DISCUSSION PAPER SERIES

DP13624
(v. 2)

THE VALUE OF UNEMPLOYMENT INSURANCE

Camille Landais and Johannes Spinnewijn

LABOUR ECONOMICS AND PUBLIC ECONOMICS

ISSN 0265-8003

THE VALUE OF UNEMPLOYMENT INSURANCE

Camille Landais and Johannes Spinnewijn

Discussion Paper DP13624
First Published 26 March 2019
This Revision 25 June 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

THE VALUE OF UNEMPLOYMENT INSURANCE

Abstract

In the absence of unemployment insurance (UI) choices, the standard approach to estimating the value of UI is to infer it from the observed consumption response to job loss in combination with some assumption on preferences. Exploiting the unique data and policy context in Sweden, we propose two alternative approaches, which we implement and compare to the standard consumption-based approach on the exact same sample of workers. We nd that the drop in consumption expenditures upon job loss is relatively small (□ 13 percent), but the marginal propensity to consume (MPC), estimated using variation in local government transfers, is around 25 percent higher when unemployed than when employed. We show that this wedge in MPCs, the focus of our rst approach, reveals a high relative price of smoothing consumption, con rmed by direct evidence on the limited consumption smoothing means available during unemployment. The estimated relative price provides a lower-bound on the value of UI, which turns out to be substantially higher than the consumption-based estimate under standard preference assumptions. Exploiting the UI choices embedded in the Swedish UI system, we also propose a Revealed-Preference approach, which con rms that the average value of UI is large in our setting, but also reveals substantial dispersion in the value of UI, above and beyond the variation in consumption drops.

JEL Classification: H20, J64

Keywords: Unemployment insurance, Consumption Smoothing, Revealed Preference, MPC

Camille Landais - c.landais@lse.ac.uk
London School of Economics and CEPR

Johannes Spinnewijn - j.spinnewijn@lse.ac.uk
London School of Economics and CEPR

Acknowledgements
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 Ten- reyro, Danny Yagan and seminar participants at Berkeley, Bristol, Bonn, Cambridge, Columbia University, IFS, IIPF, LSE, MIT, Oxford, Princeton, Stanford, Tinbergen, TSE, and Uppsala for helpful discussions and suggestions. We also thank Henriette Druba, Arnaud Dyevre, Jack Fisher, Mathilde Munoz, Will Parker and Quirin Von Blomberg for excellent research assistance. We acknowledge financial support from ERC grants 679704 and 716485.

The Value of Unemployment Insurance

Camille Landais — LSE
Johannes Spinnewijn* — LSE

June 24, 2019

Abstract

In the absence of unemployment insurance (UI) choices, the standard approach to estimating the value of UI is to infer it from the observed consumption response to job loss in combination with some assumption on preferences. Exploiting the unique data and policy context in Sweden, we propose two alternative approaches, which we implement and compare to the standard consumption-based approach on the exact same sample of workers. We find that the drop in consumption expenditures upon job loss is relatively small ($\sim$ 13 percent), but the marginal propensity to consume (MPC), estimated using variation in local government transfers, is around 25 percent higher when unemployed than when employed. We show that this wedge in MPCs, the focus of our first approach, reveals a high relative price of smoothing consumption, confirmed by direct evidence on the limited consumption smoothing means available during unemployment. The estimated relative price provides a lower-bound on the value of UI, which turns out to be substantially higher than the consumption-based estimate under standard preference assumptions. Exploiting the UI choices embedded in the Swedish UI system, we also propose a Revealed-Preference approach, which confirms that the average value of UI is large in our setting, but also reveals substantial dispersion in the value of UI, above and beyond the variation in consumption drops.

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, Jack Fisher, 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 the value of providing more insurance is easy to understand and 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 standard approach in the literature, famously implemented by Gruber [1997] in the context of unemployment, is to study consumption smoothing in response to adverse events. The estimated drop in consumption can be scaled by workers' risk aversion to get an estimate of the value of providing additional insurance. As the consumption drops are consistently found to be small in the literature, the corresponding value of insurance is deemed to be low. Given the large unemployment responses to UI, the consumption-based (CB) approach suggests that UI policies are too generous. This conclusion, however, relies on assumptions on preferences, which are difficult to relax in practice.

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. 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. 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.

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. We leverage the granularity of this measure to study the means used to smooth consumption and the relevant sources of heterogeneity in consumption smoothing. In line with prior work, the estimated consumption drops are relatively small, and translate into relatively 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 significant role, but only few unemployed workers have them. The take-up of debt in fact decreases when unemployed, suggesting that borrowing constraints become binding with the most indebted workers facing the the largest consumption drops. We do not find that the earnings of spouses and other members of the household significantly increase in response to a job loss in the household. 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 which leverages the fact that the marginal propensity to consume out of extra income is directly related to the price of increasing consumption at the margin. The relative marginal propensity to consume (MPC) when unemployed rather than employed therefore reveals the price of smoothing out unemployment shocks, which puts a bound on the value of consumption smoothing at the margin. The normalization with respect to employment and the use of MPCs rather than consumption levels makes this method robust to important challenges when converting wedges in consumption expenditures into wedges in marginal utilities (e.g., presence of home production, work-related expenditures, committed expenditures, non-durable goods, etc). The normalization makes the MPC approach also robust to behavioral frictions which affect MPCs in the same way when employed and unemployed (e.g., present bias).

To estimate the state-specific MPCs, we need comparable exogenous variation in income when unemployed and employed. 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. Estimating this in a first-difference model on the same sample of job losers used in the CB approach, we estimate a relatively high MPC while employed, of around .43. However, when unemployed, the MPC is around 25 percent higher. 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 the MRS, both non-parametrically and in a parametric choice model. We try to account for potential risk misperceptions, using salient risk shifters and belief elicitations, and other potential choice frictions, that may confound the estimates of the MRS in the RP approach. 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. The estimated mark-ups that workers are willing to pay would justify generous UI benefits, despite the high moral hazard costs. The UI policy recommendations are thus opposite to what one would conclude based on the CB approach.

To reconcile the high average value of UI with the modest consumption drops, we need state-dependence in preferences and/or higher risk aversion parameters than conventionally used in the literature. Our RP approach also reveals large dispersion in MRS, above and beyond the variation in consumption drops. We find for example important variation in MRS by age, gender and family status, which seems suggestive of significant heterogeneity in preferences. We also illustrate that preference heterogeneity, through precautionary behaviors, can in fact introduce a negative correlation between MRS and consumption drops [Chetty and Looney, 2007]. This asks for caution when using cross-sectional heterogeneity in consumption smoothing as a guide for differentiating social transfers.[1] However, similar caution is warranted when using the RP approach, since behavioral biases and other frictions will distort the value revealed through workers' insurance choices. When correcting for misperceived risks, we find a lower mean and variance in the MRS, reminiscent of the findings in Handel and Kolstad [2015]. We also show that the average cognitive ability, measured by scores from army-entry tests, is substantially lower for workers with the largest revealed values of UI.

The paper will proceed as follows. In Section 2, we set up a conceptual framework and present the three approaches. We then discuss the data and context in Section 3 before implementing the respective approaches in Sections 4-6. 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

[1] Kolsrud et al. [2018] analyze how consumption evolves during the unemployment spell to determine the optimal timing of unemployment benefits over the unemployment spell. Recommendations based on within-individual changes in consumption over the unemployment spell can arguably overcome the preference identification challenge for the CB implementation (see also Ganong and Noel [2017]).

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 model of unemployment to present the three different approaches to estimating the value of unemployment insurance. We keep the model stylized to improve the comparability of the approaches, but we discuss model extensions and provide further details in the technical Appendix A.

2.1 Setup

An agent is either employed or unemployed. The respective states are denoted by $s \in \{e, u\}$. When employed the agent has income y_e, which depends on her earnings and the taxes she pays. When unemployed the agent has income y_u, which depends on the unemployment benefits she receives.

The corresponding expected utility equals

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

where c denotes consumption and z and x_s denote the 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 efforts to find a job when unemployed and to keep a job when employed. We assume that the probability of unemployment equals $\pi(z)$ where z is the utility cost of effort.

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 an extra krona of revenue in any given state:

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

The generality of our stylized framework eases the comparison of the three approaches, and still allows to capture various models of the common resources used to smooth consumption:

Added Worker. Our framework 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,

$$-p_s \frac{\partial v_s(c_s, x_s)}{\partial x} \equiv \frac{1}{w_s} g_s'(x_s), \tag{4}$$

where $g_s(x_s)$ denotes the cost of household labor supply and $p_s = 1/w_s$ is the inverse of the household's marginal wage, which may change with a member's employment status (e.g., Fadlon and Nielsen [2018]; Hendren [2018]).

Savings. Our framework could also be extended to incorporate intertemporal decisions, where the resource cost of increasing consumption today is having to lower consumption in the future. That is,

$$-p_{s,t} \frac{\partial v_s(c_{s,t}, x_{s,t})}{\partial x} \equiv \beta R_{s,t} E_s[V_{t+1}'(x_{s,t})], \tag{5}$$

with $E_s[V_{t+1}]$ the expected continuation utility when in state s, given the assets or debt x rolled over to the next period. The price of consumption $p_{s,t}$ is the gross interest rate $R_{s,t}$, which may be higher when unemployed if liquidity or borrowing constraints become binding.

Insurance. Our framework could also be extended to incorporate insurance choice, where the resource cost of increasing consumption when unemployed is to lower consumption when employed. 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. The marginal cost of increasing state-specific resources is then equal to

$$-\frac{\partial v_s\left(c_s, x_s\right)}{\partial x} \equiv \frac{V'(A_0 - x_u - x_e)}{\pi_s}, \tag{6}$$

which is the marginal value of resources at the start of this model scaled by the state-specific probability. We fully develop the dynamic extensions with savings and insurance in Section A.2 of the technical appendix.

In this stylized model of unemployment, the optimal UI generosity is characterized by the Baily-Chetty formula (Baily [1978]; Chetty [2006]),

$$MRS = 1 + \varepsilon_{\frac{\pi}{1-\pi}}, \tag{7}$$

where $MRS = \frac{\partial v_u(c_u,x_u)}{\partial c} / \frac{\partial v_e(c_e,x_e)}{\partial c}$ is the marginal rate of substitution between consumption when unemployed and employed and $\varepsilon_{\frac{\pi}{1-\pi}}$ is the unemployment elasticity with respect to a tax-funded increase in UI. The formula relies on the envelope theorem and requires concavity and differentiability of both v and π (see Chetty [2006]). As a result, a small change in z or x has a second order impact on the agent's own welfare. The envelope conditions imply that the welfare impact of a small increase in state-specific income y_s depends only on the direct effect, captured by the state-specific marginal utility of consumption. The value of UI is therefore fully determined by the MRS, which tells us how much consumption the worker is willing to give up when employed to increase her consumption when unemployed. This mark-up simply needs to be compared to the extra fiscal cost due to the unemployment response, to evaluate whether an increase in UI generosity is desirable.

2.2 Approach I: Consumption-Based Implementation

We briefly revisit the consumption-based (CB) approach in the context of our model, which links the MRS 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 exposed to when becoming unemployed. In our model, we can state:

Proposition 1. *For state-dependence $\theta = \frac{\partial v_u/\partial c}{\partial v_e/\partial c}$ and curvature $\sigma_s^c = -\frac{\partial^2 v_s/\partial c^2}{\partial v_s/\partial c}$,*

$$MRS \cong [1 + \sigma_e^c \times [c_e - c_u]] \times \theta, \tag{8}$$

when preferences are separable and third and higher-order derivatives are small.

As is well known, the relation between the MRS and the consumption wedge can be derived from a Taylor expansion of the marginal utility of consumption when unemployed. This expansion becomes implementable when we make two types of assumptions. First, we need the higher-order derivatives of the utility functions to be small, or we would need extra information on the parameters underlying these derivatives. Second, we need preferences to be separable in consumption and resources used to increase consumption, or we would need extra information on the parameters underlying their complementarity and on the changes in resources. The resulting expression shows clearly how the MRS depends on the drop in consumption, which is scaled by the curvature of the utility function σ and any state-dependence in preferences θ.

Implementation With the right information on individuals' consumption preferences in hand, the CB approach is remarkably easy to implement. It becomes sufficient to observe the wedge in consumption between employment and unemployment to estimate the value of UI at the margin. A prime advantage of the CB approach is that it does not rely on information on the various means used for smoothing consumption. Recent work has extended the intuition of the CB approach. First, rather than looking at wedges in consumption, one can look at wedges in resources used when employed and unemployed, for example changes in spousal labor supply (Fadlon and Nielsen [2018]). Second, rather than comparing the wedge in outcomes in the realized states, one can instead consider responses to changes in the anticipated unemployment risk (Hendren [2017]).

The information on individuals' preferences, however, is essential and hard to come by. The main critiques on the CB approach have been twofold. First, the approach needs to assume how the marginal utility of consumption changes with consumption and/or employment status. The curvature of the utility function, however, may critically depend on the importance of committed expenditures [Chetty and Szeidl, 2007] and durable goods [Browning and Crossley, 2009]. The marginal utility of consumption may also be state-dependent, for example reflecting complementarity between consumption and leisure. Second, the approach needs to assume how comparable consumption expenditures are when employed and unemployed. Here, the challenge is that work- and job search-related expenditures confound the actual drop in expenditures relevant for utility [Browning and Crossley, 2001]. Moreover, complementing the same expenditures with more home production or time to shop for lower prices will increase their value [Aguiar and Hurst, 2005].

Pinning down the impact of all these confounders on the MRS or even signing the net bias is difficult. Of course, any specific challenge can be overcome with better data. For example, the original implementation in Gruber [1997] did not consider the drop in consumption upon job loss as such, but estimated the consumption response to UI benefits to then impute the consumption drop due to the net-income lost when unemployed. This imputed drop is unaffected by state-dependent changes in expenditures. More generally, a worker's smoothing of consumption always depends on both her preferences and the means she has available. A worker may smooth consumption less either because the price is high or because she cares little about the exposure. The difference is essential for inferring the value of UI and at the root of the challenge for the CB approach as information on preferences is not readily available.

2.3 Approach II: State-Specific MPC

To avoid the preference assumptions of the CB approach, we propose an alternative approach that identifies the price of consumption smoothing at the margin. This price allows to bound the value of extra consumption smoothing through UI. The key insight is that the price of smoothing consumption in a given state is linked to the marginal propensity to consume (MPC) out of state-contingent income. We can state, dropping the arguments of the partial derivatives:

Proposition 2. *For* $O_s^{mpc} = \frac{dc_s/dy_s}{1-dc_s/dy_s}$ *and curvatures* $\sigma_s^c = -\frac{\partial^2 v_s/\partial c^2}{\partial v_s/\partial c}$ *and* $\sigma_s^x = \frac{\partial^2 v_s/\partial x^2}{\partial v_s/\partial x}$,

$$MRS = \frac{O_u^{mpc}}{O_e^{mpc}} \times \frac{\sigma_u^c/\sigma_u^x}{\sigma_e^c/\sigma_e^x} \times \frac{\partial v_u/\partial x}{\partial v_e/\partial x}, \tag{9}$$

when preferences are separable.

To establish the practical relevance of Proposition 2, we will argue below that for standard preferences and models of consumption smoothing: (i) $\frac{\sigma_u^c/\sigma_u^x}{\sigma_e^c/\sigma_e^x} \geq 1$ and (ii) $\frac{\partial v_u/\partial x}{\partial v_e/\partial x} \geq 1$. As a result, the ratio of the MPC odds ratios provides a lower bound on the MRS,

$$MRS \geq \frac{O_u^{mpc}}{O_e^{mpc}}. \tag{10}$$

Now the characterization of the MRS itself (i.e., equation (9) in Proposition 2) relies on two steps:

First, optimizing workers equalize the marginal utility of consumption and the marginal cost of raising revenue, as stated in condition (3). The marginal rate of substitution between two states therefore depends on the ratio of state-specific prices:

$$MRS = \frac{p_u}{p_e} \times \frac{\partial v_u/\partial x}{\partial v_e/\partial x}. \tag{11}$$

The challenge here is to know the prices of the resources used to smooth consumption in different states.

Second, the relevant state-specific prices can be recovered from the marginal propensities to consume out of state-contingent income, dc_s/dy_s. Intuitively, the higher the shadow value 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 that the MPC depends on the state-specific price and second-order derivatives of the utility function,

$$\frac{dc_s}{dy_s} = \frac{(p_s)^2 \frac{\partial^2 v_s}{\partial x^2} + p_s \frac{\partial^2 v_s}{\partial x \partial c}}{\frac{\partial^2 v_s}{\partial c^2} + (p_s)^2 \frac{\partial^2 v_s}{\partial x^2} + 2p_s \frac{\partial^2 v_s}{\partial x \partial c}}. \tag{12}$$

Using the optimality condition (3) again and assuming separable preferences, we can re-express the

MPC as

$$\frac{dc_s}{dy_s} = \frac{p_s \frac{\partial^2 v_s}{\partial x^2} / \frac{\partial v_s}{\partial x}}{-\frac{\partial^2 v_s}{\partial c^2} / \frac{\partial v_s}{\partial c} + p_s \frac{\partial^2 v_s}{\partial x^2} / \frac{\partial v_s}{\partial x}}, \tag{13}$$

which is a simple function of price and curvature of the utility function (wrt consumption and resources). Expressed as an odds ratio and using short-hand notation for the curvature, this further simplifies to

$$O_s^{mpc} = \frac{dc_s/dy_s}{1 - dc_s/dy_s} = p_s \frac{\sigma_s^x}{\sigma_s^c}. \tag{14}$$

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, which is captured by the relative curvature $\frac{\sigma_s^x}{\sigma_s^c}$.

The MPC approach uses these two steps combined: the marginal utility of consumption depends on the shadow price of income in that state, which can be inferred from the state-specific MPC. The MRS between unemployment and employment thus depends on the difference in MPCs when unemployed and when employed respectively.

We have presented the derivation underlying the MPC approach in a stylized model. Appendix A.2 shows that the exact same relationship holds between MRS and MPCs in a dynamic model with savings, considering the MPC out of a temporary shock in state-contingent income. Appendix A.3 then further generalizes this relationship in extensions with multiple resources, multiple consumption goods, state-specific expenditures and endogenous prices.

Lower Bound Proposition 2 also identifies the two factors that can confound the relationship between the MPCs and the MRS and thus affect the potential value of the MPC approach:

A first confounding factor would be any difference in the relative curvature σ_s^c/σ_s^x across states. As argued in the context of the CB approach, obtaining information on the curvature of preferences is difficult. However, compared to the CB approach, the MPC approach does not require any information or assumptions on the curvature of preferences itself, but on how the curvature differs across states. For example, in a model with savings or insurance, it is only the difference in the preference curvature with respect to consumption (i.e., the coefficient of absolute risk aversion) when employed vs. unemployed that matters. For CARA preferences, this potential confounding factor naturally disappears. For CRRA preferences, the difference in MPCs under-estimates the difference in prices, implying that the MPC ratio provides a lower bound on the MRS. This, however, assumes no other state-dependence in the curvature of preferences.[2]

A second confounding factor would be any difference in the marginal resource cost $\partial v_s/\partial x$ across states. Like the marginal utility of consumption, this cannot be directly observed. However, since income is lower when unemployed, we expect more state-specific resources to be used when

[2] When the curvature is not constant, but depends on the consumption and resource levels, the relative curvature will indirectly depend on the prices as well. Information on the consumption and resource wedge can help quantifying this factor for different curvature parameters.

unemployed, causing the marginal cost of generating resources to be higher. This would result in a MRS that is greater than the price ratio, again implying that the MPC ratio provides a lower bound on the MRS. Stated differently, as income drops when unemployed, the MPC approach only picks up the corresponding value of insurance as revealed by workers' willingness to pay a higher price to increase their consumption when unemployed.[3]

We develop the lower bound arguments formally in Appendix A.4.

Implementation To implement the MPC approach, we need comparable exogenous variation in state-contingent income both when employed and unemployed to then estimate the differential response in consumption.[4] The MPC approach is closely related to the approaches followed by Chetty [2008] and Landais [2015], who instead consider the differential response in unemployment risk to changes in unemployment benefits relative to other income changes. In contrast with previous approaches, the MPC approach does not require assumptions on the resources used, since, by construction, the MPC reveals the shadow cost of the resource that is used at the margin.[5] The MPC approach can thus also account for precautionary means used to smooth consumption, but only when using anticipated variation in state-contingent income so that workers can adjust precautionary means in response.[6]

We also note that the MPC approach is robust to some of the important implementation challenges for the CB approach (beyond requiring information on preferences). First, work- or job search related expenditures affect the drop in consumption expenditures between employment and unemployment, but do not change the MPCs and how they relate to the state-specific prices. Second, the differential availability of home production provides another source of state-specific variation in the price of consumption, which is exactly what is picked up by the MPC approach. Third, more generally, the robustness of the MPC approach comes from the fact that potential confounders that do affect the relationship between the MPC and the state-specific price, arguably do so in the same way when employed and unemployed. Hence, these factors would cancel out when using the ratio of state-specific MPCs. For example, this holds for consumption commitments, or durable good expenses, which affect the curvature of consumption preferences, but arguably in the same way when employed as when unemployed. We develop these arguments formally in Appendix A.3.

[3] For example, in a dynamic model with no borrowing constraints, the MPCs can be the same (e.g., for CARA preferences, or when job loss implies a permanent income shock), although the MRS is greater than 1. The MPC approach does, however, capture the increase in price that a decrease in resources when unemployed may imply. We consider an extension of the model with endogenous prices in Appendix A.3.

[4] In the original implementation of the CB approach, Gruber [1997] already used estimates of the MPC when unemployed, based on variation in UI benefits, to impute the drop in consumption caused by the drop in income when unemployed. Since then, most papers have simply considered the drop in consumption when unemployed.

[5] Similar to the MPC approach, the approach in Chetty [2008] and Landais [2015] relies on implicit differentiation of an optimality condition, but now for search effort z. This would provide an exact estimate of the MRS rather than a bound, but only when UI does not affect other consumption smoothing means (e.g., savings, household labor supply).

[6] In practice, there may be a trade-off between finding anticipated and exogenous variation in income. This issue is less binding if the relevant margins of adjustment are *ex post* means of consumption smoothing (e.g., household labour supply).

Importantly, the same robustness argument applies for behavioral frictions, which do not invalidate the MPC approach as long as they affect the MPC in a similar way when employed and unemployed (e.g., present bias).[7] Still, as noted before, while the MPC approach may continue to identify the difference in prices, it only allows us to bound the MRS. This bound can be less or more informative, affecting the practical value of the MPC approach. For example, state-dependent preferences or optimistic expectations [e.g., Spinnewijn [2015]], which change the utility of unemployment consumption relative to (future) employment consumption, would make the bound less informative if they increase the relative use of resources when unemployed (and thus $\frac{\partial v_u/\partial x}{\partial v_e/\partial x}$) without being reflected in prices.[8]

2.4 Approach III: Revealed Preference

By identifying the shadow price of income, the MPC approach helps relaxing the assumptions on preferences needed for the CB approach. If prices are observable instead, we could estimate the value of insurance more directly. In particular, when workers are offered choice in UI coverage, we can use the insurance price in combination with an estimate of an individual's risk to bound her MRS:

Proposition 3. *When offered extra UI coverage at rate $dc_u/dc_e = -p_e/p_u$, the MRS for an individual (not) buying the extra coverage is bounded from below (above) by the expected price*

$$\tilde{p} \equiv \frac{p_u}{p_e} \times \frac{1-\pi}{\pi}. \tag{15}$$

The extra UI coverage allows a worker to increase her consumption when unemployed at the expense of her consumption when employed. As discussed before, this corresponds to the availability at the margin of unemployment- and employment-contingent securities at prices p_u and p_e respectively. The worker will buy the extra UI coverage only if

$$\pi \frac{\partial v_u}{\partial c} \frac{1}{p_u} + (1-\pi) \frac{\partial v_e}{\partial c} \frac{1}{p_e} \geq 0. \tag{16}$$

The willingness to take up the extra insurance depends crucially on the worker's unemployment risk π. For given terms, the expected price per unit of coverage $\tilde{p}$ is lower the higher one's unemployment risk. A worker will only take up the extra UI coverage if the MRS exceeds this expected price and vice versa. In an Arrow-Debreu market, an individual would buy coverage up to the point

[7] This robustness seems particularly relevant in light of the recent evidence on excess sensitivity of consumer spending to cash-on-hand during unemployment [e.g., Ganong and Noel [2017], Gerard and Naritomi [2019]], which for our purposes should be compared to the sensitivity of consumer spending when employed.

[8] We note two distinct issues in relation to behavioral frictions. First, like other so-called optimization approaches, our approach relies on individuals not being at a corner in their use of consumption smoothing means. If so, the MPC can be equal to 1 for small changes in income. MPC estimates based on larger income variation would overcome this issue (unless individuals are fully constrained, corresponding to a price of increasing consumption that is infinite). Second, as we can no longer invoke envelope conditions, the presence of frictions affects the use of the MRS in evaluating the desirability of UI [Spinnewijn [2015]]. The presence of frictions naturally calls for different types of interventions as well.

that the MRS equals the expected price. This corresponds to condition (11) 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]).

Implementation The RP approach is the most direct approach and in principle allows us to relax any parametric assumptions on preferences.[9] However, it adds two very strong data requirements. First, it requires the availability of insurance choices, which are often absent in the context of social insurance and motivated the CB approach.[10] Second, it requires information on individuals' risk to estimate the expected price of coverage determining their choice.[11] In fact, we would need the worker's perception of the unemployment risk at the time he or she decides to buy the insurance. The presence of biased beliefs, or behavioral frictions more generally, poses an important challenge for the implementation of the revealed preference approach as they would be wrongly attributed to workers' valuation of insurance and thus bias our estimates of the MRS. Additional data or assumptions are required to deal with these choice frictions. We come back to this in detail in Section 6.

In practice, when exogenous variation in expected prices is available, we can go beyond bounding the MRS and uncover the MRS distribution in the population. Insurance coverage choices, however, are discrete, which entails two extra caveats. First, we are no longer identifying the MRS at the margin, like in the CB approach, but an average MRS for the supplemental coverage which workers can choose to buy. Indeed, for risk-averse workers the MRS is decreasing in the coverage level so that the expected price for workers buying the supplemental coverage may no longer provide a lower bound on their MRS evaluated at the margin. Second, we can no longer invoke the envelope theorem to conclude that only the MRS and expected price determine a worker's coverage choice. Indeed, the supplemental coverage may crowd out the use of other means to smooth consumption and induce moral hazard responses as the worker reduces her effort to avoid unemployment with first-order effects on their utility. In this case, we can still bound the MRS as we show formally in Appendix A.5. The intuition is that the unemployment risk is increasing in coverage, and the corresponding expected price used to bound the MRS thus decreasing, due to moral hazard. In particular, for individual buying the extra coverage, the lower expected price for coverage, based on the predicted risk under the extra coverage, will still be a lower bound on their MRS. Similarly, for the individuals not buying the extra coverage, the expected price based on the predicted risk without the extra coverage will still be an upper bound on their MRS.

[9]Note that the MPC method is still valid when workers use insurance to smooth consumption at the margin. However, it may provide an uninformative bound on the MRS. For actuarially fair insurance of an unemployment risk $\pi < 1/2$, the MPC ratio will be below 1 as the price of unemployment-contingent consumption is lower than the price of employment-contingent consumption.

[10]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.

[11]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].

3 Data and Context

The three approaches offer complementary ways of estimating the value of insurance. The most fruitful approach will depend on the data available and the implementation assumptions required given the specific context. We exploit the unique institutional setting and data in Sweden to estimate the value of UI using the three different approaches in the exact same context on the same individuals. To recap, the CB approach 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 merge data from various registers in Sweden which allows us to implement all three approaches on the same sample of workers. We present here the institutional background and data used.

3.1 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.[12] 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.[13] 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. 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 contributing to an unemployment fund to get the comprehensive coverage.[14] Further institutional

[12] Benefits are paid per "working day", which means that there are 5 days of benefits paid per week. Benefits of 320 SEK a day therefore translate into 6960 SEK a month (≈765 USD).

[13] Apart from the level of benefits, there are no coverage differences between the basic and the comprehensive UI scheme. Individuals can receive unemployment benefits indefinitely in theory during our period of analysis. To continue receiving benefits after 60 weeks of unemployment, the unemployed must accept to participate in counselling activities and, potentially, active labor market programs set up by the Public Employment Service.

[14] 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. The

details underlying the UI choice are provided in Landais et al. [2017].

3.2 Data & Sample

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.[15] 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.[16] 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.[17]

Our measure of household consumption expenditure is linked to rich data on unemployment histories and unemployment insurance choices coming from Swedish registers covering the universe of Swedish individuals.

Data on Unemployment (Risk) Unemployment history 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, unemployment status and participation in labor market programs.[18]

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 deter-

government does not allow UI funds to charge different prices to different individuals, except for a small rebate to union members and unemployed workers.

[15] 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.

[16] All financial institutions are compelled to report this information directly to the tax administration for the purpose of the wealth tax (which was abolished in 2007). All asset holdings are reported to the tax administration at the individual level. We aggregated assets at the household level using household identifiers from the registry data.

[17] 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`

[18] 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]).

minants 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 following proposition 3. 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 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.[19] But most individuals (91%) in our sample are contributing to the comprehensive coverage prior to job loss.

4 Consumption-Based Approach

We start by revisiting the consumption-based implementation in the Swedish context, which provides us with a benchmark estimate to compare our alternative implementations to. We further exploit the nature of our registry-based expenditure measure to shed light on the means used to smooth consumption and the dimensions of heterogeneity in consumption smoothing.

4.1 Baseline Implementation

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:

$$C_{it} = \alpha_i + \nu_t + \sum_{j=-N_0}^{N_1} \beta_j \cdot \mathbb{1}[J_{it} = j] + \varepsilon_{it} \tag{17}$$

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 3.2 above. We follow the recent literature on event studies [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

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

independently of the dynamic treatment effect of the event $\{\beta_j\}_{j=-N_0}^{N_1}$.[20]

We estimate specification (17) 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 (17). 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 (17) 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 close to linear with an intercept equal to zero, as predicted by our simple parametric model. Based on this evidence, we can estimate the following specification:

$$C_{it} = \alpha_i + \nu_t + \sum_{\substack{j=-N_0 \\ j \neq 0}}^{N_1} \beta_j \cdot \mathbb{1}[J_{it} = j] + \beta_0 \cdot M_i \cdot \mathbb{1}[J_{it} = 0] + \varepsilon_{it} \tag{18}$$

In the above regression, the flow drop in consumption at unemployment $\frac{c_u - c_e}{c_e}$ is identified by $\frac{12 \cdot \hat{\beta}_0}{\widehat{C}_{-1}}$. 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.

We can now scale the estimated consumption drop by the curvature in preferences over con-

[20] 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.

sumption to get an estimate of the MRS. We follow the most standard version of the CB implementation, approximating $MRS \cong 1+\gamma \times \frac{\Delta c}{c}$ with $\gamma = \sigma c$ equal to the relative risk aversion. This approximation relies on negligible third- and higher-order derivatives of the utility function and state-independence in marginal utilities. In the absence of consensus on γ, we report estimates of the MRS for γ ranging from 1 ($MRS = 1.13$) to 4 ($MRS = 1.51$) in Figure 1. 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.

Anticipation The validity of the standard CB implementation depends on the extent to which job losses are anticipated (see Gruber [1997], Hendren [2017]). 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.[21] 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 (17) 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. 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.

4.2 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 (18) replacing consumption

[21] In comparison, the RP approach measures the insurance value at the time the insurance decision is made. See also Hendren [2018] on measuring the *ex ante* value of insurance.

by each component of total household expenditures. In Figure 2, 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 ($\frac{12\cdot\hat{\beta}_0^X}{\widehat{C}_{-1}}$). 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 2 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 (18) with bin dummies. We run the following specification:

$$C_{it} = \alpha_i + \nu_t + \sum_h \sum_{\substack{j=-N_0 \\ j\neq 0}}^{N_1} \beta_{jh} \cdot \mathbb{1}[J_{it} = j] \cdot \mathbb{1}[H_i = h] + \sum_h \beta_{0h} \cdot M_i \cdot \mathbb{1}[J_{it} = 0] \cdot \mathbb{1}[H_i = h] + \varepsilon_{it} \quad (19)$$

In Figure 3 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 (19). 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.[22] 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.

Translating these heterogeneous patterns of consumption drops into heterogeneous valuations of UI is not straightforward. This requires information on preferences and the potential correlation between consumption drops and preferences, which also affect the precautionary actions taken by households to help them smooth consumption at unemployment [Chetty and Looney [2007], Andrews and Miller [2013]]. Households with liquid assets may have smaller consumption drops but not necessarily a lower valuation of the UI. It could precisely be their risk aversion that induced them to accumulate liquid assets as precaution. Conversely, this raises the question: if many individuals have little means to smooth consumption at unemployment above what transfers are providing them, is this indicative of a low risk aversion and therefore of a low valuation of additional insurance? Or is it an indication that the shadow price of smoothing consumption is high? Taken together, the evidence from Figures 2 and 3 provides some insights already. The role of liquidity and the decrease in consumption out of debt at unemployment, the important role played by transfers, and the sensitivity of consumption drops to transfers all suggest that liquidity and borrowing constraints may bind, and that the cost of consumption smoothing may be high. In the next section, we analyze marginal propensities to consume while employed and unemployed precisely with the aim of identifying these shadow prices of consumption smoothing.

5 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.

[22] 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.

5.1 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.[23] 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.[24] 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 of characteristics $\mathbf{V}_{it} = \{\mathbf{X}_{it}, y_{it}, a_{it}\}$:

$$B_{imt} = \sum_k \tau^k_{mt} \cdot V^k_{it}, \tag{20}$$

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

[23]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.

[24]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 observable characteristics $\mathbf{V}_{it}$,

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

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_m^k from formula (20) differently (i.e., the functions $G_m(\mathbf{X})$, $\tau_m^y(y)$ and $\tau_m^a(a)$ 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.[25]

Figure 4 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 residualize welfare transfers using the following specification,

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

We then plot in a map the statistics $\mathbf{E}[\tilde{B}_{imt}|V^k = v^k] - \mathbf{E}[\tilde{B}_{imt}|V^k = v'^k]$ for each municipality

[25] 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.

m. From formula (20) which defines welfare transfers B_{imt}, we have that $\tilde{B}_{imt} = (\tau_m^k - \bar{\tau}^k)V_{it}^k$. 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 4, 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 $\tilde{B}_{imt}$ 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 $\tilde{B}_{imt}$ 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 $\tilde{B}_{imt}$ 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.

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. The map suggests significant variation within municipality over time in the generosity of welfare transfers.

5.2 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 4 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.[26]

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

$$C_{imt} = \alpha_i + \nu_t + \eta_m + \mathbf{V}'_{it}\beta + \mu_e \cdot \tilde{B}_{imt} + \mu_u \cdot \tilde{B}_{imt} \cdot \frac{M_i}{12} \cdot \mathbb{1}[J_{it} = 0] + \eta \cdot \frac{M_i}{12} \cdot \mathbb{1}[J_{it} = 0] + \upsilon_{imt}. \quad (21)$$

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 $\mathbf{V}_{it}$:

$$B_{imt} = \sum_{k \in \mathbf{K}} \tau^k_m V^k_{it} + \sum_{j \in \mathbf{K}^\mathsf{c}} \bar{\tau}^j V^j_{it} + \tilde{B}_{imt}. \quad (22)$$

This specification interacts characteristics V^k, for $k \in \mathbf{K}$, with municipality fixed-effects, while characteristics V^j for $j \in \mathbf{K}^\mathsf{c}$ are not interacted with municipality fixed effects. This residualization means that we exploit variation in the way the schedule actually varies across municipalities according to characteristics $V^j,\ j \in \mathbf{K}^\mathsf{c}$, but we shut down identification coming from differences in the schedule according to all other characteristics $V^k,\ k \in \mathbf{K}$. As a baseline, we exploit variations stemming from all characteristics $\mathbf{V}$ and from local reforms over time. That is, we do not interact any characteristics V^k with municipality fixed effects in (22), neither do we include municipality-year fixed effects. But by varying the characteristics included in $\mathbf{K}$ and/or including municipality-year fixed effects, we can shut down particular sources of variations in the schedule of welfare transfers. Our identifying assumption is that the residual variation is orthogonal to the dynamics of household consumption conditional on $\mathbf{V}$. In other words differences across municipalities in the schedule τ^j_m for characteristics V^j or over time are not correlated with any other unobserved characteristics affecting consumption. We probe into the credibility of this assumption below.

Figure 5 shows 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

[26]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]).

t.[27] 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 (21) estimated in first-differences. 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.[28] The Table confirms that the estimated MPC out of local transfers is large and significantly larger when unemployed (.55) than prior to unemployment (.44). We also report in the table the ratio of the odds ratios of the MPCs, our estimate for the lower bound on the MRS following formula (10), and its corresponding standard error, 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. Our baseline results imply a large lower bound for the MRS of 1.59 (.21).

The remainder of Table 2 explores the sensitivity of our results to exploiting different sources of underlying variation in $\tilde{B}_{imt}$. This can be done by adding different characteristics in the set $\mathbf{K}$ in residualization (22). In column (2), we interact income and asset deciles with municipality fixed effects, so as to exploit only 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 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), finally, we do the opposite exercise and add the interaction of municipality and 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 stemming from local reforms over time, in case we are concerned that these reforms are endogenous.

In all specifications, we find that the MPC out of welfare transfers is large in both states, when employed and unemployed. This is in line with evidence from section 4.2 that individuals facing unemployment have relatively limited means to smooth consumption to start with so that their

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

[28] This baseline residualization approach corresponds to the case where $\mathbf{K}^c$ includes all characteristics of vector $\mathbf{V}$ in specification (22).

consumption is expected to be sensitive to variation in the transfers they receive. 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. Only when restricting to variation due to local reforms over time, the difference is not as stark ($\sim 10\%$, see column 4).

Robustness 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 other 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.

A second concern is that $\tilde{B}_{imt}$ may be correlated with employment status. While we directly control for employment status in specification (21), 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 final concern 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.

Discussion As outlined in Section 2.3, the relative odds ratios of state-specific MPCs provide an estimate of the relative price of raising consumption when unemployed vs. employed. Our estimates of the relative price range p_u/p_e from 1.2 to 1.6, as reported in Table 2. This confirms that the price of increasing consumption at the margin is significantly higher when unemployed than when employed, which again resonates with our previous evidence of section 4.2, showing that liquidity and borrowing constraints may be binding during unemployment. Under the assumptions discussed before, this ratio translates into a lower bound on the MRS. According to our preferred specification, the lower bound on the MRS is 1.59 (.22). Overall, the MPC approach therefore suggests that the value of unemployment insurance is substantially larger than what can be inferred from the traditional CB approach. 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 (18) 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.

We further evaluate the validity of our MRS estimates using the MPC approach in three different ways:

First, we assess the internal validity of our 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: individuals receive a replacement rate of 80% of their previous daily wage up to a cap. 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]. 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 is sensitive to the identification strategy chosen in a given sample. For the sake of brevity, we report all the details of the estimation in appendix Appendix E. Appendix Figure E-1 displays the identification design and the main result. We find a significant kink into the relationship between the daily wage and the drop in consumption at unemployment. This translates into an estimate of the MPC while unemployed of .63 (.16), which is very robust across specifications. This estimate is remarkably similar to our MPC estimate while unemployed using the local welfare transfer variation of .55 (.02). This evidence provides additional internal credibility to our estimates of the MPC.

A second way to probe into the 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 compares well with 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.[29] 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.

[29] 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.

Interestingly, 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.[30] 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.

Finally, we can compare our implied MRS to other "optimization" approaches used to estimate the value of unemployment insurance. Our finding of a higher value of UI using the MPC approach compared to the CB approach is in line with prior work relying on the implicit differentiation of household's optimality conditions (e.g., Chetty [2008], Landais [2015]). Our results confirm the high value of UI using yet an alternative (and arguably more robust) "optimization" method and, importantly, our setting allows us to implement both the standard CB approach and the "optimization" approach in the same context and on the same sample.

6 Revealed Preference Approach

Instead of studying how consumption smoothing (or other behavior) changes when unemployed, this section turns to the ex ante insurance choice embedded in the Swedish UI setting. We use workers' insurance choices in combination with their predicted risk to implement the Revealed Preference approach outlined in Section 2.4.

6.1 Non-Parametric Estimates

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.

Expected Price 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.[31]

[30] 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.

[31] 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).

On average the comprehensive plan increases the net daily UI benefit by 188 SEK. As explained in section 3.1, 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.[32]

To predict the unemployment risk, we use the risk model introduced in Section 3.2 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 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, $\tilde{p}_i = \frac{p_u}{p_e} \frac{1-\pi(Z_i)}{\pi(Z_i)}$.[33] 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.

Bounds We compute the individual expected price per unit of coverage $\tilde{p}_i$ for the entire population of Swedish workers aged 25 - 55 between 2002 and 2007. We estimate the model separately on the workers buying comprehensive and basic coverage to account for the presence of moral hazard, as discussed in 2.4. Following the RP approach, the expected price $\tilde{p}_{1,i}$ 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 $\tilde{p}_{0,i}$ using the predicted unemployment risk under basic coverage provides a (conservative) upper bound on the MRS.

We start by simply averaging 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 seem to face a price that implies 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.

Heterogeneity To obtain more informative bounds on the distribution of MRS using the RP approach, we need variation in the expected price per unit of coverage $\tilde{p}_i$, which can come from variation in p_u/p_e or in the underlying risk π_i. We will leverage the fact that there is a significant amount of heterogeneity in risk π_i in our context. Figure 6 divides workers in different cells

[32] 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.

[33] 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.

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.[34] The graph illustrates that there is substantial variation in the expected price per unit of coverage across cells. On average, workers facing a higher expected price are less likely to buy the comprehensive coverage, but there is substantial variation in take-up. Of course, some determinants of the unemployment risk may also directly affect the willingness to buy insurance and affect the observed relationship, either introducing positive or negative correlation.[35] 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.

6.2 Structural Estimation

We now put additional structure on the choice model to estimate the distribution of MRS using the variation in expected prices across workers. Building on the result in Proposition 3 that individuals buy additional coverage when their MRS is larger than the expected price, we assume the following linear choice model: an individual i buys comprehensive insurance at time t if and only if

$$X_{it}\beta - \gamma\tilde{p}\left(Z_{it}\right) + \varepsilon_{it} \geq 0. \tag{23}$$

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 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 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 use in our risk model both tenure ranking and the

[34] The observables include age, gender, marital status, education, income, unionisation, region and predictors of unemployment risk (firm layoff rate, unemployment history, tenure rank). Appendix Figure D-1 shows the same non-parameteric plot, but using the predicted risk under basic coverage to calculate the expected price.

[35] 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.

full history of firm layoff notifications, as well as their interaction, as shifters of workers' risk.[36]

Panel A of 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. 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 again 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.

Risk (Mis)perception The structural estimation relies 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 for the RP approach is that workers perceive their risk differently from what is predicted by our risk model. A growing empirical literature documents biases in perceived risks using belief elicitations [e.g., Spinnewijn [2015]] and insurance choices [e.g., Sydnor [2010]]. Appendix Figure D-2 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. For comparison, using the Survey of Consumer Expectations in the US, we find a similar estimate of .27 (.08) in a regression of job loss on elicited beliefs about job loss.

We explore the sensitivity of the estimation results for our choice model 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.[37] In the second specification, we convert the predicted risk using our baseline model into an estimate of the perceived risk, assuming an imperfect correlation of .26 based on the estimates above.[38] The change in estimates is comparable for both specifications. 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 perceptions, the mean MRS is 2.13, which is now very close 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

[36] 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].

[37] 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.

[38] The low coefficient estimate when regressing true job loss on perceived can 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] separately identify the correlation and the relative variance in the context of job finding probabilities, also using the Survey of Consumer Expectations, and find a correlation coefficient that is approximately the OLS estimate in that specific context.

the 75th percentile at 3.43. Panel A of Figure 7 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. Figure 6 plots the implied demand curve from the MRS estimation using the perceived risk model, showing the average share of individuals predicted to buy the comprehensive coverage at different prices. This allows comparing the estimated demand curve to the observed shares of individuals buying the comprehensive coverage at the observed (endogenous) expected prices.

Other Sensitivity Columns (4) and (5) of Table 3 allow for extra controls affecting the MRS in the choice model (using again the perceived risk 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 or MRS distribution. 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.

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.[39] 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 7 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 structural estimation implies that the average MRS is substantially higher than the CB estimates indicate, corroborating the findings of the MPC

[39]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.

approach. Moreover, the RP estimation shows substantial heterogeneity in the value of insurance, above and beyond the heterogeneity in unemployment risk.

7 Comparison of Results

We have implemented the three approaches on the exact same set of individuals. Figure 8 puts these results together.

For the CB approach, we highlight 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). The corresponding mark-up workers would be willing to pay to transfer a krona of consumption from employment to unemployment lies between 13 and 51 percent. For the MPC approach, we report the ratio of the MPCs, expressed as odds ratios, using the within municipality variation in local welfare transfers, both across household types and over time. This ratio provides a lower bound on the MRS of 1.59 (.22) (see Table 2) and thus a corresponding mark-up of about 60%. For the RP approach, we report a mean of 2.13 (.02), together with the distribution of MRS in our baseline sample. These estimates are obtained using the parametric RP model, where we use the perception-adjusted predicted risks under comprehensive coverage (see Figure 7). The mean estimate corresponds to a mark-up of more than 100%.

We also display in Figure 8 a range of plausible values for the moral hazard (MH) costs of UI, which is equal to 1 plus the unemployment elasticity $\varepsilon_{\frac{\pi}{1-\pi}}$. Following the Baily-Chetty formula in (7), an increase in the generosity of UI is desirable only if the insurance value is higher than the MH cost. In the same Swedish context as the one used here, 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 take [1.5-2.5] as a range of credible estimates of the moral hazard cost of UI in our context.

Average UI valuation The first important conclusion from Figure 8 is that the estimated average value of insurance is substantially larger in the MPC and RP approaches compared to the CB implementation. Furthermore, 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.[40] 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.

What can explain this discrepancy between the CB approach and the two alternative approaches? An obvious way to reduce the gap is to assume even higher values of risk aversion, for example driven by committed expenditures (Chetty and Szeidl [2007]). There is little consensus about the appropriate γ, which in part motivated the alternative approaches (see Chetty and

[40] This conclusion is similar to the conclusion in Gruber [1997], who in the US context finds a low value of UI based on a moderate consumption drop at unemployment relative to the MH costs and thus concludes that UI may be too generous in the US.

Finkelstein [2013]), but then the question remains whether even larger values ($\gamma \geq 8$) that fully close the gap are plausible. An alternative explanation of the gap are state-dependent preferences. Both the MPC and RP approaches do not impose state-independence, while the standard implementation of the CB approach does. In the context of unemployment, state-dependence in preferences can be driven by for example home production opportunities, which are arguably different when unemployed. As discussed in section 2.2, to account for state-dependence in the CB approach, we need to scale up the standard estimate by $\theta = \frac{\partial v_u/\partial c}{\partial v_e/\partial c}$, capturing the worker's willingness to give up employment consumption for unemployment consumption when consumption expenditures are the same in both states. For a standard risk aversion of $\gamma = 3$, the state-dependence needed to rationalize the discrepancy between the CB approach and the other approaches would be quite large ($\theta > 1.5$). But the required state-dependence goes down as the risk aversion increases (e.g., $\theta > 1.2$ for $\gamma > 5$). Combining the consumption drop and MRS constructively to learn about the structural preference parameters is difficult and subject to the same challenges as for the CB implementation.

Heterogeneity in UI valuation The second important conclusion from our 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 D-3, 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.[41] 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 9 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 3), these three characteristics are responsible for a significant amount of the residual heterogeneity in MRS

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

conditional on consumption drops displayed in Appendix Figure D-3. 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 9 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 6 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 9 provides additional evidence showing that cognitive ability is negatively correlated with the estimated MRS. The measure of cognitive ability comes from tests administered by the Swedish Army to all enlisted individuals.[42] 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.

8 Conclusion

The challenges to the standard consumption-based approach in estimating the value of insurance have inspired us to develop two alternative approaches. The MPC approach uncovers the price of smoothing consumption to bound the value of insurance at the margin. The MRS approach uses the choice of extra insurance at a known price, but accounting for the predicted risk, to bound the value of the extra insurance. We have implemented all three approaches for the same sample of workers and in the same setting. Our analysis calls for caution when using observed consumption responses to adverse events, especially when considering the optimal allocation of social transfers across workers. More constructively, we have also shown how changes in marginal propensity to consume when hit by an adverse event may overcome some of the major obstacles. We have applied this in the context of unemployment insurance, but our methods can be generalized to any other social insurance program. Any embedded choice will allow us to further improve the evaluation of these programs, but whether the introduction of (unregulated) choice is desirable will

[42] 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.

crucially depend on the presence of choice frictions. Our analysis in Sweden has revealed large heterogeneity in willingness-to-insure, conditional on risk and exposure. However, in line with the growing evidence of choice in health insurance markets, the role of risk misperceptions and choice frictions may be sizeable. Our focus on the insurance value of social insurance programs follows the tradition in the literature, mostly ignoring the corresponding redistributive value. In order to make progress in this dimension, differences in MPCs above and beyond differences in consumption across recipients of social transfers and tax payers, may provide useful information to guide policy further.

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

Aguiar, Mark and Erik Hurst, "Consumption versus Expenditure," *Journal of Political Economy*, 2005, *113* (5), pp. 919–948. 8

Andrews, Isaiah and Conrad Miller, "Optimal social insurance with heterogeneity," *Working Paper*, 2013. 21, 36

Arrow, Kenneth, "Uncertainty and the welfare economics of medical care," *The American Economic Review*, 1963, *53* (5), 941–973. 13

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. 17

Browning, Martin and Søren Leth-Petersen, "Imputing consumption from income and wealth information," *The Economic Journal*, 2003, *113* (488), 282–301. 15

_ and Thomas F Crossley, "Unemployment insurance benefit levels and consumption changes," *Journal of Public Economics*, 2001, *80* (1), 1–23. 4, 8, 18

_ 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. 8

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. 29

Cabral, Marika and Mark Cullen, "Estimating the Value of Public Insurance Using Complementary Private Insurance," *Forthcoming American Economic Journal: Economic Policy*, 2016. 5, 13

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

_ , "Moral hazard versus liquidity and optimal unemployment insurance," *Journal of Political Economy*, 2008, *116* (2), 173–234. 4, 11, 25, 29

_ and Adam Looney, "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. 4, 21, 36

_ and Adam Szeidl, "Consumption Commitments and Risk Preferences," *The Quarterly Journal of Economics*, 2007, *122* (2), pp. 831–877. 8, 34

__ and Amy Finkelstein, "Social insurance: Connecting theory to data," in "Handbook of Public Economics," Vol. 5, Elsevier, 2013, pp. 111–193. 2, 34

Eika, Lasse, Magne Mogstad, and Ola Vestad, "What can we learn about household consumption from information on income and wealth," 2017. 5, 15

Fadlon, Itzik and Torben Heien Nielsen, "Household labor supply and the gains from social insurance," *Forthcoming Journal of Public Economics*, 2018. 5, 6, 8

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, 13

Freyaldenhoven, Simon, Christian Hansen, and Jesse M Shapiro, "Pre-event trends in the panel event-study design," *Working paper*, 2018. 17

Ganong, Peter and Pascal Noel, "Consumer spending during unemployment: Positive and normative implications," *Forthcoming American Economic Review*, 2017. 4, 5, 12, 18

__ and Simon Jaeger, "A Permutation Test for the Regression Kink Design," *Journal of the American Statistical Association*, 2018, *113* (522), 494–504. 77, 78

Gerard, Francois and Joana Naritomi, "Job Displacement Insurance and Consumption: Evidence from Brazil," *working paper*, 2019. 5, 12

Grönqvist, Hans, J Peter Nilsson, Per-Olof Robling et al., "Early lead exposure and outcomes in adulthood," *Working Paper, IIES*, 2017. 36

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. 28

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. 28

Gruber, Jonathan, "The consumption smoothing benefits of unemployment insurance," *The American Economic Review*, 1997, *87* (1), 192–205. 2, 4, 8, 11, 18, 19, 34

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

__ and J. Kolstad, "Health Insurance For Humans: Information Frictions, Plan Choice, and Consumer Welfare," *American Economic Review*, 2015, *105* (8), 2449–2500. 4, 5

Hendren, Nathaniel, "Knowledge of future job loss and implications for unemployment insurance," *American Economic Review*, 2017, *107* (7), 1778–1823. 4, 8, 19, 68

__, "Measuring ex-ante welfare in insurance markets," Working Paper, National Bureau of Economic Research 2018. 6, 19

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. 28, 29

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. 28

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239. 21

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, 15

Kolsrud, Jonas, Camille Landais, and Johannes Spinnewijn, "Studying consumption patterns using registry data: lessons from Swedish administrative data," *Working Paper*, 2017. 2, 5, 15, 17, 51

— , — , 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, 13, 15, 18, 28, 34, 76, 77

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. 4, 11, 29

— , Arash Nekoei, Peter Nilsson, David Seim, and Johannes Spinnewijn, "Risk-based selection in unemployment insurance: Evidence and implications," *Working Paper*, 2017. 4, 15, 16, 31, 32

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. 13

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. 32

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. 28

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, 34

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. 12, 32

_ , "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. 32

Figures

Figure 1: Estimated Consumption Dynamics around Start of Unemployment Spell

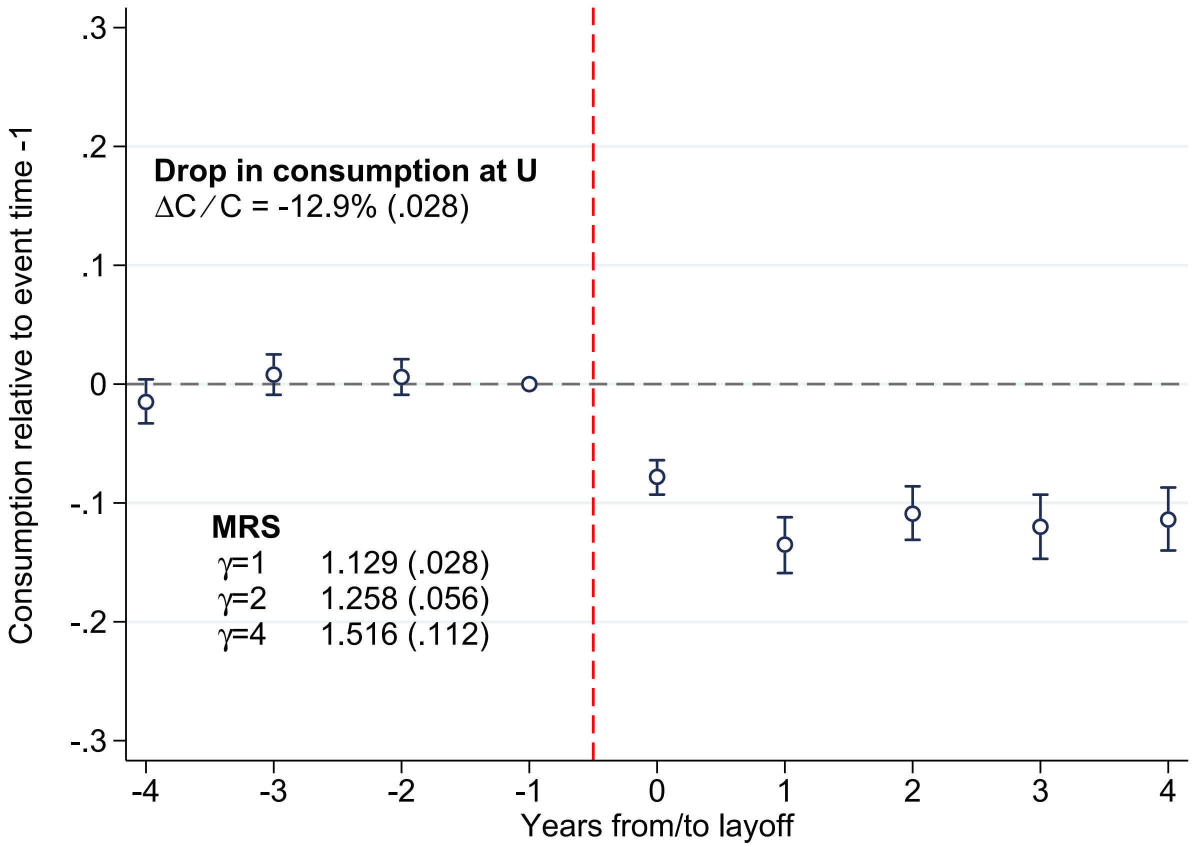

Notes: The figure reports event study estimates of household annual consumption around the time when a household member loses her job. Coefficients and confidence intervals come from specification (17) run on the sample of treated individuals and a control group of individuals obtained from nearest-neighbor 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 (18) 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: Decomposition of the Estimated Drop in Consumption at Unemployment

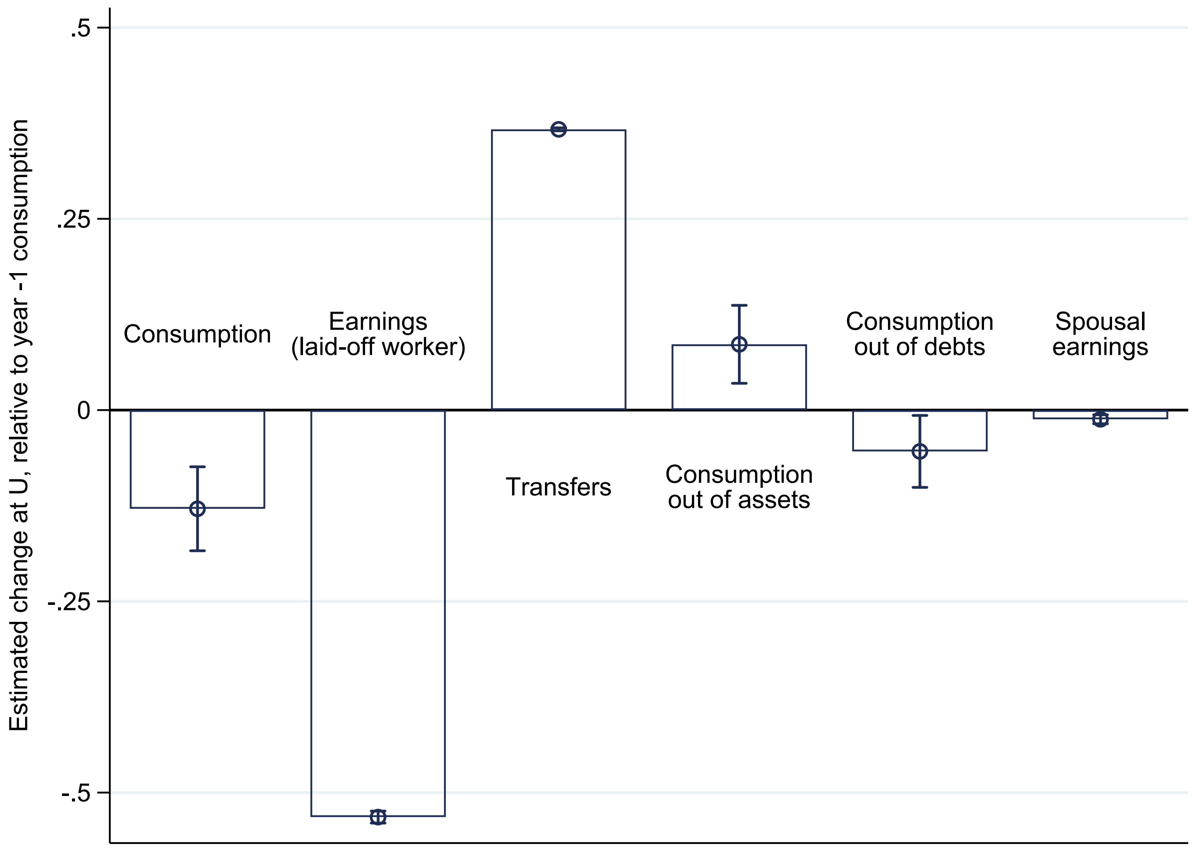

Notes: The graph decomposes the variation in consumption expenditures at unemployment, into the variations of five different components of total household expenditures: earnings of the individual subject to the unemployment shock ($E_{i,u,t}$), spousal earnings ($E_{i,-u,t}$), all transfers net of taxes 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 (18) 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 3: Heterogeneity in Estimated Drop in Consumption at Unemployment

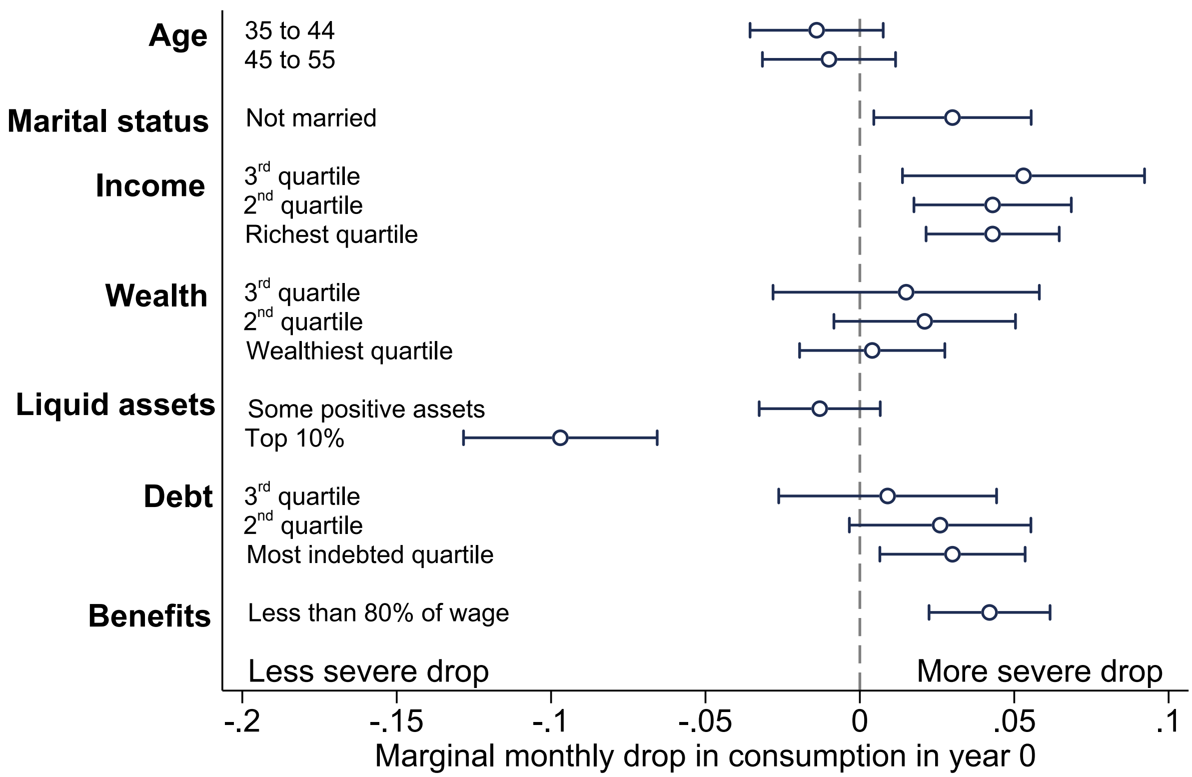

Notes: The graph analyzes heterogeneity in consumption drops at unemployment. The figure reports estimates of the effect of having characteristic $H = h$ on the drop in consumption at unemployment, following specification (19). Note that all dimensions of heterogeneity are entered simultaneously in the 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 4: Identifying Variation in Local Transfers

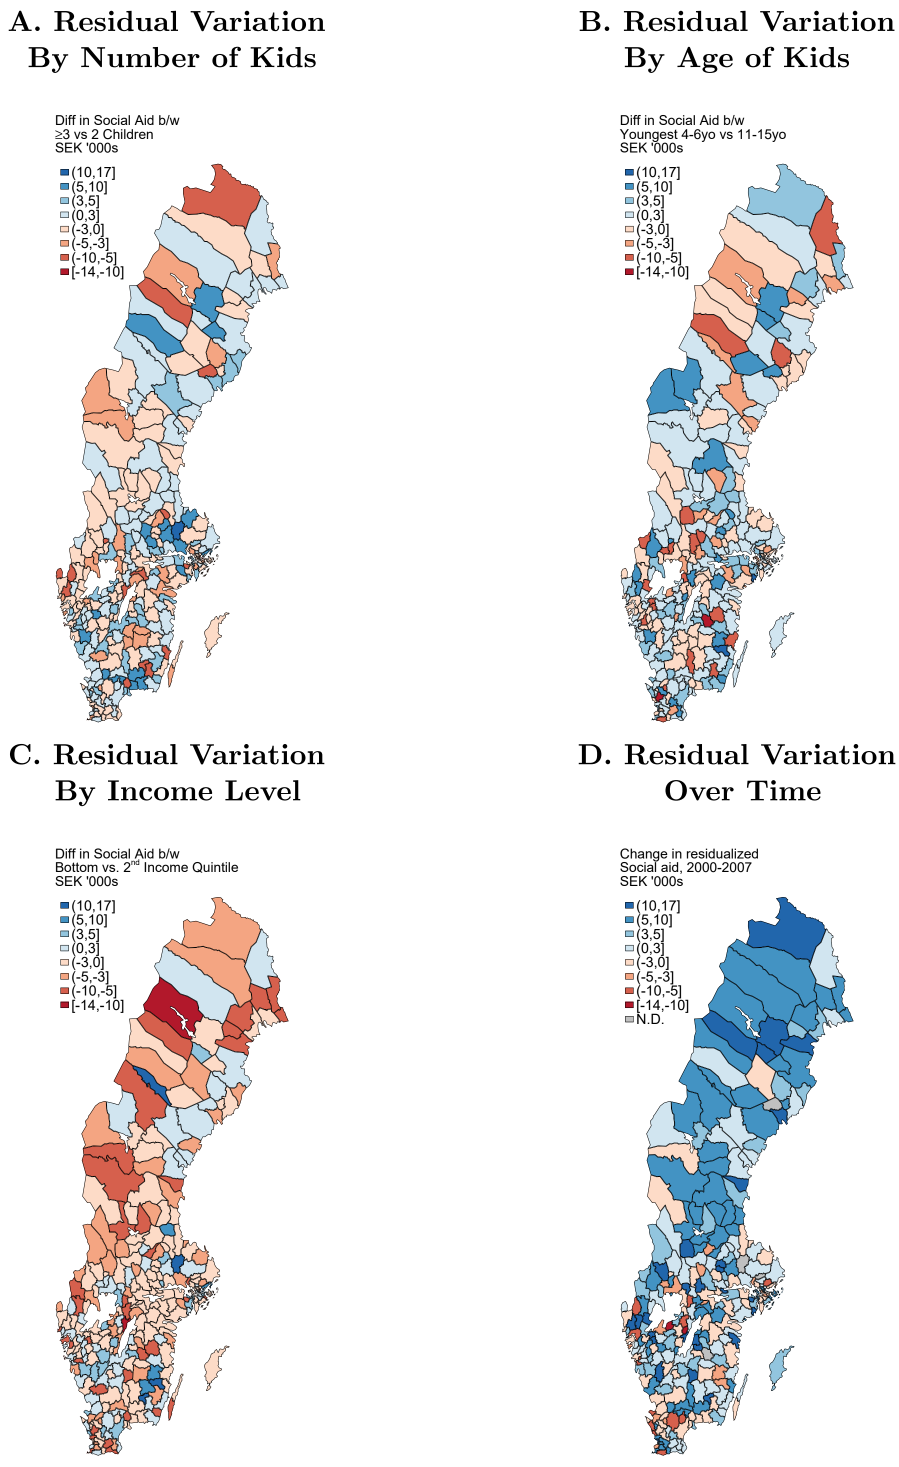

Notes: The Figure provides evidence of the variation in the way Swedish municipalities set local welfare transfers ("*social bidrag*"). By law, transfers are functions of characteristics $\mathbf{V}$, which include the number of dependents, the age of the dependent children, the liquid assets and income 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 residualize welfare transfers using specification:

$$B_{imt} = \sum_{j \neq k} \tau^j_m \cdot V^j_{it} + \bar{\tau}^k V^k_{it} + \tilde{B}_{imt}$$

We then plot $\mathbf{E}[\tilde{B}_{imt}|V^k = v^k] - \mathbf{E}[\tilde{B}_{imt}|V^k = v'^k]$ 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 $\tilde{B}_{imt}$ 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 $\tilde{B}_{imt}$ 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 5: Relationship Between First-Difference in Residual Local Transfers and First-Difference in Consumption by Employment Status

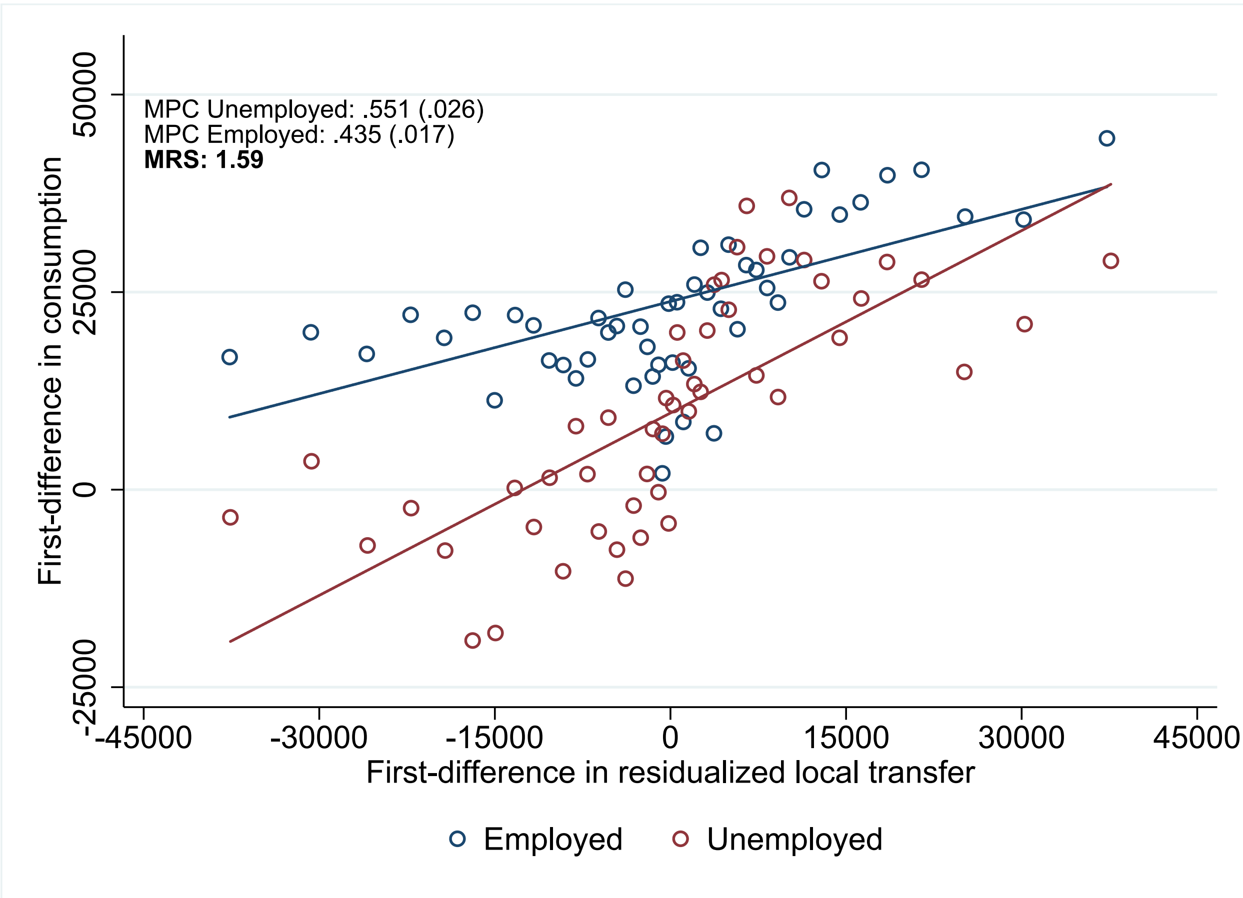

Notes: The graph is a bin-scatter plot of the relationship between the first-difference in residualized transfers $\tilde{B}_{imt}$ and the first-difference in annual household consumption, splitting the sample between households prior to the unemployment shock and households who experience unemployment in year t. The sample is the same as the one used for the CB approach in section 4 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 6: Non-parametric Relation between Expected Price and Insurance Coverage, and Estimated Demand Curve

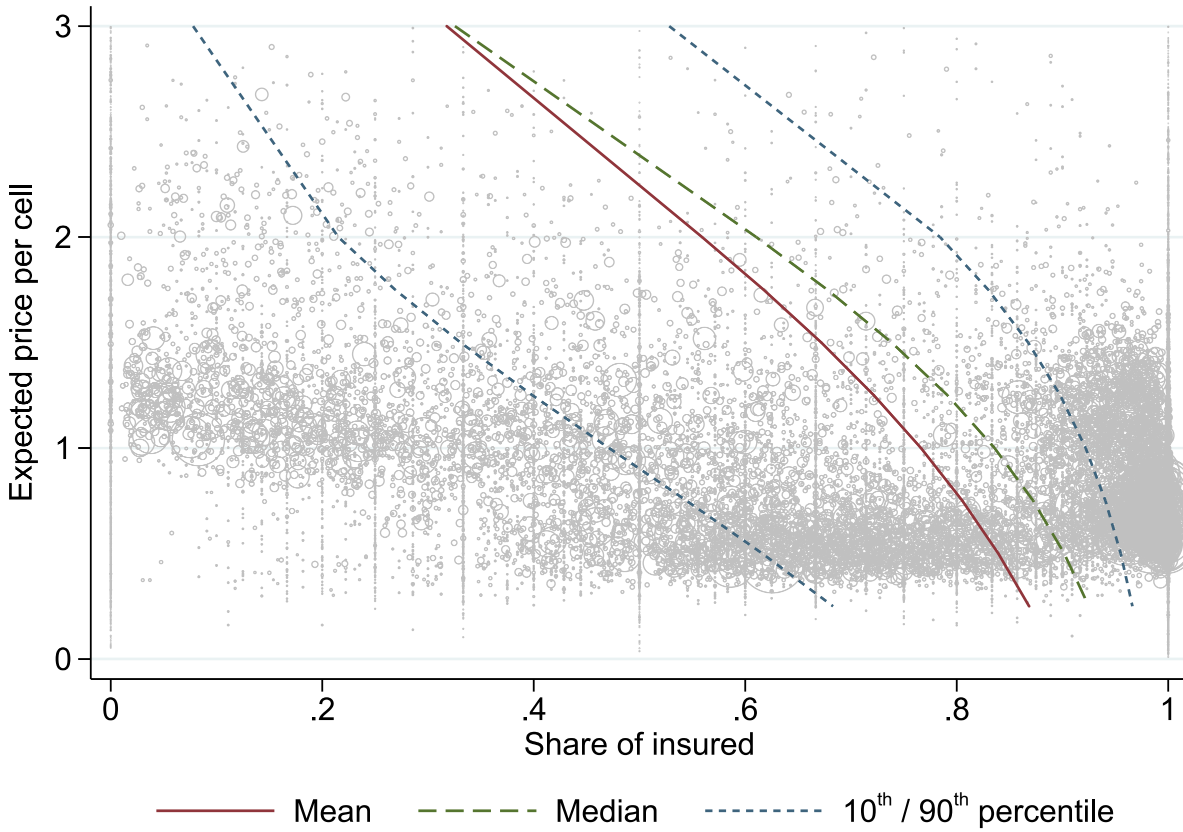

Notes: The scatter plot shows the average expected price and share buying comprehensive insurance coverage for workers grouped by cells based on a rich set of observables. The expected price is calculated given the predicted risk under comprehensive coverage. Appendix Figure D-1 shows the same plot using 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. The full line 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. The dashed lines show other quantiles of the estimated probability to buy comprehensive coverage.

Figure 7: Distributions of MRS from RP Structural Estimation

A. Risk Models

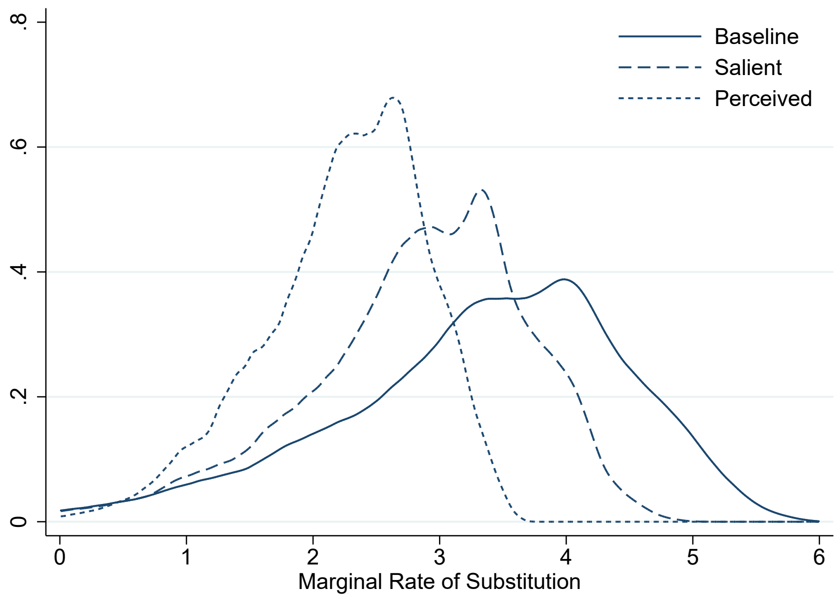

B. Lower and Upper Bound in Perceived Risk Model

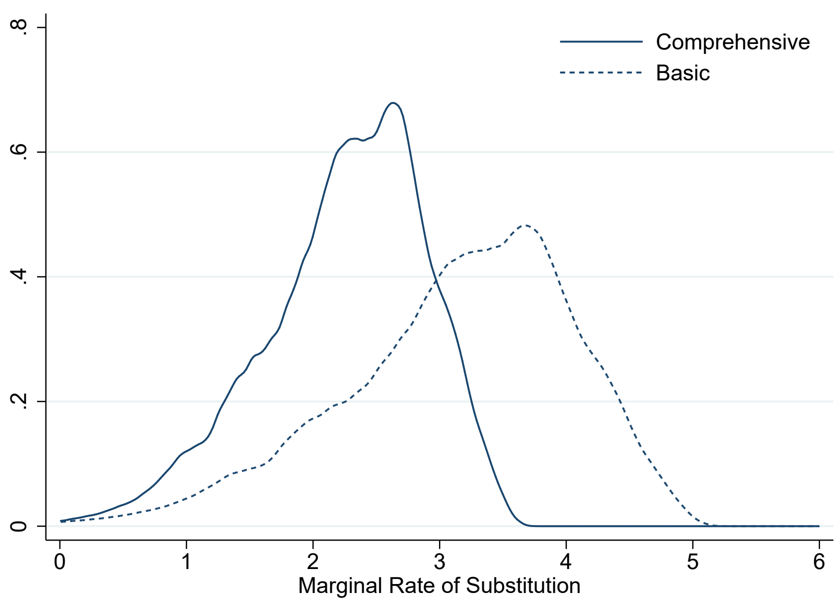

Notes: Panel A shows the distribution of MRS estimated under different risk models. The solid line is based on the baseline risk model (see column (1) of Table 3). The long dashed line is based on the salient risk shifters (see column (2) of Table 3). The short dashed line used the perceived risk model (see column (3) of Table 3). In all risk models used for Panel A, the risk is estimated for workers under comprehensive coverage. Panel B, however, 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). 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.

Figure 8: Comparison of MRS Estimates Across Different Approaches for the Baseline Sample

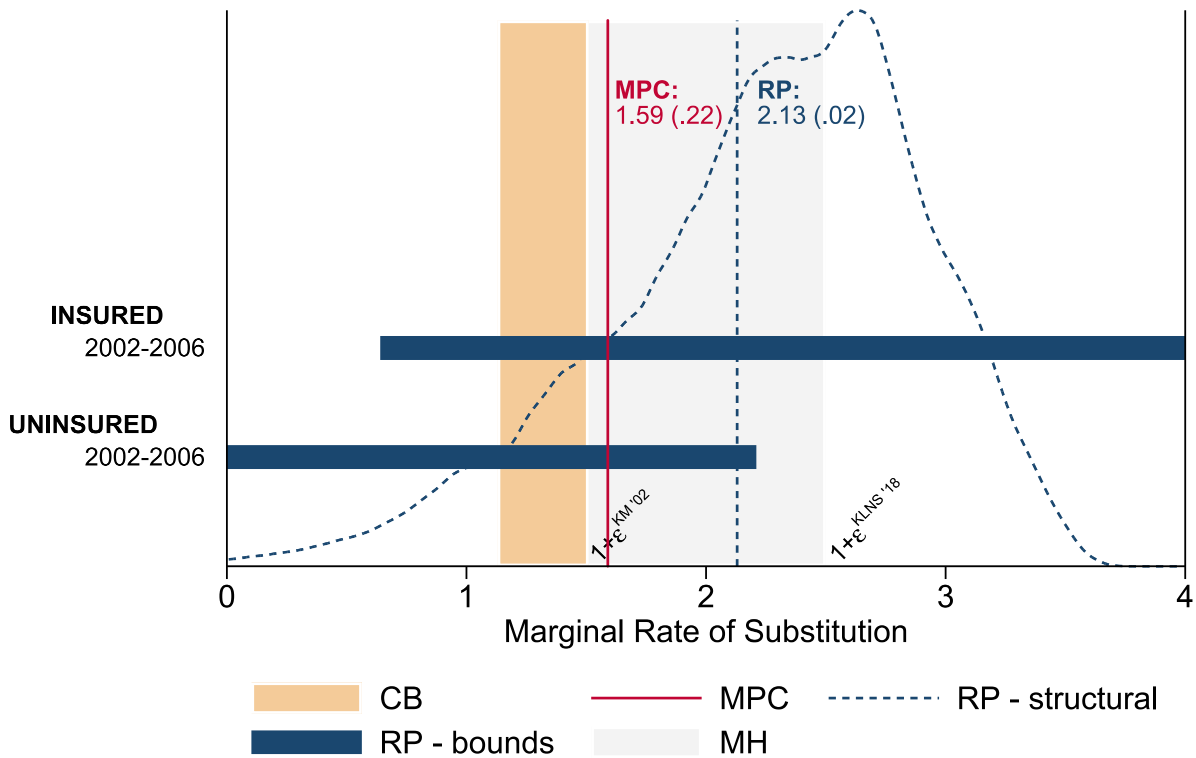

Notes: The graph compares the estimates of the MRS from the three different approaches, all implemented for the same baseline sample of workers. The region shaded in orange represents the range of MRS estimates from the CB approach, based on a consumption drop of 12.9% and relative risk aversion $\gamma \in [1; 4]$. 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 6.1, 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 7.

Figure 9: Heterogeneity in Estimated MRS using the RP Approach

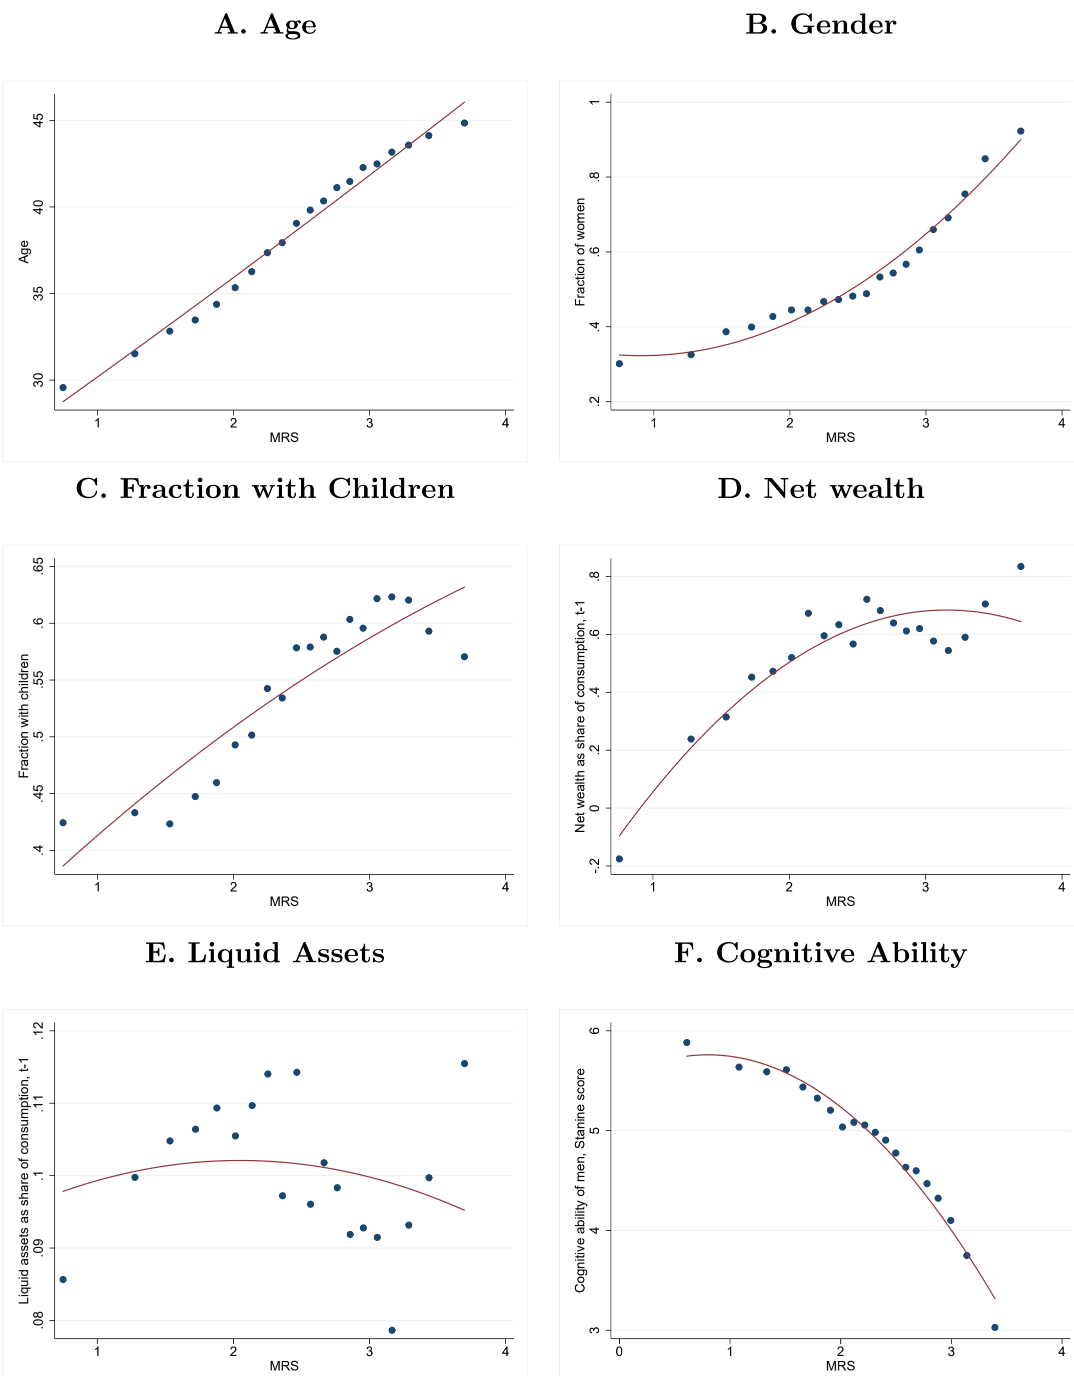

Notes: The graph correlates the estimated MRS from the RP approach with various observable characteristics, using bin scatter plots, by bins of estimated MRS. Panel A, B and C focus on age, gender and the presence of children, three types of observable characteristics that may correlate with risk preferences. 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.

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 & Implied Lower Bound on MRS

	First Difference in Household Consumption				
	(1)	(2)	(3)	(4)	(5)
MPC employment	0.435	0.439	0.414	0.413	0.431
	(0.017)	(0.017)	(0.017)	(0.018)	(0.017)
MPC unemployment	0.551	0.544	0.519	0.454	0.547
	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)
Lower Bound on MRS	1.591	1.528	1.530	1.185	1.596
SE (delta method)	(0.262)	(0.248)	(0.254)	(0.198)	(0.262)
SE (bootstrap)	(0.218)	(0.204)	(0.211)	(0.156)	(0.218)
Residualization:					
$\mathbf{V}_{it}$	×	×	×	×	×
Municipality FE	×	×	×	×	×
Income × Municipality FE		×		×	
Family type × Municipality FE			×	×	
# & age child. × Municipality FE			×	×	
Year × Municipality FE					×
Observations	89673	89673	89673	89673	89673
Adjusted R-squared	0.026	0.026	0.024	0.021	0.025

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 as the one used for the CB approach in section 4 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 (21) estimated in first-differences. We also report our estimate for the lower bound on the MRS, which, following formula (10), is equal to $\frac{\hat{\mu}_e+\hat{\mu}_u}{1-(\hat{\mu}_e+\hat{\mu}_u)} \Big/ \frac{\hat{\mu}_e}{1-\hat{\mu}_e}$, and its corresponding standard error, using two alternative approaches: the Delta-method, and a block-bootstrap computation. 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 (i.e. τ^{age}, $\tau^{children}$, τ^y, etc.) 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. In column (2), we add income and asset deciles interacted with municipality fixed effects, so as to exploit only the variation arising from how transfers differently account for the family structure of the household across municipalities. In column (3), we instead add family structure dummies interacted with municipality fixed effects: this specification only exploits variation in the phase-out rates τ_m^y and τ_m^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 only stems from changes in the average generosity of transfers within municipality over time due to local reforms. Finally, in column (5), we do the opposite exercise and add municipality fixed-effects interacted with year fixed effects in the residualization of transfers. By controlling for the year by year change in average generosity of transfers within municipality, this column only exploits within household variation stemming from the structure of τ_k.

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, $X_{it}\beta - \gamma\tilde{p}(Z_{it}) + \varepsilon_{it} \geq 0$. 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 $MRS(X_{it}) = \frac{X_{it}\beta}{\gamma}$. Column (1) to (3) show the estimates for different risk models. Column (1) uses the baseline risk model, as discussed in Section 3. 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 $\tilde{p}$. Columns (4) to (10) perform further sensitivity checks of the structural 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 9.

Appendix A Technical Appendix

This technical appendix provides the proofs of the Propositions 1-3 underlying the three implementations. We also set up the model extensions referred to in Section 2 and demonstrate the robustness of the MPC approach in more detail. We also further discuss the assumptions required for the implementation of the MPC and RP approach.

A.1 Proofs

Proof of Proposition 1: 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 $\theta = \frac{\partial v_u(c,x)/\partial c}{\partial v_e(c,x)/\partial c}$ (and thus $\theta = \frac{\partial^2 v_u(c,x)/\partial c^2}{\partial^2 v_e(c,x)/\partial c^2}$) and $\sigma_e^c = -\frac{\partial^2 v_e(c,x)/\partial c^2}{\partial v_e(c,x)/\partial c}$, we obtain

$$\frac{\frac{\partial v_u(c_u,x_u)}{\partial c}}{\frac{\partial v_e(c_e,x_e)}{\partial c}} \cong \theta + \theta \sigma_e^c [c_e - c_u],$$

which proves the Proposition.

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

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

This implies that the MRS equals

$$\frac{\frac{\partial v_u(c_u,x_u)}{\partial c}}{\frac{\partial v_e(c_e,x_e)}{\partial c}} = \frac{p_u}{p_e} \times \frac{\frac{\partial v_u(c_u,x_u)}{\partial x}}{\frac{\partial v_e(c_e,x_e)}{\partial x}}. \tag{24}$$

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,

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

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

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

and thus

$$\frac{dc_s}{dy_s}=\frac{p_s\sigma_s^x}{\sigma_s^c+p_s\sigma_s^x},$$

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

$$p_s=\frac{dc_s/dy_s}{1-dc_s/dy_s}\frac{\sigma_s^c}{\sigma_s^x}.$$

Plugging this into equation (24), we obtain

$$\frac{\frac{\partial v_u(c_u,x_u)}{\partial c}}{\frac{\partial v_e(c_e,x_e)}{\partial c}}=\frac{\frac{dc_u/dy_u}{1-dc_u/dy_u}}{\frac{dc_e/dy_e}{1-dc_e/dy_e}}\times\frac{\frac{\sigma_u^c}{\sigma_u^x}}{\frac{\sigma_e^c}{\sigma_e^x}}\times\frac{\frac{\partial v_u(c_u,x_u)}{\partial x}}{\frac{\partial v_e(c_e,x_e)}{\partial x}}.$$

Proof of Proposition 3: We consider an individual who is offered insurance at the margin at price p_u/p_e. The extra insurance allows her to marginally increase her consumption when unemployed at the expense of consumption when employed at rate $\frac{dc_e}{dc_u}=-\frac{p_u}{p_e}$. This corresponds to the combination of buying an additional unemployment-contingent security $d\tilde{x}_u$, while selling an additional employment-contingent security $d\tilde{x}_e(=-d\tilde{x}_u)$. These securities are priced at p_u and p_e respectively and could be offered in addition to the use of other means to smooth consumption, which we continue to denote by x_s.

The marginal impact of the 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. For an individual who buys that extra unit of coverage, we need

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

The opposite is true for an individual not buying that extra unit over coverage. This implies that the expected price puts a lower bound on the MRS in the former case and an upper bound in the latter case.

A.2 Model Extensions

Intertemporal Model We set up a dynamic model with a state-dependent (exogenous) probability of becoming/remaining unemployed, $\tilde{\pi}_s$. The individual's problem is summarized by,

$$V_{s,t}(A_t) = \max_{A_{t+1},c_{s,t},z_{s,t}} \left\{ v_s(c_{s,t}) + \beta \left[\tilde{\pi}_s V_{u,t+1}(A_{s,t+1}) + (1-\tilde{\pi}_s) V_{e,t+1}(A_{s,t+1}) \right] \right\},$$

$$\text{s.t. } c_{s,t} = A_t + y_{s,t} - \frac{A_{s,t+1}}{R_{s,t}} \text{ and } A_{s,t+1} \geq \bar{A},$$

for $s \in \{e,u\}$ and t. We can think of the assets (or debt) left for the next period $A_{s,t+1}$ as the endogenous resource to change consumption relative to cash-on-hand, $A_t + y_{s,t}$. We allow the interest rate to be different when employed and unemployed. As an agent may run out of assets when unemployed, and the interest rate is higher on loans than on savings, this would give rise to $R_{u,t} \geq R_{e,t}$. We endogenize this further below in subsection A.3.

At an interior optimum, the first-order condition with respect to consumption equals,

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

Using $V'_{s,t}(A_t) = \partial v_s(c_{s,t})/\partial c$, this 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

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

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,

$$\frac{\frac{dc_{s,t}}{dy_{s,t}}}{1 - \frac{dc_{s,t}}{dy_{s,t}}} = R_{s,t} \left(\frac{E_s \left[V_{\tilde{s},t+1}''(A_{s,t+1})\right]}{E_s \left[V_{\tilde{s},t+1}'(A_{s,t+1})\right]} \bigg/ \frac{v_s''(c_{s,t})}{v_s'(c_{s,t})} \right).$$

Insurance Model We allow for an insurance choice in our stylized model by adding a pre-period in which individuals can allocate an endowment to future employment and unemployment states using Arrow-Debreu securities. Individuals gain utility from the residual endowment in the pre-period. An individual faces the following problem,

$$\begin{aligned} &\max_{x_e, x_u, z} v_0(c_0) + \pi(z) v_u(c_u) + [1 - \pi(z)]\, v_e(c_e) - z, \\ &\text{s.t. } c_0 = A_0 - x_e - x_u \geq 0, \\ &\qquad c_s = y_s + \frac{1}{p_s} x_s \geq 0 \text{ for } s \in \{e, u\}. \end{aligned}$$

Note that in this setting the utility cost of the state-contingent resource is borne even if the state does not realize. At an interior optimum, the following condition holds for each state s,

$$-p_s \frac{\partial v_0(c_0)}{\partial c} + \pi_s(z) \frac{\partial v_s(c_s, x_s)}{\partial c} = 0,$$

where $\pi_u(z) = \pi(z)$ and $\pi_e(z) = 1 - \pi(z)$. The first-order condition does directly depend on the probability that the state realizes, but this probability can be reflected in the price as well. For actuarially fair insurance, we have $p_s = \pi_s(z)$. Using the optimality condition for the employment and unemployment-contingent securities, we get the following expression for the MRS

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

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

$$p_s v_0''(c_0) \frac{dx_s}{dy_s} = \pi_s v_s''(c_s) \left(1 + \frac{1}{p_s} \frac{dx_s}{dy_s} \right)$$

and using $\frac{dc_s}{dy_s} = 1 + \frac{1}{p_s}\frac{dx_s}{dy_s}$, this can be rewritten as

$$\frac{\frac{dc_s}{dy_s}}{1 - \frac{dc_s}{dy_s}} = p_s \left(\frac{v_0''(c_0)}{v_0'(c_0)} \bigg/ \frac{v_s''(c_s)}{v_s'(c_s)} \right).$$

A.3 MPC Approach: Robustness

We consider extensions of our stylized model and show how our MPC approach remains robust in the sense that the ratio of MPCs identifies the ratio of the state-specific price of increasing consumption, which then determines the MRS. We consider extensions to multiple goods and resources and address all the standard challenges to the CB approach (e.g., home production, work-related expenditures, durable goods, committed expenditures). We also consider a model where the price of consumption is endogenous, depending on the level of resources used.

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 $\Delta [c_s - \phi_s]$ 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,

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

Another reason for the 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,

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

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

$$u'_s(\phi_s c_s) = \frac{p_s}{\phi_s} v'_s(x_s),$$

from which we can derive,

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

The MPC depends on the state-specific consumption scalar ϕ. 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 $\Delta [\phi_s c_s]$ that is relevant for the MRS. The ratio of the MPC odds ratios, however, still identifies the relative prices and bounds the MRS as in our stylized model,

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

Multiple consumption goods To study the robustness of our approach when allowing for multiple consumption goods or 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

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

This assumes that we observe total consumption. If we only observe a partial measure of consumption expenditures, we would have

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

As a consequence, the ratio of partial MPC odds ratios would underestimate the price ratio p_u/p_e (unless $q_s = 1/p_s$), and thus provides a weaker lower bound on the MRS. This indicates that a more comprehensive measure of expenditures is preferable.

Multiple Resources In analogy to considering multiple consumption goods, we can introduce a second resource into our setup:

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

If both resources are used at the margin, optimality is characterized by,

$$u_s'(c_s) = q_s w_s'(z_s) = p_s v_s'(x_s).$$

From this and the budget constraint we derive,

$$\begin{aligned}\frac{dc_s}{dy_s} &= \frac{1}{1 - \frac{1}{p_s^2}\frac{u_s''(c_s)}{v_s''(x_s)} - \frac{1}{q_s^2}\frac{u_s''(c_s)}{w_s''(z_s)}}, \\ &= \frac{1}{1 + \frac{1}{p_s}\left[\frac{\sigma_s^c}{\sigma_s^x} + \frac{\sigma_s^c}{\sigma_s^z} MRS_s^{z,x}\right]},\end{aligned}$$

where $MRS_s^{z,x} = \frac{w_s'(z_s)}{v_s'(x_s)} (= \frac{p_s}{q_s})$. The MPC is thus smaller if the worker uses more resources at the margin. Translating this into the MRS, we get

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

Hence, to the extent that the marginal rate of substitution between these resources is similar (i.e., $\frac{p_u}{q_u} = \frac{p_s}{q_s}$), the ratio of MPC odds ratios still identifies a lower-bound on the MRS. Moreover, if there is a common resource for which prices are not-state specific (e.g., $q_u = q_e$), the ratio would understate the difference in prices $\frac{p_u}{p_e}$, but since workers no longer solely rely on the expensive resource, it may no longer be true that $\frac{\partial v_u(x_u)}{\partial x_u} > \frac{\partial v_e(x_e)}{\partial x_e}$.

Endogenous Prices We now consider a model where the price of using extra resources is endogenous to the level of the resource used, $p(x)$. This is a more parsimonious way to think about the availability of different means to smooth consumption. 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 $\varepsilon_s^{p,x} = \frac{p'(x_s)}{p(x_s)} x_s$. Translating this into the MRS, we get

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

Hence, as long $\varepsilon_e^{p,x}/x_e > \varepsilon_u^{p,x}/x_u$, the ratio of MPC odds ratios understates the price difference and thus provides a lower bound on the MRS. This clearly holds when the price elasticity is constant. Note also that preferences can now be quasi-linear in the resources used for consumption.

A.4 MPC Approach: Implementation

The MPC approach provides a lower bound in the stylized model under two assumptions. The first assumption is that the relative curvature does not change across states, or, if it does, that $(\sigma_u^c/\sigma_u^x)/(\sigma_e^c/\sigma_e^u) \geq 1$. The second assumption is that the marginal resource cost is higher when unemployed than when employed, $(\partial v_u/\partial x)/(\partial v_e/\partial x) \geq 1$. We argue that these assumptions hold for standard preferences and the common means used to smooth consumption.

CARA/CRRA We illustrate how the MRS relates to the ratio of MPCs for CARA and CRRA preferences, both over consumption and the resources used.

For CARA preferences, $U(c, x) = \frac{e^{-\sigma^c c}}{-\sigma^c} - \frac{e^{\sigma^x x}}{\sigma^x}$, the price ratio is identified exactly by the ratio of MPC odds ratios for $\frac{\sigma_u^c}{\sigma_u^x} = \frac{\sigma_e^c}{\sigma_e^x}$. Hence, we have

$$MRS = \frac{O_u^{mpc}}{O_e^{mpc}} \times \frac{\frac{\sigma_u^c}{\sigma_u^x}}{\frac{\sigma_e^c}{\sigma_e^x}} \times e^{\sigma_u^x x_u - \sigma_e^x x_e}.$$

The price ratio provides a lower bound if $\frac{\sigma_u^c}{\sigma_u^x} \geq \frac{\sigma_e^c}{\sigma_e^x}$ and $\sigma_u^x x_u \geq \sigma_e^x x_e$.

For CRRA preferences, $U(c,x) = \frac{c^{1-\gamma^c}}{1-\gamma^c} - \frac{x^{1+\gamma^x}}{1+\gamma^x}$, we have

$$MRS = \frac{O_u^{mpc}}{O_e^{mpc}} \times \frac{\frac{\gamma_u^c}{\gamma_u^x}\frac{x_u}{c_u}}{\frac{\gamma_e^c}{\sigma_e^x}\frac{x_e}{c_e}} \times \frac{(x_u)^{\gamma_u^x}}{(x_e)^{\gamma_e^x}}.$$

When $x_u > x_e$ and $c_u > c_e$, the ratio of MPC odds ratios understates the price ratio for $\frac{\gamma_u^c}{\gamma_u^x} = \frac{\gamma_e^c}{\gamma_e^x}$. Hence, it then provides a lower bound on the MRS (at least when $(x_u)^{\gamma_u^x} > (x_e)^{\gamma_e^x}$). So in the absence of state-dependent curvature, for both CARA and CRRA preferences, the ratio provides a lower bound on the MRS.

Marginal resource cost When unemployed, a worker has lower income y and, ceteris paribus, will use more of her resources x to smooth her consumption. The marginal utility cost from further using those resources will be higher if the disutility of using more resources is convex. That is, $(\partial v_u/\partial x)/(\partial v_e/\partial x) \geq 1$. This naturally holds in the model with endogenous prices and quasi-linear preferences over the resources used. This also holds in the intertemporal consumption model. An unemployed agent will increase consumption more at the expense of future assets than an employed agent. That is, $A_{u,t+1} < A_{e,t+1}$. For risk-averse preferences, the continuation utility is concave, so we have $V'_{s,t+1}(A_{u,t+1}) > V'_{s,t+1}(A_{e,t+1})$ and $V'_{u,t+1}(A_{s,t+1}) > V'_{e,t+1}(A_{s,t+1})$. With the risk of remaining unemployed higher than the risk of becoming unemployed, we therefore have $E_u(V'_{s,t+1}(A_{u,t+1})) > E_e(V'_{s,t+1}(A_{e,t+1}))$. The result also holds in a model with spousal labor supply, where a worker's spouse increases his or her labor supply when the worker becomes unemployed. Given convex cost of labor, the spouse faces a higher marginal cost of labor when a worker becomes unemployed, assuming separability in a worker's employment status and the cost of spousal labor. For the insurance model, the condition continues to hold, but the lower bound can well be uninformative. The argument is different as we start from $MRS = \frac{p_u}{p_e} \times \frac{1-\pi(z)}{\pi(z)}$. When the ex ante unemployment risk $\pi(z)$ is smaller than one half, we have $MRS > \frac{p_u}{p_e}$. The bound will become more uninformative as prices become more actuarially fair and the unemployment risk becomes smaller. In fact, with actuarially fair prices and $\pi(z) < 1/2$, the ratio of MPCs should be smaller than 1.

State-dependent Preferences In general, the lower-bound argument assumes a preference for smoothing consumption relative to the income loss when unemployed. In principle, with state-dependent 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. Hence, the assumption that the ratio of MPC odds ratios provides a lower bound seems weak. This doesn't mean that the lower bound will be informative. It will be less informative as workers have stronger preferences to smooth consumption, for given prices of smoothing consumption, implying that the difference in the utility of cost of using extra resources goes up at the optimum.

Income Variation As discussed, an important challenge for the implementation of the MPC approach is to find exogenous variation in state-contingent income that is otherwise comparable across states. We note that the implementation does not require unanticipated variation in income. In fact, the extensions with multiple resources and endogenous prices indicate that if ex ante insurance is used at the margin to smooth consumption, the income variation used for implementation should be anticipated in order to allow this margin to be adjusted as well. We also note that our analysis suggests a trade-off between using variation in state-contingent income (e.g., wages vs. UI benefits) vs. shocks in income in different states (e.g., other social transfers). The issue with using exogenous income shocks is that workers can adjust future consumption, in addition to other resources. The extensions with multiple resources or consumption goods suggest that we would then under-estimate the difference in prices. The challenge with variation in state-contingent income, however, is that the MPC may be different because of mental accounting or other behavioral frictions that are unrelated to the price of consumption smoothing.

A.5 RP Approach: Implementation

We show that the RP approach generalizes to a model with a discrete insurance choice. Using the notation of the insurance model, consider the 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

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

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

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

and thus

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

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.

Appendix B Additional Figures: Consumption-Based Approach

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

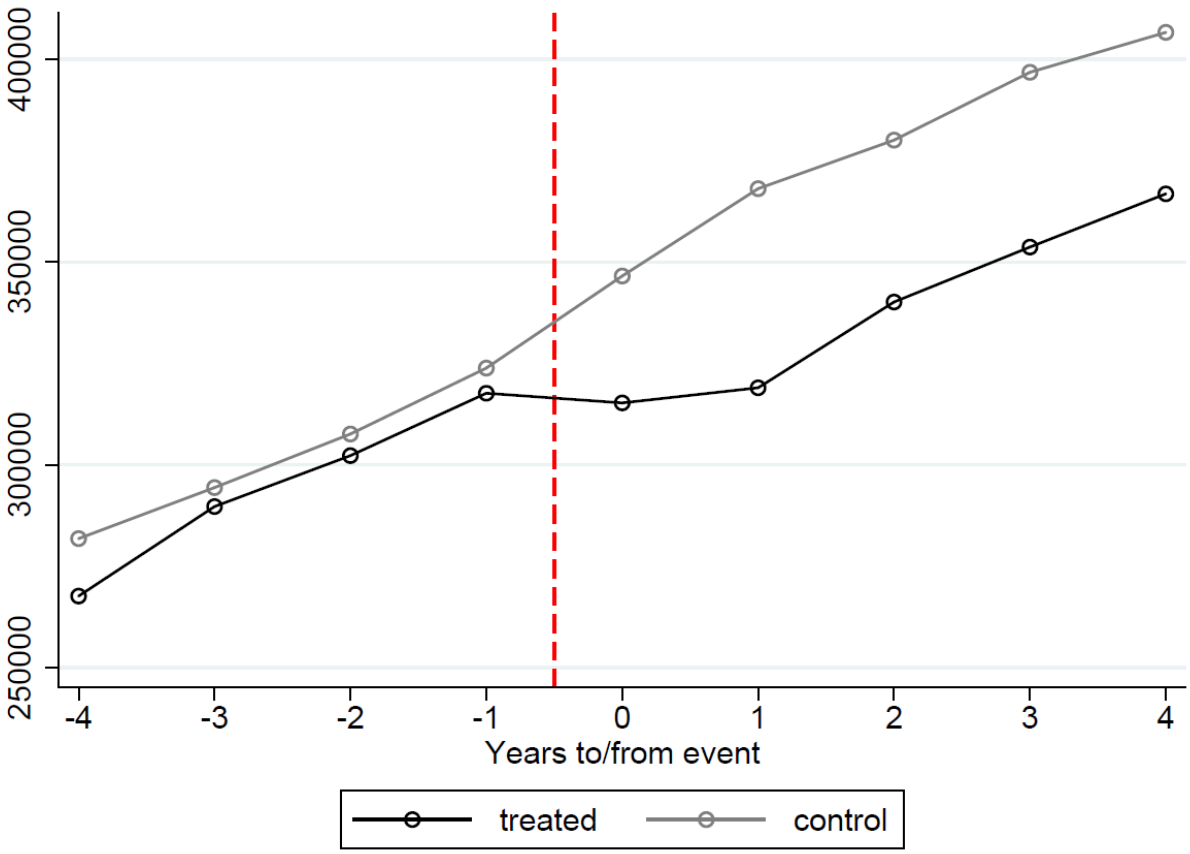

Notes: The figure reports the evolution of average household annual consumption in constant SEK2003 around the time when a household member loses her job. Event time is the time in years relative to the occurrence of the first job loss. The treatment group is composed of all the households from our baseline sample described in section 3.2 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

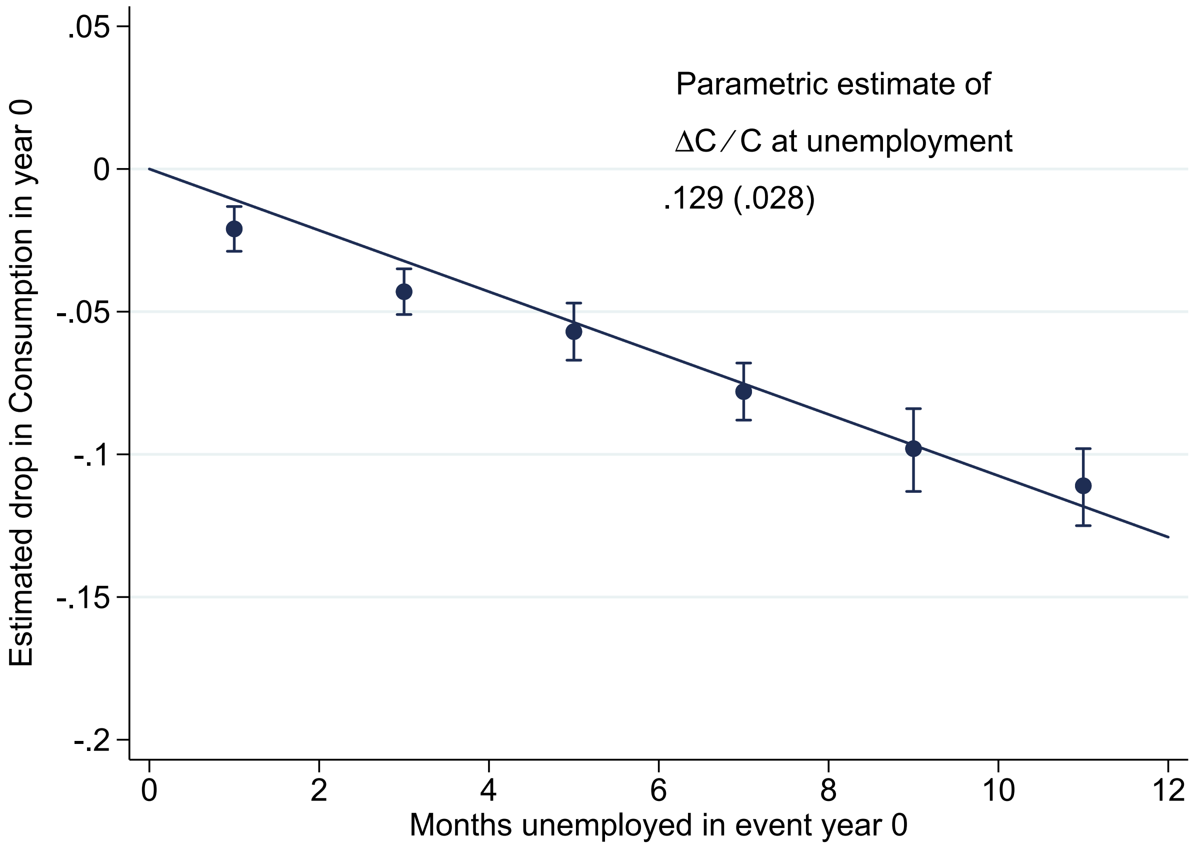

Notes: The graph displays the relationship between the drop in annual household consumption in event year 0 and the number of months spent unemployed in event year 0. In our sample, in event year 0, individuals are all observed unemployed in December, but differ in the time in months M_i they have spent unemployed in that year. We split the sample in 6 bins of M_i, and estimate specification (17) 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 (18) to identify the *flow* drop in consumption when unemployed $c_u - c_e$ from our annual consumption measure. See text for details.

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

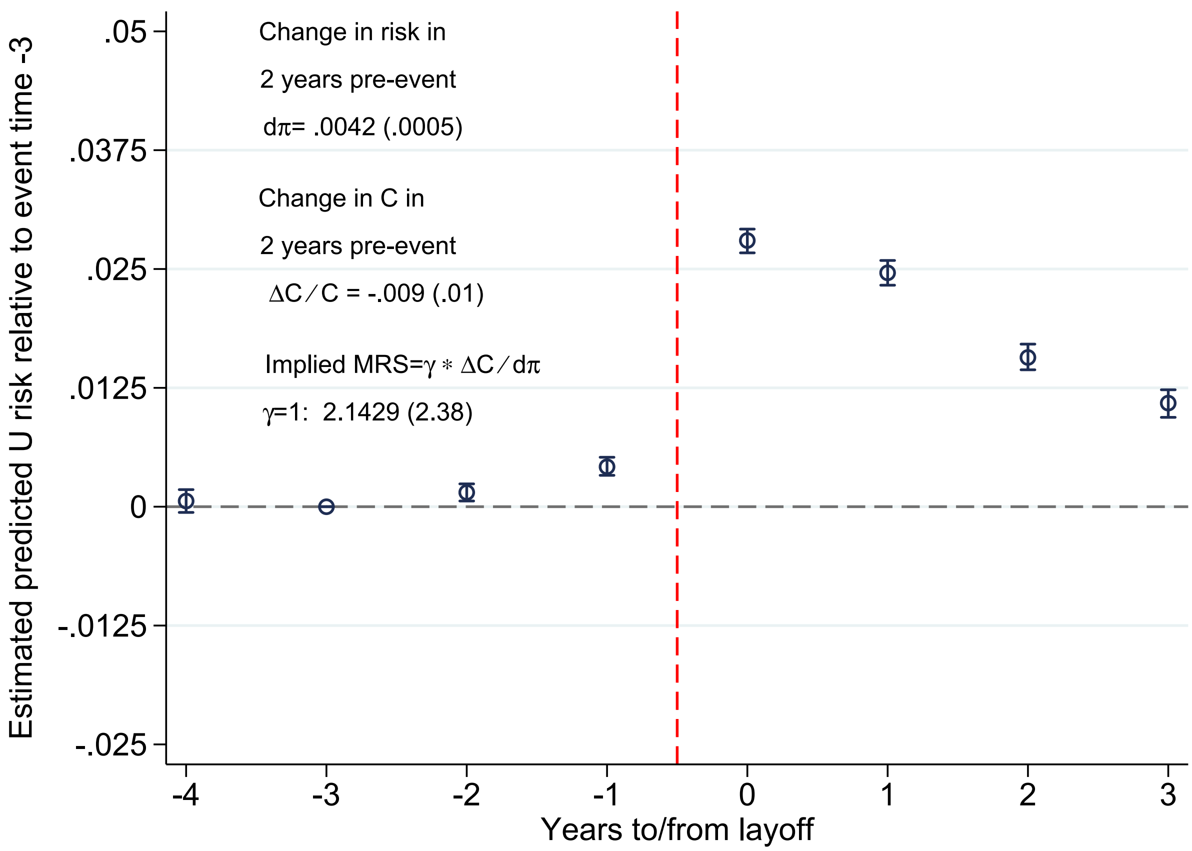

Notes: The graph studies how much individual unemployment risk gets revealed through changes in observables in the years around job loss. It reports estimates from specification (17) where we use as the outcome our measure of predicted unemployment risk, based on our rich model of observable determinants of unemployment risk in Sweden. See section 3.2 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.

Appendix C Additional Figures: MPC Approach

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

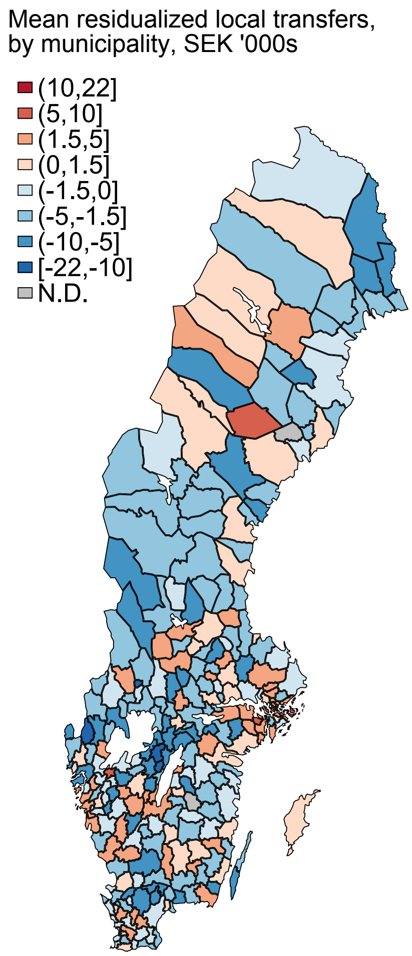

Notes: The Figure provides evidence of the variation in the way Swedish municipalities set local welfare transfers ("*social bidrag*"). By law, transfers are functions of characteristics $\mathbf{V}$, which include the number of dependents, the age of the dependent children, the liquid assets and income 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

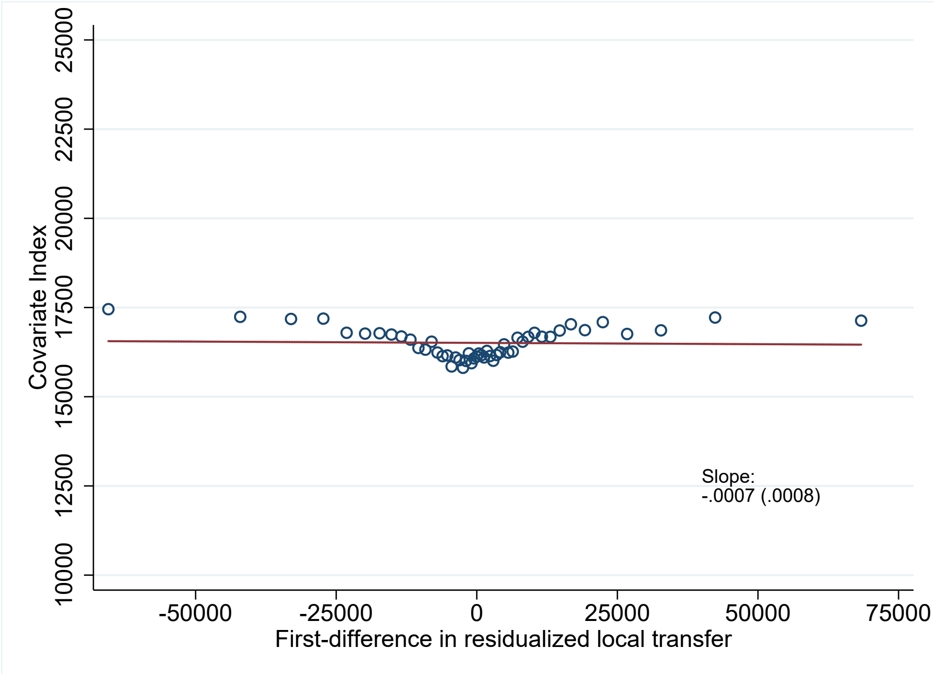

B. By Employment Status

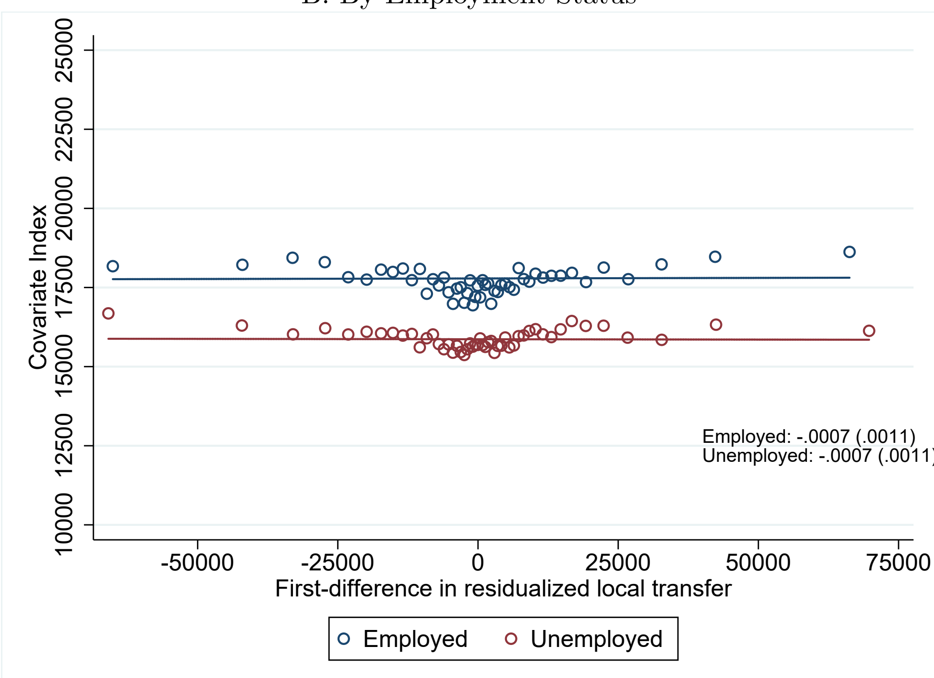

Notes: The Figure probes into the validity of our identifying assumption that the residual variation in welfare benefits $\tilde{B}_{imt}$ is orthogonal to the dynamics of household consumption. We use observables characteristics available in the registry data, that correlate with consumption, and that 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

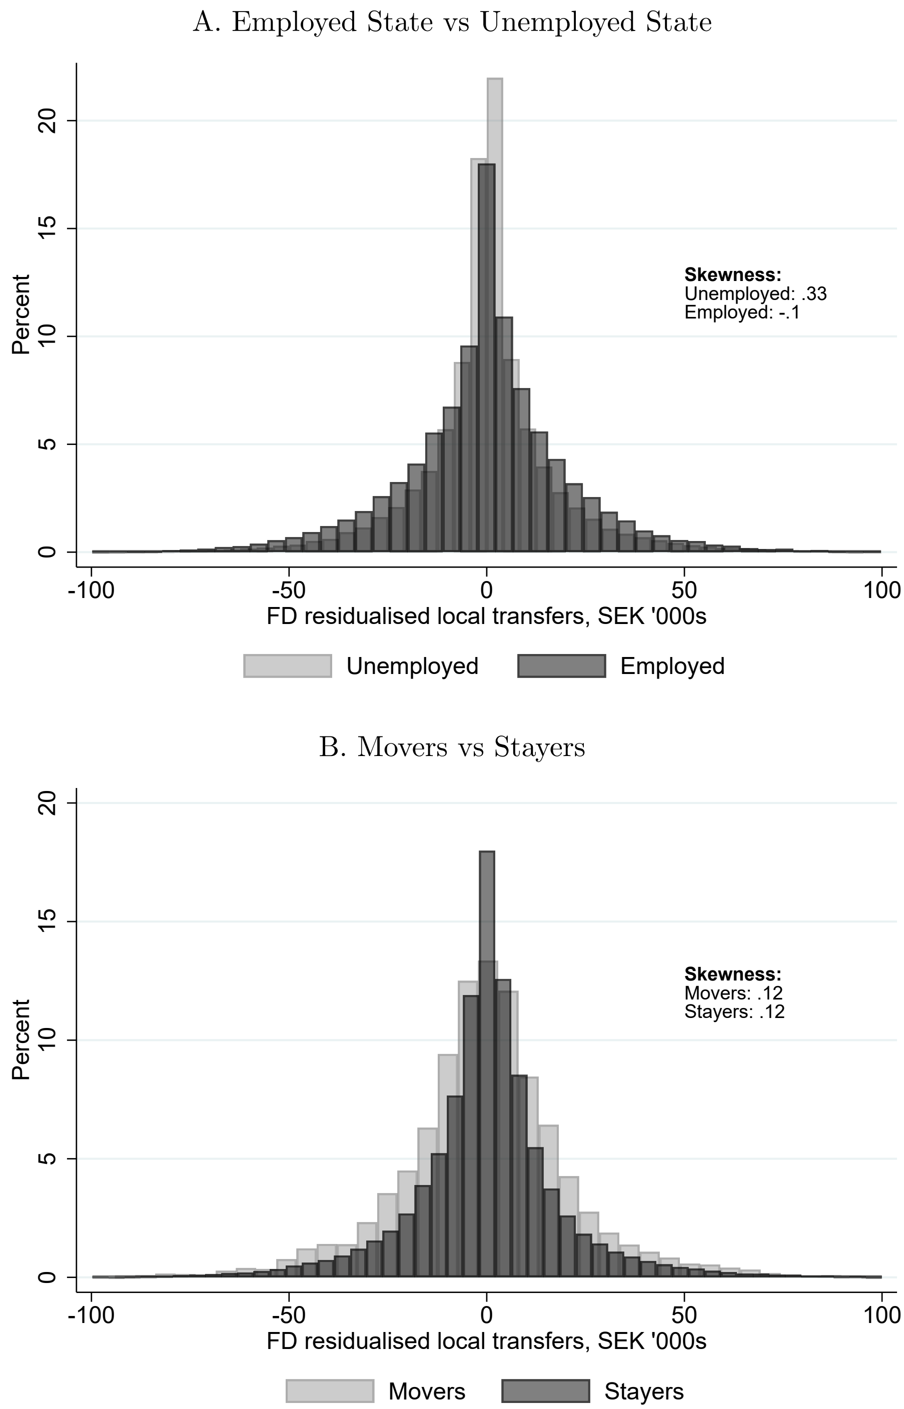

Notes: Panel A explores whether $\tilde{B}_{imt}$ is correlated with employment status. We plot the distribution of our baseline residual variation $\tilde{B}_{imt}$ by employment status. The figure shows that the distribution of our identifying variation in welfare transfers is very similar across employment 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.

Appendix D Additional Figures and Tables: Revealed-Preference Approach Approach

Figure D-1: Non-parametric Relation between Expected Price (under Basic Coverage) and Insurance Coverage

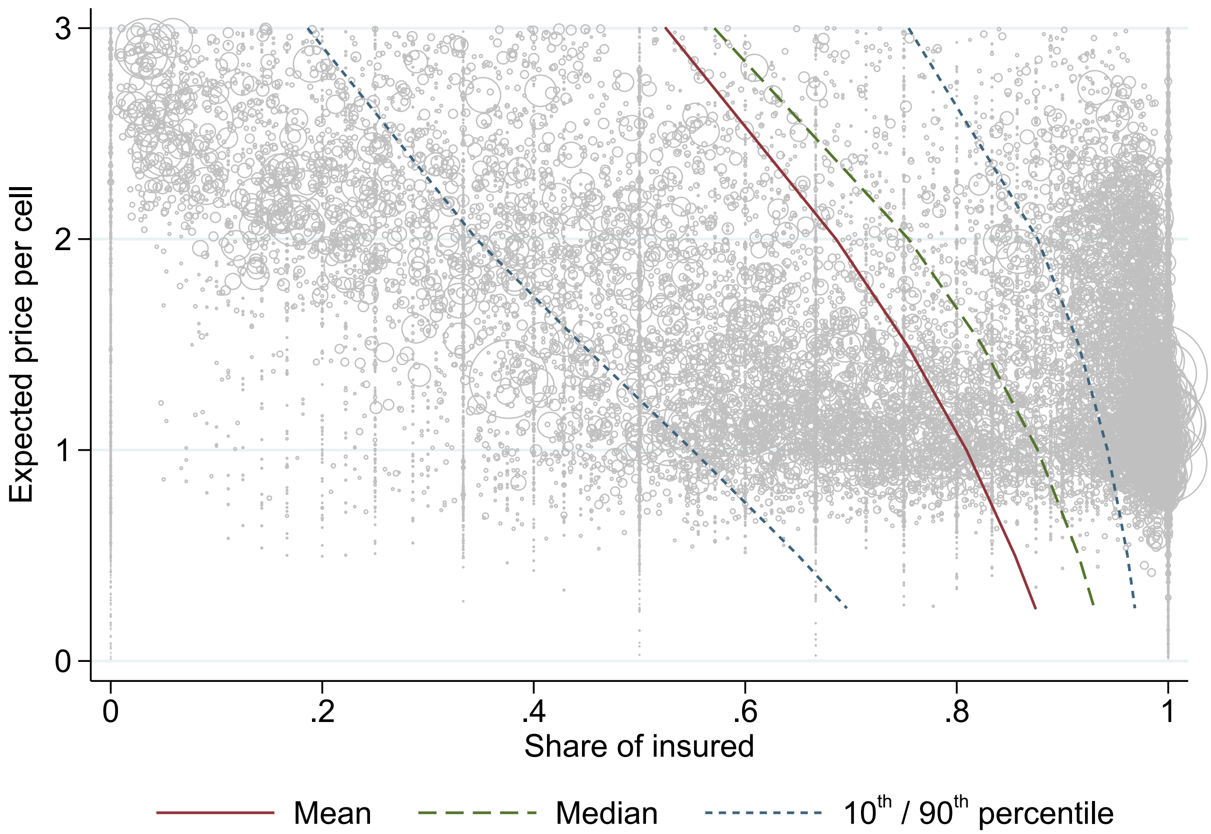

Notes: This Figure complements Figure 6 showing the average expected price and share buying comprehensive insurance coverage for workers grouped by cells based on a rich set of observables, but calculating the expected price using the predicted risk under basic coverage rather than comprehensive coverage. 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. The lines superimpose the estimated demand from the parametric MRS estimation.

Figure D-2: Realized vs. Perceived Job Loss Probabilities in HUS Survey

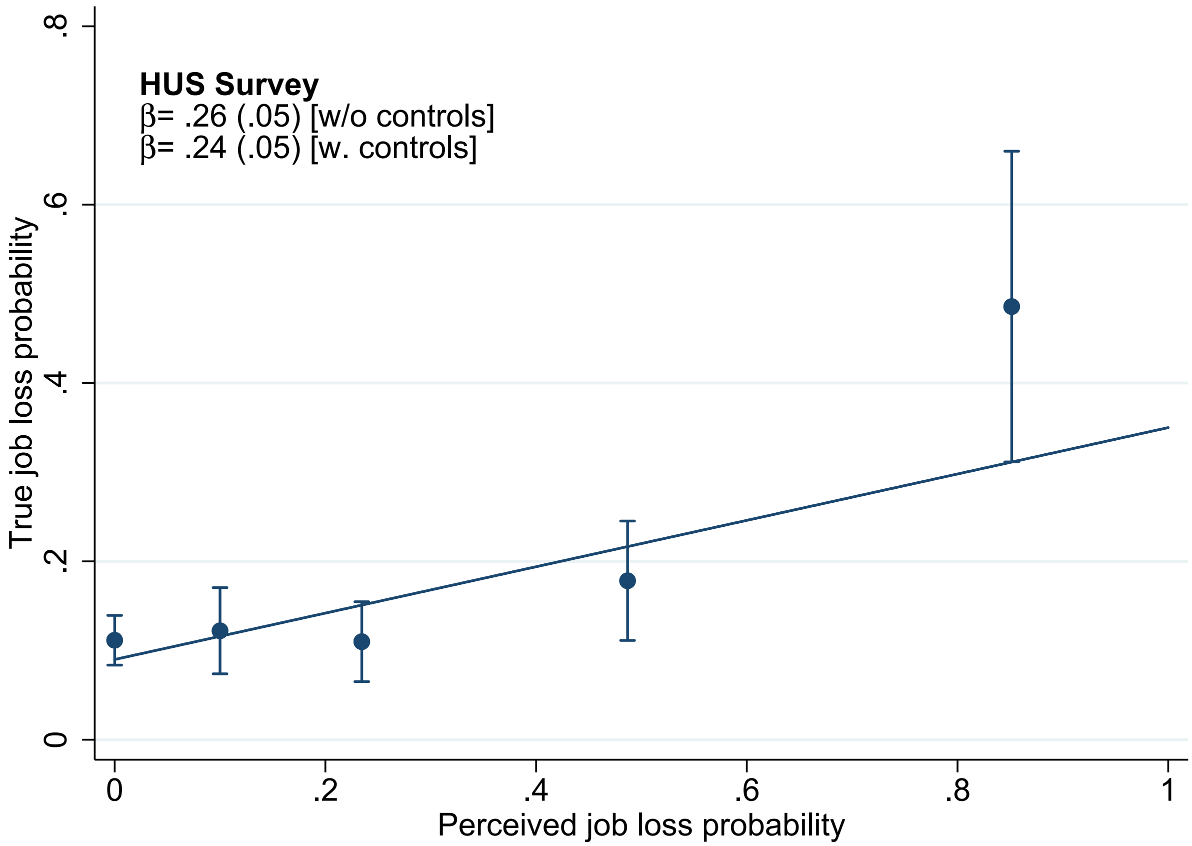

Notes: The figure compares the true and perceived probability to keep one's job, interpreting the complement of these probabilities as job loss probabilities. We use the responses to the question, "How likely is it that you will keep your current job next year?" in the HUS survey. For different bins 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. 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.

Figure D-3: MRS vs. Consumption Drop Estimates in Baseline Sample

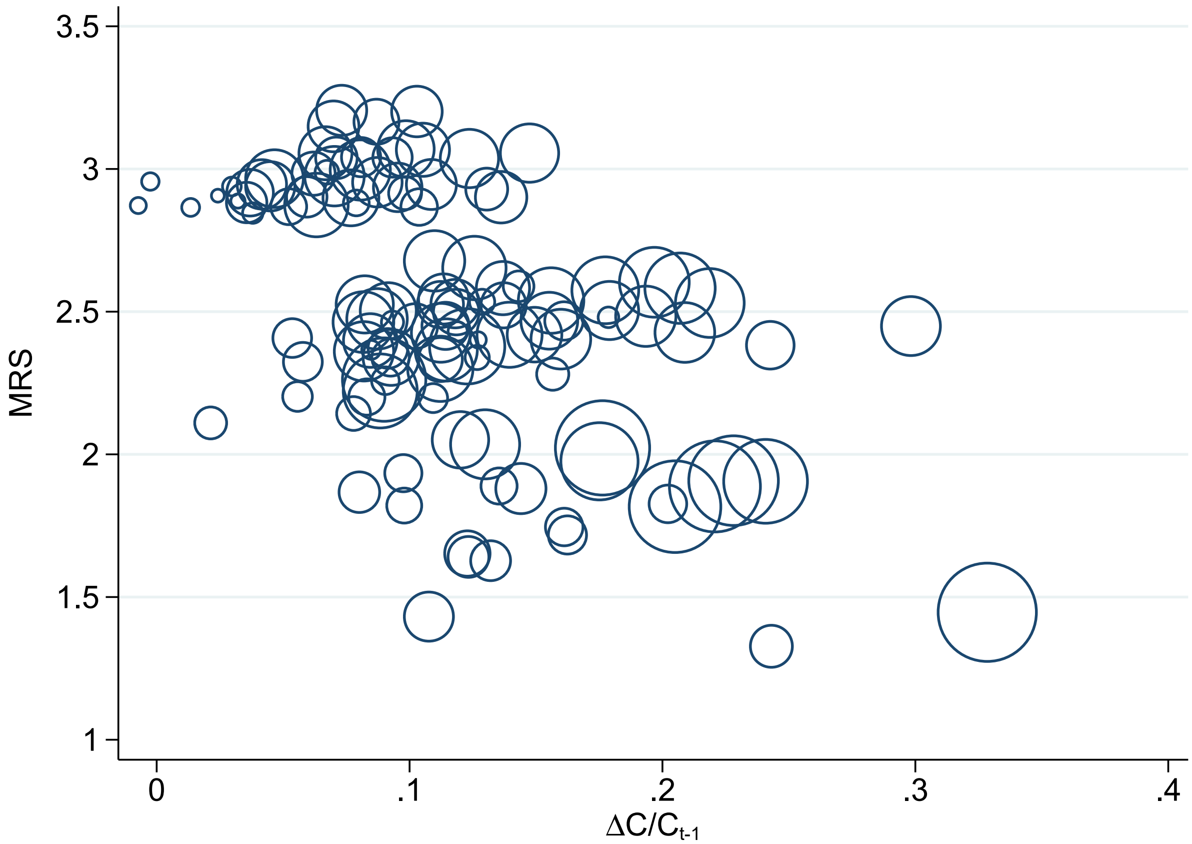

Notes: The graph correlates estimated MRS from the RP approach with estimated drops in consumption at job loss used in the CB approach for the individuals in our baseline sample split in cells of observables. We consider 120 cells of observable characteristics, using three age groups, income deciles, 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 (18). 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 conditional on consumption drops, there is still a very large amount of residual variation in MRS left in the data.

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 Income Quartile						-1.052				
						(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 3. 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 structural estimation using the perceived risk model. See Table 3 for details.

Appendix E 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.

E.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 $w = 850SEK$.[43]

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 E-1 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.[44] 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:

$$MPC = \frac{\lim_{w^-} dE[\Delta C|w]/dw - \lim_{w^+} dE[\Delta C|w]/dw}{\lim_{w^-} dE[b|w]/dw - \lim_{w^+} dE[b|w]/dw} \tag{25}$$

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 (2), as discussed in section A.4.

We estimate the numerator of the estimand in (25), in the same baseline sample of analysis used throughout the paper, based on the following RK specification:

$$\Delta C_i = \beta_0 \cdot (w-k) + \beta_1 \cdot (w-k) \cdot \mathbf{1}[w > k] + \sum_j \gamma_j \mathbf{1}[M_i = j] + \mathbf{X}'\gamma \tag{26}$$

[43] 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.

[44] 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 (25) 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 E-2 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 **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 E-2 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 **X** in specification (26) 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.

E.2 Results

Figure E-1 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 **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 E-3 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 E-3 panel B, we report how MPC estimates change as we start including cumulatively the characteristics of vector $\mathbf{X}$ in specification (26). 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 E-1: Regression kink design: Effect of UI benefits variation on consumption at unemployment

A. UI Benefit Replacement Rate

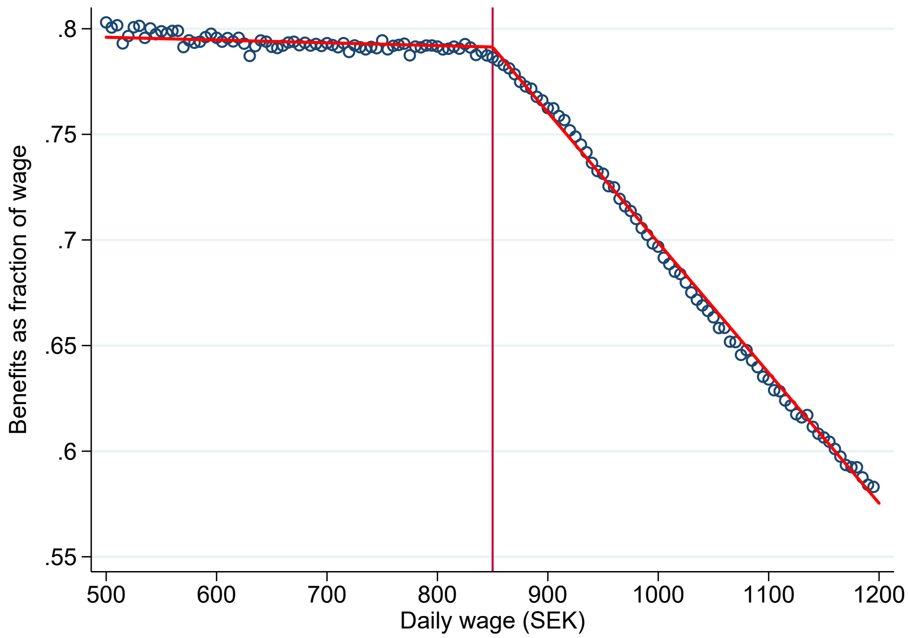

B. Consumption Drop at Unemployment

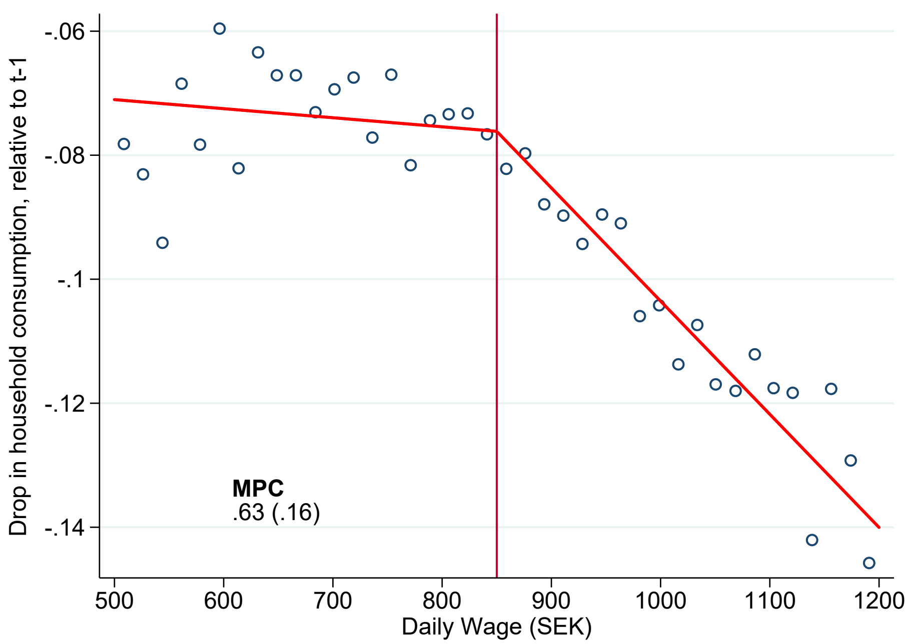

Notes: The Figure presents our RK design and main result. The design relies on the presence of a cap in UI benefits: the replacement rate is 80% of previous daily wage, up to a cap, when daily wages = 850SEK. Panel A plots, in our main sample over the period 2002 to 2007, a binscatter of the relationship 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 E-2: Regression kink design: Robustness

A. Pdf of Assignment Variable

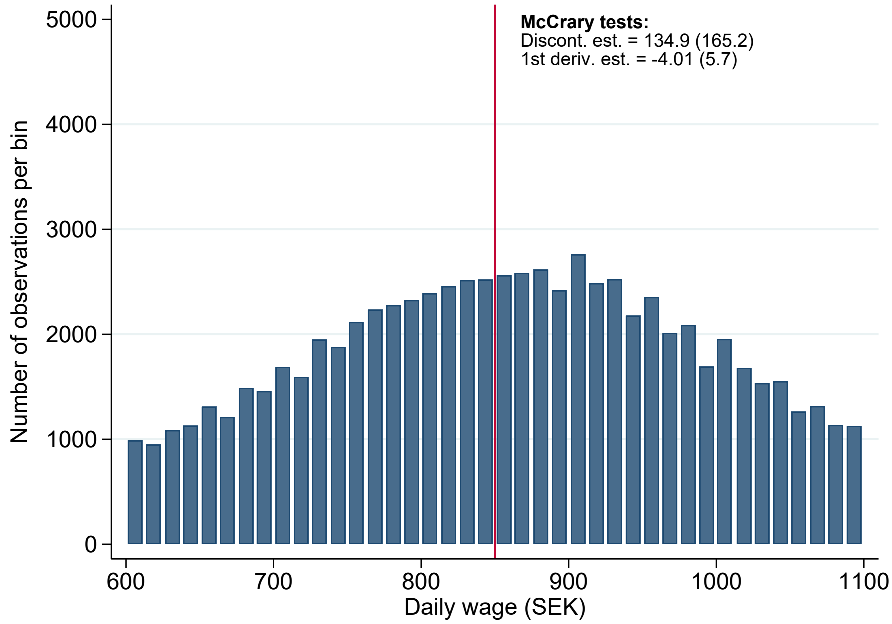

B. Covariate Index vs Assignment Variable

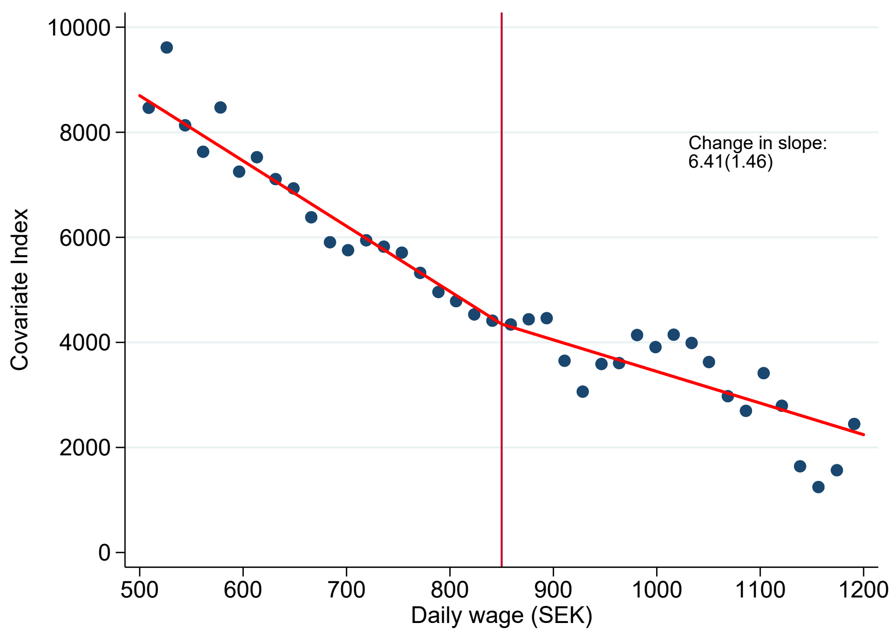

Notes: Panel A displays the probability density function of daily wage. We also report on the graph formal McCrary tests for the existence of a discontinuity nor lack of smoothness at the 850SEK threshold. Panel B investigates the presence of potential selection along observable characteristics around 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 E-3: Regression kink design: Sensitivity

A. Estimates by Bandwidth

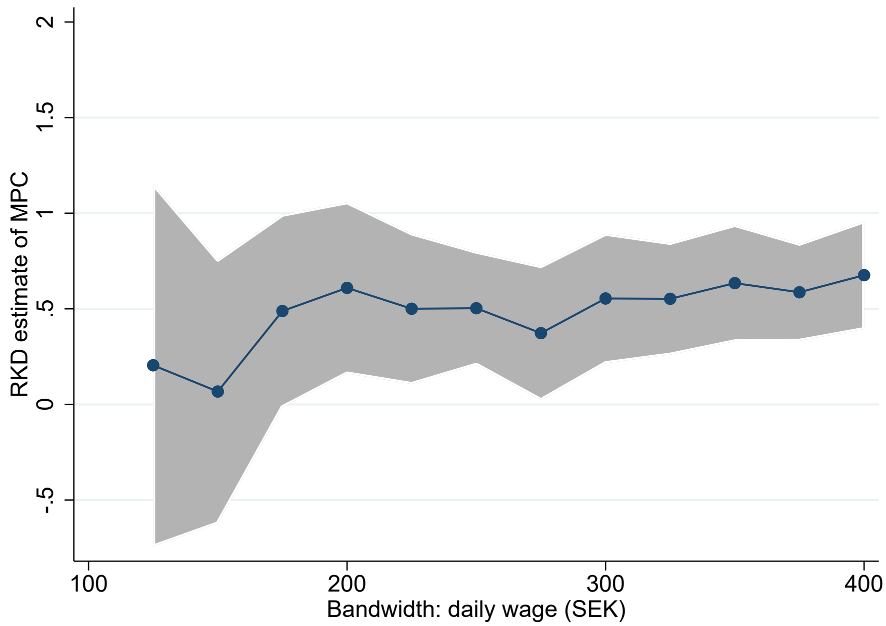

B. Sensitivity to Inclusion of Controls

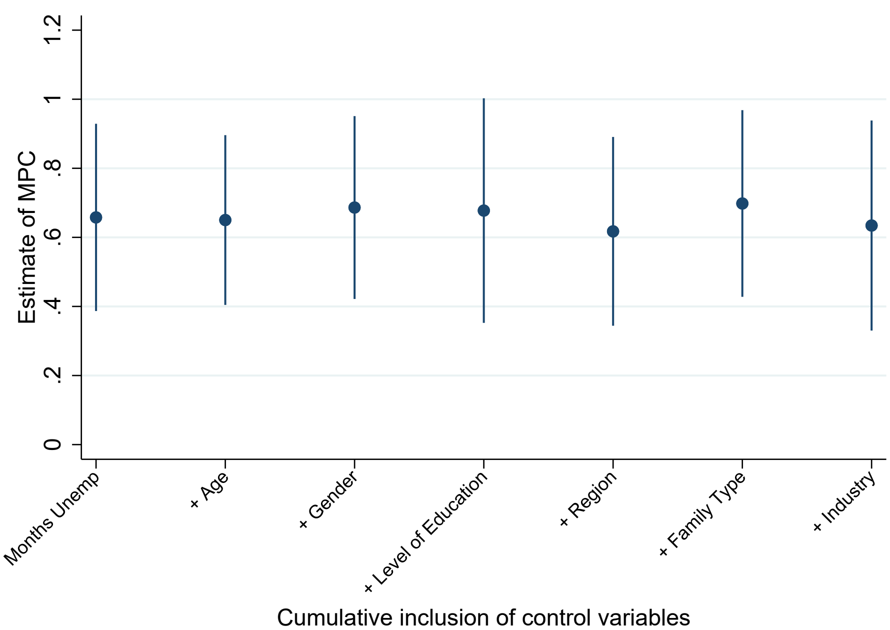

Notes: The Figure investigates the sensitivity of our estimates of the MPC while unemployed to specification assumptions and implementation choices. Panel A shows the sensitivity of our MPC estimates to the choice of bandwidth for the RK estimation. Our baseline bandwidth is 300. The optimal bandwidth 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 (26). 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.