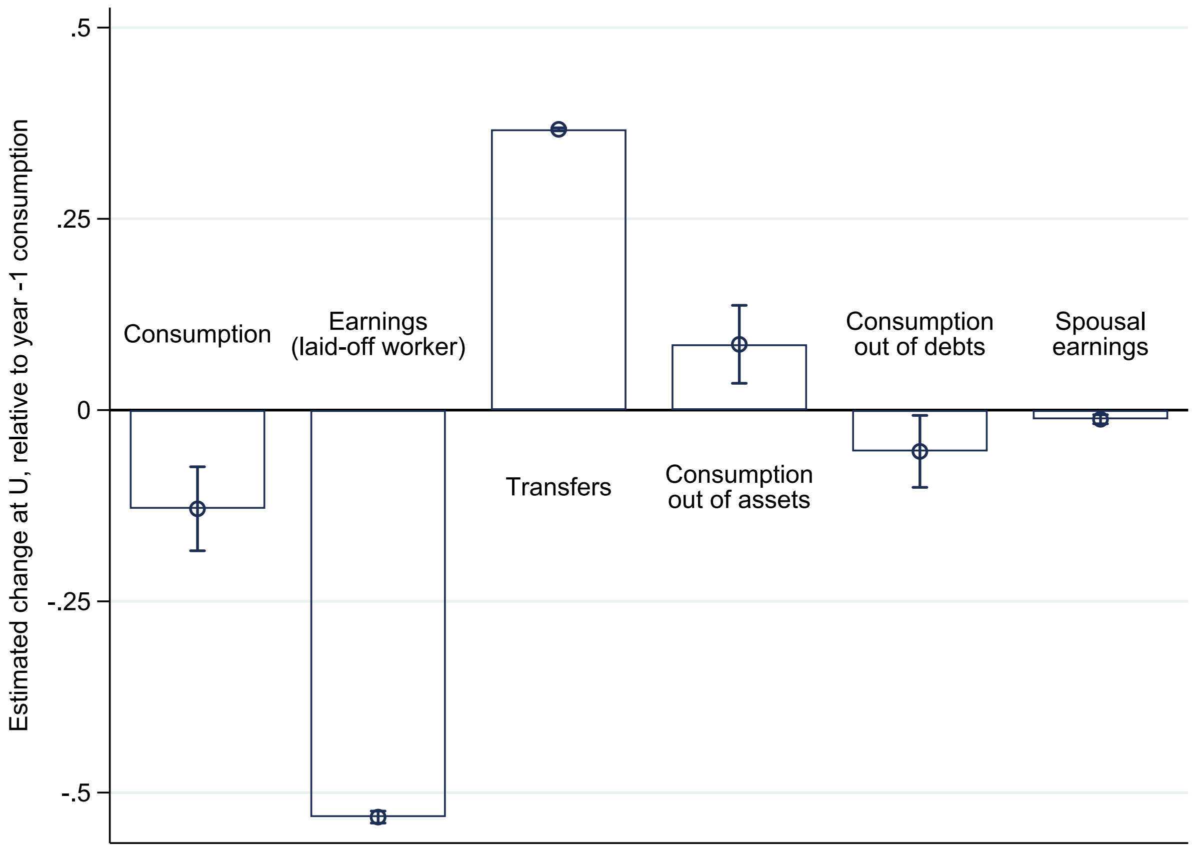

Notes: The graph decomposes the variation in consumption expenditures at unemployment, into the variations of five different components of total household expenditures: earnings of the individual subject to the unemployment shock ($E_{i,u,t}$), spousal earnings ($E_{i,-u,t}$), all transfers net of taxes paid (T_{it}), consumption out of assets ($-\Delta A_{it}$), consumption out of debt (ΔD_{it}). To document the respective role of each margin in smoothing consumption at job loss, we estimate specification (20) replacing consumption by each component of total household expenditures. We report on the graph, for each component, the estimate $\frac{12 \cdot \hat{\beta}_0}{\hat{C}_{-1}}$, of the change in this component at job loss, scaled by the consumption level prior to unemployment. The figure shows for instance that upon unemployment, individuals experience a loss of earnings amounting to more than 50% of their pre-unemployment total household expenditures. See text for details.

Figure B-5: Heterogeneity in Estimated Drop in Consumption at Unemployment

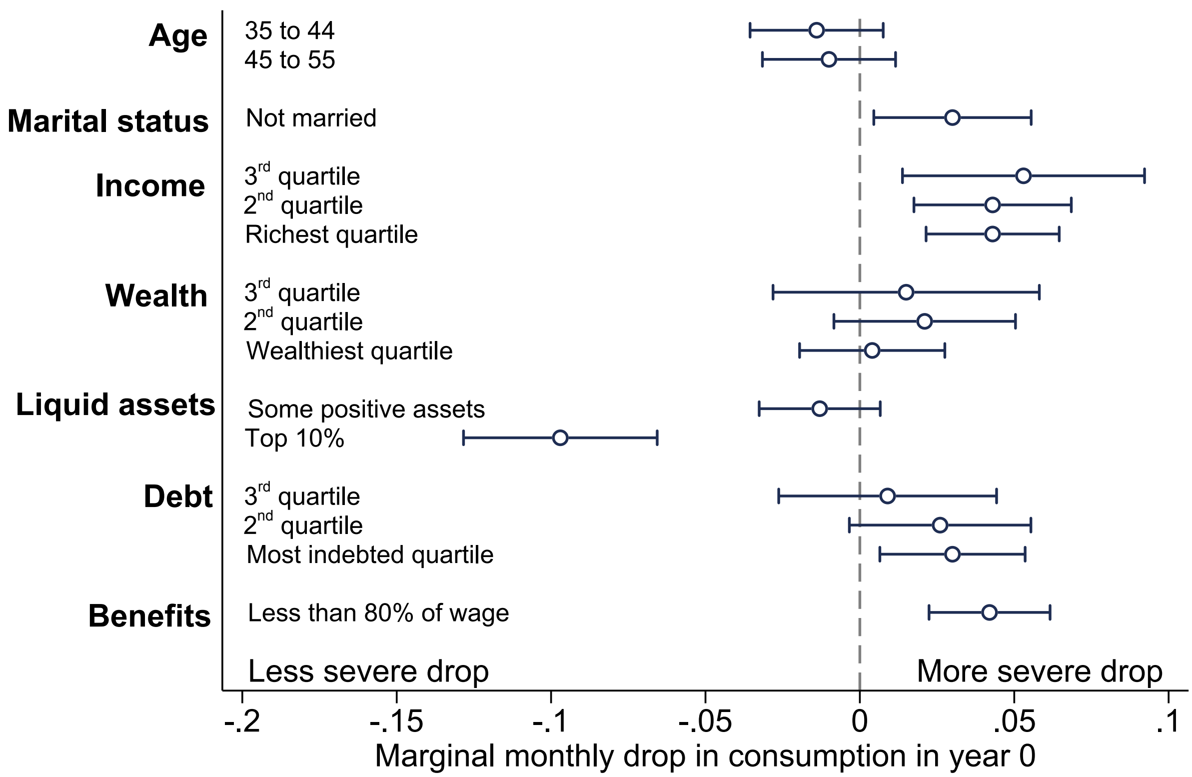

Notes: The graph analyzes heterogeneity in consumption drops at unemployment. The figure reports estimates of the effect of having characteristic $H = h$ on the drop in consumption at unemployment, following specification (29). Note that all dimensions of heterogeneity are entered simultaneously in the regression. We focus on demographic characteristics (age, marital status), as well as characteristics affecting the ability to smooth consumption over the spell (wealth, and portfolio composition at the start of the spell, UI replacement rate, etc.). All estimates are relative to the baseline category for each dimension. For age, the baseline is being less than 35 at the start of the spell. For marital status, the baseline is being married. For wealth, income and debt, results are relative to the bottom quartile of the distribution pre-unemployment. For liquid assets, the baseline is having no liquid assets at the start of the spell. For UI benefits, the baseline is having a replacement rate of 80%, which is the maximum replacement rate under the comprehensive coverage. See text for details.

Figure B-6: Change in Consumption Following Job Loss, by Expenditure Category

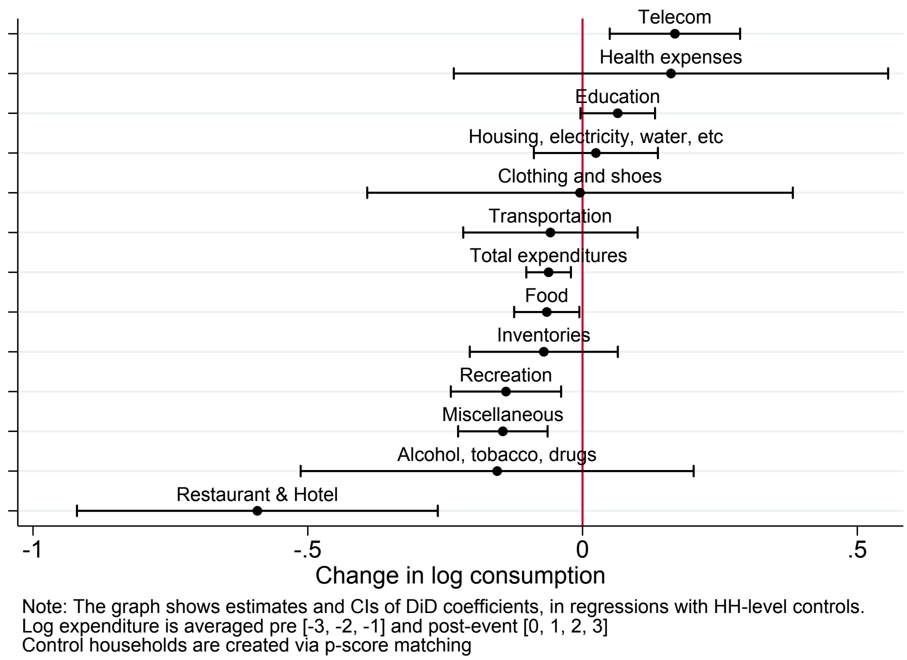

Notes: The Figure explores how the consumption responses to job loss differ by expenditure category. We use the HUT consumption survey, and compute the drop in consumption at job loss from an event study specification for different categories of expenditure observable in the survey. The graph confirms that during unemployment, expenditures that are complement to spending time home actually do go up (housing, telecom, etc.) while expenditures that are complement to spending time away from home decrease sharply.

Figure B-7: Probability to Receive a Transfer from Family/Friends Around Job Loss

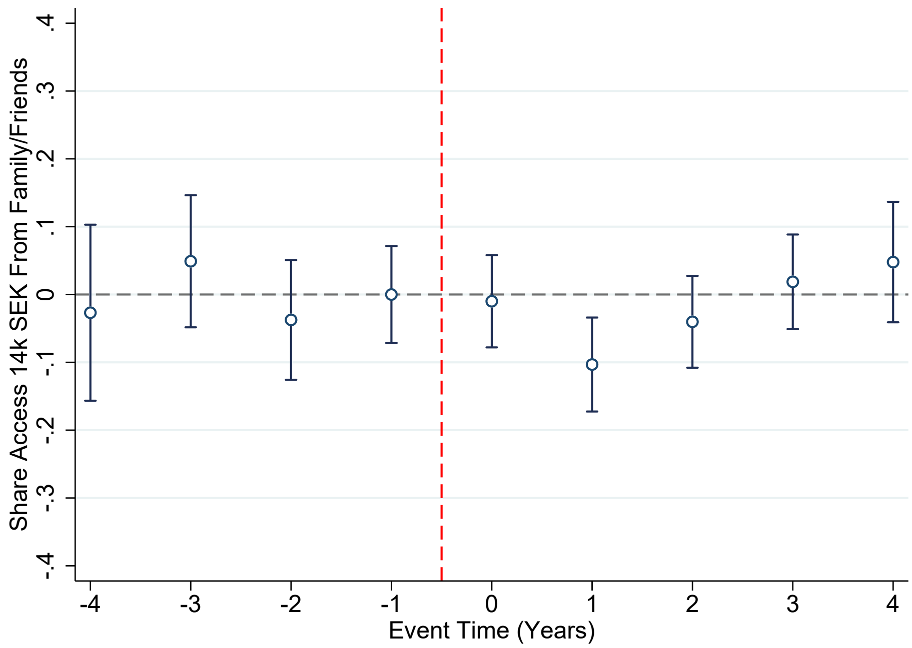

Notes: The Figure explores the evolution of the fraction of individuals who report borrowing from family and friends as a function of time since/until job loss. In the ULF survey, which we matched to our registry data, respondents are asked whether they borrowed at least 14,000 SEK last year. 5% of respondents indicated that they did borrow at least 14,000 SEK. Then individuals are asked where they borrowed from. Informal transfers are captured by borrowing from family and friends. We use a control group of individuals who never experience job loss, as in the event study strategy underlying Figure 1 in the paper. The graph reports the estimated coefficient $\hat{\beta}_j$ from the following event study specification:

$$D_{it} = \alpha_C + \alpha_T \cdot \mathbb{1}[\text{Treatment} = 1] + \nu_t + \sum_{j \neq -1} \beta_j \cdot \mathbb{1}[J_{it} = j] + \varepsilon_{it} \tag{30}$$

where D_{it} is an indicator for borrowing from family and friends in year t.

Appendix C Additional Figures: MPC Approach

Figure C-1: Average Residual Variation in Local Transfers Conditional on $\mathbf{V}$

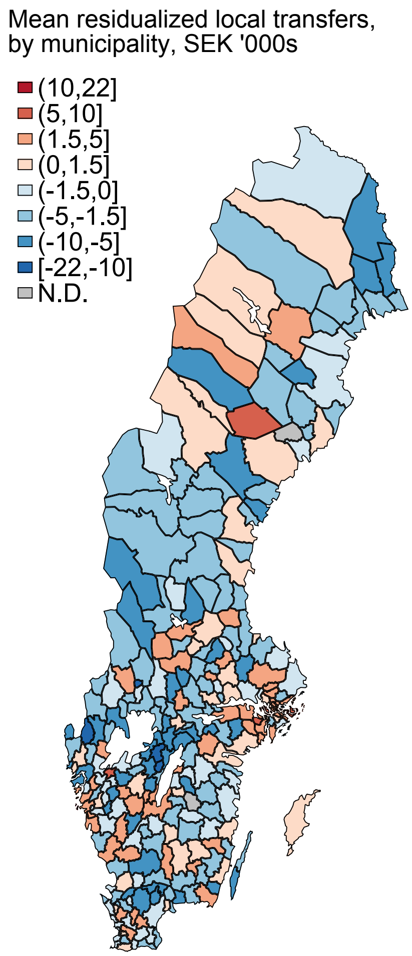

Notes: The Figure provides evidence of the variation in the way Swedish municipalities set local welfare transfers ("*social bidrag*"). By law, transfers are functions of characteristics $\mathbf{V}$, which include the number of dependents, the age of the dependent children, the liquid assets and income of the household: $B_{imt} = \sum_k \tau^k_{mt} \cdot V^k_{it}$. Because of the discretion left to municipalities, there is, after controlling for characteristics $\mathbf{V}$, a significant amount of variation left in the generosity of local welfare transfers across municipalities. To provides an illustration of this sizeable variation, we residualize transfers B_{imt} received by household i in year t on the vector of observable characteristics $\mathbf{V}_{it}$, which by law determine B. We include in $\mathbf{V}_{it}$ marital and cohabitation status of the household head, dummies for the number of adults in the household, dummies for the number of children in the household and their age, and dummies for the decile of disposable income (excluding local transfers) and for the decile of net liquid assets of the household. The figure plots the average residualized transfer $\tilde{B}_m$ in each municipality over the period 2000-2007. The map shows that there is a large amount of variation in the average residual generosity of welfare transfers between municipalities. For example, the urban municipalities in Stockholm, Gothenburg or Malmö in the South, but also some less populated municipalities in the North are significantly more generous.

Figure C-2: Robustness: Relationship Between Residualized Transfers & Covariate Index of Observed Heterogeneity

A. Whole Sample

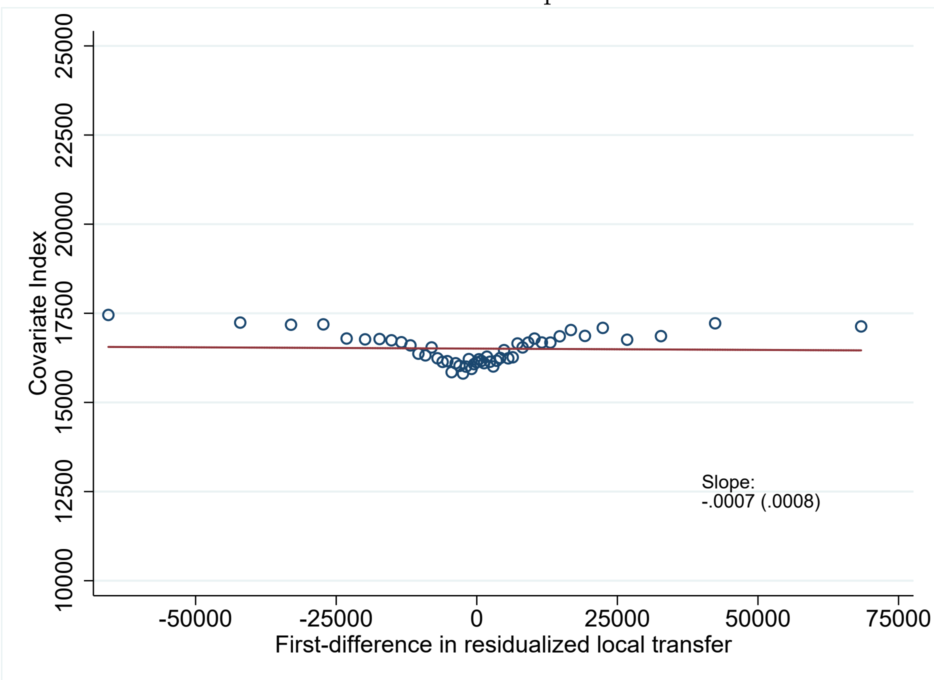

B. By Employment Status

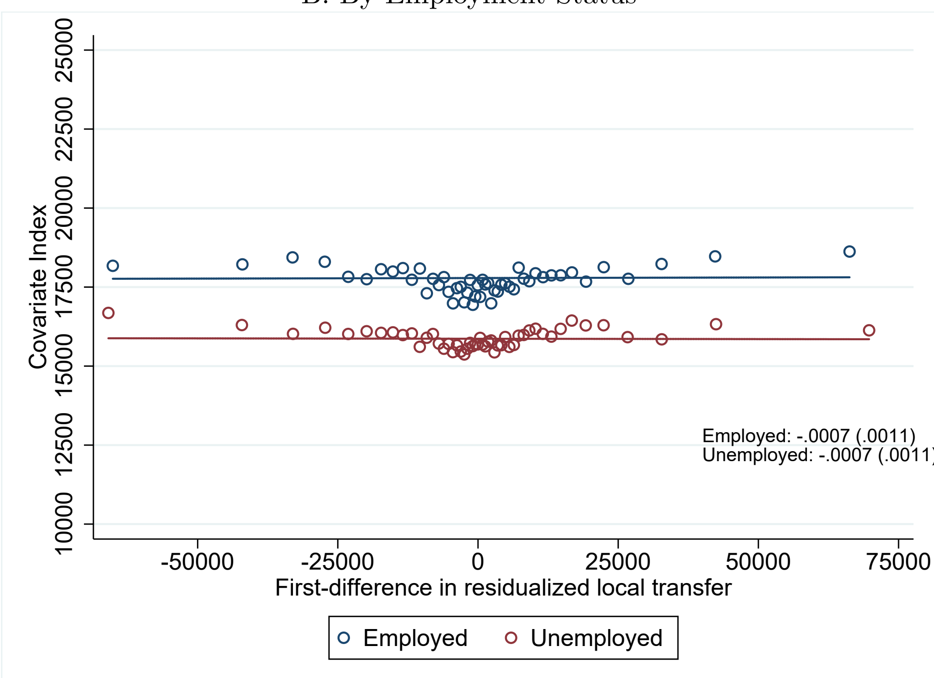

Notes: The Figure probes into the validity of our identifying assumption that the residual variation in welfare benefits $\tilde{B}_{imt}$ is orthogonal to the dynamics of household consumption. We use observables characteristics available in the registry data, that correlate with consumption, and that do not enter the benefit formula of welfare transfers. We use as covariates: the education level of the household members, the age of the head of the household, the total amount of real estate wealth of the household, the lagged value of total household debt, and the industry of the head of the household. We build a covariate index, which is a linear combination of these variables where the coefficients are obtained by regressing consumption on these covariates. We then test for the presence of a significant correlation between $\tilde{B}_{imt}$ and this covariate index. The graph is a binscatter of the relationship between the residual $\tilde{B}_{imt}$ obtained from our baseline residualization and the covariate index. Panel A shows this relationship in our whole sample. Panel B splits the sample by employment status. In each panel, we report the estimated correlation between the covariate index and $\tilde{B}_{imt}$, and find no statistically significant correlation.

Figure C-3: Distribution of Residualized Transfers

A. Employed State vs Unemployed State

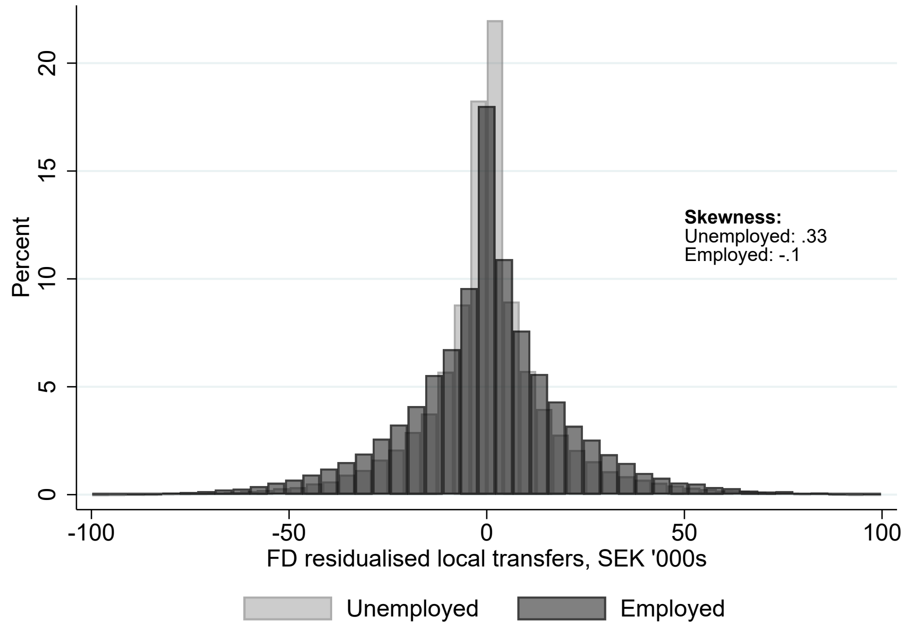

B. Movers vs Stayers

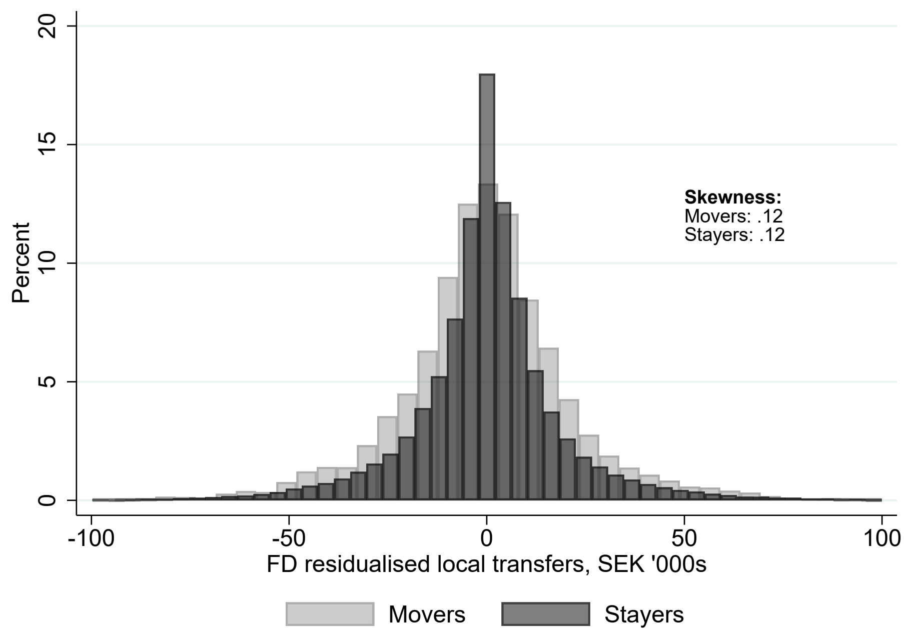

Notes: Panel A explores whether $\tilde{B}_{imt}$ is correlated with employment status. We plot the distribution of our baseline residual variation $\tilde{B}_{imt}$ by employment status. The figure shows that the distribution of our identifying variation in welfare transfers is very similar across employment status. This alleviates the concern that the difference in our MPC estimates while employed and unemployed are simply driven by different distributions of underlying variation in transfer. Panel B displays the distribution of $\tilde{B}_{imt}$, splitting the sample between movers (households who moved municipality in year t) and stayers. We find no significant correlation between $\tilde{B}_{imt}$ and the probability of moving, which indicates that our identifying variation in transfers is immune to the bias of selective migration.

Figure C-4: Evolution of Benefits and Consumption Around the Time of a Large Change in an Individual's Residualized Welfare Transfers

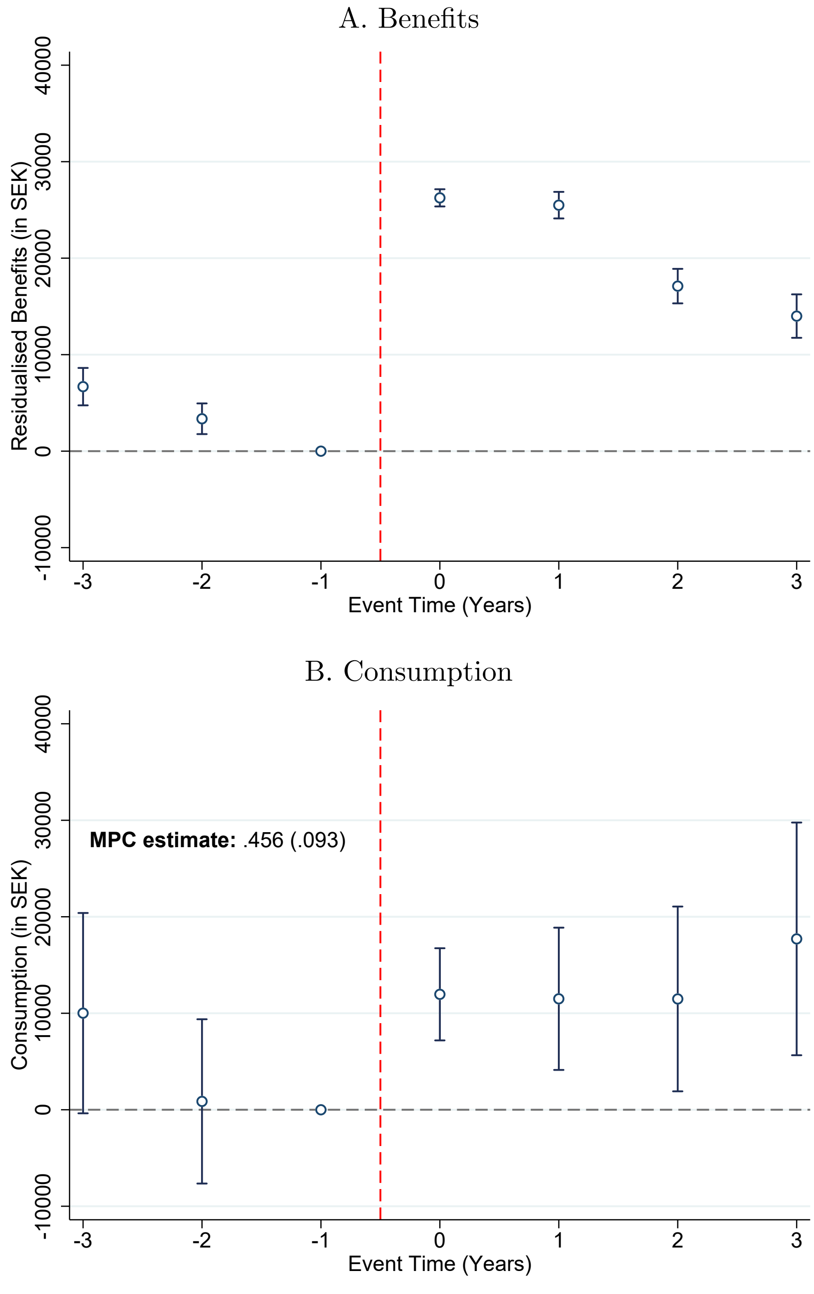

Notes: The Figure explores the evolution of benefits and consumption around sudden increases in an individual's welfare transfers. We follow an an event study design. We define an event as a year in which the residual transfer $\tilde{B}_{imt}$ received by an individual experiences a sudden increase of more than 12,500SEK. The Figure shows the evolution of benefits (panel A) and of consumption around the time of the event, following an event study specification with year and individual fixed effects. Both panels A and B suggests that the identifying variation brought about by variation in residual welfare transfers is not strongly correlated with the past dynamics of individual benefits nor with the past dynamics of household consumption. We compute an implied MPC corresponding to the estimated change in consumption in year 0, divided by the estimated change in benefits in year 0. We find an MPC of .456 (.093), which is very similar to our baseline estimated MPCs in Table 2.

Figure C-5: Evolution of Benefits and Consumption Around the Time of a Large Change in Average Residualized Welfare Transfers At the Municipality Level

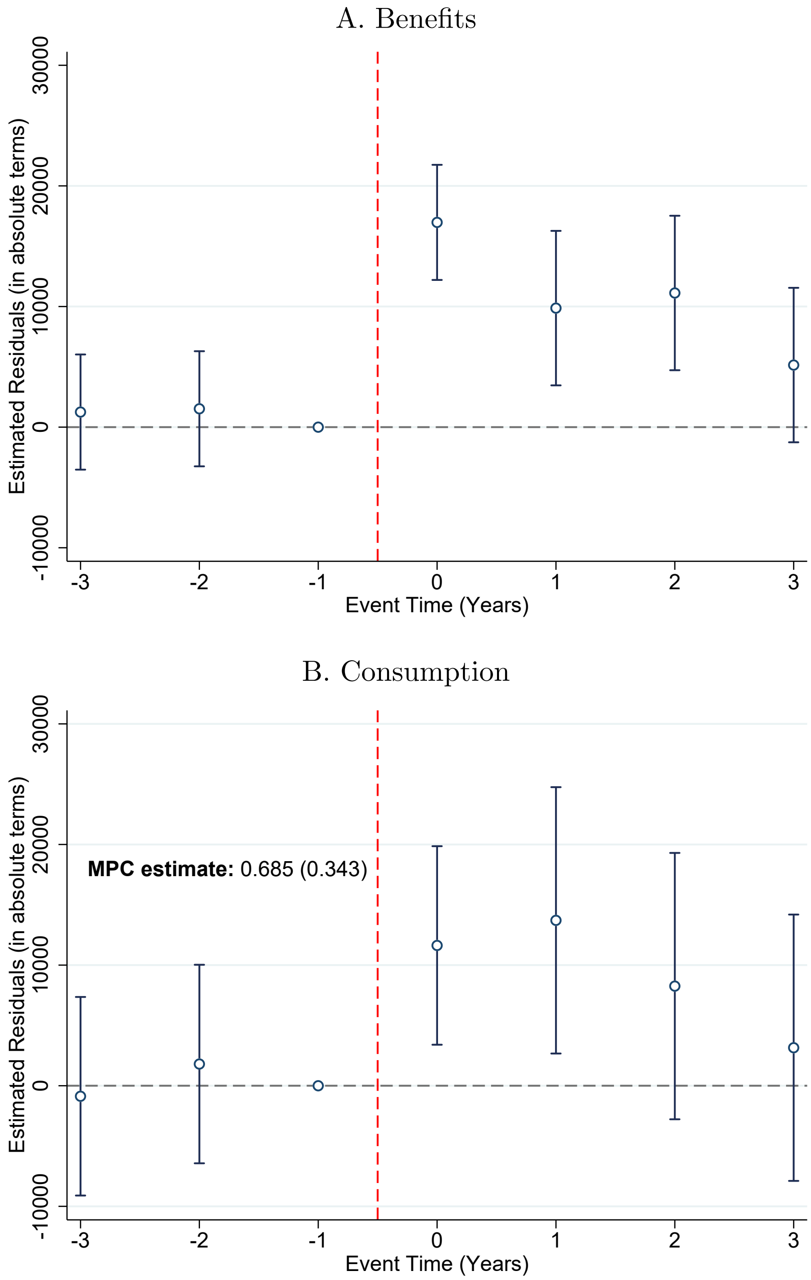

Notes: The Figure explores the evolution of benefits and consumption around sudden increases in the average generosity of welfare transfers at the municipal level. We follow an an event study design. We define an event as a year in which the average residual transfer $\tilde{B}_{imt}$ in municipality m experiences a sudden increase of more than 12,500SEK. We found 8 municipalities experiencing such events over our sample period. The Figure shows the evolution of average benefits (panel A) and of average consumption at the municipality level around the time of the event, following an event study specification with year and municipality fixed effects. The absence of pre-trends in both panel A and B confirms that the identifying variation brought about by these reforms is not endogenous to the past dynamics of benefits nor to the past dynamics of consumption.

Figure C-6: Decomposition of Consumption Responses to a Marginal Increase in Welfare Transfers, by Employment Status

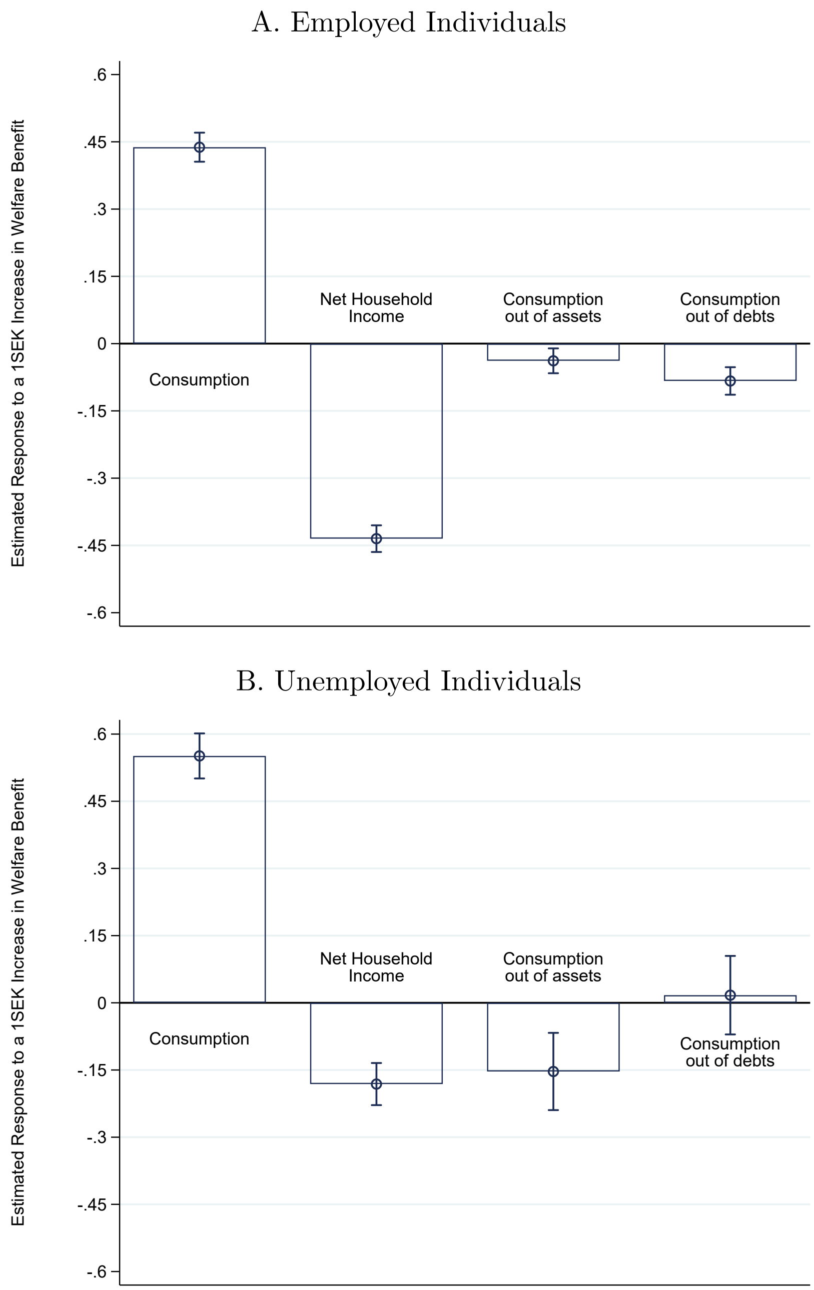

Notes: The Figure decomposes the response of consumption to a 1SEK increase in welfare benefit payments into the responses of various components of consumption. We decompose consumption into 5 components:

$$C = Y + B + C_{\text{Assets}} + C_{\text{Debt}} + C_{\text{Residual}}$$

where Y is total household labor income net of taxes, B are local welfare transfers, C_{Assets} is consumption out of assets, C_{Debt} is consumption out of debt, and C_{Residual} is the residual part of consumption not captured by the previous aggregates, including, among other things, some other transfers or taxes. The Figure reports the estimated change of each of these consumption components in response to a change in welfare benefits, using our baseline specification of Table 2 column (1).

Appendix D Alternative Identification of MPC While Unemployed

We assess the internal validity of our baseline MPC estimates by using an alternative identification strategy in the same sample to estimate the MPC. For this purpose we take advantage of the existence of a kink in the Swedish UI benefit schedule. This offers a credible source of exogenous variation in income that can be exploited in a regression kink design, as discussed in detail in Kolsrud et al. [2018]. While this source of variation is only valid to identify the MPC in the unemployment state, it is useful to gauge whether the magnitude of our MPC estimates are sensitive to the identification strategy chosen in a given sample.

D.1 Identification Strategy: RK Design

In Sweden the schedule of UI benefits is a kinked function of pre-unemployment earnings. Eligible workers receive daily unemployment benefits equal to 80% of their daily wage prior to unemployment, up to a cap. Over the period 2002 to 2007, the cap in daily UI benefits was fixed at 680SEK, meaning that the relationship between UI benefits and daily wage w exhibited a kink at $w = 850SEK$.[47]

We identify the effect of unemployment benefits on consumption using a RK design, taking advantage of the kink in the schedule of UI benefits as a function of the daily wage. Our identifying variation is displayed in Figure 4 panel A, which plots, in our main sample over the period 2002 to 2007, a binscatter of the relationship between the daily wage and the average replacement rate. The latter is computed as the average benefit received during unemployment from the IAF data divided by the daily wage.

The graph shows first that the replacement rate is close to exactly 80% on the left hand side.[48] The graph also displays a clear kink at $w = 850SEK$, with the replacement rate declining sharply, as benefits are capped. We use this kinked relationship and treat it as a fuzzy RKD around the 850SEK threshold. Our RK estimand of the MPC in the unemployment state is given by:

$$MPC = \frac{\lim_{w^-} dE[\Delta C|w]/dw - \lim_{w^+} dE[\Delta C|w]/dw}{\lim_{w^-} dE[b|w]/dw - \lim_{w^+} dE[b|w]/dw} \tag{31}$$

Importantly, the MPC from this RK design is identified out of an anticipated change in state-contingent income while unemployed, which is the relevant MPC concept from the point of view of Proposition (1).

We estimate the numerator of the estimand in (31), in the same baseline sample of analysis used throughout the paper, based on the following RK specification:

$$\Delta C_i = \beta_0 \cdot (w-k) + \beta_1 \cdot (w-k) \cdot \mathbf{1}[w > k] + \sum_j \gamma_j \mathbf{1}[M_i = j] + \mathbf{X}'\gamma \tag{32}$$

[47] A daily wage of 850SEK corresponds to about 468USD a week using the average exchange rate over the period 2002 to 2007 of 1SEK $\approx$ 0.11USD.

[48] Note that the reason why the replacement rate is a bit below 80% is that some workers have their UI benefits reduced due to sanctions.

where ΔC is the drop in annual household consumption at unemployment. We control non parametrically for the time spent unemployed during the year by adding a set of dummies for the number of months M_i spent in unemployment. We estimate the denominator in (31) using

$$\Delta b_i = \eta_0 \cdot (w-k) + \eta_1 \cdot (w-k) \cdot \mathbf{1}[w > k] + \sum_j \zeta_j \mathbf{1}[D = j] + X'\zeta$$

where b_i are UI benefits received.

Our fuzzy RK estimate of the marginal propensity to consume in unemployment is $MPC = \frac{\hat{\beta}_1}{\hat{\eta}_1}$. As far as inference is concerned, we provide robust standard errors, bootstrapped standard errors, as well as a permutation test analysis a la Ganong and Jaeger [2018].

An important assumption of the RK design is the existence of a smooth relationship at the threshold $w = 850SEK$ between the assignment variable and any heterogeneity affecting the outcome. To assess the credibility of this assumption, we conduct two types of analysis (see also Kolsrud et al. [2018]). First, we focus on the probability density function of the assignment variable, to detect manipulation or lack of smoothness around the kink that could indicate the presence of selection. Figure D-1 panel A shows that the pdf of daily wage does not exhibit a discontinuity nor lack of smoothness at the kink, which is confirmed by the results of formal McCrary tests. Second, we investigate the presence of potential selection along observable characteristics around the kink. For this purpose, instead of looking at each characteristics in isolation, we aggregate them in a covariate index. The index is a linear combination of a vector of characteristics $\mathbf{X}$ that correlate with consumption, which includes age, gender, level of education, region, family type and industry. The coefficients in the linear combination are obtained from a regression of the outcome variable ΔC_i on these covariates. In Figure D-1 panel B, we display the relationship between this covariate index and the assignment variable. The relationship between the index and daily wage appears smooth around the 850SEK threshold. Yet, formal tests of non-linearity suggest the presence of a signifiant (although economically small) kink at the threshold. Furthermore, the graph also reveals some volatility in the index on the right hand side of the threshold. We therefore include the vector of characteristics $\mathbf{X}$ in specification (32) to control for the small lack of smoothness, and increase precision of our RK MPC estimates. We explore below the sensitivity of our results to the inclusion of these controls.

D.2 Results

Figure 4 panel B plots the graphical representation of our baseline result. It shows the average change in consumption ΔC_i between the year the individual is unemployed and the year prior to the start of the spell by bins of daily wages. For the purpose of the plot, the change in household consumption ΔC_i is first residualized on a set of dummies for the number of months spent unemployed M_i and the vector of characteristics $\mathbf{X}$ which includes year, age gender, education, region, family structure, and industry fixed effects. To make the magnitude of the results interpretable, we scale consumption change by the average consumption in the year prior to unemployment in each bin.

The graph shows evidence of a large non-linearity in the relationship between daily wage and the consumption drop at unemployment. There is a sharp and significant change in the slope of this relationship at the 850SEK threshold. We also report on the panel our baseline RK estimates, using a bandwidth of 300, of the $MPC = \frac{\hat{\beta}_1}{\hat{\eta}_1} = .63$ (.16).

This MPC estimate is remarkably similar to our estimates of the MPC while unemployed from the local transfer variation in Table 2.

Sensitivity Tests We investigate the sensitivity of our estimates of the MPC while unemployed to our various specification assumptions and implementation choices. We start by analysing the sensitivity of our MPC estimates to the choice of bandwidth for the RK estimation. Figure D-2 panel A shows that our estimates are very stable across all bandwidth sizes. Our baseline bandwidth is 300. The optimal bandwidth from Calonico, Cattaneo, & Titiunik (2014) is 244.

Second, we investigate the sensitivity of our MPC estimates to the inclusion of the vector of controls $\mathbf{X}$. In Figure D-2 panel B, we report how MPC estimates change as we start including cumulatively the characteristics of vector $\mathbf{X}$ in specification (32). The graph shows that the cumulative inclusion of controls has very limited effect on our MPC estimates, which are very stable, lying between .6 and .7 for all specifications.

Standard errors on our MPC estimate are obtained from a bootstrap procedure. But we also assessed the sensitivity of our estimates to potential non-linearities in the relationship between consumption drops and the daily wage. To this effect, we produced placebo estimates at 1,000 placebo kinks and followed a permutation approach to inference a la Ganong and Jaeger [2018]. Our baseline estimate lies in the upper tail of the distribution of these placebo estimates. The p-value from a one-sided test is .046, indicating that the probability of finding an MPC estimate of .63 at random in at these placebo kinks is less than 5%. The 95% confidence interval for our MPC estimate obtained from this permutation procedure is [.459; .673].

FIGURE D-1: REGRESSION KINK DESIGN: ROBUSTNESS

A. Pdf of Assignment Variable

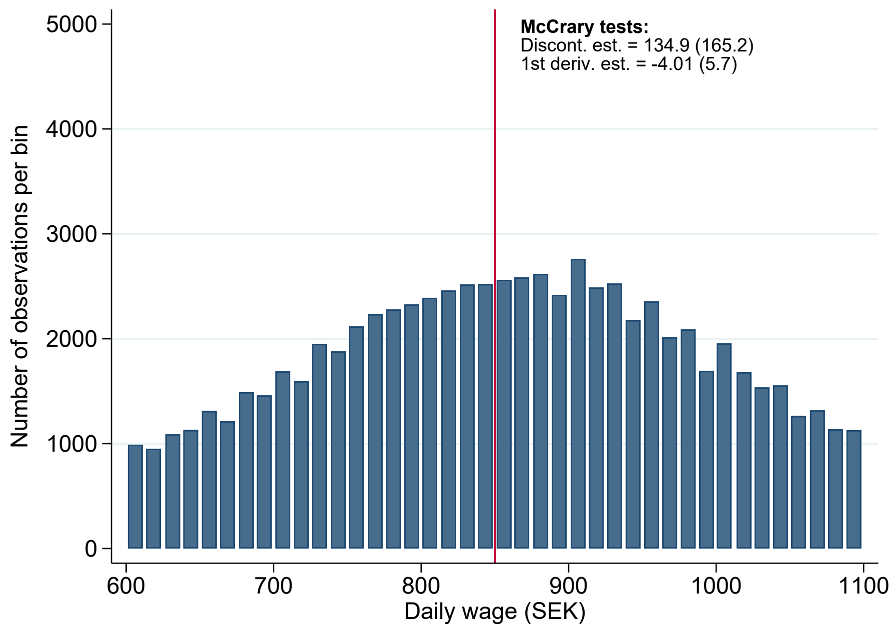

B. Covariate Index vs Assignment Variable

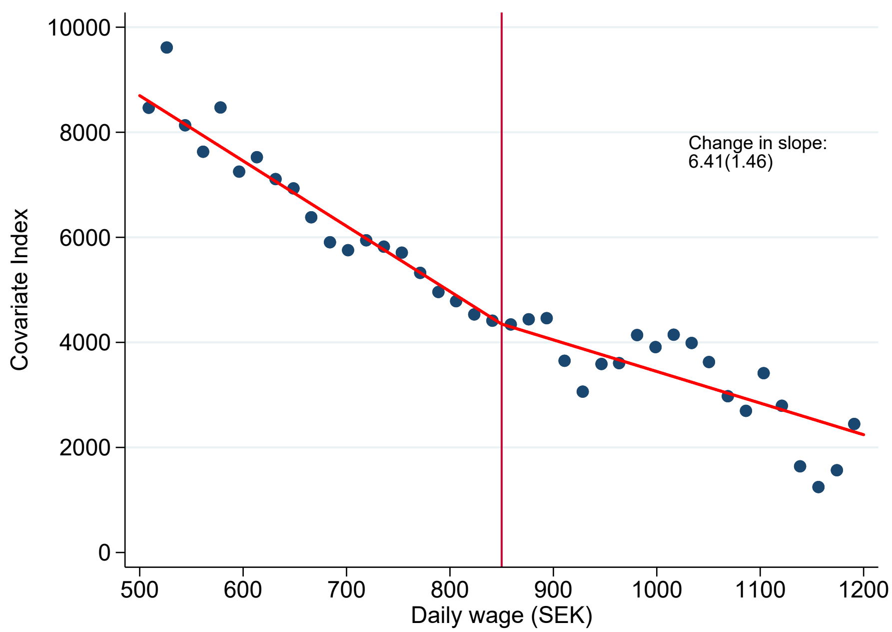

Notes: Panel A displays the probability density function of daily wage. We also report on the graph formal McCrary tests for the existence of a discontinuity nor lack of smoothness at the 850SEK threshold. Panel B investigates the presence of potential selection along observable characteristics around the kink. For this purpose, we aggregate observable characteristics into a covariate index. The index is a linear combination of a vector of characteristics $\mathbf{X}$ that correlate with consumption, which includes age, gender, level of education, region, family type and industry. The coefficients in the linear combination are obtained from a regression of the outcome variable ΔC_i on these covariates. The panel displays the relationship between this covariate index and the assignment variable. The relationship between the index and daily wage appears smooth around the 850SEK threshold. Yet, formal tests of non-linearity suggest the presence of a signifiant (although economically small) kink at the threshold.

Figure D-2: Regression kink design: Sensitivity

A. Estimates by Bandwidth

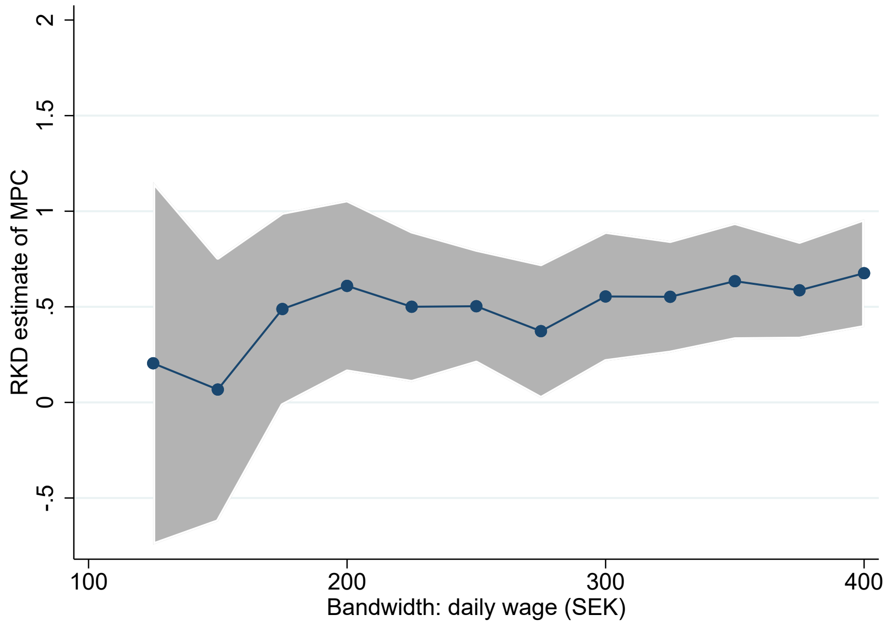

B. Sensitivity to Inclusion of Controls

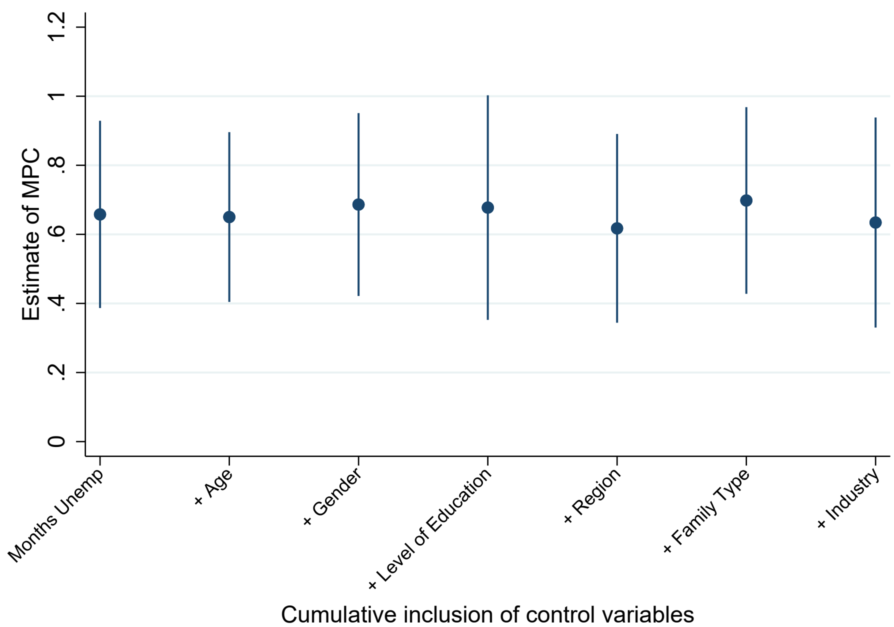

Notes: The Figure investigates the sensitivity of our estimates of the MPC while unemployed to specification assumptions and implementation choices. Panel A shows the sensitivity of our MPC estimates to the choice of bandwidth for the RK estimation. Our baseline bandwidth is 300. The optimal bandwidth from Calonico, Cattaneo, & Titiunik (2014) is 244. Panel B investigates the sensitivity of our MPC estimates to the inclusion of the vector of controls **X**. We report our MPC estimates when including cumulatively the characteristics of vector **X** in specification (32). The graph shows that the cumulative inclusion of controls has very limited effect on our MPC estimates, which are very stable, lying between .6 and .7 for all specifications.

Appendix E Revealed-Preference Approach: Additional Results & Figures

Figure E-1: Realized vs. Perceived Job Loss Probabilities in HUS Survey

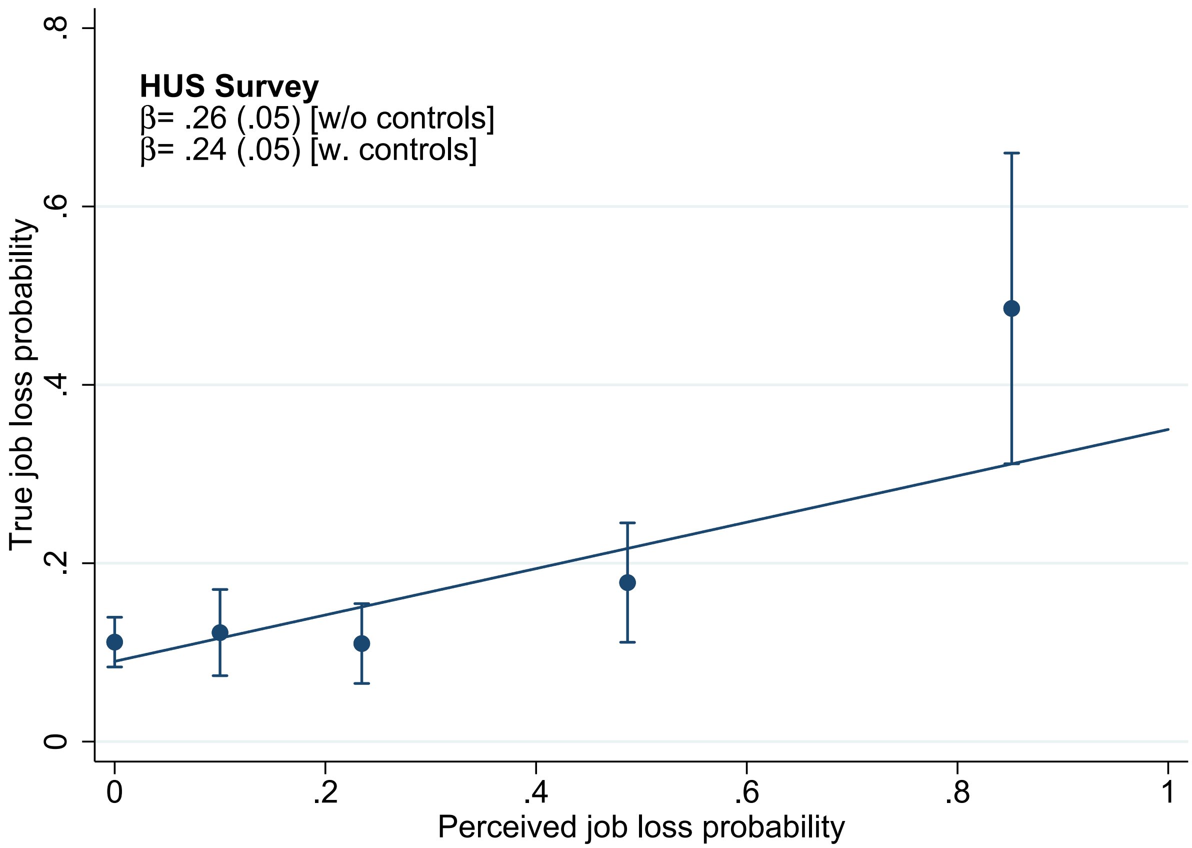

Notes: The figure compares the true and perceived probability to keep one's job, interpreting the complement of these probabilities as job loss probabilities. We use the responses to the question, "How likely is it that you will keep your current job next year?" in the HUS survey. For different bins of reported probabilities in the 1996 wave, we calculate the share of workers who lost their job between the 1996 wave and the 1998 wave, where job loss is defined as an individual reporting 1) being unemployed in 1998, 2) reporting a starting date for the job held in 1998 that lies after 1996, or 3) if the job held in 1998 was started in 1996, a starting month that is later than the starting month reported for the job held in 1996. While the perceived and the actual job loss probability are similar on average, workers who report a 1 percent higher job loss probability are only .26 (.05) percent more likely to lose their job. We use this estimate to correct the perceived risks in our RP estimation in columns 3-10 in Table 3. Note, however, that the two-year interval between the waves does not allow us to evaluate the perceived and actual job loss probability at the same horizon. Moreover, our interpretation of job loss includes workers who have switched jobs (and have not drawn any unemployment benefits). Using the perceived job loss question in the Survey of Consumer Expectations in the US, we find a very similar estimate of .27 (.08) when regressing actual job loss on perceived job loss. The average perceived and actual job loss are almost the same as well.

Moral Hazard In the presence of moral hazard, our estimate of the MRS distribution, using the predicted risk under comprehensive coverage, will be biased downward. When a worker considers to get basic coverage, she could gain from increasing her effort as well, which we ignore in the choice model. The magnitude of the bias will depend on the size of the ignored utility gain workers get from changing their effort.[49] Our estimate can thus be interpreted as a lower bound on the MRS, just like our estimate from the MPC implementation (albeit for different reasons). To gauge the potential magnitude of the bias due to moral hazard, Panel B of Figure E-2 compares the distributions of the estimated MRS using the (perception-corrected) risk under basic and comprehensive coverage respectively. These can be interpreted as an upper- and lower bound on the MRS. The entire distribution of the MRS is shifted upward with a mean of 2.13, using risk under comprehensive coverage, compared to a mean of 2.98, using risk under basic coverage. For completeness, Appendix Table 3 shows all earlier estimation results, but using workers' predicted risk under basic coverage instead.

While correctly accounting for perceived risks and other potential frictions remains an important challenge for the RP approach, overall, the RP estimation implies that the average MRS is substantially higher than the CB estimates indicate, corroborating the findings of the MPC approach. Moreover, the RP estimation shows substantial heterogeneity in the value of insurance, above and beyond the heterogeneity in unemployment risk.

[49] Note that when this omitted utility gain is uncorrelated with the expected price and observables determining the MRS, the coefficients β and γ and thus the dispersion in MRS are estimated consistently.

Figure E-2: Non-parametric and parametric RP Estimation under Basic Coverage

A. Expected Price vs. Insurance Coverage

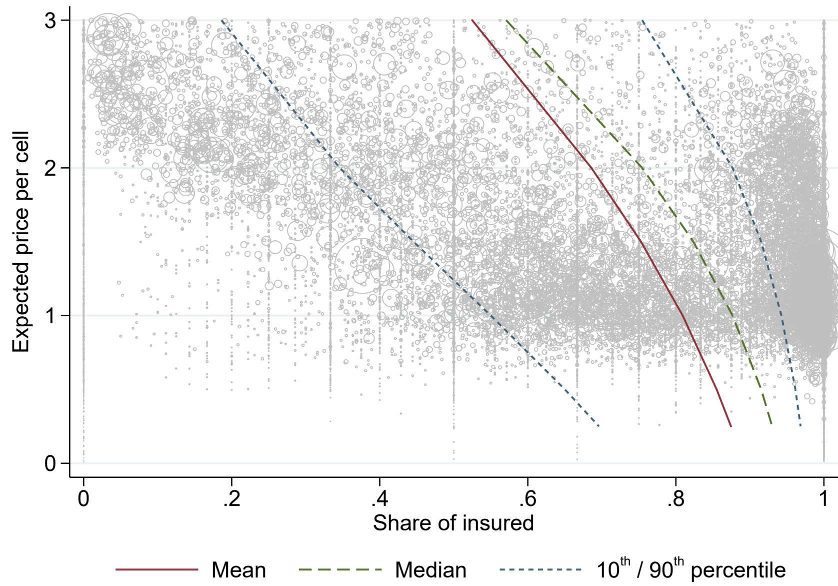

B. Lower and Upper Bound of MRS Distribution

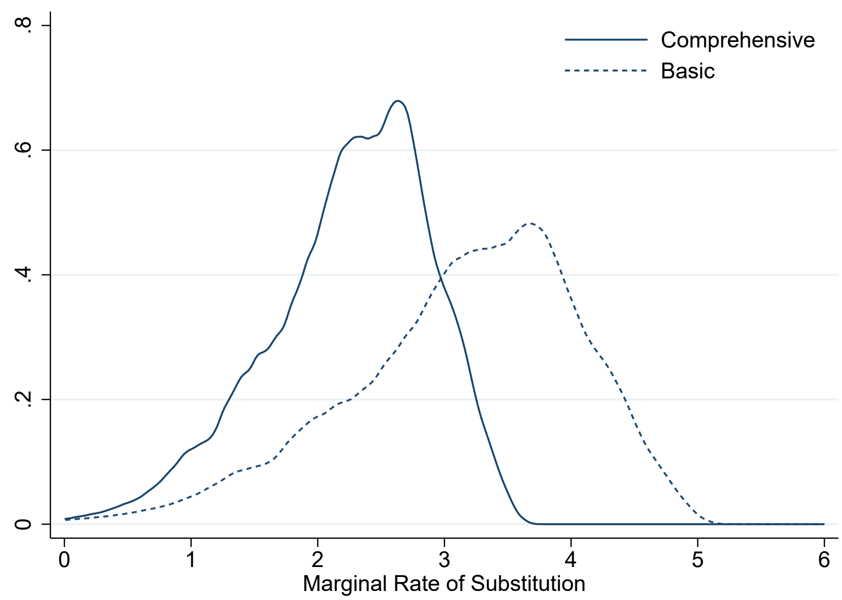

Notes: This Figure complements Figure 5 with Panel A showing the average expected price and share buying comprehensive insurance coverage for workers grouped by cells based on a rich set of observables, but calculating the expected price using the predicted risk under basic coverage rather than comprehensive coverage. Panel B contrasts the estimated distribution of MRS when the unemployment risk is estimated under the comprehensive coverage (solid line) vs. when it is estimated under the basic coverage (dashed line). In both cases depicted in Panel B, the perceived risk model is used (see column (3) of Table 3 and Appendix Table 4 respectively).

Table 4: Robustness of RP Approach using Risk under Basic Coverage

	(1) Baseline	(2) Salient	(3)	(4)	(5)	(6) Perceived Risk	(7)	(8)	(9)	(10)
Coefficient on price (γ)	-0.495	-0.673	-0.855	-0.841	-0.785		-0.853	-0.805	-0.653	-0.813
	(0.002)	(0.003)	(0.005)	(0.005)	(0.005)		(0.005)	(0.01)	(0.019)	(0.006)
Bottom						-1.052				
Income Quartile						(0.006)				
2^{nd}						-0.797				
						(0.005)				
3^{rd}						-0.352				
						(0.008)				
Top						-0.684				
						(0.01)				
Financial Variables				×						×
Region and Industry FE					×					
UI Choice persistence							×			
Cognitive ability										×
Observations	1,052,294	1,205,844	1,052,294	1,052,294	1,034,364	1,052,294	310,316	97,381	1,052,294	862,100
MRS Distribution										
Mean	4.82	3.42	2.98	3.07	3.31	4.40	2.99	3.02	3.41	3.34
10^{th}	1.78	0.69	1.22	1.2	1.12	1.24	1.22	1.34	1.21	1.53
25^{th}	3.99	2.52	2.51	2.41	2.51	2.76	2.51	2.59	2.94	2.63
50^{th}	5.32	3.93	3.27	3.35	3.60	3.87	3.27	3.32	3.84	3.59
75^{th}	6.28	4.73	3.82	4.06	4.45	5.78	3.83	3.83	4.45	4.31
90^{th}	7.01	5.32	4.24	4.55	5.06	8.73	4.25	4.2	4.88	4.81

Notes: The table mirrors the estimates of the choice model reported in Table 4, but with risk predicted under basic coverage rather than under comprehensive coverage. The specifications in Columns (1) to (10) are otherwise identical. Column (1) to (3) show the estimates for different risk models. Column (1) uses the baseline risk model, as discussed in Section 4. Column (2) uses only salient risk shifters to predict unemployment risk. Columns (3) uses the perceived risk model. Columns (4) to (10) perform further sensitivity checks of the RP estimation using the perceived risk model. See Table 3 for details.

Heterogeneity in UI valuation One important conclusion from our revealed-preference analysis is the existence of remarkable heterogeneity in the revealed value of UI, conditional on unemployment risk. While 10% of workers in our baseline sample would not be willing to pay any mark-up to transfer an extra krona to unemployment, 50% of them would be willing to pay more than a 100% mark-up, holding unemployment risk constant. To understand if the CB approach is a good guide to capture this heterogeneity in UI valuation, we examine for our baseline sample how much the MRS heterogeneity from the RP approach correlates with realized drops in consumption at job loss. In Appendix Figure E-4, we split our baseline sample in cells of observable characteristics and report the estimated average drop in consumption at job loss for households in that cell against the average MRS in the cell estimated from the RP approach in the year prior to job loss.[50] The graph shows that conditional on consumption drops, there is still a very large amount of residual variation in MRS left in the data.

To draw welfare conclusions from this residual heterogeneity, it is critical to understand where it stems from: is it capturing deep structural heterogeneity in risk preferences, or some form of heterogeneous frictions? To better understand the source driving this heterogeneity in UI valuation, in Figure E-3 we correlate our estimated MRS from the RP approach (see column 10 in Table 3) with various observable characteristics. Panel A, B and C focus on age, gender and the presence of children, three types of observable characteristics that potentially correlate with risk (or state-dependent) preferences. The three panels show evidence of a strong correlation between these characteristics and the MRS: older people, women, and individuals with children all have a significantly larger revealed-preference value of UI conditional on risk. Interestingly, since age, gender and the presence of children hardly affect the drop in consumption at job loss (see Figure B-5), these three characteristics are responsible for a significant amount of the residual heterogeneity in MRS conditional on consumption drops displayed in Appendix Figure E-4. In Panel D and E, we correlate our estimates of the individual MRS with asset holdings. Individuals with higher MRS have more wealth on average, which is again consistent with heterogeneity in preferences underlying the insurance choice and wealth accumulation. The relationship with the share of liquid assets is less clear. Workers with higher risk aversion may invest more in liquid assets, which in turn reduces the need to insure unemployment risk. Overall, risk preferences may thus well be negatively correlated with consumption drops at job loss, suggesting that heterogeneity in consumption drops can be a rather poor guide to infer heterogeneity in the value of UI (e.g., Chetty and Looney [2007], Andrews and Miller [2013]).

While evidence from Figure E-3 panels A to E is consistent with substantial heterogeneneity in preferences, there are also clear indications that part of the variance in the estimated MRS can be due to heterogeneity in frictions. We have already shown in Section 7 that correcting for risk misperceptions has a significant impact on the estimated distribution of MRS in the RP approach, reducing both the average and the variance of our estimates of the MRS. Panel F of Appendix Figure E-3 provides additional evidence showing that cognitive ability is negatively correlated with

[50] We create 120 cells using as observable characteristics three age bins, income deciles, family type and gender.

the estimated MRS. The measure of cognitive ability comes from tests administered by the Swedish Army to all enlisted individuals.[51] The graph shows that average cognitive ability score for the workers with rather extreme MRS, willing to pay a mark-up of more than 200%, is nearly half of the score for workers with MRS close to 1, who are not willing to pay a significant mark-up. This may indicate that choice frictions, rather than preferences, may be partly responsible for the high mean and variance in the MRS revealed by workers' choices.

[51] Until the late 1990s, enlistment was compulsory for men, and over 90 percent of all men in each cohort went through the whole enlistment procedure when turning 18. We use the measure of cognitive ability ranging from 1 to 9. This variable follows a Stanine scale that approximates a normal distribution. The score is standardized within each cohort of draftees to account for any minor changes in the tests over time. See for instance Grönqvist et al. [2017] for details on these measures.

Figure E-3: Heterogeneity in Estimated MRS using the RP Approach

A. Age

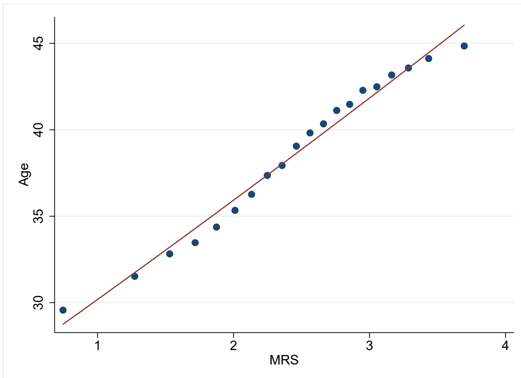

B. Gender

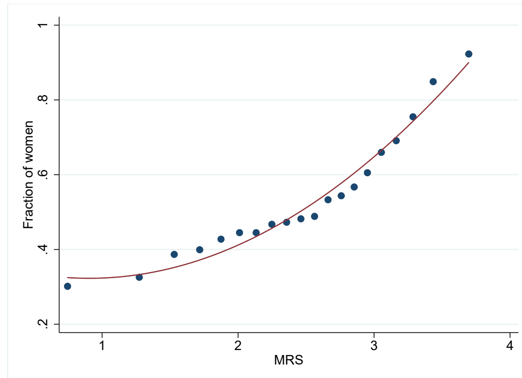

C. Fraction with Children

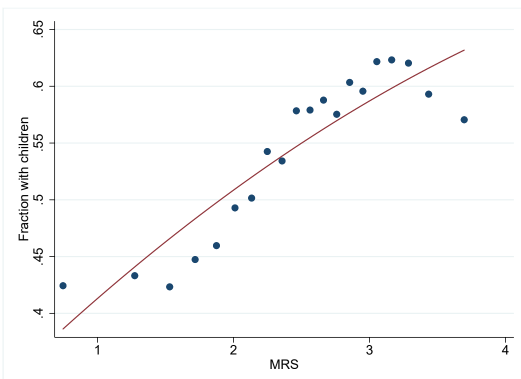

D. Net wealth

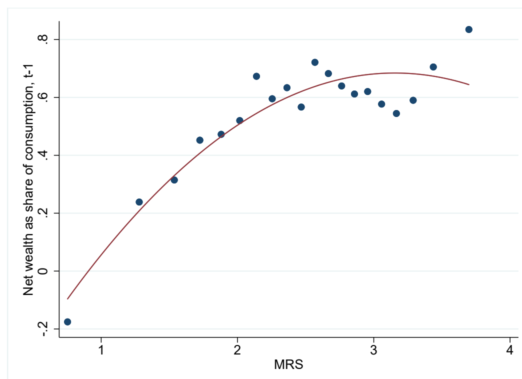

E. Liquid Assets

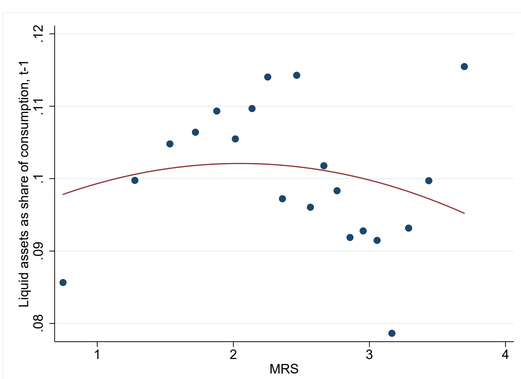

F. Cognitive Ability

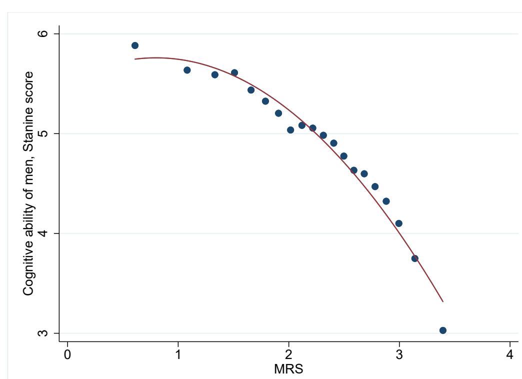

Notes: The graph correlates the estimated MRS from the RP approach with various observable characteristics, using bin scatter plots, by bins of estimated MRS. Panel A, B and C focus on age, gender and the presence of children, three types of observable characteristics that may correlate with risk preferences. In Panel D, we correlate our estimates of the individual MRS with net household wealth as a fraction of total household consumption in the year prior to job loss. Panel E looks at the amount of total household liquid assets in bank accounts as a fraction of total household consumption in the year prior to job loss. Panel F correlates the estimates of the MRS with a direct measure of cognitive ability from tests administered by the Swedish Army to all enlisted individuals. The measure of cognitive ability is ranging from 1 to 9 and follows a Stanine scale that approximates a normal distribution. The specification of the choice model underlying this exercise is reported in Table 3 Column (10), using the perceived risk model with risk estimated under comprehensive coverage.

Figure E-4: MRS vs. Consumption Drop Estimates in Baseline Sample

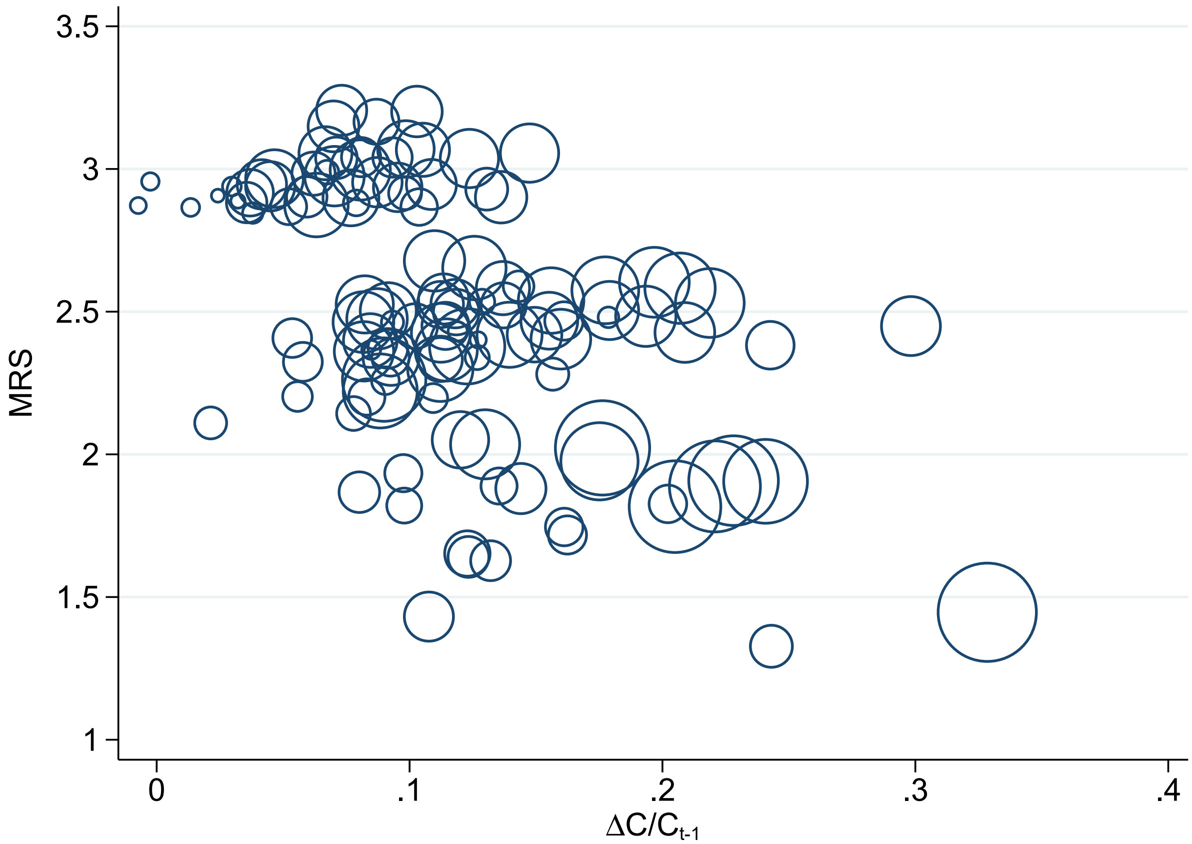

Notes: The graph correlates estimated MRS from the RP approach with estimated drops in consumption at job loss used in the CB approach for the individuals in our baseline sample split in cells of observables. We consider 120 cells of observable characteristics, using three age groups, income deciles, civil status and gender. For each cell, we calculate the average of the estimated drops in consumption at job loss for households in that cell, from specification (20). We then plot this estimate against the average MRS in the cell estimated using the RP model in the year prior to job loss. Underlying the MRS predictions is the choice model estimated using the perceived risk model, with risks predicted under comprehensive coverage. The size of each dot is proportional to the number of individuals in that cell. The graph shows that the MRS and consumption drops are negatively correlated. Conditional on consumption drops, there is also a large amount of residual variation in MRS left in the data.