

The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden

J Kolsrud
Uppsala

C Landais
LSE

P Nilsson
IIES

J Spinnewijn*
LSE

July 2, 2015

Abstract

This paper provides a simple, yet general framework to analyze the optimal time profile of benefits during the unemployment spell. We derive simple sufficient-statistics formulae capturing the insurance value and incentive costs of unemployment benefits paid at different times during the unemployment spell. Our general approach allows to revisit and evaluate in a transparent way the separate arguments for inclining or declining profiles put forward in the theoretical literature. We then estimate our sufficient statistics using administrative data on unemployment, income and wealth in Sweden. First, we exploit duration-dependent kinks in the replacement rate and find that the moral hazard cost of benefits is larger when paid earlier in the spell. Second, we find that the drop in consumption determining the insurance value of benefits is large from the start of the spell, but further increases throughout the spell. On average, savings and credit play a limited role in smoothing consumption. Our evidence therefore indicates that the recent change from a flat to a declining benefit profile in Sweden has decreased welfare. In fact, the local welfare gains push towards an increasing rather than decreasing benefit profile over the spell.

Keywords: Unemployment, Dynamic Policy, Sufficient Statistics, Consumption Smoothing

JEL codes: H20, J64

*We thank Tony Atkinson, Richard Blundell, Raj Chetty, Liran Einav, Hugo Hopenhayn, Henrik Kleven, Alan Manning, Nicola Pavoni, Torsten Persson, Emmanuel Saez, Florian Scheuer, Robert Shimer, Frans Spinnewyn, Ivan Werning, Gabriel Zucman and seminar participants at the NBER PF Spring Meeting, SED Warsaw, DIW Berlin, Kiel, Zurich, Helsinki, Stanford, Leuven and LSE for helpful discussions and suggestions. We also thank Iain Bamford and Panos Mavrokonstantis for excellent research assistance. We acknowledge financial support from the Sloan Foundation (NBER grant number 22-2382-15-1-33-003), STICERD and the CEP.

1 Introduction

How should unemployment benefits be paid over time? Should benefits be cut after 6 months of unemployment like in the US? Should the unemployed receive the same benefits level for ever? Substantial pressure in the policy debate has indeed pushed countries like Belgium and Sweden where flat-benefit profiles were in place until recently, to reduce the benefit generosity for the long-term unemployed relative to the short-term unemployed. But why should we prefer declining profiles of benefits over the unemployment spell rather than flat or even increasing profiles?

Social insurance programs aim to provide insurance while maintaining incentives. However, solving the dynamic problem of balancing incentives and insurance during unemployment can prove daunting, especially when adding important features such as heterogeneity and non-stationarity. We critically lack a simple, general, and evidence-based framework, tightly integrating theory and empirics, to evaluate dynamic social insurance policies like unemployment insurance (UI). While there is a growing empirical literature that evaluates social insurance design using a *sufficient-statistics* approach, this literature has ignored the dynamic features of social insurance policies. At the same time, there is an influential theoretical literature on optimal dynamic policies, but derived in stylized models that are hard to connect to the data.

This paper tries to fill this important gap in the literature. In the spirit of the sufficient-statistics approach we derive a characterization of the optimal dynamic policy based on a limited set of high-level statistics. This general, but simple characterization shows how the trade-off between insurance and incentives evolves over the unemployment spell and identifies the relevant behavioral responses in this dynamic context. We then use Swedish administrative data on unemployment, income and wealth to estimate the full set of statistics required to evaluate the benefit profile in Sweden. Surprisingly, both from the insurance and the incentives side, we find no evidence to support the recent policy change in Sweden to a declining benefit profile.

Our analysis starts by considering a general dynamic model of unemployment incorporating job search behavior and consumption decisions, allowing for unobservable heterogeneity, duration dependence in job finding rates, etc. Using dynamic envelope conditions, we show that the Baily-Chetty intuition (Baily [1978], Chetty [2006]) generalizes for UI benefits paid at any unemployment duration t . The optimal benefit paid at unemployment duration t should balance the corresponding insurance value with the implied moral hazard cost at the margin. Importantly, both the insurance value and moral hazard cost can be expressed as a function of identifiable and estimable statistics. The insurance value of benefits at time t of the unemployment spell will depend on the drop in consumption at time t , while the incentive cost will depend on the effect of these benefits on the government expenditures through agents' unemployment responses. These are simply captured by survival rate responses throughout the unemployment spell weighted by the benefit levels paid.

Through identifying the relevant set of statistics to evaluate welfare, our approach provides a simple, yet robust guide to evaluate the time profile of unemployment benefits. While a large literature has analyzed average duration responses to particular benefit changes or spikes in exit rates at benefit exhaustion, our framework requires responses in unemployment survival throughout

the spell to variations in the benefit profile (i.e., variations in UI benefits at different points of an unemployment spell). Similarly, while some previous work has analyzed the average consumption drop at unemployment, the evaluation of the benefit profile depends on the ability of the unemployed to smooth consumption throughout the unemployment spell. We analyze and provide important contributions on these two missing pieces in the empirical part of this paper.

The empirical analysis uses a unique administrative dataset in Sweden that combines unemployment registers and tax registers with comprehensive information on income and wealth for the universe of Swedish individuals from 1999 until 2007. We first exploit duration-dependent caps on unemployment benefits using a regression-kink design. These caps have been affected by several policy reforms, allowing us to estimate non-parametrically how unemployment survival responds to different variations in the benefit profile. The policy variation also offers compelling placebo settings that confirm the robustness of our approach. We then take advantage of the exhaustive information on income, transfers and wealth in Sweden to construct a residual measure of yearly expenditures, accounting for all income sources and changes in assets. Linking this measure to administrative data with precise information on the duration of unemployment, we can identify how consumption expenditures at yearly, but also at higher frequency, react to unemployment and the length of an unemployment spell in particular. We also confirm the robustness of our conclusions using surveyed consumption measures.

Our empirical analysis provides the following main results;

First, unemployment durations are very responsive overall to changes in the benefit level. Surprisingly, the response to changes in benefits early in the spell is more pronounced than changes later on. The implied moral hazard cost of increasing the benefit level is 21 percent higher for benefits paid in the first 20 weeks of unemployment than for benefits paid after 20 weeks. This result may seem surprising. All else equal, the incentive cost from increasing benefits for the long-term unemployed is expected to be larger as it also discourages the short-term unemployed from leaving unemployment when they are forward-looking. Using the same regression-kink design, we do provide clear evidence that exit rates early in the spell respond to benefit changes applying later in the spell, but also that agents become less responsive later in the spell to comparable changes in the policy. Importantly, such non-stationary forces, which may be driven by duration dependence or dynamic selection over the unemployment spell, are large enough to offset the significant effect of forward-looking incentives.

Second, expenditures drop substantially and early in the spell. We find that expenditures drop on average by 19% in the first 20 weeks of unemployment, compared to their pre-unemployment level. This drop deepens to on average 27% for those who are unemployed for longer. Our data allows us to provide a detailed account of the different smoothing mechanisms. Overall, we find a limited ability to smooth consumption in addition to the government transfers, as most unemployed individuals have few assets prior to becoming unemployed. Individuals who have access to assets do use them to smooth consumption, but mostly by depleting liquid assets from their bank accounts. Consumption through increases in debt declines over the unemployment spell, reflecting both a

reduction in debt-financed purchases of real estate and in non-mortgage related credit. Evidence from the consumption surveys offers complementary insights into the types of consumption goods that individuals adjust over the spell. It is non-committed, non-durable consumption that is more likely to decrease strongly over the spell. Interestingly, the drop in consumption of durable goods is most pronounced early on in the spell and reduces later on in the spell. The panel structure of our administrative data allows us to provide evidence that almost all of the drop in average consumption over the spell happens at the within-individual level rather than between individuals. We also aggregate our expenditure measures at the household level and demonstrate that households offer only little additional consumption smoothing opportunities, either through their earnings or assets.

Our empirical analysis thus indicates that the insurance value of unemployment benefits increases over the unemployment spell, while the incentive cost decreases slightly over the unemployment spell. Putting together our estimates, which are evaluated for the flat-benefit profile in Sweden before 2007, we find that welfare could be increased by having an inclining profile of benefits over the spell. Our estimates suggest a welfare gain of 80% for the first kroner transferred from the short-term unemployed to the long-term unemployed, accounting for the unemployment responses. Instead Sweden has introduced a declining benefit profile in 2007.

Our paper contributes to several literatures. First, the sufficient statistics approach has a long tradition in UI starting with Baily [1978], implemented by Gruber [1997] and generalized by Chetty [2006]. To date, this literature has focused almost entirely on the optimal average generosity of the system.¹ Conversely, the theoretical literature on the optimal time profile of UI has generated results in stationary, representative-agent models, which are hard to take to the data. Our analysis shows how the previously identified forces (e.g., in Hopenhayn and Nicolini [1997] and Shimer and Werning [2008]) come together, but also integrates heterogeneity and duration-dependence (see for example Shimer and Werning [2006], Pavoni [2009]). Second, our empirical analysis of unemployment responses contributes to a long literature on labor supply effects of social insurance (see Krueger and Meyer [2002]) by explicitly using duration-dependent variation in benefits and identifying the welfare-relevant unemployment responses. Our analysis also indicates that differences in the timing of the benefit variation could explain different estimates of unemployment responses in the literature. Finally, a large literature has used consumption surveys to analyze consumption smoothing of income shocks and unemployment in particular (e.g., Gruber [1997]). The use of an administrative expenditure measure is a powerful alternative, which allows to exploit important advantages from registry data and seems a promising instrument to evaluate policy more generally².

The remainder of the paper proceeds as follows. Section 2 identifies the sufficient statistics for policy design in a dynamic model of unemployment. Section 3 describes our data and the

¹Recent extensions of the Baily-Chetty formula can be found in Schmieder et al. [2012], analyzing the potential benefit duration, but for a given benefit level, or also in Gerard and Gonzaga [2014]. Spinnewijn [2015] provides a formula for the optimal intercept and slope of a linear benefit profile.

²See for instance Mogstad and Kostol [2015], Kreiner et al. [2014] and Pistaferri [2015] for a survey of these recent developments.

policy context in Sweden. Section 4 describes our regression kink design and provides estimates of the relevant unemployment elasticities. Section 5 describes the construction of our consumption measure and analyzes how this evolves during the unemployment spell. Section 6 analyzes welfare linking the theory and the empirical estimates and providing a back-of-the-envelope calibration of the optimal benefit profile. Section 7 concludes.

2 Model

We characterize the optimal unemployment policy in a stylized job search model closely related to Chetty [2008]. In the spirit of the “sufficient statistics” literature, our approach consists in identifying the minimal level of information necessary to determine the optimal time profile of unemployment benefits in a broad class of models. For that matter, we avoid specifying the primitives of the environment and the underlying agent’s behavior when possible. We focus instead on observable variables that are relevant for policy. As a consequence, our setup can account for various non-stationary forces (e.g., duration-dependence in exit rates) and various forms of heterogeneity (e.g., heterogeneity in savings or credit).

We start by presenting our “sufficient statistics” characterization of the optimal UI policy. We then revisit the insights from previous work on the optimal timing of benefits. In particular, we show how the different forces identified in more specific models affect the “sufficient statistics” in our encompassing model.

2.1 Setup

We consider a dynamic model of unemployment in discrete time ending at T . We first present our setup and results for a representative agent, but allow for heterogeneity later on. The agent starts unemployed and remains unemployed until she finds work. Once she has found work, she remains employed until T . The government designs an unemployment policy (\mathbf{b}, τ) to provide insurance against the unemployment risk. The policy pays out unemployment benefits b_t depending on the unemployment duration t and is funded by a uniform tax τ paid when employed.

Agent’s behavior. An agent makes two types of choices: job search choices \mathbf{s} and consumption choices \mathbf{c} . The first type of decisions concern her general job search strategy: each period of unemployment the agent decides how hard to search, what jobs to search for, how to set her reservation wage, etc. While we do not model these search decisions explicitly, our model accounts for their dependence on the unemployment policy and their impact on the probability that the agent leaves unemployment. That is, we focus on the resulting exit rate out of unemployment denoted by h_t . This hazard rate changes during the unemployment spell if the policy incentives or the environment do. For example, the hazard rate is allowed to evolve over the unemployment spell due to skill-depreciation or any other form of duration dependence. The agent’s probability to be still unemployed after t periods equals the survival rate $S_t = \prod_{s=1}^t (1 - h_s)$ (with $S_0 = 1$).

The expected duration of the unemployment spell simply equals the sum of the survival rates at each duration, $D = \sum_{t=0}^T S_t$.

The second type of decisions concern the agent's consumption. She decides how much to save or borrow, or how much to use any other means of smoothing consumption over time (e.g., family transfers, household production). These decisions determine her consumption level throughout the unemployment spell and when employed. We denote her consumption level by c_t^u and c_t^e for when unemployed and employed respectively at time t . Like her job search strategy, the agent's consumption allocation will depend on the unemployment policy and her unemployment history.

The agent makes her search and consumption decisions to maximize her expected utility taking the unemployment policy as given, $V(\mathbf{b}, \tau) \equiv \max_{\mathbf{s}, \mathbf{c}} U(\mathbf{s}, \mathbf{c} | \mathbf{b}, \tau)$. We assume time-separable preferences with instantaneous utility $u(c_t^u, s_t)$ and utility $u(c_t^e)$ when unemployed and employed respectively.

Unemployment Policy. We consider an n -part policy that pays benefit b_1 for the first B_1 periods, b_2 for the next $B_2 - B_1$ periods, and so on.³ We take the potential benefit durations as given, but our characterization generalizes for a fully flexible policy with $n = T$.

The government runs a balanced budget. That is, the expected benefit payments of the unemployment policy are covered by the expected tax revenues,

$$\underbrace{\sum_0^{B_1} S_t b_1}_{\equiv D_1} + \underbrace{\sum_{B_1}^{B_2} S_t b_2}_{\equiv D_2} + \dots + \underbrace{\sum_{B_{n-1}}^T S_t b_n}_{\equiv D_n} = (T - D) \tau.$$

We use D_k to denote the expected time spent unemployed while receiving b_k . The total expected unemployment duration equals the sum of the durations spent on the different parts of the unemployment policy, $D = \sum_{k=1}^n D_k$. We illustrate this for a two-part policy in Figure 1, a simple scheme that is commonly used in practice. As shown in Panel A, the scheme pays a benefit level b_1 when *short-term* unemployed, up to B periods, and a benefit level b_2 afterwards when *long-term* unemployed. Panel C illustrates the corresponding survival function S_t . The areas under the survival function before and after B denote the expected time spent receiving the benefits b_1 and b_2 respectively.

Note that we ignore time discounting for notational convenience only. The analysis also naturally generalizes when considering more than one unemployment spell or overlapping cohorts which move in and out of unemployment.⁴

³As it will be clear which case applies, we use (with some abuse of notation) the subindex to either refer to time t or to part k of the unemployment policy.

⁴Note first that when the discount and interest rates are equal, the survival rates used in the agent's expected utility and the government's budget constraint could simply be re-interpreted as time-weighted probabilities. Second, when employment is not an absorbing state, the end time T could be re-interpreted as the time between the start of two consecutive employment spells, which then would be endogenous to the policy. Finally, in the steady state of a model with overlapping cohorts, the share D_k/D_l would simply capture the share of unemployed in the k -th part relative to the l -th part of the unemployment spell at a given point in time.

2.2 Optimal Unemployment Policy

We characterize the budget-balanced policy that maximizes the agent's expected utility. The optimal policy trades off the provision of insurance and incentives, but both the value of UI and its incentive cost may change over the unemployment spell. Baily [1978] and Chetty [2006] have shown that the marginal insurance value of a flat benefit profile is determined by the difference in marginal utilities when employed and unemployed and that the incentive cost is determined by the unemployment response to the coverage provided. We generalize these insights in a dynamic setting for benefits paid at different times during the unemployment spell.

Consumption Smoothing. An agent loses her labor earnings when she loses her job. The utility impact of this loss depends on the resulting drop in her consumption. When the unemployed have access to unemployment benefits, liquid assets or any other means, they can smoothen the impact of unemployment and thus experience a smaller drop in consumption as a result. When these means change while unemployed, their consumption will vary as well over the spell.

The utility gain from increasing benefits at duration t equals $\partial V / \partial b_t = S_t \partial u(c_t^u, s_t) / \partial c$ and thus only depends on the agent's marginal utility of consumption at that duration. While the agent will change her consumption allocation dc in response to changes in the unemployment policy, such behavioral changes have only a second order impact on the agent's expected utility by the envelope condition $\partial V / \partial c_t = 0$ (see Chetty 2006, 2009). As a consequence, the insurance gain from a small tax-funded increase in unemployment benefits depends only on the difference in marginal utilities when receiving the benefit and paying the tax respectively. That is, the consumption smoothing benefit of an increase in the benefit level b_k received during the k -th part of the unemployment policy paid by an increase in the tax τ equals,

$$CS_k \equiv \frac{E_k u'(c^u) - Eu'(c^e)}{Eu'(c^e)}, \quad (1)$$

where $E_k u'(c^u) = \left[\sum_{B_{k-1}}^{B_k} S_t \partial u(c_t^u, s_t) / \partial c \right] / D_k$ and $Eu'(c^e) = [\sum_0^T (1 - S_t) E[u'(c_t^e)]] / [T - D]$.

Interestingly, when preferences are time-separable and consumption utility is separable from other variables, the consumption smoothing gain of increasing unemployment benefits at time t will simply depend on the consumption level of unemployed individuals *at time t* . In practice, we can then approximate the consumption smoothing gain by

$$CS_k \cong \gamma \times \frac{Ec^e - E_k c^u}{Ec^e},$$

where γ is the relative risk aversion $-cu''(c) / u'(c)$.⁵ To evaluate the desirability of the policy's benefit profile from a consumption smoothing perspective, it is sufficient to compare the agent's average consumption level during different parts of the unemployment spell. If consumption levels

⁵Note that this relies on the Taylor expansion of the marginal utility of consumption to provide a good approximation. Our approximation ignores third-order utility terms in the utility function (see Chetty 2006).

drop during the unemployment spell, increases in the unemployment benefits are more desirable when timed later in the spell.

Moral Hazard. An agent remains unemployed for longer when receiving higher unemployment benefits. As for the consumption response, the job search response ds to changes in unemployment benefits has only a second-order impact on the agent's expected utility. However, this response has a first-order impact on the government's budget constraint since the agent does not take the fiscal externalities of her job search strategy into account. The moral hazard cost of an increase in unemployment benefits depends entirely on the impact on the government's budget constraint. This impact is captured by the changes in the time D_l spent on the different parts of the unemployment policy, weighted by the benefits b_l paid. In particular, the moral hazard cost of a budget-balanced increase in benefit b_k and tax τ equals

$$MH_k = \sum_{l=1}^n \frac{D_l b_l}{D_k b_k} \times \varepsilon_{D_l, b_k}, \quad (2)$$

where ε_{D_l, b_k} equals the elasticity of the benefit duration D_l , relative to the time spent employed $T - D$, with respect to the benefit b_k , adjusting the tax τ to keep the budget balanced. That is,

$$\varepsilon_{D_l, b_k} = \frac{d[D_l/(T-D)]}{db_k} \frac{b_k}{D_l/(T-D)} \cong \frac{\partial D_l}{\partial b_k} \frac{b_k}{D_l},$$

where the approximation follows as the employment response and the budget-balancing tax change are small for standard unemployment rates.

We can thus distinguish between three types of responses throughout the unemployment spell to an increase in benefits db_k paid between time B_{k-1} and B_k of the spell. First, a higher b_k lowers the instantaneous exit rate out of unemployment and thus increases the time spent receiving the increased benefits b_k ($\varepsilon_{D_k, b_k} \geq 0$). Second, by increasing the share of people still unemployed at time B_k , a higher b_k also increases the time spent receiving benefits later in the spell ($\varepsilon_{D_l, b_k} \geq 0$ for $l > k$). Finally, a higher b_k also reduces the incentives to leave unemployment earlier in the spell as agent anticipates the more generous benefit received later in the spell ($\varepsilon_{D_l, b_k} \geq 0$ for $l < k$).

All duration responses to a change in b_k are weighted by the relative budget shares of the respective parts of the unemployment policy, as shown in equation (2). When starting from a flat profile, however, it is the total duration response that captures the fiscal externality. The moral hazard cost of increasing benefits b_k then simply depends on the elasticity of the average unemployment duration ε_{D, b_k} scaled by the relative time spent receiving b_k . That is,

$$MH_k = \varepsilon_{D, b_k} \frac{D}{D_k} \text{ when } b_t = b \text{ for all } t. \quad (3)$$

To evaluate the desirability of the policy's benefit profile, it is sufficient to compare the incentive cost from increasing the benefit level of the different parts of the unemployment policy. Specific models of unemployment may have different implications for the values of these elasticities and how

they depend on the timing of the benefits, but the crucial insight is that the elasticities themselves are sufficient to evaluate the moral hazard cost of the policy and can be estimated empirically.

Dynamic Baily Formula. For a benefit payment during the unemployment spell to be set optimally, the consumption smoothing benefit from a further increase must equal its moral hazard cost. For an n -part benefit policy to be optimal, this needs to hold for all benefit levels. We can state the following result:

Proposition 1. *A budget-balanced unemployment policy (b_1, \dots, b_n, τ) for given potential duration thresholds (B_1, \dots, B_n) is optimal only if*

$$\frac{E_k u'(c^u) - E u'(c^e)}{E u'(c^e)} = \sum_{l=1}^n \frac{D_l b_l}{D_k b_k} \times \varepsilon_{D_l, b_k} \text{ for each } k.$$

Proof. *The policy maker solves*

$$\max_{\mathbf{b}, \tau} V(\mathbf{b}, \tau) \text{ s.t. } (T - D) \tau = \sum_k \left[\left[\sum_{t=B_{k-1}}^{B_k} S_t \right] b_k \right]$$

where V denotes the agent's indirect utility maximized over her search choices \mathbf{s} and consumption choices \mathbf{c} , while S_t and D denote the survival rate and unemployment duration implied by the agent's search behavior. For the policy to be optimal, a budget-balanced increase in b_k and τ (accounting for the behavioral responses) cannot increase welfare and this for any $k \leq n$. That is,

$$\sum_{t=B_{k-1}}^{B_k} S_t \frac{\partial}{\partial c} u(c_t^u, s_t) - \sum_{t=0}^T [1 - S_t] E[u'(c_t^e)] \frac{d\tau}{db_k} = 0.$$

The impact of any search or consumption response to the change in the policy on the agent's utility is of second order by the envelope theorem ($\partial V / \partial x = 0$ for $x \in s$ or $x \in c$). The direct impact of an increase in the benefit level when unemployed at time $t \in [B_{k-1}, B_k]$ equals $\partial u(c_t^u, s_t) / \partial c$, given the unemployed's optimal consumption and search choices. The direct impact of an increase in the tax rate when employed at time t equals $E[u'(c_t^e)]$, given the employed's optimal consumption choice accounting for the variation in the wage offer the agent accepted when leaving unemployment. The marginal utilities of consumption in the states affected by the policy change are weighted by the probability to be unemployed and employed respectively.

The budget-balancing change in the tax $d\tau / db_k$ does depend on the behavioral responses through the change in the survival rates. Accounting for these responses wrt this budget-balanced change,

we can write

$$\begin{aligned}
\frac{d\tau}{db_k} &= \frac{D_k}{T-D} + \frac{d\frac{D_k}{T-D}}{db_k} b_k + \Sigma_{l \neq k} \frac{d\frac{D_l}{T-D}}{db_k} b_l \\
&= \frac{D_k}{T-D} \left\{ 1 + \frac{d\frac{D_k}{T-D}}{db_k} \frac{b_k}{\frac{D_k}{T-D}} + \Sigma_{l \neq k} \frac{d\frac{D_l}{T-D}}{db_k} \frac{b_k}{\frac{D_l}{T-D}} \frac{b_l D_l}{b_k D_k} \right\} \\
&= \frac{D_k}{T-D} \left\{ 1 + \Sigma_{l=1}^n \frac{D_l b_l}{D_k b_k} \times \varepsilon_{D_l, b_k} \right\}.
\end{aligned}$$

Substituting this expression into the original first-order condition and using the expected value notation introduced before, we find

$$D_k E_k u'(c^u) - (T-D) E u'(c^e) \frac{D_k}{T-D} \left\{ 1 + \Sigma_{l=1}^n \frac{D_l b_l}{D_k b_k} \times \varepsilon_{D_l, b_k} \right\} = 0.$$

Hence, the formula in the Proposition immediately follows. ■

Our characterization simplifies to the Baily-Chetty formula for a flat profile ($b_t = b$),

$$\frac{E u'(c^u) - E u'(c^e)}{E u'(c^e)} = \varepsilon_{D,b}. \quad (4)$$

To evaluate the optimality of a flat policy, the average wedge in marginal utilities and the average duration elasticity with respect to an overall change in benefit level are sufficient.

To evaluate the optimal benefit profile, more information is required. For a simple two-part policy - paying benefit b_1 until time B and b_2 thereafter - the characterization becomes:

$$\frac{E_1 u'(c^u) - E u'(c^e)}{E u'(c^e)} = \varepsilon_{D_1, b_1} + \frac{D_2 b_2}{D_1 b_1} \varepsilon_{D_2, b_1}, \quad (5)$$

$$\frac{E_2 u'(c^u) - E u'(c^e)}{E u'(c^e)} = \frac{D_1 b_1}{D_2 b_2} \varepsilon_{D_1, b_2} + \varepsilon_{D_2, b_2}. \quad (6)$$

For the wedge in marginal utilities, we simply need to decompose the average wedge into a wedge for the short-term unemployed and a wedge for the long-term unemployed, as illustrated in Panel B of Figure 1. No policy variation is needed to estimate these wedges. For the moral hazard cost, we need, as illustrated in Panel C and Panel D of Figure 1, to decompose duration responses into the responses to variations in short-term benefits b_1 and the responses to variations in long-term benefits b_2 . To identify such responses, duration-dependent policy variation is required. That is, rather than having benefits change throughout the spell, we need changes in benefits paid only to the short-term db_1 or to the long-term unemployed db_2 as illustrated in Panel A of Figure 1.

Comparing the consumption smoothing gain and moral hazard cost not only allows for a characterization of the optimal policy, but also provides a simple, yet robust guide to evaluate local policy changes. When, for some part of the policy, the moral hazard cost MH_k exceeds the consumption smoothing gain CS_k , evaluated at the current policy, welfare could be increased by decreasing the

generosity of that part of the policy and vice versa. We can also directly evaluate local changes in the benefit profile. For example, the relative consumption smoothing gain CS_1/CS_2 and moral hazard cost MH_1/MH_2 will determine in which direction the profile of a two-part policy should be adjusted to increase welfare.

Heterogeneous Agents. We briefly extend the characterization in the presence of heterogeneous agents to evaluate when heterogeneity matters. Consider J types of unemployed agents who differ in preferences, productivity, search skills, assets, etc. An agent of type j chooses her job search strategy given her (type-specific) employment prospects and decides how much to consume maximizing her (type-specific) expected utility. Her optimized behavior results in a type-dependent exit rate $h_{j,t}$ and consumption level $c_{j,t}$. We denote the share of agents of type j by α_j and the corresponding social welfare weight by λ_j . Note that the type-weighted average of exit rates and consumption levels equals the empirically observed averages. The following extension of Proposition 1 follows immediately:

Corollary 1. *In a model with heterogeneous types $j \in J$, a budget-balanced unemployment policy (b_1, \dots, b_n, τ) for given potential duration thresholds (B_1, \dots, B_n) is optimal only if*

$$\frac{E_k(\lambda u'(c^u)) - E(\lambda u'(c^e))}{E\lambda u'(c^e)} = \sum_{l=1}^n \frac{D_l b_l}{D_k b_k} \times \varepsilon_{D_l, b_k} \text{ for each } k,$$

where

$$\begin{aligned} E_k \lambda u'(c^u) &= \left[\sum_1^J \lambda_j \alpha_j \sum_{B_{k-1}}^{B_k} S_{j,t} \partial u_j(c_{j,t}^u, s_{j,t}) / \partial c \right] / D_k, \\ E \lambda u'(c^e) &= \left[\sum_1^J \lambda_j \alpha_j \sum_0^T (1 - S_{j,t}) u_j'(c_{j,t}^e) \right] / [T - D], \\ D_k &= \sum_1^J \alpha_j \sum_{B_{k-1}}^{B_k} S_{j,t}. \end{aligned}$$

The first thing to notice is that the sufficient statistics capturing the moral hazard cost remain unchanged. While striking, this is an immediate consequence of the fact that the moral hazard cost only depends on the fiscal externality associated with the unemployed's behavior. Different agents may respond very differently to changes in benefits, but it is the resulting change in the average benefit durations that captures the welfare-relevant impact of these responses.

The consumption smoothing gain CS_k of benefits paid in the k -th part of the policy now equals the wedge in Pareto-weighted marginal utilities. In practice, we can approximate this by

$$CS_k \cong \tilde{\lambda}_k \times \tilde{\gamma}_k \times \frac{Ec^e - E_k c^u}{Ec^e},$$

for some representative Pareto-weight $\tilde{\lambda}_k$ and risk-aversion $\tilde{\gamma}_k$ applicable to the individuals who are unemployed during the k -th part of the spell. For a utilitarian social welfare function and homogeneous risk preferences, the consumption profile is again sufficient to evaluate the benefit profile. If Pareto-weights and preferences are different across types, the representative parameters

$\tilde{\lambda}_k$ and $\tilde{\gamma}_k$ should account for the selection of types throughout the unemployment spell.

2.3 Optimal Timing of Benefits

A large literature has analyzed the optimal dynamics of the unemployment insurance policy. We revisit some of the forces for decreasing or increasing unemployment benefits as identified in previous work through the lens of our sufficient statistics. In particular, we study how the consumption smoothing benefits and the moral hazard cost evolve during the unemployment spell.

2.3.1 Stationary Environment

Consider as a benchmark case a representative agent in a stationary environment. We assume that the agent's exit rate h_t only depends on the continuation policy $(b_t, b_{t+1}, \dots, \tau)$. This implies that the exit rate would remain constant during unemployment when the benefit profile is flat for $T \rightarrow \infty$.⁶

Moral Hazard. We first analyze how the moral hazard cost depends on the timing of the UI benefits, evaluating changes relative to a flat benefit profile. When benefits are increased later in the spell, an agent not just reduces her search effort (or increases her reservation wage) when receiving the higher benefits, but already responds from the start of the spell, $\frac{\partial h_{t-s}}{\partial b_t} \leq 0$. This force depends on the agent's forward-looking behavior and tends to raise the incentive cost of benefits timed later in the spell relative to benefits timed earlier.

To illustrate this, consider an increase in benefit level b_2 of a two-part policy which an agent receives when unemployed for longer than time B . The time spent receiving this benefit level depends on the probability S_B to be still unemployed at time B and the expected time she remains unemployed afterwards $\tilde{D}_B \equiv \sum_{s=0}^{\infty} S_{B+s}/S_B$. We can thus rewrite

$$\varepsilon_{D_2, b_2} = \varepsilon_{S_B, b_2} + \varepsilon_{\tilde{D}_B, b_2}.$$

In our stationary environment, the latter only depends on the continuation policy. In fact, the response in the *remaining duration* of unemployment \tilde{D}_B , conditional on still being unemployed at time B , is independent of the time B at which the benefits change. Starting from a flat profile ($b_1 = b_2$), we therefore have that $\varepsilon_{\tilde{D}_B, b_2} = \varepsilon_{D, b}$. The immediate consequence is that the moral hazard cost of increasing benefits later in the spell, which exceeds ε_{D_2, b_2} , also exceeds the moral hazard cost of increasing benefits throughout, which equals $\varepsilon_{D, b}$. This result generalizes as follows:

Proposition 2. *Starting from a flat benefit profile ($b_t = b$ for all t) in a stationary environment with equal discount and interest rate, the moral hazard cost increases during the unemployment spell,*

$$MH_t \leq MH_{t'} \text{ for } t < t'.$$

⁶Note that this assumption is satisfied in a McCall search model with CARA preferences (e.g., Shimer and Werning [2008]) or in any stationary model without savings (e.g., Hopenhayn and Nicolini [1997]).

Proof. See appendix.

The forward-looking behavior of the unemployed causes the incentive cost of providing UI benefits to be higher when timed later in the spell. This provides a force towards declining unemployment benefits and underlies the declining optimal benefit profile in models without savings as considered in Shavell and Weiss [1979] and Hopenhayn and Nicolini [1997].

Consumption Smoothing. We now turn to the evolution of the consumption smoothing benefits during the unemployment spell. This is rather straightforward as standard models of intertemporal consumption predict that consumption is weakly decreasing during unemployment. For equal discount and interest rate, the Euler equation states

$$\frac{\partial}{\partial c} u(c_t^u, s_t) \geq h_t Eu'(c_{t+1}^e) + (1 - h_t) \frac{\partial}{\partial c} u(c_{t+1}^u, s_{t+1}).$$

When the marginal utility is higher when unemployed than when employed (e.g., separable preferences), an agent prefers to consume out of her savings or borrow against her future earnings to increase her unemployment c_t^u at time t at the expense of future consumption. When facing a flat benefit profile ($b_t = b$), an agent thus runs down her assets and reduces her consumption while unemployed. This continues until she becomes liquidity-constrained and starts consuming "hand-to-mouth".

Proposition 3. *Starting from a flat benefit profile ($b_t = b$ for all t) in a stationary model with $\partial u(c_t^u, s_t) / \partial c > Eu'(c_t^e)$ and equal discount and interest rate, the consumption smoothing gain increases during the unemployment spell,*

$$CS_t \leq CS_{t'} \text{ for } t < t'.$$

Proof. See appendix.

Long-term unemployment implies a larger shock than short-term unemployment and this larger shock requires more savings or credit to be smoothed. All else equal, this implies that social insurance against long-term unemployment is valued more. This pushes the optimal benefit profile to be increasing over the unemployment spell and underlies the inclining optimal benefit profile in a model without search considered in Shavell and Weiss [1979].

Propositions 2 and 3 state two forces with opposing effects on the optimal profile of the unemployment benefits. Werning [2002] and Shimer and Werning [2008] analyze these two opposing forces in models with search and savings and show that these exactly cancel out in case of CARA preferences; a flat benefit profile is optimal conditional on the unemployed having access to liquidity. In general, it is difficult to provide conditions on the primitives of the model to determine whether the optimal profile is decreasing or increasing. Our sufficient statistics approach, to the

contrary, provides a simple way to evaluate the dynamic profile of the UI policy, while accounting for the empirical importance of the two forces.

2.3.2 Duration-Dependence and Heterogeneity

While the theoretical literature has considered stationary environments to identify the forces for inclining and declining benefit profiles, a large empirical literature documents important non-stationarities during the unemployment spell: different job seekers leave unemployment at different rates and job market opportunities change over the unemployment spell.⁷ Two major challenges have arisen in this context. Empirically, it is a major challenge to disentangle and identify the role of selection and duration-dependence. Theoretically, it has proven very difficult to characterize the policy impact of different forms of non-stationarities during unemployment. Our analysis shows that a sufficient-statistics approach can overcome these two challenges.

Moral Hazard. The moral hazard cost may no longer be larger for benefits timed later in the spell when non-stationary forces reduce the responsiveness to incentives over the unemployment spell. Examples of such forces include “negative duration dependence”, referring to a decline in the exit rate probability over the unemployment spell. This decline can be due to skill-depreciation, stock-flow sampling of vacancies, or discrimination by employers against long-term unemployed. Dynamic selection may also reduce the responsiveness to incentives over the unemployment spell when less employable job seekers select into longer unemployment spells. When these forces are strong enough, they may offset the forward-looking channel, and push towards an inclining time profile of benefits.

Importantly, our approach can characterize the optimal benefit profile without having to disentangle negative duration dependence and selection effects. As discussed before, it is the budgetary impact of the behavioral responses to the benefit profile that is sufficient to evaluate the welfare cost, regardless of the forces underlying the differential responses.⁸

Consumption Smoothing. We established for a representative agent that the consumption smoothing gains of UI are weakly increasing over the unemployment spell. Heterogeneity will typically affect the gradient of consumption smoothing gains over the spell. Since it depends on how the average marginal utility of consumption evolves during the unemployment spell, selection effects may increase or decrease (and even revert) the gradient depending on the correlation between consumption levels and unemployment durations.

⁷See Machin and Manning [1999] for a review. See Schmieder et al. [2012] and Kroft et al. [2013] for recent examples.

⁸This does not mean that distinguishing between duration-dependence and heterogeneity is irrelevant for policy. First, how the unemployment elasticities change when changing the unemployment policy will crucially depend on the underlying forces. For example, more generous unemployment benefits paid to the short-term unemployed discourage the more elastic job seekers from leaving unemployment and thus increase the benefit elasticity of the long-term unemployed if this type of selection is important. Second, heterogeneity in job seekers’ responses becomes directly policy-relevant if the unemployment policy can condition on the source of heterogeneity.

Note that when there are selection effects in consumption, the average consumption drop will differ from the average “within”-individual consumption drop at different unemployment lengths. If the policy maker only cares about the insurance value provided by the unemployment policy and not about the redistributive value, one needs to identify the average “within”-individual consumption drop at each time t , which therefore requires controlling for selection effects in consumption. More generally, identifying selection effects underlying the consumption profile becomes important when individuals receive different policy weights or have different risk preferences.

3 Empirical Implementation: Context and Data

To implement our sufficient statistics approach and determine the optimal profile of UI benefits, two important pieces of empirical evidence are needed. First, one needs to identify and estimate responses of unemployment durations to variations in the benefit profile, i.e., variations in UI benefits at different points of an unemployment spell. Second, one needs to estimate the time profile of consumption to identify how consumption (relative to employment) drops over an unemployment spell.

Our empirical analysis offers contributions on both dimensions by using a unique administrative dataset that we created in Sweden combining unemployment registers and tax registers with exhaustive information on income and wealth. We present here the institutional background and data used in our empirical implementation.

3.1 Institutional background

In Sweden, displaced workers who have worked for at least 6 months prior to being laid-off are eligible to unemployment benefits, replacing 80% of their earnings up to a cap. In practice, the level of the cap is quite low relative to the earnings distribution and applies to about 50% of unemployed workers. Individuals can receive unemployment benefits indefinitely. 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.⁹

The time profile of benefits has changed during the period we study. Before 2001, the time profile of UI benefits was flat for all unemployed workers. Full-time workers would get daily benefits of 80% of their pre-unemployment daily wage throughout the spell, with daily benefits capped at 580SEK a day. The cap thus applies for daily wages above 725SEK.¹⁰ In July 2001, a system of duration-dependent caps was introduced, which created a decreasing time profile of benefits for the unemployed above the threshold wage. The cap for the benefits received during the first 20 weeks of unemployment was increased to 680SEK (daily wage above 850SEK) while the cap for benefits

⁹Like in other Scandinavian countries, UI in Sweden is administered by different unemployment funds (of which most are affiliated with a labor union) and contributions to the funds are voluntary in principle. In 2006, 90% of all workers were contributing to an unemployment fund.

¹⁰The daily wage is computed as gross monthly earnings divided by number of days worked in the last month prior to becoming unemployed.

received after the first 20 weeks was kept unchanged at 580SEK. In July 2002, the cap for benefits received during the first 20 weeks of unemployment was increased to 730SEK (daily wage above 912.5SEK) and the cap for benefits received after the first 20 weeks was increased to 680SEK.¹¹

The 2001 and 2002 reforms introduce variation in the benefit profile which makes it possible to estimate the causal impact of benefits received at different times during the unemployment spell on survival in unemployment. We explain in Section 4 how the 2001 and 2002 variations in the time-dependent caps can be used in a regression kink design to identify the effects on unemployment durations of UI benefits given in the first 20 weeks of a spell and of benefits given after 20 weeks. Importantly, because the 2001 and 2002 variations in UI benefits are introduced starting from a flat profile, the Swedish context gives us the possibility to evaluate the welfare consequences of local departures from a flat benefit profile. In other words, our empirical implementation can directly assess whether welfare is increased by switching towards a declining or inclining benefit profile.

3.2 Data

Unemployment history data come from the HÄNDEL register of the Public Employment Service (PES, Arbetsförmedlingen) and were merged with the ASTAT register from the UI administration (IAF, Inspektionen för Arbetslösheftsförsäkringen) in Sweden. The data contain information from 1999 to 2007 on the date the unemployed registered with the PES (which is a pre-requisite to start receiving UI benefits), eligibility to receive UI benefits, earnings used to determine UI benefits, weekly information on benefits received, unemployment status and participation in labor market programs. To define unemployment spells, we use the registration date at the PES as the start date and focus on individuals with no earnings who report to be searching for a full-time work. The end of a spell is defined as finding any employment (part-time or full-time employment, entering a PES program with subsidized work or training, etc.) or leaving the PES (labor force exit, exit to another social insurance program such as disability insurance, etc.).¹²

These data are linked with the longitudinal dataset LISA which merges several administrative and tax registers for the universe of Swedish individuals aged 16 and above. In addition to socio-demographic information (such as age, family situation, education, county of residence, etc.), LISA contains exhaustive information on earnings, taxes and transfer and capital income on an annual basis. Data on wealth comes from the wealth tax register (Förmögenhetsregistret), which covers the asset portfolio's for the universe of Swedish individuals from 1999 to 2007. The register contains detailed information on all financial assets (including debt) and real assets.¹³ For the

¹¹Some unions have launched their own complementary UI-schemes which further increased the cap (by up to 3 times the cap on regular UI) by topping up the regular UI-benefit to 80 percent of the previous wage. Importantly, our regression-kink design analysis focuses on the effect of the 725SEK-kink in the UI schedule, which was removed in 2002 before the introduction of the top-ups, so that all unemployed had to comply to the same kinked schedule of benefits.

¹²To deal with a few observations without any end date, we censor the duration of spells at two years.

¹³All financial institutions are compelled to report this information directly to the tax administration for the purpose of the wealth tax, which ensures quality and exhaustiveness of the data. The wealth tax was abolished in Sweden in 2007, after which the government collected only limited information on the stock of assets.

financial assets, we have information on all savings by asset class (bank accounts, bonds, stocks, mutual funds, private retirement accounts, etc.). The dataset also contains information on total outstanding debt including mortgage debt, consumer credit, student debt, etc. For real estate, we have information on all asset holdings at market value as used for the wealth tax assessment. All asset holdings are reported at the individual level. The comprehensiveness and detailed nature of both the income and wealth data in Sweden is exceptional, providing a unique opportunity to construct a registry-based consumption measure at the yearly level and to investigate what means individuals use to smooth consumption (transfers, asset rebalancing, increase in debt, etc.).

Finally, additional information on consumption is available through the yearly household budget survey (HUT, Hushållens Utgifter). From 2003 to 2009, individuals sampled in the HUT can be matched to the registry data, which allows us to compare registry-imputed and survey-based measures of consumption. The sample size of the HUT ($\approx 3,000$ households per year) is limited and HUT does not have a panel structure, but it has the advantage of offering direct measures of the flow of consumption on different items at the moment the household is surveyed. The HUT also offers complementary insights into what types of consumption (durables, non-durables, etc.) individuals adjust over an unemployment spell.

In Table 1, we provide summary statistics on unemployment, demographics, income and wealth for individuals observed in December of year n and about to become unemployed, i.e., who are observed starting an unemployment spell in the first quarter of year $n + 1$. From 1999 to 2007, the average unemployment spell is 26 weeks. The average time spent unemployed during the first twenty weeks of the spell (when the cap on benefits is potentially relaxed) equals $D1 = 12.6$ weeks. The average replacement rate is 72%. Roughly half of all unemployed individuals have a wage above the threshold at which the benefit cap applies.

Socio-demographic characteristics reported in Panel B of Table 1 show that the unemployed are relatively young (34 years old on average) and relatively few of the unemployed are married (23% on average). This young and unmarried group of the population may have more limited consumption smoothing opportunities on average.

Prior to the onset of a spell, the average unemployed has yearly gross earnings (before any tax or payroll contribution) of 151,000 SEK.¹⁴ Panel C of Table 1 also reveals that individuals about to become unemployed have almost zero capital income on average.¹⁵ The redistribution through the tax and transfer system is significant. The average disposable income (defined as earnings plus capital income minus all taxes paid plus all transfers received) for individuals about to become unemployed is 140,000 SEK, and much more equally distributed than gross earnings, with the median disposable income equal to 136,000 SEK, and the 25th percentile equal to 89,000 SEK.

The last panel of Table 1 provides information about individuals' asset positions prior to becoming unemployed. The average unemployed has a net worth at the onset of a spell equivalent to 124% of her yearly disposable income. Financial assets represent 66% of yearly disposable income,

¹⁴All figures are expressed in constant SEK2003.

¹⁵Capital income is here defined as all dividends and earned interests, plus realized capital gains on all private properties (real estate, financial assets, valuable commodities such as art or cars, etc.)

among which liquid bank account holdings represent only 21% of yearly disposable income. Other financial assets are predominantly stocks and assets held in mutual funds. Real estate (gross of mortgage debt) held by the average unemployed represents 228% of yearly disposable income on average. Total debt, which mostly comprises mortgage, student loans and credit card debt, is fairly large in Sweden and represents on average 170% of the yearly disposable income of an unemployed at the onset of her spell.

Importantly, Panel D reveals that the distribution of wealth at the onset of an unemployment spell is extremely skewed. Within the different asset classes, between 50% and 70% of individuals have no asset at all. In particular, more than 50% of individuals have no liquid assets on their bank accounts at the start of an unemployment spell. Overall, more than 50% of unemployed start an unemployment spell with no positive net wealth, and more than 25% even start a spell with negative net wealth. For a majority of individuals, existing assets thus offer very limited ability to smooth consumption over an unemployment spell.

4 Duration Responses

In this section we analyze unemployment responses to changes in the benefit profile. The presence of duration-dependent caps in the Swedish UI system provides compelling variation in UI benefits at different points in time during the unemployment spell. We exploit this variation using a regression kink (RK) design.

4.1 Regression-Kink Design: Strategy & Results

The time-dependent caps introduce kinks in the schedule of UI benefits given during the first 20 weeks of unemployment and after 20 weeks of unemployment. Figure 2 shows UI benefits as a function of daily pre-unemployment wages for spells starting in 1999 and up to July 2001 (panel A.1), for spells starting from July 2001 to July 2002 (panel B.1) and for spells starting after July 2002 (panel C.1). For spells starting before July 2001, the same cap applies to unemployment benefits given in the first 20 weeks of unemployment (b_1) and after 20 weeks of unemployment (b_2). The schedule of both b_1 and b_2 thus exhibits a kink at a daily wage of 725 SEK. This gives the opportunity to identify the effect of a joint change in b_1 and b_2 on unemployment durations. For spells starting after July 2001 and before July 2002, the cap for b_1 is increased, while the cap in b_2 remains unchanged. The relationship between b_1 and previous wages therefore becomes linear around the 725 SEK threshold, where the schedule of b_2 still exhibits a kink at 725 SEK. This makes it possible to identify the effect on unemployment duration of a change in b_2 only. Finally, the cap in b_2 is also increased for spells starting after July 2002, so that kinks in the schedule of both b_1 and b_2 disappear at the 725 SEK threshold. This offers a placebo setting to test for the robustness of our approach at the 725 threshold.

Our identification strategy relies on a RK design. Formally, we consider the general model:

$$Y = y(b_1, b_2, w, \mu),$$

where Y is the duration outcome of interest, μ is (unrestricted) unobserved heterogeneity, and b_1 , b_2 and w (previous daily wage) are endogenous regressors. We are interested in identifying the marginal effect of benefits given during part k of the spell on the duration outcome Y , $\alpha_k = \frac{\partial Y}{\partial b_k}$. The RK design consists in exploiting the fact that b_k is a deterministic, continuous function of the wage w , kinked at $w = \bar{w}_k$. The RK design relies on two identifying assumptions. First, the direct marginal effect of w on Y should be smooth around the kink point \bar{w}_k . Second, the density of the “forcing variable” w should also be smooth around \bar{w}_k . This second assumption implies imperfect sorting around the threshold \bar{w}_k , i.e., individuals do not have perfect control over their assignment in the schedule. We provide below several formal tests for the validity of these identifying assumptions and the robustness of our RK design.

Under these two identifying assumptions, α_k can be identified as:

$$\alpha_k = \frac{\lim_{w \rightarrow \bar{w}_k^+} \frac{\partial E[Y|w]}{\partial w} - \lim_{w \rightarrow \bar{w}_k^-} \frac{\partial E[Y|w]}{\partial w}}{\lim_{w \rightarrow \bar{w}_k^+} \frac{\partial b_k}{\partial w} - \lim_{w \rightarrow \bar{w}_k^-} \frac{\partial b_k}{\partial w}}$$

In practice, we provide estimates $\hat{\alpha}_k = \frac{\hat{\delta}_k}{\nu_k}$ where $\hat{\delta}_k$ is the estimated change in slope between Y and w at \bar{w}_k and ν_k is the deterministic change in slope between b_k and w at \bar{w}_k . We estimate the former using the following regression model:

$$E[Y|w] = \beta_0 + \beta_1(w - \bar{w}_k) + \delta_k(w - \bar{w}_k) \cdot \mathbf{1}[w \geq \bar{w}_k]. \quad (7)$$

This model is estimated for $|w - \bar{w}_k| \leq h$, where h is the bandwidth size.

Preliminary graphical evidence of a change in slope in the relationship between (total) duration of the unemployment spell and previous daily wage in response to the kink in UI benefits is provided in the right-hand side panels of Figure 2. They plot average unemployment duration in bins of previous daily wage for the three periods of interest. Panel A.II shows a significant change in the relationship between wage and unemployment duration around the 725SEK threshold for spells starting up to July 2001. In this period the schedule of UI benefits exhibits kinks in both b_1 and b_2 at 725SEK (as shown in Panel A.1). In panel B.II, a significant yet smaller change in slope can be detected at the 725SEK threshold for spells starting between July 2001 and July 2002 when the schedule at 725SEK exhibits a kink in b_2 only. Finally, panel C.II shows evidence of perfect linearity in the relationship between wage and unemployment duration around the 725SEK threshold for spells starting after July 2002, when kinks in the schedule at 725SEK are eliminated for both b_1 and b_2 .

RK estimates for the effects of benefits on unemployment duration D are shown in Figure 4, where we report for each policy period the point estimate and 95% robust confidence interval of

the change in slope $\hat{\delta}_k$ for the regression model in (7) using a bandwidth $h = 90\text{SEK}$ around the 725SEK threshold.¹⁶ We also report the implied benefit elasticity of unemployment duration. In line with the evidence presented in Figure 2, the estimated change in slope is large and significant for spells starting before July 2001. This implies an elasticity of unemployment duration w.r.t. an overall change in the benefit level (both b_1 and b_2) of $\varepsilon_{D,b} = 1.53$ (.13). The change in slope for spells starting between July 2001 and July 2002 is smaller but precisely estimated, and implies an elasticity of unemployment duration w.r.t. b_2 of $\varepsilon_{D,b_2} = .69$ (.14).

The same approach can be used to estimate the effect of benefits on the survival rate in unemployment at any spell length. Figure 5 shows the RK estimates for the effect of the benefit changes on the benefit durations D_1 and D_2 , where $D_1 = \sum_{t < 20wks} S_t$ is the time spent receiving benefit b_1 and $D_2 = \sum_{t \geq 20wks} S_t$ is the time spent receiving benefit b_2 . For both D_1 and D_2 , we find that the change in slope of the relation between wage and benefit duration is significant, but substantially smaller at the kink in b_2 than at the kink in both b_1 and b_2 . Note that the effect of benefits b_k on total duration D should be equal to the sum of the effect of b_k on durations D_1 and D_2 . Reassuringly, without imposing this restriction, we indeed find that the estimated change in slope in D equals the sum of the estimated slope changes for D_1 and D_2 .

The most common threat to identification and inference in the RK design is the presence of non-linearity that underlies the relationship between the assignment variable and the outcome, but is unrelated to the effect of the kinked policy schedule. To deal with this threat, one can implement a difference-in-difference RK design, as suggested in Landais [2015], using spells for which the UI schedule is linear at the threshold. We use the spells starting after July 2002 as a placebo and reject the presence of non-linearity around the 725SEK threshold. In line with the evidence presented in Panel C.II of Figure 2, the estimated change in slope is very close to zero and not statistically significant. This precisely estimated zero effect alleviates the concern that our estimates are spuriously capturing some non-linear functional dependence between wages and unemployment duration around the 725SEK threshold.

We provide several additional tests to assess the validity of the RK design and the robustness of the RK estimates in Appendix B. We start by providing supporting evidence in favor of the local random assignment assumption underlying the RK design. Figure 14 displays the probability density function of wages and reports tests in the spirit of McCrary [2008] that confirm smoothness of the pdf around the 725SEK threshold. Figure 15 provides evidence of smoothness in the relationship between observable characteristics of unemployed workers and wages at the 725SEK threshold. The sensitivity of the RK estimates to the size of the bandwidth is explored in Figure 16. The stability of the RK estimates across bandwidth sizes further alleviates the concern that the RK estimates pick up some underlying non-linearity in the relationship between wages and unemployment duration. We also explore the sensitivity of the results to the order of the polynomial used to fit the data around the 725SEK threshold. In Table 2, we report RK estimates where

¹⁶Because the change in caps applies to ongoing spells as well, we censor spell duration at their duration as of July 2001 and July 2002, and report here estimates from censored duration models. We checked on fully uncensored data that our censoring strategy does not affect the point estimates.

we estimate quadratic and cubic (rather than linear) models for equation 7. Estimates are stable across the different specifications, but get less precise as the order of the polynomial increases.

We report White robust standard errors for inference and explore sensitivity to alternative strategies in Table 2. In particular, we report 95% confidence interval based on permutation tests as in Ganong and Jaeger [2014]. Interestingly, due to the linearity in the relationship between unemployment duration and wages across the whole support of the assignment variable, such confidence intervals often appear much tighter than those based on robust standard errors. Finally, in Figure 18 we perform tests aimed at detecting non-parametrically the presence and location of a kink point in the relationship between unemployment duration and wages, as suggested in Landais [2015]. All these tests strongly support the conclusion that there is a change in slope that occurs right at the actual kink point in the UI schedule.

4.2 Implications for Moral Hazard Costs

Our results carry important implications for the moral hazard costs of modifying the time profile of UI benefits. Importantly, because our estimates are based on local departures from an initially flat benefit profile ($b_1 = b_2$), this gives us the possibility to evaluate the moral hazard costs of switching towards a declining or inclining benefit profile.

First, the moral hazard cost of the Swedish unemployment policy is large overall. For a flat profile, the moral hazard cost from increasing the benefit level throughout the unemployment spell (i.e., increasing both b_1 and b_2) simply equals the elasticity of the average duration. That is, $MH = \varepsilon_{D,b} = 1.53 (.13)$. Because of behavioral responses when increasing all benefits by 1%, the planner would need to levy from the employed 1.53 times more resources to balance its budget than the implied static cost absent behavioral responses. While these estimated elasticities with respect to the UI benefit level imply a large incentive cost of the UI system, to gauge the magnitude of the implied behavioral response it is more appropriate to re-express these elasticities with respect to the net of replacement rate $1 - r$, as advocated in Landais et al. [2010].¹⁷ Changing benefits b by 1% percent is equivalent to changing the net of replacement rate $1 - r$ by $r/(1 - r)$ percent. Given the high replacement rate $r = .8$ at the kink, our estimated elasticity $\varepsilon_{D,b} = 1.53$ translates into a much lower estimated elasticity w.r.t. $1 - r$. That is, $\varepsilon_{D,1-r} = -\frac{1-r}{r}\varepsilon_{D,b} = -.38$.

Second, our results suggest that the moral hazard cost of an increase in benefits is lower when timed later in the spell. As discussed before, for a flat profile, the moral hazard cost MH_k simplifies to the elasticity of the average duration scaled by the relative time spent on the policy b_k . That is,

$$MH_k = \varepsilon_{D,b_k} \frac{D}{D_k} \text{ when } b_1 = b_2.$$

The moral hazard cost of increasing benefits b_2 after 20 weeks of unemployment is therefore $MH_2 = \frac{D}{D_2}\varepsilon_{D,b_2} = 1.38 (.27)$. Compared to $MH = \varepsilon_{D,b} = 1.53 (.13)$, this implies that the incentive cost

¹⁷First, $1 - r$ is the relevant wedge due to UI from a theoretical standpoint. Second, this makes unemployment elasticities comparable to labor supply elasticities estimated in the taxation literature.

of increasing benefits after 20 weeks is smaller than the incentive cost from increasing benefits throughout the unemployment spell. A second implication is that the incentive cost is larger for increasing benefits in the first 20 weeks than after. In fact, we can use our estimates to back out the elasticity of unemployment duration with respect to a change in b_1 only: $\varepsilon_{D,b_1} = .83$ (.18).¹⁸ This then implies that $MH_1 = 1.67$ (.37).

Our estimates thus indicate that the moral hazard cost of increasing benefits for the first 20 weeks is larger than that of increasing benefits after 20 weeks of unemployment. That is, to increase benefits by 1% for the first 20 weeks of unemployment while balancing its budget, the planner would need to levy from the employed 1.67 times more resources than the implied static cost because of behavioral responses. This is 21 percent larger than the moral hazard cost of increasing benefits after 20 weeks of unemployment. While our way of estimating the response to changes in benefits early in the spell is somewhat indirect, we provide additional evidence that the effect of b_1 on unemployment duration is larger if anything in Figure 20 in appendix. We exploit the fact that for spells starting between July 2001 and July 2002, there is a kink in b_1 (and not b_2) at the 850SEK threshold that identifies the effect of b_1 on unemployment duration.¹⁹

The fact that the moral hazard cost of unemployment benefits decreases with the spell length is somewhat surprising. As discussed in Proposition 2, in a stationary environment, one would expect forward-looking incentives to play in the opposite direction. Indeed, our results unambiguously show that unemployed individuals are forward-looking. The estimated elasticity of D_1 , the duration spent on the first-part of the profile, with respect to benefits b_2 received in the second part of the profile, reported in Figure 5, is positive and significant ($\varepsilon_{D_1,b_2} = .79$ (.13)). Unemployed individuals are not fully myopic: their instantaneous hazard rate out of unemployment does react to variation in benefits later in the spell.

This implies that there is substantial non-stationarity as well. Evidence of such non-stationarity can directly be found by comparing, for different durations t , the elasticity of \tilde{D}_t with respect to a flat benefit level b , where \tilde{D}_t is the remaining duration of unemployment at t , conditional on surviving until t . In a stationary environment, starting from a flat profile, one would expect these elasticities to be constant as a function of t , and thus equal to $\varepsilon_{D,b}$, the elasticity of the total unemployment duration with respect to benefits throughout the spell, as discussed in Section 2.3.1.²⁰ Results, reported in Figure 6, show that $\varepsilon_{\tilde{D}_t,b}$ strongly declines as a function of t , which means that the responsiveness of exit rates decreases substantially over the unemployment spell.

¹⁸We simply use the fact that $\varepsilon_{D,b} = \varepsilon_{D,b_1} \frac{b}{b_1} + \varepsilon_{D,b_2} \frac{b}{b_2} = \varepsilon_{D,b_1} + \varepsilon_{D,b_2}$ for $b_1 = b_2$.

¹⁹Although this kink in the schedule of b_1 offers a more direct source of variation to identify MH_1 , it does not evaluate MH_1 for the same sample of unemployed, nor does it evaluate MH_1 at a flat profile as individuals at the 850SEK face a declining time profile of benefits. Note also that for the period 1999-2000, when the schedule of both b_1 and b_2 is linear at the 850SEK threshold, the placebo test does not pass at the 850SEK kink. We therefore implement a strict DD-RKD in this case, where the effect of b_1 is identified from the change between 2000 and 2001 in the change in slope of unemployment duration at the 850SEK wage threshold following the introduction of the

kink in the schedule of b_1 . Formally, this implies $\alpha^{DD} = \frac{\Delta_{2000,2001} \delta_k}{\Delta_{2000,2001} \nu_k}$.

²⁰Note that in the stationary environment considered in Section 2.3.1, exit rates only depend on the continuation policy, regardless of the unemployment duration. This implies $\varepsilon_{\tilde{D}_t,b} = \varepsilon_{\tilde{D}_t,b_2}$.

Such non-stationarity, which can be the result of dynamic selection and/or duration dependence, is large enough to offset the significant effect of forward-looking incentives, so that the incentive cost of increasing the generosity of long-term benefits is still somewhat smaller than that of increasing the generosity of short-term benefits.

Our results thus suggest that for the estimation of behavioral responses to variation in UI and evaluation of associated MH costs, the timing of the benefit variation critically matters. The duration responses to a 1% change in UI benefits throughout the spell, or to a 1% change in UI benefits for the first 26 weeks only, or to a 1% extension in the potential duration of benefits, are likely to be significantly different. This point is important to keep in mind when trying to compare different estimates of the behavioral responses to UI found in the existing literature. Interestingly, our results are consistent with earlier literature, and may also help reconcile some apparent discrepancies in earlier findings. Like Meyer [1990] and Landais [2015], who focus on variation in benefit level in the US for the first 26 weeks of the unemployment spell, we find that the moral hazard cost of changing UI benefits early on in the spell is large. But like Schmieder et al. [2012], Rothstein [2011] or Valla and Farber [2011], who focus on extensions in the potential duration of benefits, we find that the effect of increasing benefits later on in the unemployment spell is somewhat smaller.²¹

5 Consumption Smoothing Over the Unemployment Spell

In this section we discuss the construction of our registry-based consumption measure and use it to analyze how consumption evolves during the unemployment spell.

5.1 Residual Consumption Measure

We start from the accounting identity that expenditures in period n are the sum of all income and transfers received in period n , minus the change in assets between period $n - 1$ and period n ,

$$\text{expenditures}_n = \text{income}_n - \Delta \text{assets}_n.$$

As a result of the comprehensiveness of the longitudinal administrative dataset that we assembled including all earnings, income, taxes, transfers and wealth, we have precise third-party reported information on all the components needed to construct such residual measure of yearly expenditures for the universe of Swedish individuals and households for years 1999 to 2007. Our approach is closely related to Koijen et al. [2014] who constructed a similar measure in Sweden for years 2003 to 2007 using a smaller subset of individuals, and confirmed its consistency with HUT data.²² In

²¹Note however that in a recent paper, Mas and Jonhston [2015] find relatively large effect from potential duration reductions. Consistent with what we find is that the effect is mostly driven by exit rate responses early in the spell rather than late in the spell.

²²The registry-based data provides detailed information on the sources of expenditures, but only limited information on the type of consumption goods (e.g., durables vs. non-durables) that money is spent on. To address this limitation, we exploit information from the consumption surveys (HUT) where we do observe the type of goods that the household

practice, we compute consumption in year n as:

$$C_n = y_n + T_n + \tilde{C}_n^b + \tilde{C}_n^d + \tilde{C}_n^v + \tilde{C}_n^h, \quad (8)$$

where:

- y_n represents all earnings and is computed from the tax registers, which contain third-party reported earnings for all employment contracts, including all fringe benefits and severance payments.²³
- T_n accounts for all income taxes and transfers, including unemployment insurance, disability insurance, sick pay, housing and parental benefits, etc.
- $\tilde{C}_n^b = y_n^b - \Delta b_n$ equals consumption out of bank holdings. It is equal to interests earned on these bank holdings during year n , y_n^b , minus the change in the value of bank holdings between year $n-1$ and year n , $\Delta b_n = b_n - b_{n-1}$. An increase in the value of bank holdings net of earned interests implies positive savings or conversely negative consumption out of bank holdings.²⁴
- $\tilde{C}_n^d = -y_n^d + \Delta d_n$ is consumption out of debt, which includes student loans, credit card debt, mortgages, etc., and is third-party reported by financial institutions to the tax authority. It is equal to the change in the stock of debt Δd_n , minus all interests paid on the existing stock of debt y_n^d . An increase in debt net of interests paid contributes positively to consumption expenditures, while deleveraging contributes negatively to consumption.
- $\tilde{C}_n^v = y_n^v - \Delta v_n$ is consumption out of financial assets (other than liquid holdings in bank accounts). It is equal to all income from financial assets y_n^v minus the change in the value of the portfolio of financial assets Δv_n . Each year financial assets need to be reported by financial institutions at their market value as of December 31, which makes it possible to compute v_n very accurately.²⁵ The return on financial assets y_n^v includes interests, dividends

consumes, but for a very small sample of the population. We discuss in more detail in Section 6.2 how the distinction between expenditures and consumption can be relevant for our welfare conclusions, in particular when the type of consumption changes over the unemployment spell.

²³Note that self-employed, for whom earnings are in large part self-reported, are excluded from the analysis as they are part of a different UI system.

²⁴Information on bank holdings and earned interests are reported directly by banks to the tax authority. Until 2005, positive balances are reported if the interest income during that year was greater than 100SEK (roughly \$15). After 2005, the balance of bank accounts is reported if it is greater than 10kSEK (roughly \$1,500), implying a comparable censoring level (given an average interest rate of about 1%). This may introduce a small bias in our estimates of consumption out of bank holdings, but the sign of this bias is a priori unknown. For individuals starting a spell with censored bank holdings (b_n below the censoring point), we may slightly overestimate the drop in consumption if they use some of this to smooth consumption. To the contrary, for individuals with positive observed holdings (b_n above the censoring point), but b_n dropping below the censoring point during the spell, we may slightly underestimate the drop in consumption if they do not consume all of their holdings.

²⁵It is possible that our measure misses some financial wealth held in foreign banks to the extent that these banks do not comply with the requirement to transmit information on the financial wealth of their Swedish customers to the Swedish tax authority. Yet, the fraction of foreign-held assets is low (about 3% of household assets according to

and any price change $\Delta p_n^v \times q_n^v$. Such price change would be exactly offset by a change in the value of assets, included in Δv_n , unless the return is realized by selling the asset.²⁶ In practice, we observe all interests paid and dividends received. Computing exact price changes requires precise disaggregated information on all prices and quantities for all assets held. Our approach consists in imputing price changes by asset class using the average price change for each class of asset.²⁷ This strategy of imputing prices by asset class rather than at the individual asset level may introduce a bias in our estimates of consumption profiles if price changes at the individual level are correlated with the duration of unemployment, or in other words if individuals who stay unemployed longer have, conditional on asset class, lower returns on their assets. Since only a tiny fraction of unemployed hold such financial assets, this bias is likely to be very small.

- $\tilde{C}_n^h = y_n^h - \Delta h_n$ constitutes consumption out of real estate wealth. It is equal to all income derived from holding real estate assets y_n^h minus the change in the value of real estates Δh_n . Detailed information on the stock of real estate wealth, estimated at market value as of December 31 of each year, is available from the tax authority. The return to holding real estate y_n^h includes rents, but also imputed rents for homeowners, as well as price changes in the value of real estates. Such price changes constitute a return to holding real estate, and should be counted as income. However, like for financial assets, only by selling the asset, this return generates positive consumption. We impute price changes based on the average price change of each class of real estate assets, using data on real estate prices available from Statistics Sweden.²⁸ For imputed rents, our approach consists in using the average yearly rental yield in Sweden to compute a counter-factual rent for homeowners living in their own property.²⁹ Because the large majority of unemployed individuals are renters and only few have real estate wealth, these assumptions are unlikely to significantly affect our estimated consumption profiles. Yet, to assess the robustness of our approach, we also explore different

the Savings Barometer of Statistics Sweden), and likely to be held by much wealthier households than our sample of unemployed.

²⁶Conceptually, the change in the price of any financial asset held constitutes a return to holding this asset, whether this price change generates a realized income flow or is left unrealized. However, the price change will generate an actual consumption flow only if this return is realized by selling the asset. If this return is not realized, then this unrealized income is “saved”, which translates into an increase in the market value of the asset equivalent to the value of the unrealized income. The overall effect on consumption expenditure will thus be exactly zero.

²⁷Data on prices of various asset classes comes from Statistics Sweden (see <http://www.statistikdatabasen.scb.se/pxweb/en/ssd/>) For stocks, we also used information from Ekonomifakta.se, a website run by The Confederation of Swedish Enterprise (Svenskt Näringsliv), Sweden’s largest and most influential business federation, which compiles detailed information on stock prices (see <http://www.ekonomifakta.se/sv/Fakta/Ekonomi/Finansiell-utveckling/Borsutveckling-i-Sverige-och-USA/>). For the evolution of prices of assets held in mutual funds, we used data from Fondbolagen.se, a website run by the Swedish Investment Fund Association (Fondbolagen), covering approximately 90 per cent of the net fund assets held in the Swedish market. For the so-called capital insurance accounts, we do not observe the asset composition and use a weighted index of stocks, bonds and mutual funds. For other private pension savings, we do not need to impute price changes as we observe contributions and withdrawals, which directly determine the consumption flow.

²⁸<http://www.statistikdatabasen.scb.se/pxweb/en/ssd/>

²⁹Again, this may introduce a small potential bias in our estimates of consumption profiles if rental yields are correlated with unemployment duration.

sample restrictions, and consider consumption profiles of renters only, as well as consumption profiles of individuals without any real wealth at the start of a spell.

5.2 Unemployment Duration and Yearly Consumption

In order to analyze the effect of unemployment duration on consumption, we start by constructing a sample of all individuals unemployed as of December of year n , or who will become unemployed in year $n + 1$, for all years n from 1999 up to 2007. We then correlate our yearly consumption measure, which records total yearly expenditures between December of year $n - 1$ and December of year n , with the time t (in quarters) spent unemployed since the onset of the unemployment spell. Negative values for t represent the time remaining until the start of the unemployment spell.³⁰ In practice, we run the following regression:

$$C_{it} = \sum_{t=-3}^{+8} \beta_t \cdot \mathbf{1}[T = t] + X_i' \gamma + \varepsilon_{it} \quad (9)$$

where $\mathbf{1}[T = t]$ is an indicator for being observed in the t -th quarter of the ongoing unemployment spell. We include in this regression a set of controls X , which consists of year dummies, age dummies, a dummy for being just out of school in year n , and a set of dummies for family status.

We report in Figure 7 the estimated coefficients β_t from regression model (9). These coefficients represent the average yearly consumption levels in constant SEK for individuals observed in their t -th quarter of unemployment, relative to the yearly average consumption level of individuals observed just one quarter prior to becoming unemployed. The graph provides clear evidence that unemployment causes a substantial drop in consumption. The average yearly consumption of individuals who have been unemployed for a full year ($t = 5$ quarters) is almost 40kSEK lower than the average consumption of individuals at the start of their spell. Given that the yearly consumption of individuals observed just prior to becoming unemployed is 137kSEK, this means that yearly consumption levels drop by 26% on average after a year of unemployment. The second interesting finding is the lack of anticipation prior to becoming unemployed. Consumption is perfectly flat in the four quarters preceding the onset of the unemployment spell. This suggests that unemployment shocks are relatively unanticipated, or that individuals have little ability to change consumption prior to becoming unemployed to smooth out the upcoming earnings shock. Finally, the profile of yearly consumption after the start of the unemployment spell reveals important information about the dynamics of flow consumption during the spell. For the first four quarters, yearly consumption drops significantly and almost linearly. As yearly consumption for individuals observed in the first four quarters of unemployment mixes consumption realized while still employed with consumption realized while unemployed, this sharp decline indicates that the flow consumption level while unemployed must be significantly lower than the flow consumption level while employed. From

³⁰To ensure that yearly consumption observed in quarters $t = -3$ to $t = 0$ only aggregates consumption flows while employed, we impose the further restriction that individuals must not have experienced any unemployment spell in the past two years prior to becoming unemployed.

the fifth quarter on, individuals have spent their last four quarters unemployed, and their yearly consumption therefore only aggregates flow consumption levels while unemployed. The fact that yearly consumption levels are relatively flat after the fifth quarter suggests that flow consumption decreases only slowly later on in the unemployment spell.

Decomposition. In Figure 8 we provide further insights into the means used by individuals to smooth consumption over the unemployment spell by decomposing yearly consumption into four core components, and by correlating the evolution of these components with unemployment duration. We follow the same methodology as in Figure 7, and plot the estimated coefficients from regression model (9) for each of the four outcomes. In Panel A, we plot the evolution of yearly gross earnings as a function of unemployment duration. Unemployed individuals lose all their earnings so that their yearly earnings are exactly 100% lower after one year of unemployment compared to their pre-unemployment level. In panel B, we add all transfers (UI, parental and housing benefits, social assistance, etc.). Transfers are quintessential in smoothing consumption: the drop in earnings plus transfers is just 27% after a year. In panel C, we add other income, including capital income, and subtract taxes. In panel D, we add changes in assets and debt, which gets at our residual consumption measure. Both additions have almost no effect on the average consumption profile. Figure 8 shows unambiguously that capital income and changes in assets have but a marginal role in smoothing consumption over the unemployment spell. Most of this role is played by transfers, and in particular by UI benefits. Overall, these consumption patterns indicate that unemployed individuals are on average close to being hand-to-mouth.

Asset Rebalancing. The minority of individuals who have some assets prior to becoming unemployed do use them to smooth consumption over the unemployment spell. To document this, we report in Figure 9 the evolution of consumption out of various assets and debt over the unemployment spell, following the same methodology as in Figures 7 and 8. Panel A shows that workers are increasing their bank account savings by 5kSEk in the year before unemployment, but these average contributions drop to zero during the first year of unemployment. This reduction in savings increases consumption and thus reduces the consumption drop relative to the pre-unemployment level by nearly 5% on average. Panel B indicates that the other financial assets provide limited consumption smoothing benefits in the first year of unemployment. The reduction in the consumption drop after one year from all financial assets (including bank accounts) is only 4%, but increases slightly for longer spells. Panel C shows that the unemployed also stop investing in real assets after displacement. The resulting increase in consumption, however, is more than offset by the reduction in debt increases, as shown in Panel D. The unemployed indeed consume substantially less from debt increases. Interestingly, this is also true for individuals not owning real estate, suggesting that the reduction in debt increases is not only driven by a decline in mortgage-financed purchases of real estate but also by reduced access to borrowing. In fact, as shown in appendix Figure 22, the consumption of individuals not owning any real estate falls by 5% over the first year of unemployment due to reduced access to non mortgage-related credit.

Duration Dependence. The estimated drop in consumption in our baseline Figure 7 does not separate the relative contribution of dynamic selection vs. within-individual duration dependence. Dynamic selection may occur if individuals who stay unemployed longer consume on average at significantly different levels compared to individuals who exit unemployment quickly. Within-individual duration dependence corresponds to the within-individual drop in consumption, controlling for individual consumption fixed effects. Even though, as discussed in Section 2.3.2, the relative contribution of dynamic selection vs. within-individual duration dependence is not needed for designing the optimal benefit profile, it is nevertheless interesting to disentangle the relative magnitude of these two forces. To do so, we take advantage of the panel structure of the data, and estimate, using a within-estimator, fixed-effect models of the form:

$$C_{it} = \sum_{t=-3}^{+8} \tilde{\beta}_t \cdot \mathbf{1}[T = t] + X_i' \gamma + \alpha_i + \mu_{it} \quad (10)$$

Estimates of the within-individual drop in yearly consumption $\tilde{\beta}_t$ in quarter t of the unemployment spell are reported in Figure 10. The pattern of yearly consumption during the unemployment spell controlling for selection is very similar to that of Figure 7. The average within-individual drop in consumption after one year of unemployment equals 25.3% (3.2), which accounts for almost all of the drop in the average consumption level and suggests that the role of selection effects on the profile of consumption levels is fairly limited on average.

Household Consumption. So far our analysis has focused on individual level consumption as all earnings, income and asset level information are reported at the individual level. It is nevertheless interesting to explore how much consumption insurance unemployed individuals receive within their household. To do so, we use household identifiers constructed by Statistics Sweden for the universe of the Swedish population, and compute individual residual measures of consumption for all members aged above 18 of the household in which an unemployed individual lives.³¹ We then sum these measures at the household level to obtain residual household consumption measures for all unemployed in our sample. Using the same methodology as for the individual consumption charts, we report in Figure 11 the estimated evolution of yearly household consumption over the unemployment spell based on equation (9). Two important findings emerge from the results.

First, in absolute terms, the yearly household consumption drop is extremely similar to the drop in yearly individual consumption. If anything the drop in consumption after a year in absolute terms for the household as a whole ($\approx 43\text{kSEK}$) is slightly larger than for individual consumption ($\approx 37\text{kSEK}$). This suggests that, compared to single individuals, the household as a whole offers little additional means to smooth consumption during unemployment. In Figure 21 in appendix, we show further evidence that other members of the household do not provide extra consumption opportunities over the unemployment spell. In Panel A, we report for all unemployed individuals the evolution of total disposable income of all other members of their household as a function of

³¹Statistics Sweden identifies couples as part of the same household if they are married or have children together.

the time they spent unemployed. Interestingly, this declines slightly as a function of time spent unemployed, suggesting that individuals who stay unemployed longer, live with individuals who have slightly lower disposable income compared to the average pool of unemployed. In Panel B, we control for this selection effect by estimating a fixed effect model. Results show that after a year, the within-household change in the sum of disposable income of all other members of the household is not significantly different from zero. This suggests that in our context, the added-worker effect is not playing any significant role in increasing household consumption in response to an unemployment shock.

The second lesson of Figure 11 is that the household consumption drop in percentage terms is smaller than the individual consumption drop, as the drop in absolute term is roughly similar but the denominator is much larger (household consumption prior to unemployment is 271 kSEK). After a year spent unemployed, yearly household consumption drops by 15.8% compared to 26.3% for individual consumption. Conceptually, the individual estimate offers an upper bound on the drop in consumption of unemployed individuals, which is tight when they do not receive any consumption insurance from other members of their household. Conversely, the drop in household consumption offers a lower bound on the individual drop in consumption, assuming perfect income pooling within the household. Note of course that when there is consumption insurance within the household, other members of the household will also experience a drop in consumption over the unemployment spell. This is likely to be reflected in higher risk aversion parameters, and by consumption externalities at the household level that the social planner should take into account, an issue we come back to in section 6.

5.3 Recovering High-Frequency Consumption

The analysis so far has focused on yearly consumption. As explained in Section 2, in order to evaluate the consumption smoothing benefits of benefits b_1 and b_2 , we need to recover higher frequency consumption measures, such as monthly consumption flows. To do so, we use the fact that C_t , the yearly consumption measured at month t of the unemployment spell, is simply the sum of the monthly consumption flows between month $t - 11$ and month t :

$$C_t = \sum_{k=t-11}^t c_k$$

Our strategy then consists in specifying a (somewhat flexible) functional form for the monthly consumption flows c_t as a function of t . Once a parametric model for c_t has been specified, one can use the information from yearly consumption C_t to identify the parameters of c_t . We use standard distance minimization techniques to recover the parameters from different flow consumption models.³²

In Panel A of Figure 12, we show the monthly consumption flow as a function of time spent

³²All details regarding the implementation of this strategy are reported in Appendix C.

unemployed for three different models: linear consumption with a drop at unemployment, quadratic consumption with a drop at unemployment, and a step function. All models deliver very similar results: consumption drops significantly right at the onset of the unemployment spell, and then continues to decline, but at a relatively slow rate, throughout the spell. Using the step-function as a baseline, this translates into an estimated consumption wedge $\Delta C_1 = 20\%$ between the average consumption of unemployed individuals in the first part of the Swedish UI schedule (up to 20 weeks) and the average consumption of individuals prior to becoming unemployed. For unemployed individuals in the second part of the UI profile (after 20 weeks), the estimated wedge between their average consumption and the average consumption prior to becoming unemployed is somewhat larger: $\Delta C_2 = 27\%$. This confirms earlier evidence showing that unemployed individuals have very little means of smoothing consumption over the unemployment spell, and therefore experience quite significant drops very early on in the spell. After this sharp initial drop, the flatness of the UI profile, and the quite generous replacement rate of the UI system in Sweden ensures that consumption does not drop much further over the unemployment spell.

Identification in Panel A of Figure 12 assumes that the consumption profile $\{c_t\}$ is identical for all unemployed, irrespective of the duration t at which their yearly consumption is observed. This is equivalent to assuming that consumption profiles $\{c_t\}$ are not selected on unemployment duration. In Panel B of Figure 12, we relax this assumption by estimating different consumption profiles for individuals with different realized unemployment durations. We estimate separately the step function profile for individuals whose total realized unemployment duration is less than 5 months and for individuals whose total realized unemployment duration is larger than 5 months.³³ While the profile for individuals who remain unemployed less than 5 months is not very precisely estimated, Panel B suggests that there is a little bit of dynamic selection, and that the drop in consumption is slightly milder for individuals who experience short spells than for individuals who experience longer spells. Using the profiles from Panel B, we can re-weight each profile by the fraction of individuals observed at each duration t who will be unemployed less than 5 months and more than 5 months to recover aggregate consumption drops $\Delta C_1 = .19$ and $\Delta C_2 = .27$. Reassuringly, these aggregate consumption drops are very similar to that obtained in Panel A.

In Panels C and D of Figure 12 we investigate heterogeneity in consumption profiles across age groups, and across individuals with different initial wealth level. In Panel C, we report estimated monthly consumption profiles for individuals who are less and more than 35 years old at the start of a spell. Younger individuals experience a significantly larger drop in consumption at the start of the spell compared to older unemployed. This suggests that older individuals have more means to smooth consumption early in the unemployment spell. Interestingly, although the initial drop at the start of the spell is much attenuated for older unemployed, their consumption profile quickly converges with that of younger unemployed, suggesting limited ability to smooth consumption over spells longer than 5 months. Panel D plots the estimated monthly consumption profiles for

³³This procedure is made possible by the fact that at each duration t we observe the consumption levels of individuals who will in practice have different total realized unemployment duration.

individuals who start their unemployment spell with positive net wealth versus zero or negative net wealth. It confirms that individuals in the latter group have extremely limited access to credit in order to smooth consumption and are essentially hand-to-mouth. Their consumption drops immediately at the start of the spell and remains flat afterwards. Individuals who have some positive net wealth draw down their assets to smooth consumption, which greatly attenuates the drop in consumption early in the spell. Yet, the average wealth level is not sufficient to cover spells longer than 5 months so that consumption drops quickly afterwards.

We reach similar conclusions regarding the consumption profile when using data from the consumption surveys (HUT) instead of our registry-imputed measure. Since the HUT surveys collect information on household consumption expenditures at the time of the interview, it is possible to directly recover flow (bi-weekly) measures of consumption c_t , and avoid the indirect method used above.³⁴ By matching HUT data with PES information, we are able to precisely determine the labor market situation of all individuals in households surveyed in the HUT.³⁵ Given the limited number of observations in HUT surveys, we run the simple regression:

$$c_{it} = \eta_1 \cdot \mathbf{1}[0 < t \leq 5] + \eta_2 \cdot \mathbf{1}[t > 5] + X_i' \gamma + \varepsilon_{it}, \quad (11)$$

where η_1 (resp. η_2) captures the effect of having a member of the household unemployed for less than 5 months (resp. more than 5 months) at the time of the interview, relative to households where one individual will become unemployed in the next 12 months.

Results, reported in column (1) of Table 3, suggest that, relative to pre-unemployment levels, total average surveyed household consumption drops by 5% in the first 5 months of unemployment, and by 13% after 5 months of unemployment. Although not very precisely estimated due to the small sample size, these results prove very similar to results on household-level consumption obtained using our registry measure of expenditures. Using our method to recover monthly household-level consumption from the registry measure of expenditures, we find that the average drop in household consumption in the first 5 months of unemployment is equal to 9%, while the drop after 5 months is equal to 18%.

The HUT also offers insights into the type of consumption goods that households adjust over the spell. First, committed expenditures do not seem to vary significantly over the unemployment spell. In column (3) of Table 3, we show that housing rents paid by renters do not seem to decline significantly, neither early, nor later in the unemployment spell. To the contrary, consumption of non-committed goods, such as food, recreation, transportation, or restaurants, drops significantly and is likely to further decrease over the spell (columns (2),(6),(7) and (8)).³⁶ Finally, expenditures

³⁴Information on consumption in HUT comes from shopping books where surveyed households are asked to report all their expenditures during the 2 weeks preceding their interview date. Note however that for a few durable goods, extra information is added from answers to interview questions, where individuals are asked to recollect some expenditures for up to 12 months prior to the interview date.

³⁵We restrict the sample to households where, at the date of the interview, one (and only one) individual is unemployed, or where, at the date of the interview, one (and only one) individual will become unemployed in the following 12 months.

³⁶Gruber [1997] considers only food consumption using data from the PSID for the period 1968-1987 and finds an

on durable goods such as the purchase of new vehicles or the purchase of furniture and home appliances, exhibit an interesting pattern (columns (4) and (5)). They decline very strongly early during the spell, but increase later during the spell, yet remaining largely below their pre-unemployment level. This strong substitution away from expenditures on durable goods early on in the spell has potential consequences for our welfare analysis, an issue we come back to in the next section.

6 Welfare Analysis

In this section we bring our theoretical and empirical analysis together to evaluate the optimality of the benefit profile in Sweden. We estimate the welfare impact of changes in the policy profile and discuss the robustness of the implied policy recommendations.

6.1 Unemployment Benefit Profile

The welfare consequences of an increase in benefits during the first 20 weeks of unemployment, after 20 weeks of unemployment and throughout the unemployment spell are reported in the different rows of Table 4. The different components of the welfare effect are shown in columns.

The first column repeats the moral hazard cost estimates obtained from the RKD benefit elasticities estimated in Section 4. As discussed before, our estimates imply that the moral hazard cost is not higher, but slightly lower for benefits paid later in the spell. In particular, the marginal cost of making the policy more generous is 21 percent higher during the first 20 weeks of unemployment than after 20 weeks of unemployment.

We then compare the moral hazard costs to the consumption smoothing gains. In column (2), we use estimates from our preferred specification of the consumption profile (Panel B of Figure 12) and report the average consumption drop over the full unemployment spell, the average consumption drop during the first 20 weeks of unemployment and the average consumption drop for individuals unemployed after 20 weeks. To evaluate the welfare impact of the consumption drop and be comparable to the moral hazard cost, the drop needs to be scaled by $\tilde{\lambda}_k \times \tilde{\gamma}_k$. Remember that in a representative agent model this scalar simply equals the relative risk aversion.

Putting the estimates of the CS gains and the MH costs together, we find that the MH costs are substantially larger than the CS gains at conventional levels of relative risk aversion ($\gamma < 2$). This is true regardless of the timing and thus suggests that the unemployment policy is too generous throughout the unemployment spell.³⁷ Such conclusion about the benefit levels, however, are sensitive to the assumptions on the scalars $\tilde{\lambda}_k$ and $\tilde{\gamma}_k$ and on the risk preference parameter γ in particular, which is known to vary substantially across context (see Chetty [2008], Chetty [2009]). Importantly, our estimates allow us to evaluate changes in the benefit profile as well, asking whether

average drop of 6.8% in the first year of unemployment, which is very similar to our estimates.

³⁷Note that consumption drops are estimated relative to the average consumption level just before the start of the unemployment spell. The average consumption level of all employed individuals in Sweden during our estimation period is about 33% higher than the average consumption level of individuals about to become unemployed, which would further increase estimates of the consumption wedge between employment and unemployment.

some part of the policy should be made more generous relative to the other part. Rather than on the average risk aversion, such evaluation depends on potential differences in risk aversion (or Pareto weights) at different unemployment lengths.

Column (3) of Table 4 shows that the value of a (tax-funded) kroner spent on benefits during the first 20 weeks of unemployment equals $CS_1/MH_1 = \tilde{\lambda}_1 \times \tilde{\gamma}_1 \times .11$. When spent after 20 weeks the value equals $CS_2/MH_2 = \tilde{\lambda}_2 \times \tilde{\gamma}_2 \times .20$. This is 81% percent higher than CS_1/MH_1 assuming that the relevant scalar is the same for short-term and long-term unemployed. Under that assumption, an increase in the benefits after 20 weeks implemented jointly with a decrease in the benefits before 20 weeks would thus create a welfare surplus of 81 percent at the margin. Since we evaluate these local policy changes starting from a flat benefit profile, this indicates that welfare can be increased by making the benefit profile increasing.

Because the moral hazard costs and consumption smoothing gains are endogenous to the policy parameters, moving from local recommendations to the characterization of the global optimum requires either more information or some assumptions on the way the sufficient statistics vary with the policy parameters. We perform a simple calibration of the optimum in line with the analysis of the optimal flat benefit level in Gruber [1997], but extended to our dynamic context. Figure 13 shows how our estimates of the relative consumption smoothing gains and moral hazard costs extrapolate under alternative assumptions. A benefit profile is characterized by the ratio of replacement rates in the first 20 weeks, and after 20 weeks of unemployment r_1/r_2 . The replacement rates are such that the expected total expenditures unchanged under the assumption that the duration elasticities with respect to the net of replacement rate are constant.³⁸ For the actual policy (flat profile, $r_1 = r_2 = .8$), we report the estimated relative consumption smoothing gains CS_1/CS_2 and relative moral hazard costs MH_1/MH_2 . The plain black curve shows how the relative moral hazard costs MH_1/MH_2 extrapolate for different profiles r_1/r_2 under the constant elasticities assumption.³⁹ The plain grey line shows how the relative consumption smoothing gains CS_1/CS_2 change with r_1/r_2 , under the assumption that the marginal propensity to consume out of the change in benefits is equal to one.⁴⁰ Leaving the expected benefit expenditures unchanged, the calibration suggests that the marginal value of tax money spent would be equalized (i.e., $CS_1/MH_1 = CS_2/MH_2$, point A in Figure 13) when benefits received after 20 weeks are about

³⁸That is, we consider replacement rates r_1 and r_2 that solve the system of equations consisting of

$$\log(D_k) = \alpha_k + \bar{\varepsilon}_{D_k, 1-r_k} \log(1-r_k) + \bar{\varepsilon}_{D_k, 1-r_l} \log(1-r_l) \text{ for } k = 1, 2 \text{ (and } l = 2, 1\text{),}$$

and the constraint that expected expenditures remain the same. We use our estimates of the duration elasticities w.r.t. changes in b_1 and b_2 (estimated for $r_1 = r_2 = .8$) to calibrate the (constant) elasticities w.r.t. the net-of-replacement rate, i.e., $\bar{\varepsilon}_{D_k, 1-r_l} = -\frac{1-.8}{.8} \varepsilon_{D_k, b_l}$.

³⁹That is, we re-calculate the moral hazard costs based on equations (5) and (6), using the calibrated benefit durations D_k and elasticities $\varepsilon_{D_k, b_l} = -\frac{r_l}{1-r_l} \bar{\varepsilon}_{D_k, 1-r_l}$.

⁴⁰More precisely, we compute the consumption smoothing gains assuming that the consumption drop is linear in the replacement rate,

$$\frac{Ec^e - E_k c^u}{Ec^e} (r_k) = \frac{Ec^e - E_k c^u}{Ec^e} (r_k^0) - \beta [r_k - r_k^0],$$

where r_k^0 denotes the baseline replacement rate and $\beta = 1$ implies hand-to-mouth consumption of the benefit changes.

8% more generous than benefits received during the first 20 weeks.⁴¹ Our recommendation for the policy profile is very similar when assuming a marginal propensity to consume out of the change of benefits of .28 as in Gruber [1997] (point B). The optimal profile would be even more inclining if relative moral hazard costs were assumed constant across benefit profiles (point C). While this back-of-the-envelope calculation relies on strong assumptions, it is important to emphasize that the latest policy changes in Sweden have changed the profile in the exact opposite direction, making benefits up to 23% more generous at the start compared to later in the spell.

6.2 Consumption Implementation

Attempts at quantifying the consumption smoothing gains of UI policies have been scarce, as it requires to estimate differences in marginal utility levels across employed and unemployed individuals, which proves very difficult in practice. We have followed a “consumption implementation” approach, which relies on a Taylor approximation of the marginal utility wedge, and on the estimation of the wedge in consumption expenditures across employed and unemployed individuals. Yet, this approach still requires information on individuals’ preferences, with respect to risk in particular.⁴² While the “consumption implementation” approach has some obvious limitations which may bias our estimates of the average CS gains, our focus on the *profile* of unemployment benefits changes the nature of the challenge we face. Our recommendations on the policy profile will be affected only to the extent that the bias in the estimates of the CS gains varies over the unemployment spell. We briefly assess the robustness of our policy recommendations to the limitations of the “consumption implementation”.

Heterogeneous Agents As discussed before, our (local) expressions of the optimal profile are fully robust to heterogeneity underlying unemployment responses. Regarding the CS gains, the expression in the left-hand side of Corollary 1, based on the wedges in Pareto-weighted marginal utilities, is also fully robust to unrestricted heterogeneity. And the “consumption implementation” approach to estimating these wedges in Pareto-weighted marginal utilities is robust to heterogeneity as well (Chetty [2006], Andrews and Miller [2013]). Yet, this approach still requires information on how individuals’ risk preferences and Pareto-weights evolve over the unemployment spell.⁴³ To

⁴¹Rather than adjusting the replacement rates until the marginal value of tax money spent is equalized (as based on equations (5) and (6)), we could alternatively evaluate the welfare gain of inclining the benefit profile, keeping the tax rate fixed. Such change no longer increases welfare if

$$\frac{1 + CS_1}{1 + CS_2} = \frac{1 + MH_1}{1 + MH_2}. \quad (12)$$

Note that the scalars $\tilde{\lambda}_k \times \tilde{\gamma}_k$ no longer drop out of condition (12), but our recommended profile also satisfies this alternative condition for values of $\tilde{\lambda}_k \times \tilde{\gamma}_k$ around 1.

⁴²This limitation has inspired the use of other choice data to recover the marginal utility gap: Chetty [2008] considers liquidity and substitution effects in effort, (Shimer and Werning [2007]) analyze reservation wages.

⁴³Note that our implementation is not accounting for the correlation between risk preferences and the consumption drop at a given time in the unemployment spell (Andrews and Miller [2013]). This assumption may introduce a bias in the estimated consumption smoothing gains. But our conclusions with respect to the optimal profile will only be affected to the extent that the covariance term capturing this correlation is also correlated with unemployment

assess the robustness of our welfare conclusions to dynamic selection on risk preferences and/or Pareto weights over the unemployment spell, we investigate in appendix Table 5 how various observable characteristics are distributed across short term and long term unemployed. We find that the probability of experiencing long unemployment spells is strongly and monotonically correlated with age. Other patterns of selection are of much smaller magnitude (e.g., married, number of children) or less clear. The effect of net wealth levels and of portfolio characteristics correlated with risk preferences (i.e., fraction of portfolio wealth invested in stocks, leverage defined as total debt divided by gross assets) are all small and non-monotonic. Note that our estimates imply that the scalar $\tilde{\lambda}_k \times \tilde{\gamma}_k$ should be 80 percent larger for the short-term unemployed compared to the long-term unemployed to justify a declining profile. Overall, we find little evidence of dynamic selection patterns providing such justification, unless the Pareto-weights or risk aversion of the elderly are substantially lower than average.

Expenditures vs. Consumption Our empirical analysis considers changes in expenditures. An important question is to what extent expenditures translate into consumption and thus capture the potential welfare value of unemployment benefits. In particular, unemployed workers may try to re-allocate expenditures to smooth the shocks in their consumption. A first example is the substitution towards household production, which has been analyzed extensively in the context of retirement (e.g., Aguiar and Hurst [2005]). An increase in household production would reduce the utility drop when unemployed. A second example is the substitution away from expenditures on durable goods that provide a consumption flow for future periods as well. A reduction in durable good expenditures would have a smaller impact on the immediate consumption flow.

The analysis based on the consumption surveys in Table 3 provides suggestive evidence that these substitution effects may be important. For example, we find a larger drop in restaurant expenditures than in food expenditures, consistent with substitution towards household production. We also estimate a very substantial drop in the purchase of durable goods like vehicles, furniture and house appliances. Two issues are particularly relevant when translating such findings to policy. First, the sign of the impact of substitution patterns on the optimal unemployment policy is ambiguous. For example, the smaller drop in consumption due to household production decreases the marginal utility of unemployment benefits, but complementarity between expenditures and household production increases the marginal utility of benefits. Second, it is the differential substitution over the unemployment spell that is relevant for the evaluation of the benefit profile. For example, the reduction in durable good expenditures is particularly pronounced during the first months of the spell, indicating that the cost of deferring investments in durable goods and, hence, the marginal value of unemployment benefits may increase more over the unemployment spell than we have estimated.⁴⁴ These interesting issues could be further explored and integrated in our framework to improve the design of the unemployment policy.

duration.

⁴⁴On the other hand, various types of expenditures (e.g., housing) are committed in the short term, but can be optimally adjusted in the longer term (Chetty and Szeidl [2007]).

Individuals vs. Households Our welfare analysis has focused on consumption smoothing estimates at the individual level. As analyzed earlier, households can use the income and assets of the entire household to smooth consumption when one of its members becomes unemployed and we indeed find that the relative drop in consumption is substantially smaller at the household level (16% compared to 26% after one year of unemployment). The same two caveats again apply when translating this to policy. First, it is unclear how much the smaller consumption drop decreases the marginal utility of UI benefits as the relevant risk aversion (or Pareto-weight) to weigh this consumption drop may well be higher at the household level. Second, we find no evidence that consumption levels evolve differently over the unemployment spell when aggregated at the household level, which is again what matters for evaluating the benefit profile.

7 Conclusion

This paper has offered a simple, general and empirically implementable framework to evaluate the optimal time profile of unemployment benefits. Our theoretical approach proves that, independent of the underlying primitives of the model, the dynamic problem of balancing insurance value and incentive costs can be characterized in a transparent way as a series of simple trade-offs involving just a few estimable statistics. Putting this simple characterization to the data, our empirical implementation has shown that it is not at all obvious that declining benefit profiles are always optimal. Despite the forward-looking behaviors of job seekers, important forces such as duration dependence and heterogeneity can, as in our context, make the moral hazard costs of long term benefits smaller than that of benefits offered early in the spell. The limited access to consumption smoothing opportunities that we document among the unemployed in Sweden also makes cutting benefits particularly costly for the long-term unemployed.

We have presented a framework that is easily replicable and our hope is that it will trigger new empirical work that analyzes the relevant statistics for policy evaluation in other contexts where labor market conditions, access to credit or the unemployment policy in place may be very different. Our analysis has shown that the empirical analysis of labor supply responses to UI should pay particular attention to the timing of benefits in order to produce estimates that can be meaningful from a welfare perspective. In terms of assessing the value of UI benefits, our analysis shows that fruitful avenues of research are being opened by administrative and/or proprietary data on wealth and expenditures matched with UI records. The replication of our empirical exercise for different contexts or subcategories of individuals can provide the necessary estimates in order to assess the optimality of contingent profiles such as age-dependent or business-cycle dependent profiles.

Most importantly, the tools developed in this paper can be applied to other dynamic contexts. An important area for future work will be to develop such simple, yet robust characterization of various other dynamic policies, including the design of retirement pensions or parental leave policies.

References

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

Andrews, Isaiah and Conrad Miller, "Optimal Social Insurance with Heterogeneity," *Mimeo, MIT*, 2013.

Baily, Martin N., "Some Aspects of Optimal Unemployment Insurance," *Journal of Public Economics*, 1978, 10 (3), 379–402.

Chetty, Raj, "A General Formula for the Optimal Level of Social Insurance," *Journal of Public Economics*, 2006, 90 (10-11), 1879–1901.

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

— , "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 05 2009, 1 (1), 451–488.

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

Ganong, Peter and Simon Jaeger, "A Permutation Test and Estimation Alternatives for the Regression Kink Design," *Harvard University Discussion Paper*, 2014.

Gerard, Francois and Gustavo Gonzaga, "Informal Labor and the Efficiency Cost of Social Programs: Evidence from 15 Years of Unemployment Insurance in Brazil," *Mimeo, Columbia University*, 2014.

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

Hopenhayn, Hugo A. and Juan Pablo Nicolini, "Optimal Unemployment Insurance," *Journal of Political Economy*, 1997, 105 (2), 412–438.

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*, University of Chicago Press, July 2014

Kreiner, Claus T., David Dreyer Lassen, and Søren Leth-Petersen, "Liquidity Constraint Tightness and Consumer Responses to Fiscal Stimulus Policy," *Mimeo, Copenhagen University*, 2014.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo, "Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment," *The Quarterly Journal of Economics*, 2013.

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

Landais, Camille, "Assessing the Welfare Effects of Unemployment Benefits Using the Regression Kink Design," *American Economic Journal: Economic Policy*, 2015.

— , Pascal Michaillat, and Emmanuel Saez, "Optimal Unemployment Insurance over the Business Cycle," Working Paper 16526, National Bureau of Economic Research 2010.

Machin, Stephen and Alan Manning, "The Causes and Consequences of Longterm Unemployment in Europe," in Orley C. Ashenfelter and David Card, eds., *Orley C. Ashenfelter and David Card, eds.*, Vol. 3, Part C of *Handbook of Labor Economics*, Elsevier, 1999, pp. 3085 – 3139.

Mas, Alexandre and Andrew Jonhston, "Potential Unemployment Insurance Duration and Labor Supply: the Individual and Market-Level Response to a Benefit Cut," *Mimeo, Princeton University*, 2015.

McCrary, Justin, "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, February 2008, 142 (2), 698–714.

Meyer, Bruce, "Unemployment Insurance and Unemployment Spells," *Econometrica*, 1990, 58(4), 757–782.

Mogstad, Magne and Andreas Kostol, "Earnings, Income, and Consumption of Alowed and Rejected Disability Insurance Applicants," *American Economic Review, Papers and Proceedings*, 2015, 105 (5), 137–141.

Pavoni, Nicola, "Optimal Unemployment Insurance, With Human Capital Depreciation, and Duration Dependence," *International Economic Review*, 2009, 50 (2), 323–362.

Pistaferri, Luigi, "Household Consumption: Research Questions, Measurement Issues, and Data Collection Strategies," *Mimeo Stanford University*, 2015.

Rothstein, Jesse, "Unemployment Insurance and Job Search in the Great Recession," Working Paper 17534, National Bureau of Economic Research October 2011.

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

Shavell, Steven and Laurence Weiss, "The Optimal Payment of Unemployment Insurance Benefits over Time," *Journal of Political Economy*, 1979, 87 (6), 1347–1362.

Shimer, Robert and Iván Werning, "On the Optimal Timing of Benefits with Heterogeneous Workers and Human Capital Depreciation ,," 2006.

— and —, "Reservation Wages and Unemployment Insurance," *Quarterly Journal of Economics*, 2007, 122 (3), 1145–1185.

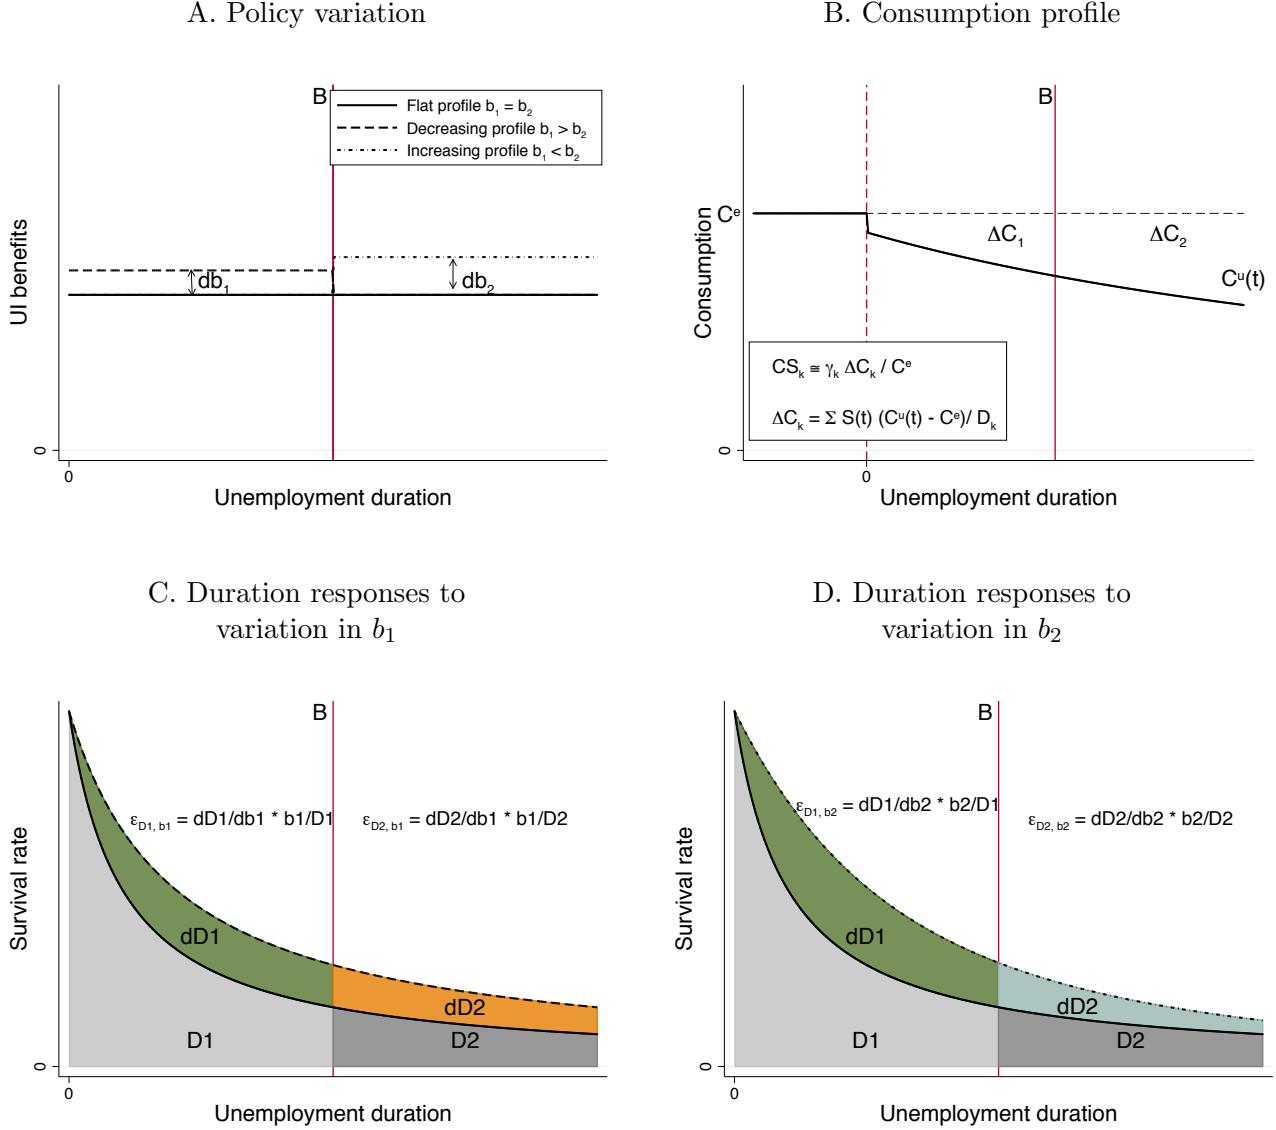
— and —, "Liquidity and Insurance for the Unemployed," *American Economic Review*, 2008, 98 (5), 1922–42.

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

Valletta, Robert and Henry S. Farber, "Extended Unemployment Insurance and Unemployment Duration in The Great Recession: The U.S. Experience," *Mimeo, Princeton University*, 2011.

Werning, Iván, "Optimal Unemployment Insurance with Unobservable Savings," *University of Chicago and UTDT*, 2002.

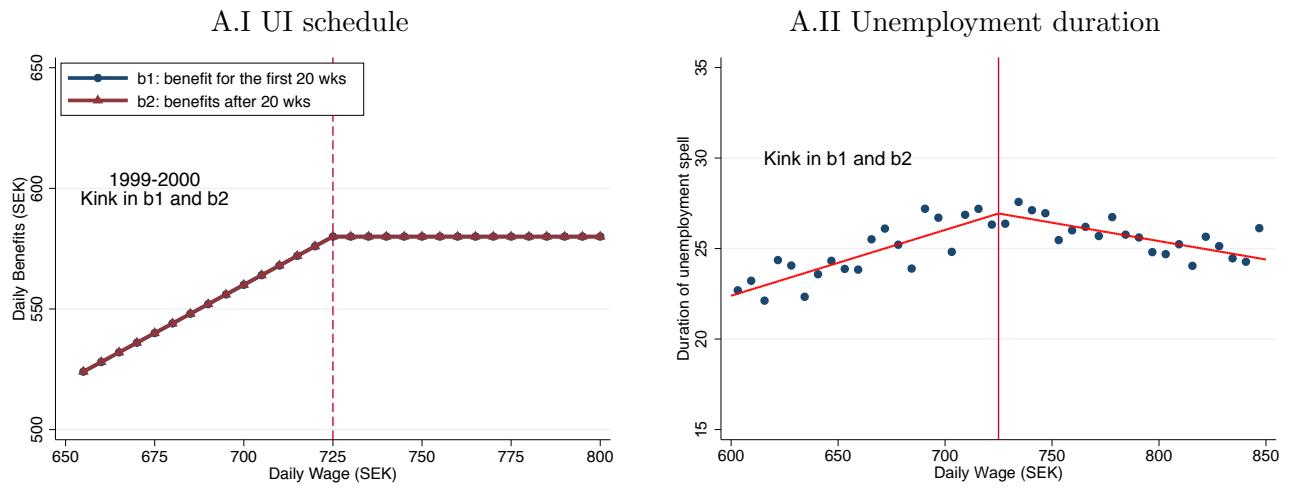
Figure 1: SUFFICIENT STATISTICS FOR WELFARE ANALYSIS OF TWO-PART POLICY



Notes: The figure summarizes the policy variation and statistics needed to characterize an optimal two-part profile giving b_1 for the first B weeks and b_2 afterwards. In Panel A, we display policy variation that allows to evaluate the welfare of a two-part profile. We start from a flat profile as is the case in our empirical application in Sweden. Variation db_1 in benefits given for the first B weeks and variation db_2 in benefits given after B weeks are both needed to evaluate the moral hazard costs of providing the respective benefit levels. The moral hazard costs are captured by the responses of the duration spent in the first part of the profile D_1 and in the second part of the profile D_2 to variations in both b_1 and b_2 , as exemplified in panels C and D. These responses enable the identification of all cross-duration elasticities $\epsilon_{D1, b1}$, $\epsilon_{D2, b1}$, $\epsilon_{D1, b2}$ and $\epsilon_{D2, b2}$ entering the RHS of dynamic Baily-Chetty formulas (5) and (6). To evaluate the consumption smoothing gains of the two-part policy, the planner requires the average drop in consumption ΔC_1 for individuals in the first part of the profile receiving b_1 , and the average drop in consumption ΔC_2 for individuals in the second part of the profile receiving b_2 . This can be calculated based on the profile of consumption as a function of time spent unemployed as depicted in Panel B. Note that these consumption statistics need to be evaluated at the current profile, and do not require any policy variation.

Figure 2: UI BENEFITS AND UNEMPLOYMENT DURATION AS A FUNCTION OF DAILY WAGE AROUND THE 725SEK KINK

A. 1999 - 2000



B. 2001

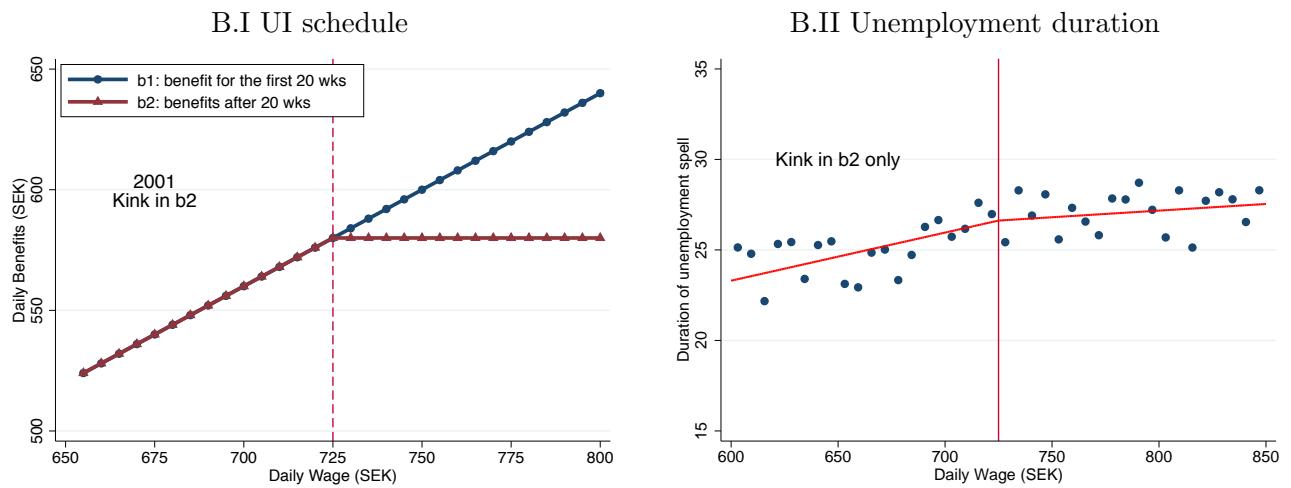
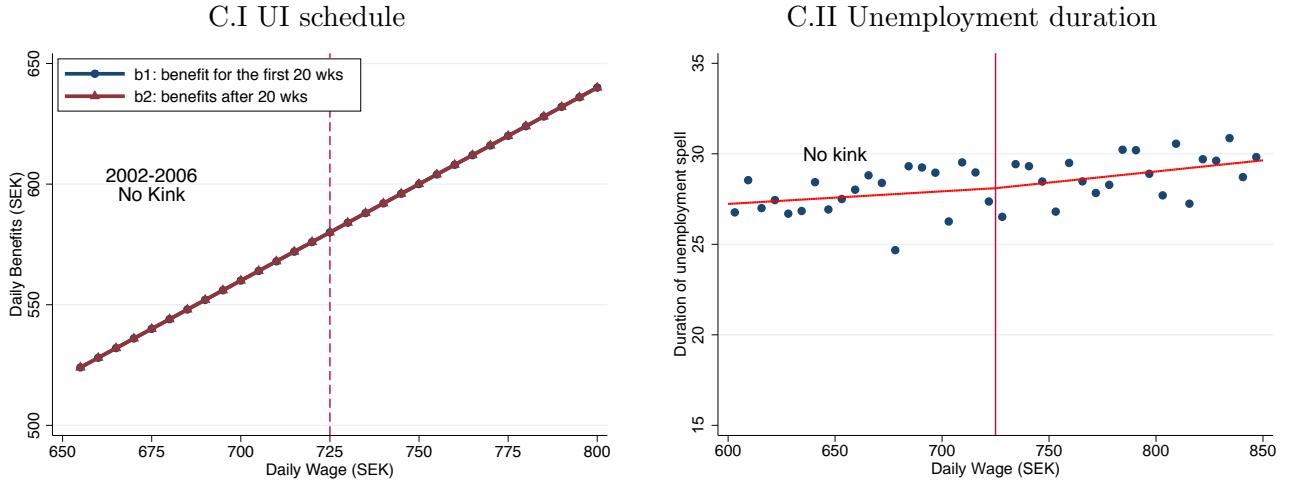


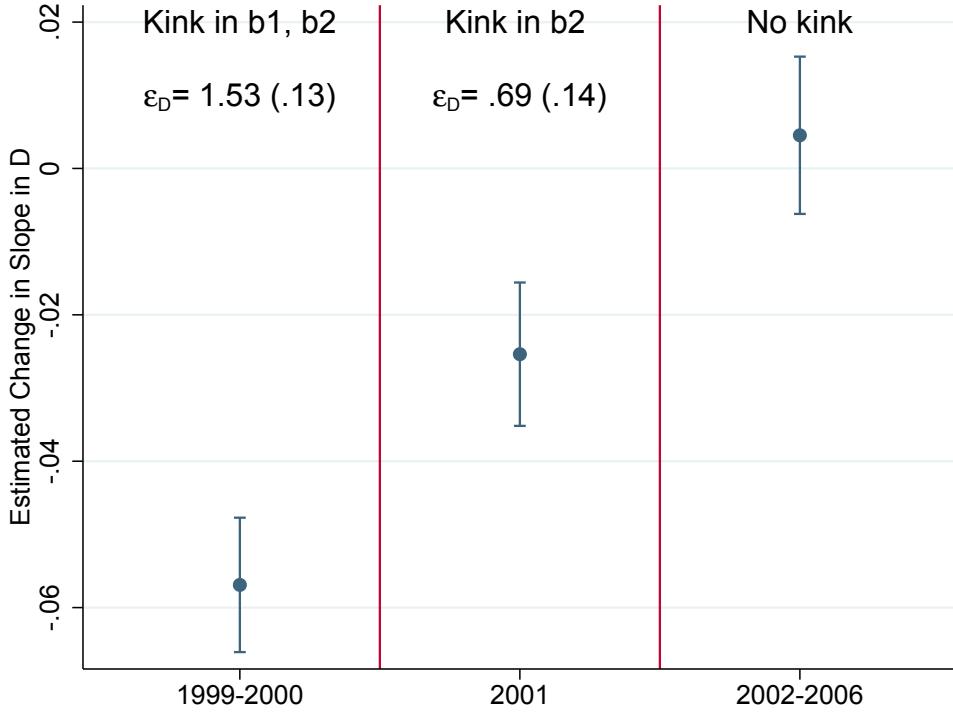
Figure 3: UI BENEFITS AND UNEMPLOYMENT DURATION AS A FUNCTION OF DAILY WAGE AROUND THE 725SEK KINK (*continued*)

A. 2002-2007



Notes: The left panels display the UI benefit level received during the first 20 weeks of unemployment (b_1) and after 20 weeks of unemployment (b_2) as a function of daily wage prior to becoming unemployed. For spells starting before July 2001 (A.I), the schedule exhibits a kink in both b_1 and b_2 at the 725SEK threshold, which can be used to identify the effect of both b_1 and b_2 on unemployment duration. For spells starting between July 2001 and July 2002 (B.I), the schedule exhibits a kink in b_2 only at the 725SEK threshold, which can be used to identify the effect of b_2 on unemployment duration. Finally, for spells starting after July 2002 (C.I), the schedule is linear for both b_1 and b_2 at the 725SEK threshold, which offers a placebo setting to assess the validity of the RK design at the 725SEK threshold. The right panels plot average unemployment duration in bins of previous daily wage for the three periods of interest. Unemployment duration is defined as the number of weeks between registration at the PES and exiting the PES or finding any employment (part-time or full-time employment, entering a PES program with subsidized work or training, etc.). Unemployment duration is capped at two years. Sample is restricted to unemployed individuals with no earnings who report being searching for full-time employment. The graphs provide graphical evidence of a change in slope in the relationship between unemployment duration and previous daily wage in response to the kink in UI benefits. The change in slope is larger for spells starting before July 2001, when both b_1 and b_2 are capped at the 725SEK threshold (A.II). The magnitude of the change in slope decreases for spells starting between July 2001 and July 2002 when only b_2 is capped at the 725SEK threshold. Finally, there is no significant change in slope for spells after July 2002, when the schedule is linear for both b_1 and b_2 at the threshold, which is supportive of the identifying assumptions of the RK design. Formal estimates of the change in slope using polynomial regressions of the form of equation (7) are displayed in Table 2. The red lines display predicted values of the regressions in the linear case.

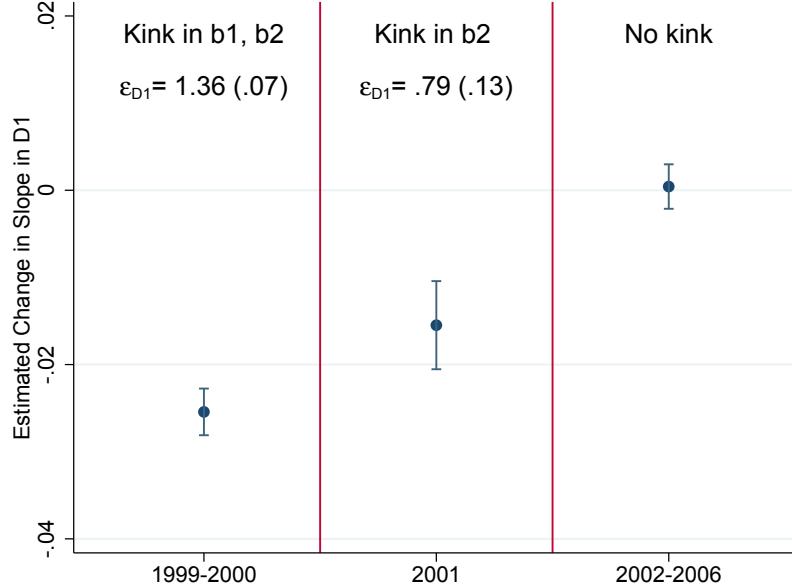
Figure 4: RKD ESTIMATES AT THE 725SEK THRESHOLD



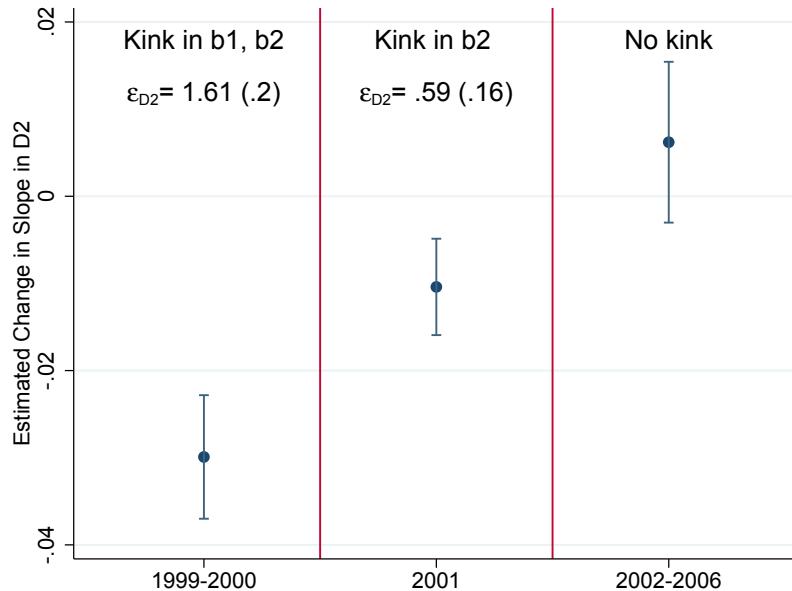
Notes: The figure reports estimates of the change in slope with 95% robust confidence interval in the relationship between unemployment duration and daily wage at the 725SEK threshold using polynomial regressions of the form of equation (7) with a bandwidth size $h = 90\text{SEK}$. These estimates are reported for three periods of interest: 1999-2000 (i.e., spells starting before July 2001), 2001 (i.e., spells starting after July 2001 and before July 2002) and 2002 (i.e., spells starting after July 2002). Unemployment duration is defined as the number of weeks between registration at the PES and exiting the PES or finding any employment (part-time or full-time employment, entering a PES program with subsidized work or training, etc.). Unemployment duration is capped at two years. Sample is restricted to unemployed individuals with no earnings who report being searching for full-time employment. The figure also reports the corresponding elasticities of unemployment duration with respect to b_1 and b_2 (for period 1999-2000) and with respect to b_2 only (period 2001). Robust standard errors computed using the Delta-method are in parentheses.

Figure 5: RKD ESTIMATES AT THE 725SEK THRESHOLD

A. Time spent on first part of the profile (D_1):
unemployment duration up to 20 weeks

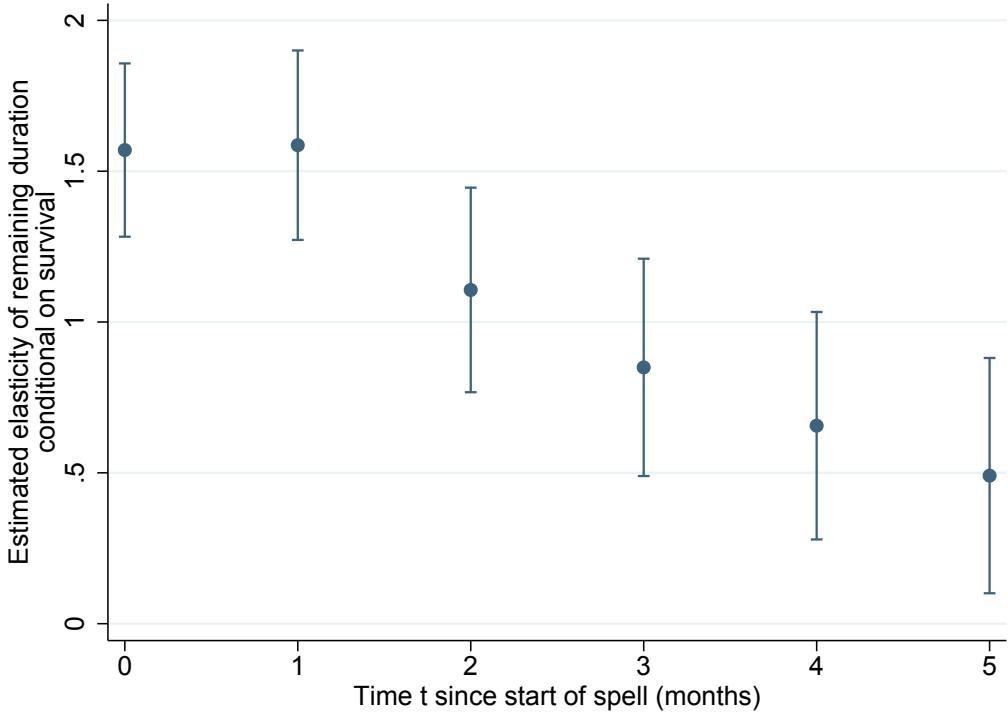


B. Time spent on second part of the profile (D_2):
unemployment duration after 20 weeks



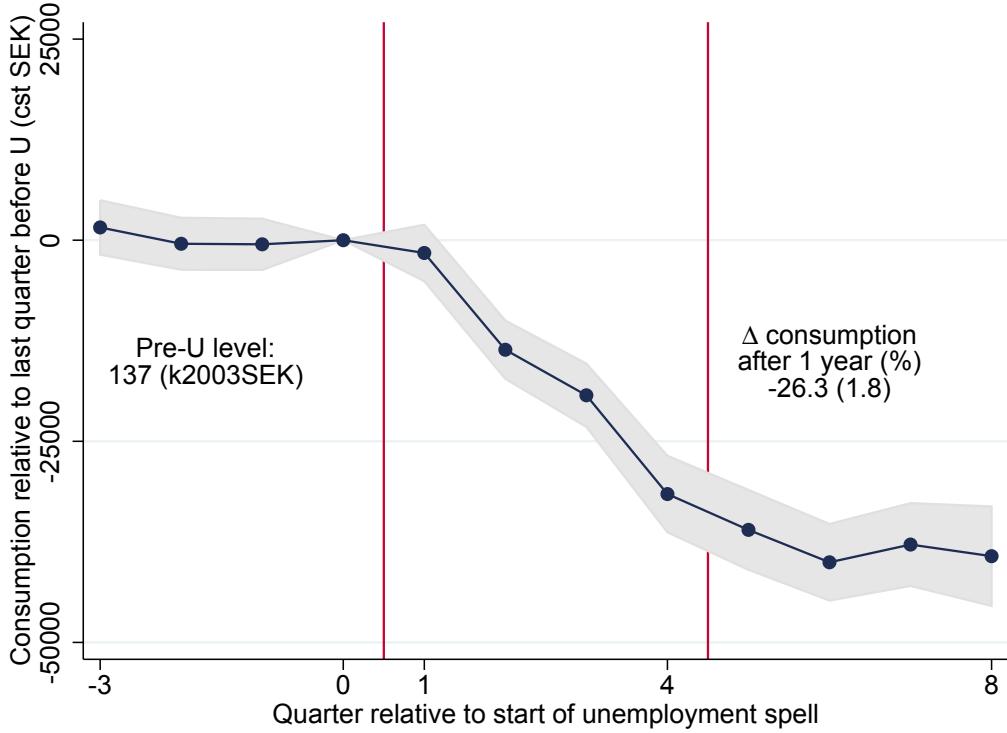
Notes: The figure reports estimates of the change in slope with 95% robust confidence interval in the relationship between daily wage and time spent in the first part of the benefit profile D_1 and time spent in the second part of the profile D_2 , at the 725SEK wage threshold, using polynomial regressions of the form of equation (7) with a bandwidth size $h = 90\text{SEK}$. These estimates are reported for three periods of interest: 1999-2000 (i.e., spells starting before July 2001), 2001 (i.e., spells starting after July 2001 and before July 2002) and 2002- (i.e., spells starting after July 2002). Unemployment duration is defined as the number of weeks between registration at the PES and exiting the PES or finding any employment (part-time or full-time employment, entering a PES program with subsidized work or training, etc.). $D_1 = \sum_{t < 20wks} S_t$ corresponds to duration censored at 20 weeks of unemployment. $D_2 = \sum_{t \geq 20wks} S_t$ corresponds to unconditional duration spent unemployed after 20 weeks of unemployment (i.e., not conditional on having survived up to 20 weeks). Total duration of unemployment is capped at 2 years. Sample is restricted to unemployed individuals with no earnings who report being searching for full-time employment. The figure also reports the corresponding elasticities of unemployment duration with respect to b_1 and b_2 (for period 1999-2000) and with respect to b_2 only (period 2001). Robust standard errors computed using the Delta-method are in parentheses.

Figure 6: TESTING FOR STATIONARITY: ELASTICITY OF THE REMAINING DURATION OF UNEMPLOYMENT, CONDITIONAL ON SURVIVING UNTIL t , WITH RESPECT TO CHANGES IN THE FLAT BENEFIT LEVEL



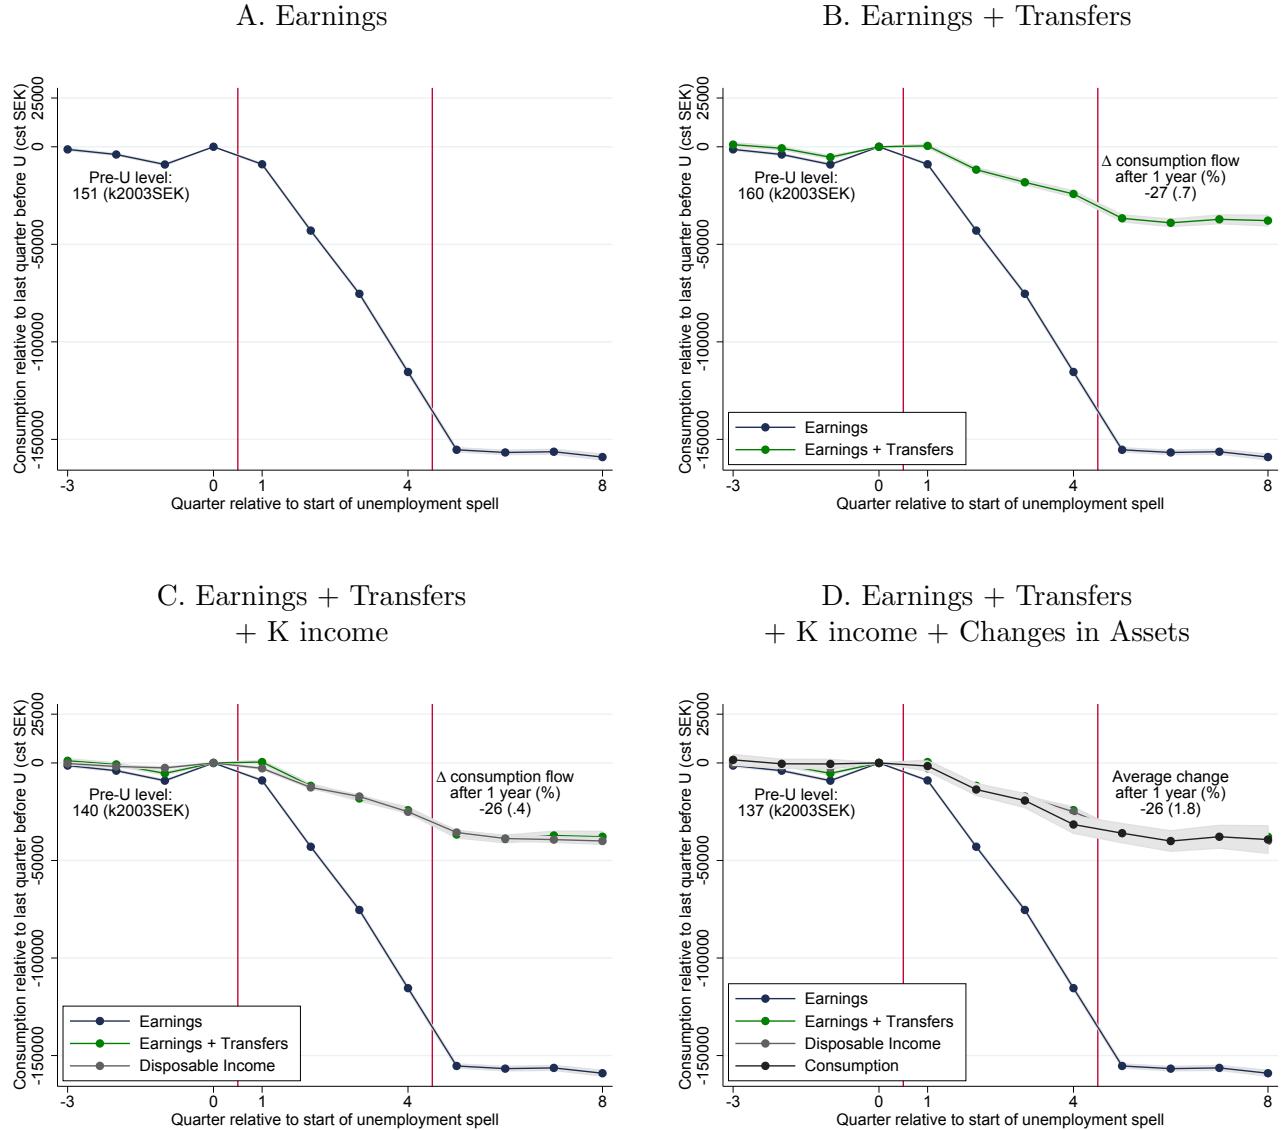
Notes: The figure reports RRD estimates (with 95% robust confidence interval) of the elasticity of the remaining duration of unemployment conditional on surviving until t with respect to changes in the flat benefit level b . Estimates use the presence, for spells starting before July 2001, of a kink in the benefit schedule of the flat benefit b at the 725SEK wage threshold. We use polynomial regressions of the form of equation (7) with a bandwidth size $h = 100$ SEK. The remaining duration \tilde{D}_t is the unemployment duration D minus t , conditional on being still unemployed after t months. Unemployment duration is defined as the number of weeks between registration at the PES and exiting the PES or finding any employment (part-time or full-time employment, entering a PES program with subsidized work or training, etc.). Sample is restricted to unemployed individuals with no earnings who report being searching for full-time employment. In a stationary environment, the elasticity of \tilde{D}_t with respect to the flat benefit b should be constant with t . As the estimated elasticities strongly decline with t , our results suggest the presence of strong non-stationary forces (i.e., dynamic selection, duration-dependence, etc.).

Figure 7: YEARLY INDIVIDUAL CONSUMPTION AS A FUNCTION OF UNEMPLOYMENT DURATION



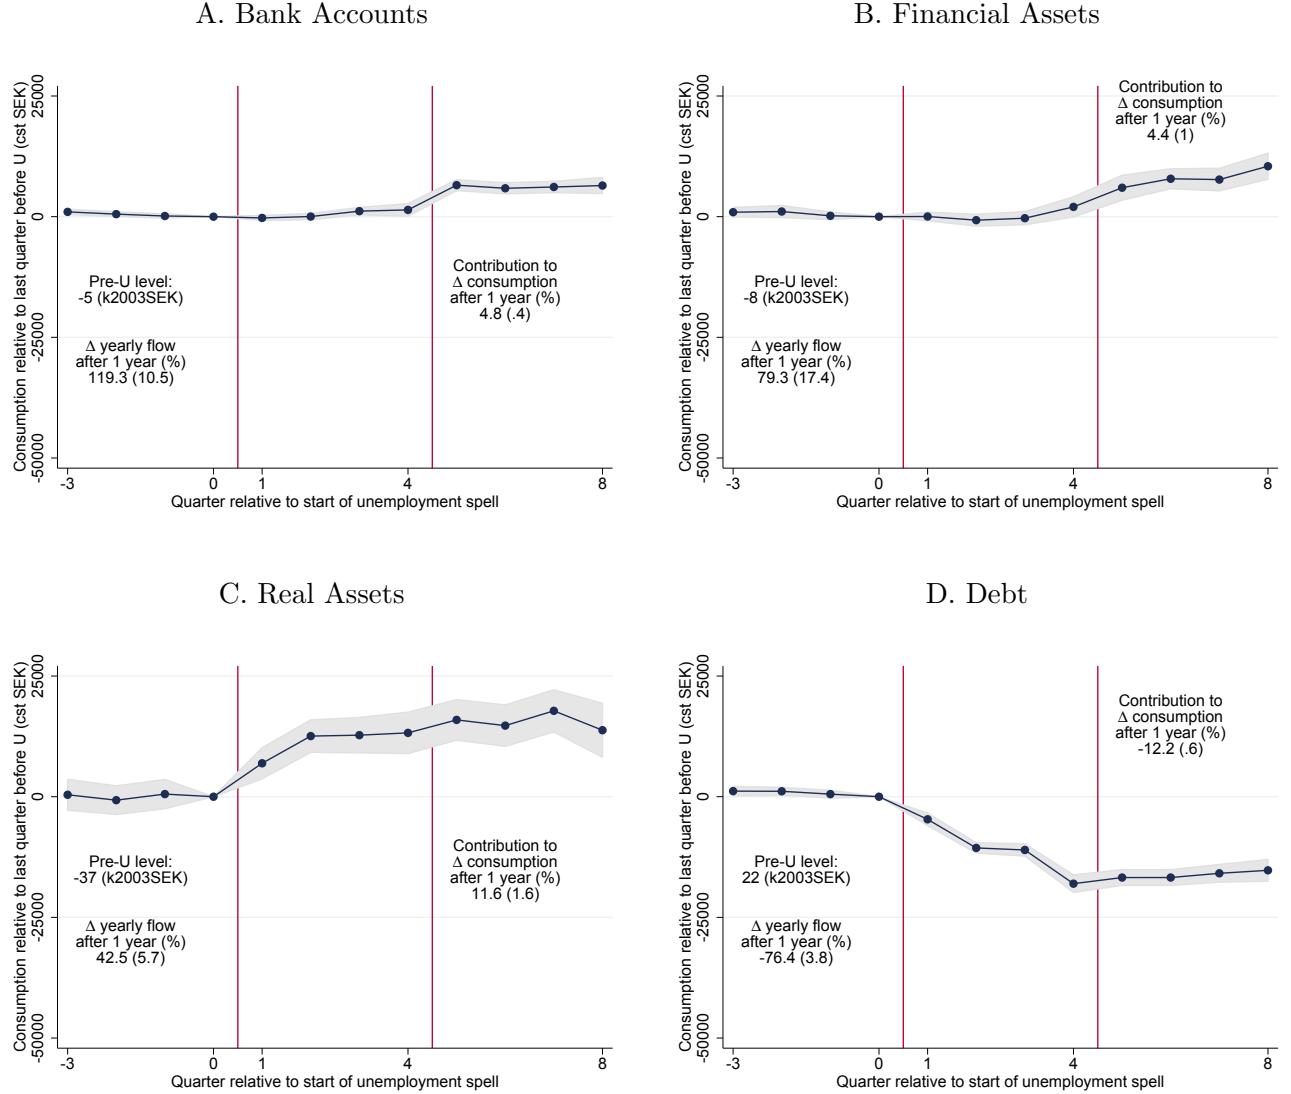
Notes: The figure correlates our yearly consumption measure with the time t (in quarters) since (or until) the onset of the unemployment spell. Details on the construction of the yearly measure of consumption expenditures are given in Section 5.1. The sample consists of all individuals unemployed as of December of year n , or who will become unemployed in year $n+1$, for all years n from 1999 up to 2007. We impose the further restriction that individuals must not have experienced any unemployment spell in the past two years prior to becoming unemployed. The figure follows from regression model (9) and plots the estimated coefficients β_t for the set of indicators $\mathbf{1}[T = t]$ for being observed, as of December of year n , in the t -th quarter since the onset of one's spell. Regression includes year dummies, age dummies a dummy for being just out of school in year n , and a set of dummies for family status. We also plot the 95% confidence interval from robust standard errors. The graph provides clear evidence that unemployment causes a substantial drop in consumption. The average yearly consumption of individuals who have been unemployed for a full year (quarter 5) is more than 40kSEK lower than the average consumption of individuals at the start of their spell. Given that the yearly consumption of individuals observed just prior to becoming unemployed is 137kSEK, this means that yearly consumption levels drop by 26.3% after a year of unemployment.

Figure 8: DECOMPOSITION OF YEARLY INDIVIDUAL CONSUMPTION AS A FUNCTION OF UNEMPLOYMENT DURATION



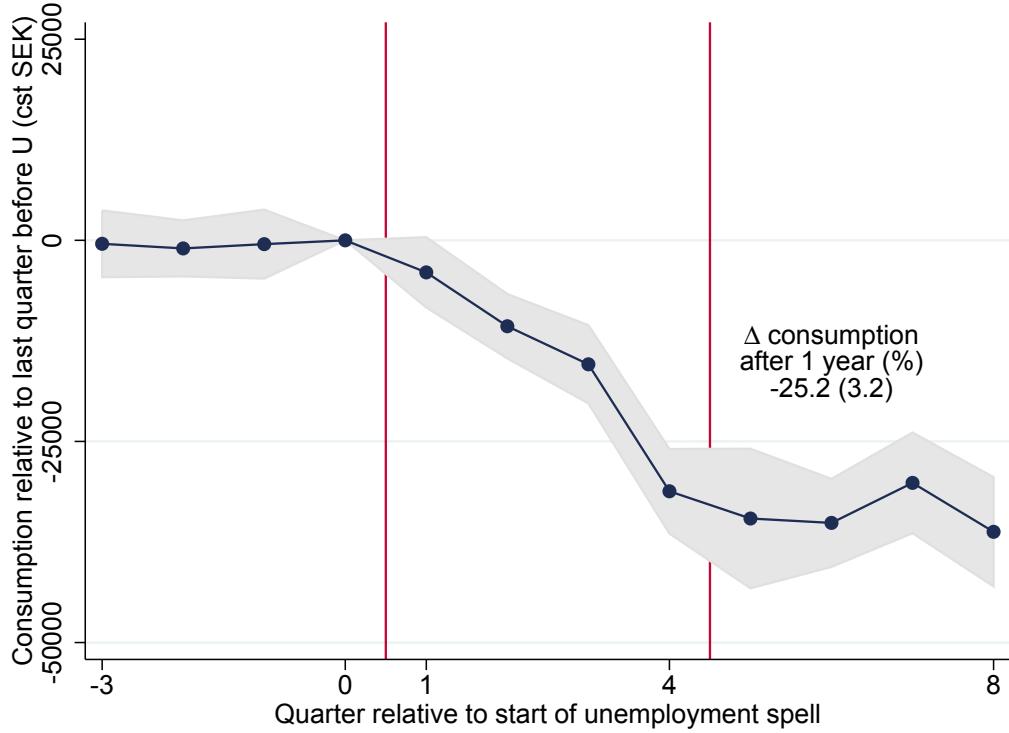
Notes: In this figure, we decompose yearly consumption into earnings, transfers, capital income and changes in assets and correlate these with the time t (in quarters) since (or until) the onset of the unemployment spell. Details on the construction of the yearly measure of consumption expenditures are given in section 5.1. We consider the same sample as in Figure 7. The figure follows from regression model (9) and plots for the various outcome variables the estimated coefficients β_t for the set of indicators $\mathbb{1}[T = t]$ for being observed, as of December of year n , in the t -th quarter since the onset of one's spell. Regression includes year dummies, age dummies, a dummy for being just out of school in year n , and a set of dummies for family status. We also plot the 95% confidence interval from robust standard errors. Panel A plots the evolution of earnings. In panel B, we add transfers and plot the evolution of earnings + transfers. In panel C, we add other income including capital income. Finally, panel D gets at yearly consumption by adding changes in assets and debt. The graph provides evidence that most of the drop in earnings is covered by transfers, while assets play a negligible role in smoothing consumption on average.

Figure 9: CHANGES IN ASSETS OVER THE UNEMPLOYMENT SPELL



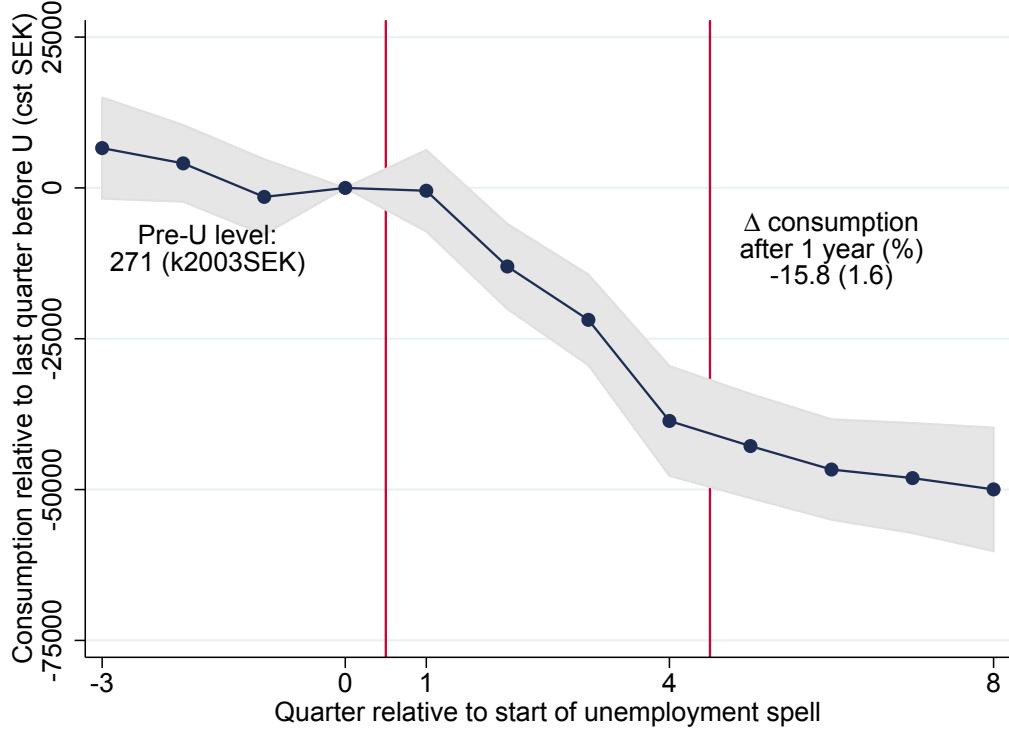
Notes: In this figure, we explore the role of various assets in smoothing consumption over the unemployment spell. For each asset type, we show how the yearly consumption out of this type of asset correlates with the time t (in quarters) since (or until) the onset of the unemployment spell. We consider the same sample as in Figure 7. The figure follows from regression model (9) and plots for the various outcomes the estimated coefficients β_t for the set of indicators $\mathbf{1}[T = t]$ for being observed, as of December of year n , in the t -th quarter since the onset of one's spell. Regression includes year dummies, age dummies, a dummy for being just out of school in year n , and a set of dummies for family status. We also plot the 95% confidence interval from robust standard errors. For each type of asset, we show the yearly consumption level out of this type of asset prior to becoming unemployed. We also report the percentage change in these yearly consumption levels after a full year of unemployment ($t = 5$). We finally report the contribution of this change in yearly consumption levels out of this type of asset to the overall change in consumption after a full year of unemployment.

Figure 10: YEARLY CONSUMPTION AS A FUNCTION OF UNEMPLOYMENT DURATION: FIXED-EFFECT MODEL



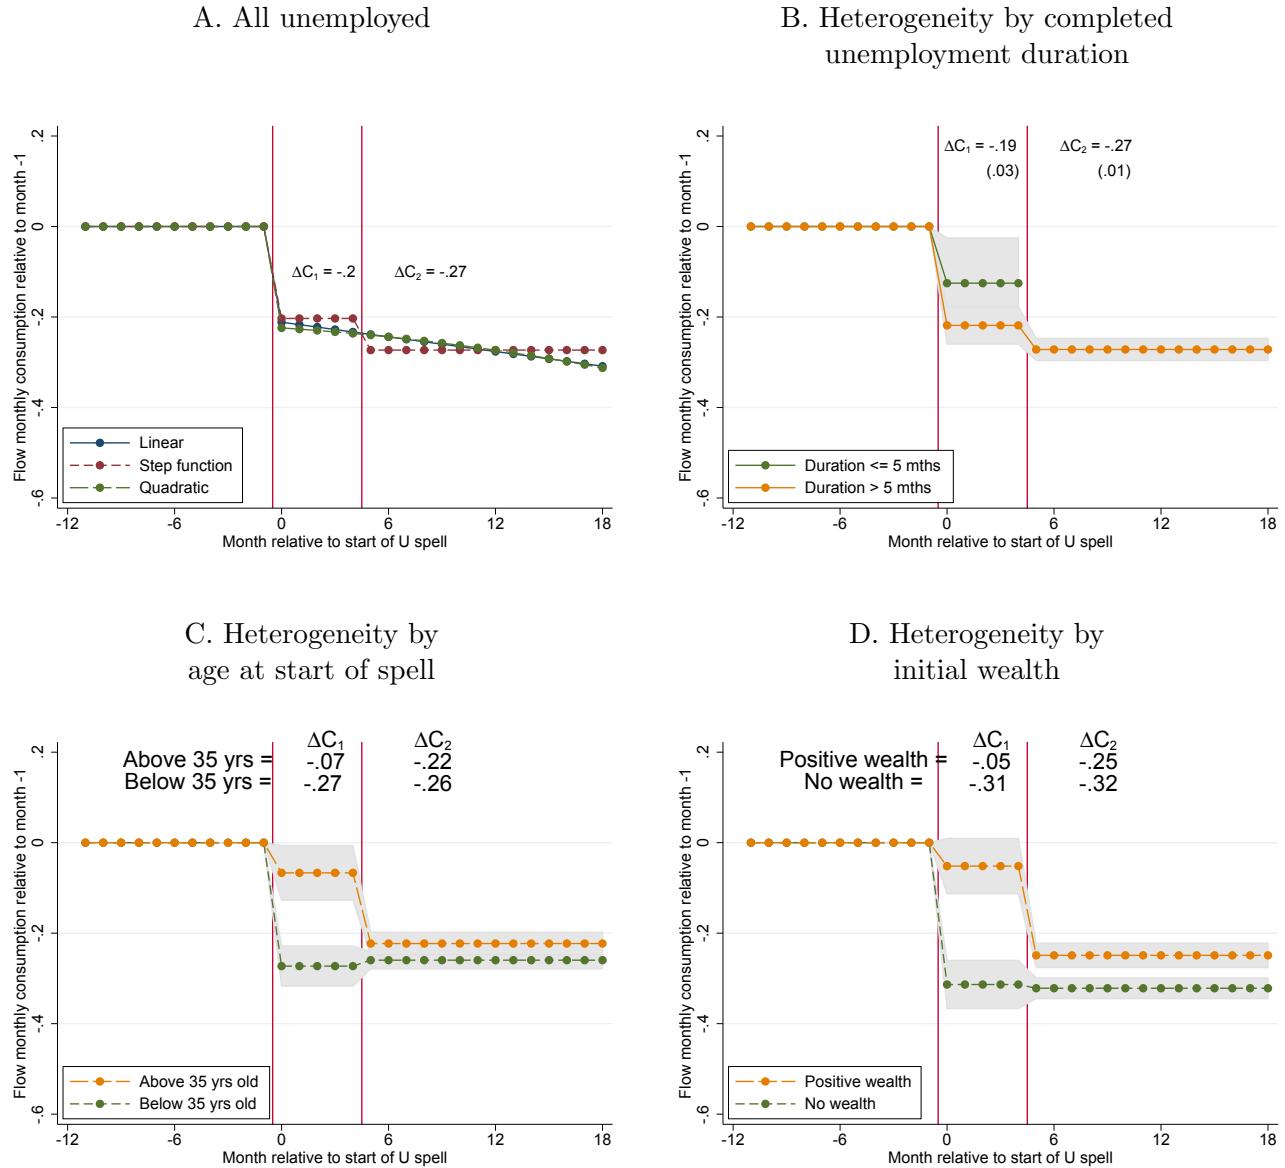
Notes: The figure correlates yearly consumption measure, with the time t (in quarters) since (or until) the onset of the unemployment spell, controlling for potential dynamic selection. We consider the same sample as in Figure 7. Taking advantage of the panel structure of the data, we estimate a model with individual fixed effects using a within estimator. The figure follows from regression model (10) and plots the estimated coefficients $\hat{\beta}_t$ for the set of indicators $\mathbf{1}[T = t]$ for being observed, as of December of year n , in the t -th quarter since the onset of one's spell. Regression includes year dummies, age dummies, a dummy for being just out of school in year n , and a set of dummies for family status. We also plot the 95% confidence interval from robust standard errors.

Figure 11: YEARLY HOUSEHOLD CONSUMPTION AS A FUNCTION OF UNEMPLOYMENT DURATION



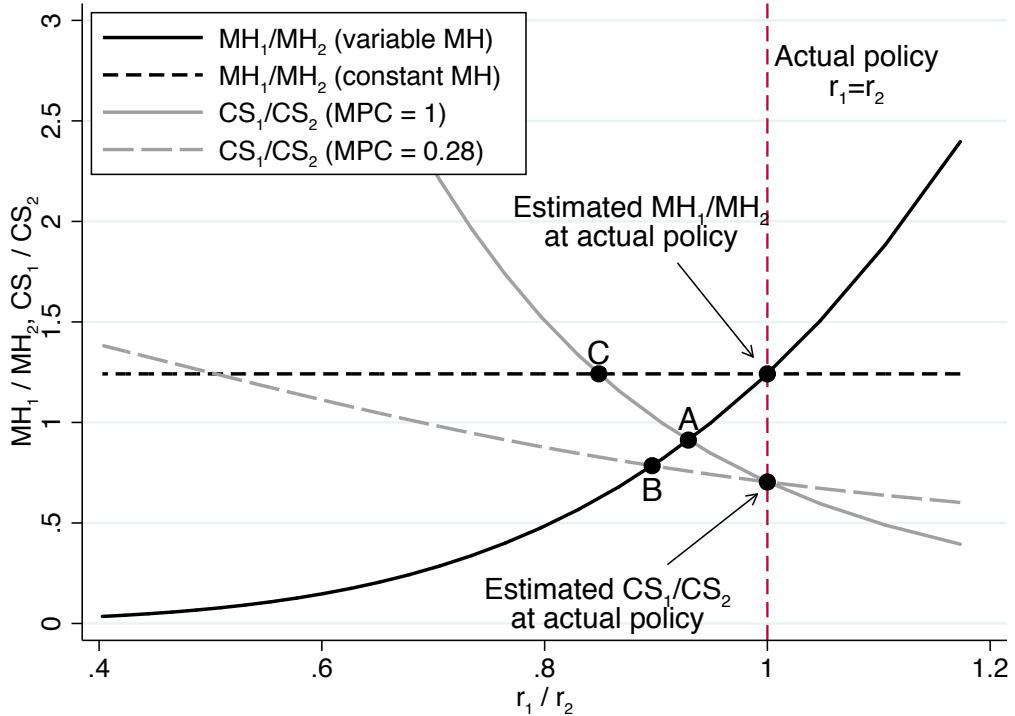
Notes: The figure correlates yearly household consumption with the time t (in quarters) since (or until) the onset of the unemployment spell. To compute household consumption, we compute an individual residual measure of consumption for all members aged above 18 of the household in which the unemployed individual lives, and then sum these measures at the household level. We consider the same sample as in Figure 7. The figure follows from regression model (9) and plots the estimated coefficients β_t for the set of indicators $\mathbf{1}[T = t]$ for being observed, as of December of year n , in the t -th quarter since the onset of one's spell. Regression includes year dummies, age dummies, a dummy for being just out of school in year n , and a set of dummies for family status. We also plot the 95% confidence interval from robust standard errors.

Figure 12: ESTIMATED MONTHLY CONSUMPTION AS A FUNCTION OF UNEMPLOYMENT DURATION



Notes: The figure reports estimated monthly consumption as a function of time relative to the onset of the unemployment spell. Parameters of the monthly consumption profiles are estimated using the information contained in yearly consumption, following the methodology explained in Appendix C. In Panel A, we report the estimated profile of monthly consumption for three different parametric specifications of the monthly consumption profiles: linear with a drop at unemployment, quadratic with a drop at unemployment, and a step-function. In Panel B, we account for selection by estimating two different consumption profiles, one for individuals whose realized duration is 5 months or less, and one for individuals whose realized duration is longer than 5 months. We compute the aggregate drops ΔC_1 and ΔC_2 by reweighting the two profiles by the fraction of individuals in each profile at each duration t . In Panel C, we estimate separate profiles depending on age at the start of the spell. In panel D, we report estimated profiles for individuals with no or negative net wealth at the start of a spell, and individuals with positive net wealth at the start of a spell. We report robust 95% confidence intervals for the step-function estimates.

Figure 13: CALIBRATION OF WELFARE EFFECTS OF DIFFERENT BENEFIT PROFILES



Notes: The figure performs a simple calibration of the optimal two-part profile of UI. A benefit profile is characterised by the ratio of replacement rates in the first 20 weeks, and after 20 weeks of unemployment r_1/r_2 . For the actual policy (flat profile, $r_1 = r_2$), we report the estimated relative consumption smoothing gains CS_1/CS_2 and relative moral hazard costs MH_1/MH_2 . The plain black curve shows how the relative moral hazard costs MH_1/MH_2 extrapolate for different profiles r_1/r_2 under the assumption of constant elasticities of duration with respect to the net of replacement rate (see text for details). Alternatively, the dashed black curve assumes a constant ratio of moral hazard costs. The plain grey line shows how the relative consumption smoothing gains CS_1/CS_2 change with r_1/r_2 , under the assumption that the marginal propensity to consume out of the change in benefits is equal to one. Alternatively, the dashed grey-line assumes a marginal propensity to consume out of the change of benefits of .28 as in Gruber [1997]. Point A shows the calibrated optimal profile ($CS_1/CS_2 = MH_1/MH_2$) in our preferred scenario, which suggests making benefits received after 20 weeks about 8% more generous than benefits received during the first 20 weeks. The optimal policy profile is very similar when assuming a marginal propensity to consume out of the change of benefits of $\beta = .28$ (point B). The optimal profile would be even more inclining when assuming that the relative moral hazard costs would remain unchanged (point C).

Table 1: SUMMARY STATISTICS: PRE-UNEMPLOYMENT

	Mean	P25	P50	P75	P90
I. Unemployment					
Duration of spell (wks)	26.0	5.71	13.29	33.14	78.14
Duration on b_1 (wks)	12.6	5.71	13.29	20	20
Duration on b_2 (wks)	13.4	0	0	13.14	58.14
Replacement rate	.72	.67	.78	.8	.8
II. Demographics					
Age	33.7	23	30	42	54
Fraction men	.49	0	0	1	1
Fraction married	.23	0	0	0	1
Number of children	1.01	0	1	2	3
III. Income, SEK 2003(K)					
Gross earnings	151	44	135	227	295
Capital Income	.8	0	0	0	3
Disposable Income	140	89	136	180	230
IV. Wealth, SEK 2003(K)					
Net worth (A+B-C)	174	-65	0	157	688
<i>As a fraction of disp. income</i>	1.24	-.49	0	1.21	4.34
Financial assets (A)	83	0	4	52	191
<i>As a fraction of disp. income</i>	.66	0	.03	.41	1.44
Bank holdings	29	0	0	14	70
<i>As a fraction of disp. income</i>	.21	0	0	.1	.49
Mutual funds	26	0	0	8	56
<i>As a fraction of disp. income</i>	.23	0	0	.06	.47
Stocks	17	0	0	0	11
<i>As a fraction of disp. income</i>	.12	0	0	0	.08
Real Estate (B)	281	0	0	321	907
<i>As a fraction of disp. income</i>	2.28	0	0	1.94	5.25
Debt (C)	190	0	71	254	514
<i>As a fraction of disp. income</i>	1.7	0	.53	1.77	3.36

Notes: The table provides summary statistics for our sample of unemployed individuals in Sweden from 1999 to 2007. All earnings, income and asset level measures are yearly measures where we focus on individuals just about to start an unemployment spell, i.e., observed in december of year n and starting unemployment spell in the first quarter of year $n+1$. Capital income is defined as all dividends and earned interests, plus realized capital gains on all private properties (real estate, financial assets, valuable commodities such as art or cars, etc.). Disposable income is gross earnings, plus capital income minus all taxes plus transfers received. Transfers include unemployment insurance, disability insurance, sick pay, and all housing and parental benefits. All financial assets are estimated at their market value. Real estate is gross of debt and assessed at market value. Debt includes student loans, mortgage, credit card debt, etc. For each asset, the values reported as a fraction of disposable income are from the distribution of the individual asset as a fraction of disposable income. Hence, the average fraction and the ratio of the average level of the asset to the average disposable income can be different.

Table 2: RKD ESTIMATES AT THE 725SEK THRESHOLD

	(1) Unemployment Duration D	(2) Duration D_1 (< 20 weeks)	(3) Duration D_2 (\geq 20 weeks)
I. 1999-2000: Kink in b_1 and b_2			
δ_k	-.0580 (.0046)	-.0254 (.0014)	-.0296 (.0039)
Robust s.e.			
95% CI - permutation test	[-.0595 ; -.0566]	[-.0319 ; -.0189]	[-.0402 ; -.019]
Quadratic model	-.0526 (.0184)	-.0299 (.0053)	-.0153 (.0156)
Cubic model	-.0526 (.0184)	-.0294 (.0053)	-.0154 (.0156)
N	177158	177158	177158
II. 2001: Kink in b_2 only			
δ_k	-.025 (.0049)	-.0155 (.0026)	-.0104 (.0028)
Robust s.e.			
95% CI - permutation test	[-.0325 ; -.0190]	[-.0167 ; -.0143]	[-.0115 ; -.0091]
Quadratic model	-.0631 (.0341)	-.0351 (.0087)	-.0196 (.0142)
Cubic model	-.0698 (.0341)	-.0366 (.0087)	-.02 (.0142)
N	46453	46453	46453
III. 2002- : Placebo			
δ_k	.0046 (.0048)	-.0004 (.0011)	.0062 (.0041)
Robust s.e.			
95% CI - permutation test	[.0017 ; .0075]	[-.0021 ; -.0011]	[.0053 ; .0071]
Quadratic model	.0212 (.0186)	.0002 (.0044)	.021 (.0159)
Cubic model	.0198 (.0188)	-.0006 (.0044)	.0204 (.0161)
N	189566	189566	189566

Notes: The table reports estimates of the change in slope δ_k , at the 725SEK threshold, in the relationship between daily wage and the total duration of unemployment D (col. (1)), the time D_1 spent on the first part of the Swedish UI profile (col. (2)) and the time D_2 spent on the second part of the Swedish UI profile (col. (3)). $D_1 = \sum_{t < 20wks} S(t)$ corresponds to duration censored at 20 weeks of unemployment. $D_2 = \sum_{t \geq 20wks} S(t)$ corresponds to unconditional duration spent unemployed after 20 weeks of unemployment (i.e., not conditional on having survived up to 20 weeks). Estimates are obtained from polynomial regressions of the form of equation (7) with a bandwidth size $h = 90$ SEK. These estimates are reported for three periods of interest. Panel I reports estimates for spells starting before July 2001. Panel II reports estimates for spells starting after July 2001 and before July 2002. Panel III reports estimates for spells starting after July 2002. Unemployment duration is capped at two years. We report for each estimates δ_k the White robust standard errors as well as 95% confidence intervals using the permutation test method of Ganong and Jaeger [2014].

Table 3: CONSUMPTION AS A FUNCTION OF TIME SPENT UNEMPLOYED RELATIVE TO PRE-UNEMPLOYMENT CONSUMPTION:
CONSUMPTION SURVEY ESTIMATES

	(1) Total expenditures	(2) Food	(3) Rents	(4) Purchase of new vehicles	(5) Furniture & house appliances	(6) Trans- portation	(7) Recre- ation	(8) Restau- rant
$\mathbb{1}[0 < t \leq 5 \text{ months}]$	-0.0447 (0.0325)	-0.0378 (0.0422)	-0.0344 (0.0413)	-0.422** (0.184)	-0.160* (0.0922)	-0.0726 (0.0737)	-0.105 (0.0672)	-0.106 (0.0837)
$\mathbb{1}[t > 5 \text{ months}]$	-0.130*** (0.0348)	-0.0751* (0.0453)	0.0119 (0.0411)	-0.172 (0.194)	-0.0570 (0.0958)	-0.326*** (0.0794)	-0.165** (0.0720)	-0.212** (0.0928)
Year fixed effects	×	×	×	×	×	×	×	×
Marital status	×	×	×	×	×	×	×	×
Family size	×	×	×	×	×	×	×	×
R^2	0.0331	0.0622	0.0148	0.0198	0.00991	0.0152	0.0109	0.0104
N	2558	2550	1128	2550	2388	2445	2551	1893

Notes: The Table reports estimates of the drop in household consumption over the unemployment spell in the HUT surveys, following model of equation (11). Because the HUT surveys collect information on household consumption expenditures at the time of the interview, HUT estimates can directly recover flow (bi-weekly) measures of consumption c_t , and avoid the parametric method of Figure 12. We restrict the sample to households where, at the date of the interview, one (and only one) individual is unemployed, or where, at the date of the interview, one (and only one) individual will become unemployed in the following 12 months. $\mathbb{1}[0 < t \leq 5]$ is an indicator for having a member of the household unemployed since less than 5 months at the time of the interview. $\mathbb{1}[t > 5]$ is an indicator for having a member of the household unemployed for more than 5 months at the time of the interview. Robust standard errors in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 4: WELFARE EVALUATION OF ACTUAL PROFILE USING ESTIMATED SUFFICIENT STATISTICS

	(1) Moral hazard costs MH_x	(2) Average consumption drop ΔC_x	(3) Value of tax-funded kroner spent CS_x/MH_x
Benefits given throughout the spell: b	1.53 (.13)	.23 (.01)	$\tilde{\lambda} \times \tilde{\gamma} \times .15$
Benefits given for the first 20 weeks: b_1	1.67 (.37)	.19 (.03)	$\tilde{\lambda}_1 \times \tilde{\gamma}_1 \times .11$
Benefits given after first 20 weeks: b_2	1.38 (.27)	.27 (.01)	$\tilde{\lambda}_2 \times \tilde{\gamma}_2 \times .20$

Notes: The Table reports estimates of the sufficient statistics needed to perform welfare evaluation of the profile of benefits in Sweden. The first row analyzes the welfare consequences of an increase in benefits b throughout the unemployment spell. The second row analyzes the welfare consequences of an increase in benefits b_1 during the first 20 weeks of the unemployment spell. The third row analyzes the welfare consequences of an increase in benefits b_2 after 20 weeks of unemployment. In each case, column (1) reports the moral hazard costs of an increase in benefits, which depend on the elasticity of unemployment duration as shown in equation (3), and which have been estimated in section 4. Standard errors computed using the Delta-method are in parenthesis. The second column reports the average consumption drop, using estimates from our preferred specification of the consumption profile (Panel B of Figure 12). To evaluate the welfare impact of the consumption drop and be comparable to the moral hazard cost, it needs to be scaled by $\tilde{\lambda}^k \times \tilde{\gamma}^k$, assuming the Taylor approximation of the relative difference in marginal utilities is accurate. Column (3) shows the value of a (tax-funded) kroner spent on benefits throughout the spell, during the first 20 weeks of unemployment, or after 20 weeks of unemployment. It is equal to the ratio of consumption smoothing to moral hazard costs, CS^k/MH^k . These values only depend on potential differences in risk aversion or Pareto weights for different unemployment lengths. Their comparison therefore enables evaluation of the benefit profile, keeping the average generosity fixed.

Appendix A: Proofs

Proof of Proposition 2

For convenience, we use the same notation as before, but we need time discounting as we consider $T \rightarrow \infty$. That is, we now have $S_t = \beta^t \prod_{s=1}^t (1 - h_s)$, we still have $D = \sum_{t=0}^T S_t$, and, with some abuse of notation, we write $T - D = \sum_{t=0}^T \beta^t - \sum_{t=0}^T S_t$. We assume $\beta(1 + r) = 1$, where β is the discount factor and r the interest rate.

Starting from a flat profile, consider an increase in the benefit level at duration t and duration $t + 1$ respectively. In our stationary environment, the impact of the respective changes on the remaining duration is the same at time 0 and 1 respectively. For an increase in b_t , the required increase in the tax rate equals

$$\begin{aligned} \frac{d\tau}{db_t} &= \frac{S_t}{T - D} + \sum_{j=0}^T \frac{d\frac{S_j}{T-D}}{db_t} b \\ &= \frac{S_t}{T - D} \left[1 + \frac{d\frac{D}{T-D}}{db_t} \frac{b}{\frac{D}{T-D}} \frac{D}{S_t} \right] \\ &= \frac{S_t}{T - D} \left[1 + \varepsilon_{D,b_t} \frac{D}{S_t} \right]. \end{aligned}$$

For an increase in b_{t+1} , the required increase in the tax rate equals

$$\frac{d\tau}{db_{t+1}} = \frac{S_{t+1}}{T - D} \left[1 + \varepsilon_{D,b_{t+1}} \frac{D}{S_{t+1}} \right].$$

Using $D = \sum_0^T S_t = S_0 + S_1 [\Sigma_1^T S_t / S_1] = 1 + S_1 \tilde{D}_1$, we find

$$\varepsilon_{D,b_{t+1}} = \left[\varepsilon_{S_1,b_{t+1}} + \varepsilon_{\tilde{D}_1,b_{t+1}} \right] \frac{D_1}{D},$$

where $\varepsilon_{\tilde{D}_1,b_{t+1}} = \varepsilon_{D,b_t}$ in our stationary environment with $T \rightarrow \infty$. That is, starting from a flat profile, an increase in b_t has the same impact on the continuation profile evaluated at time 0 as an increase in b_{t+1} has on the continuation profile evaluated at time 1. Note also that in a stationary environment with constant exit rate $h_t = h$, $D = 1/h$, while $D_1 = 1/h - 1$ and $S_{t+1} = S_t(1 - h)$. This implies

$$\frac{D_1}{S_{t+1}} = \frac{D}{S_t}.$$

Hence,

$$\frac{d\tau}{db_{t+1}} = \frac{S_{t+1}}{T - D} \left[1 + [\varepsilon_{S_1,b_{t+1}} + \varepsilon_{D,b_t}] \frac{D}{S_t} \right].$$

Using the same steps as in Proposition 1, we then find the respective moral hazard costs

$$MH_t = \varepsilon_{D,b_t} \frac{D}{S_t}$$

$$MH_{t+1} = [\varepsilon_{S_1,b_{t+1}} + \varepsilon_{D,b_t}] \frac{D}{S_t},$$

and thus $MH_{t+1} \geq MH_t$ since $\varepsilon_{S_1,b_{t+1}} \geq 0$. The Proposition immediately follows. \blacksquare

Proof of Proposition 3

We explicitly introduce a (risk-free) asset a_0 that the agent owns at the start of the unemployment spell. The agent receives a flat benefit level b and chooses each period t how much to consume out of her assets, which we denote by \tilde{c}_t . The agent faces the borrowing constraint $a_t \geq \bar{a}$. Since the environment and the policy are stationary, the continuation values during employment and unemployment are independent of the time spent unemployed, conditional on the asset level. We assume $\beta(1+r) = 1$, where β is the discount factor and r the interest rate.

When unemployed, the agent solves

$$V^u(a_t) = \max_{\tilde{c}_t, s_t} u(b_t + \tilde{c}_t, s_t) + \beta [h_t EV^e((1+r)[a_t - \tilde{c}_t]) + (1 - h_t) V^u((1+r)[a_t - \tilde{c}_t])].$$

We denote by EV^e the expected continuation value when becoming employed, which may depend on her search choice s_t through for example the reservation wage she sets. With $\beta(1+r) = 1$, the agent sets \tilde{c}_t such that

$$\partial u(b_t + \tilde{c}_t, s_t) / \partial c \geq h_t EV^{e'}((1+r)[a_t - \tilde{c}_t]) + (1 - h_t) V^{u'}((1+r)[a_t - \tilde{c}_t]).$$

A strict inequality holds if the agent is borrowing constrained $a_t = \bar{a}$. Using the envelope conditions $V^{u'}(a_t) = \partial u(c_t^u, s_t) / \partial c$ and $EV^{e'}(a_t) = Eu'(c_t^e)$, we find the Euler condition

$$\frac{\partial}{\partial c} u(c_t^u, s_t) \geq h_t Eu'(c_{t+1}^e) + (1 - h_t) \frac{\partial}{\partial c} u(c_{t+1}^u, s_{t+1}).$$

We now invoke the condition that $Eu'(c_{t+1}^e) < \partial u(c_{t+1}^u, s_{t+1}) / \partial c$, which is trivially satisfied when the reservation wage exceeds the benefit level and preferences are separable or job search efforts increase the marginal utility of consumption. The Euler condition then implies $\partial u(c_t^u, s_t) / \partial c < \partial u(c_{t+1}^u, s_{t+1}) / \partial c$ when the borrow constraint is not binding. This implies that consumption is decreasing during unemployment, since $s_t = s_{t+1}$ by stationarity. Moreover, in a stationary environment, the borrowing constraint will be binding when unemployed at $t+1$ if it was already binding at t . In that case, $\partial u(c_t^u, s_t) / \partial c = \partial u(c_{t+1}^u, s_{t+1}) / \partial c = u'(b)$. This proves the Proposition. \blacksquare

Appendix B: Robustness of the RK design

This appendix presents various robustness checks of the RK design.

Smoothness in unobserved heterogeneity The RK design relies on the assumption of smooth unobserved heterogeneity across the kink point in the schedule. In our context, this amounts to saying that individuals should not be able to locate their wage strategically around the kink point. A few tests can help assess the robustness of this assumption. First, Figure 14 plots the density of the daily wage and shows graphically the smoothness of the distribution of the assignment variable at the kink point in the UI schedules. Second, the panels in Figure 15 show how the mean values of different covariates (age, fraction of men, highly educated and foreigners) evolve smoothly with the daily wage. We do not find any non-linearity around the kink. Third, the panels in Figure 16 report our RKD estimates for different bandwidth sizes. For all periods we consider, the estimates remain stable for bandwidths above $h = 60\text{SEK}$.

Permutation tests Ganong and Jaeger [2014] suggest that it can be helpful to assess whether the true coefficient estimate is larger than those at “placebo” kinks placed away from the true kink. The idea behind their permutation test is that, if the counter-factual relationship between the assignment variable and the outcome (i.e., in the absence of the kink in the budget set) is non-linear, then the curvature in this relationship will result in many of the placebo estimates being large and statistically significant. In Table 2, we report 95% confidence interval based on this permutation procedure.

Non-parametric detection of kink point Figure 18 shows the R-squared when we run the RKD regression in (7) for “placebo” kinks placed in 10SEK increments from the true location of the threshold. This procedure, proposed in Landais [2015], and inspired from the time series literature on detection of trend breaks, enables to non-parametrically detect where a true kink is the most likely to be located in the data, by looking at the placebo kink where the R-squared is maximized. The figure shows that the R-squared is maximized in the exact location of the actual kink point, again supporting the evidence that there is in fact a change in slope that occurs at the actual kink point. In both panels A and B, the preferred location of the kink is extremely close to the true kink and the relationship between the placebo kink location and R-squared of the model exhibits a clear concave shape. In panel C, reassuringly, when there is no true kink at 725SEK, this relationship is perfectly flat.

Appendix C: Recovering high-frequency consumption from yearly consumption measure

To recover high-frequency consumption flows from our yearly consumption measure, we use the fact that C_t , the yearly consumption measured at month t of the unemployment spell, is simply

the sum of the monthly consumption flows between month $t - 11$ and month t :

$$C_t = \sum_{k=t-11}^t c_k$$

Our strategy then consists in specifying a functional form for the monthly consumption flows c_t as a function of time t . Once a parametric model for c_t has been specified, one can use the information from yearly consumption C_t to identify the parameters of c_t , using standard distance minimization techniques. Identification of the parameters for the c_t function relies on the fact that the C_t 's of individuals observed at different t 's will differ in their sequence of underlying c_t 's.

In practice, we focus on three type of functional forms for c_t : linear with a drop at unemployment, quadratic with a drop at unemployment, and step-functions.

Linear with a drop We specify the average monthly consumption flow at time t , c_t to be:

$$c_t = \begin{cases} \gamma_0 & \text{if } t < 0 \\ \gamma_1 + \alpha t & \text{if } t \geq 0 \end{cases}$$

In this case, we can rewrite yearly consumption levels $C_t = \sum_{k=t-11}^t c_k$ at each time t :

$$\begin{aligned} t = 0 : \quad C_0 &= 12\gamma_0 + (\gamma_1 - \gamma_0) \\ t = 1 : \quad C_1 &= 12\gamma_0 + 2(\gamma_1 - \gamma_0) + \alpha \\ t = 2 : \quad C_2 &= 12\gamma_0 + 3(\gamma_1 - \gamma_0) + 3\alpha \\ &\text{etc.} \end{aligned}$$

It follows that we can generally rewrite, $\forall t$:

$$C_t = 12\gamma_0 + (\gamma_1 - \gamma_0) \cdot A(t) + \alpha \cdot B(t)$$

where $A(t) = \min(12, t + 1)$ and $B(t) = \frac{1}{2} \cdot (t \cdot (t + 1) - \max(t - 11, 0) \cdot \max(0, t - 10))$.

Various distance minimization techniques can be used to estimate the parameters of the c_t function. In practice, we identify the drop in consumption at unemployment ($\gamma_1 - \gamma_0$) and the drift parameter α by regressing for all unemployed individuals in our sample their yearly consumption C_t on the two variables $A(t)$ and $B(t)$.

Quadratic with a drop The same method applies for the quadratic case, in which we specify the average monthly consumption flow at time t , c_t to be:

$$c_t = \begin{cases} \gamma_0 & \text{if } t < 0 \\ \gamma_1 + \alpha t + \beta t^2 & \text{if } t \geq 0 \end{cases}$$

In this case, we can generally rewrite, $\forall t$:

$$C_t = 12\gamma_0 + (\gamma_1 - \gamma_0) \cdot A(t) + \alpha \cdot B(t) + \beta \cdot D(t)$$

where, using the fact that $\sum_{i=0}^n i^2 = n(n+1)(2n+1)/6$, we have

$$D(t) = \frac{1}{6} \cdot (t(t+1)(2t+1) - \max(0, t-11) \max(0, t-10)(2 \max(0, t-11) + 1)).$$

We can again identify the drop in consumption at unemployment ($\gamma_1 - \gamma_0$) and the parameters α and β by regressing for all unemployed individuals in our sample their yearly consumption C_t on the three variables $A(t)$, $B(t)$ and $D(t)$.

Step-functions Step-functions take the following form (3-steps example):

$$c_t = \begin{cases} \gamma_0 & \text{if } t < 0 \\ \gamma_1 & \text{if } k > t \geq 0 \\ \gamma_2 & \text{if } t \geq k \end{cases}$$

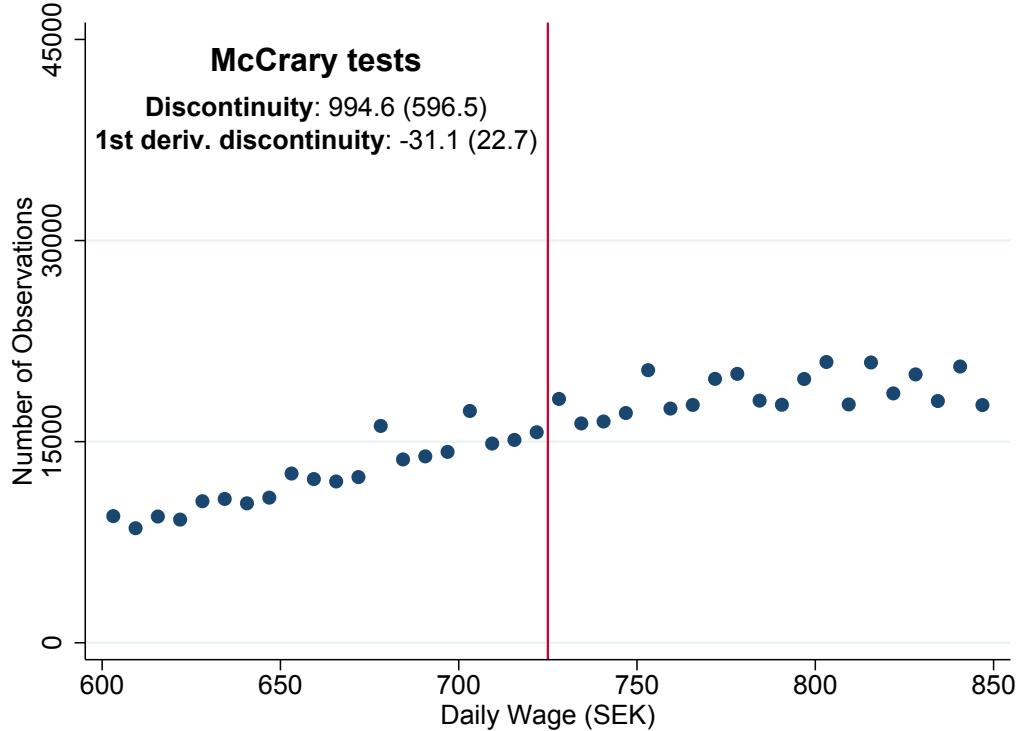
In practice, we focus on the three-step function, with steps at time $t = 0$ and $t = 5$ and use the same technique as previously to identify the parameters $\gamma_1 - \gamma_0$ and $\gamma_2 - \gamma_0$. Note that, for the sake of simplicity, we treat k as given, but we could also treat k (or the number of steps) as parameters to be estimated.

Selection on profile The key identification assumption so far is that the sequence of c_k 's is the same for individuals, irrespective of the duration t at which their yearly consumption is observed. But individuals observed at t and $t+1$ can be different because only individuals who have duration at least as long as $t+1$ will be observed in $t+1$, while not all individuals observed in t will survive to $t+1$. If there is correlation between unemployment duration and consumption profile, our estimate will suffer from dynamic selection bias. To remedy this, we take advantage of the fact that for all individuals observed at time t , we know their total ex-post duration. So that we can condition our exercise on total realized unemployment duration, and estimate separate profiles for individuals based on their unemployment duration.

In Panel B of Figure 12, we therefore estimate separately the step function profile for individuals whose total realized unemployment duration is less than 5 months and for individuals whose total realized unemployment duration is larger than 5 months.

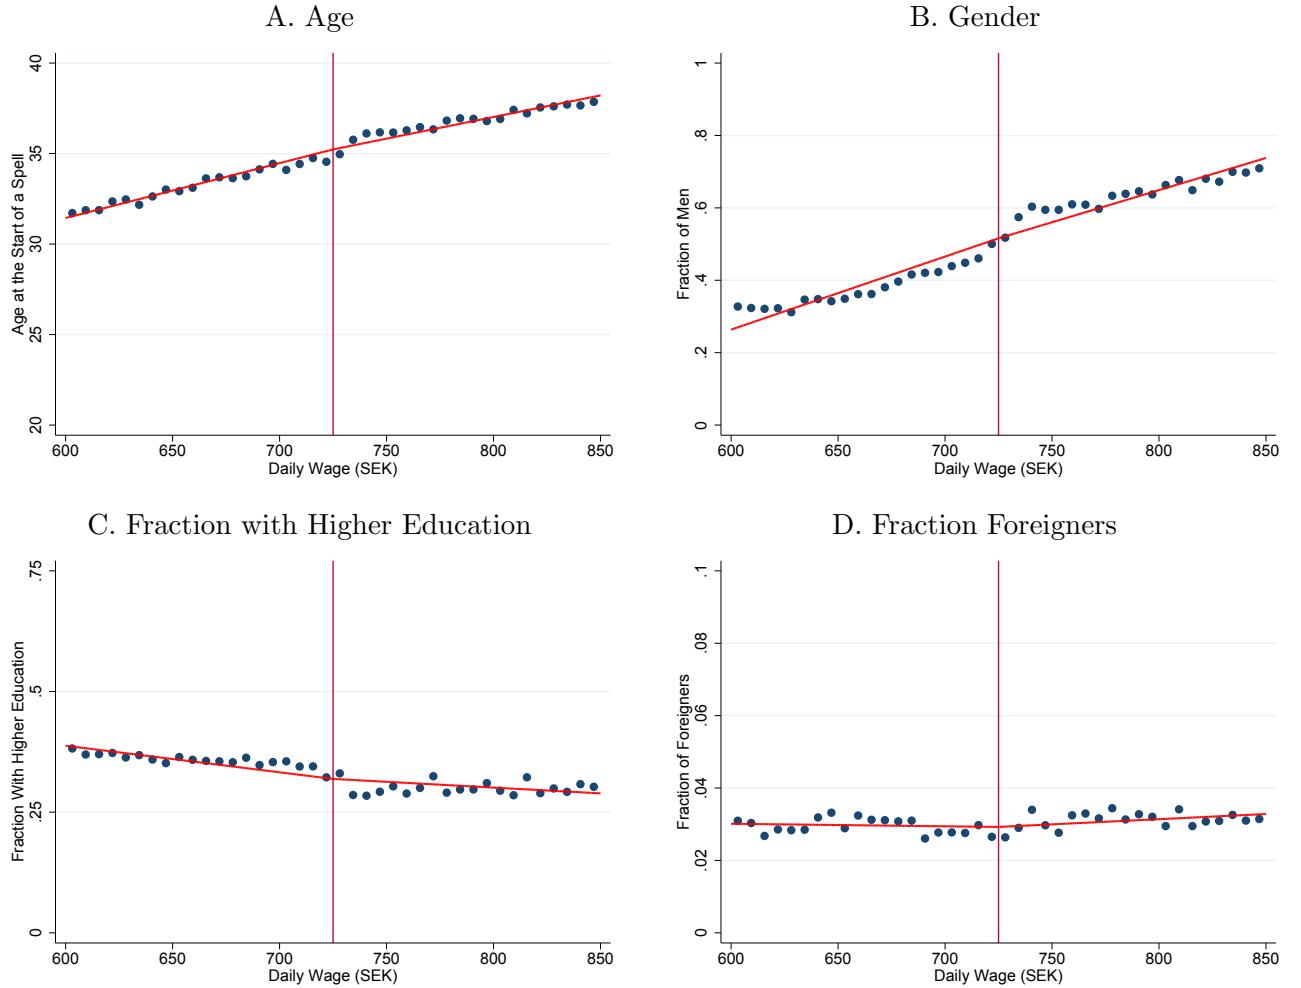
Appendix D: Figures

Figure 14: ROBUSTNESS OF THE RK DESIGN: P.D.F OF DAILY WAGE



Notes: The figure tests graphically the smoothness of the distribution of the assignment variable at the kink point in the UI schedules to assess the validity of the local random assignment assumption underlying the RK design. The Panel shows the probability density function of the daily wage around the 725SEK threshold. We also display two formal tests of the identifying assumptions of the RKD. The first is a standard McCrary test of the discontinuity of the p.d.f of the assignment variable. We report the difference in height of the p.d.f at the threshold. The second is a test for the continuity of the first derivative of the p.d.f. We report the coefficient estimate of the change in slope of the p.d.f in a regression of the number of individuals in each bin on polynomials of the assignment variable interacted with a dummy for being above the threshold. Both tests suggest smoothness of the assignment variable around the threshold, in support of the identifying assumptions of the RK design.

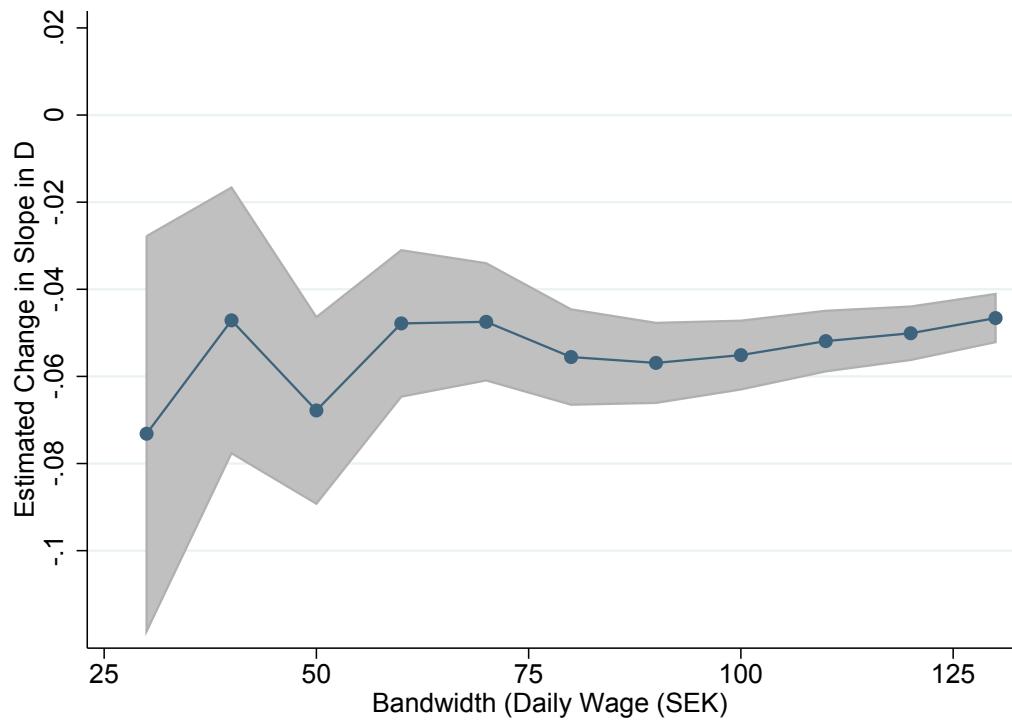
Figure 15: ROBUSTNESS OF THE RK DESIGN: COVARIATES



Notes: The figure tests the validity of the smoothness assumptions of the RK design. Each panel shows the mean values of a different covariate in bins of the assignment variable around the 725SEK threshold. The red lines display predicted values of polynomial regressions of the form of equation (7) in order to detect potential non-linearity around the threshold. The sample is restricted to all spells starting before July 2002, when kinks in the UI schedule are active at the 725SEK threshold. The graphs show evidence of smoothness in the evolution of all covariates at the kink, in support of the RKD identification assumptions.

Figure 16: RKD ESTIMATES AS A FUNCTION OF BANDWIDTH SIZE

A. 1999 - 2000



B. 2001

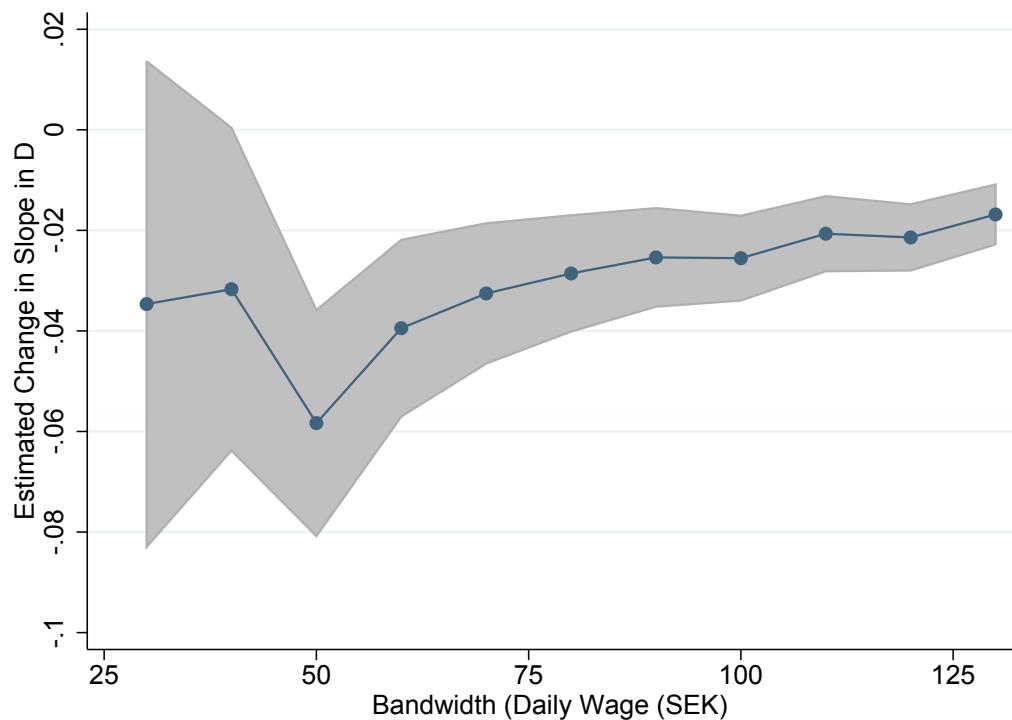
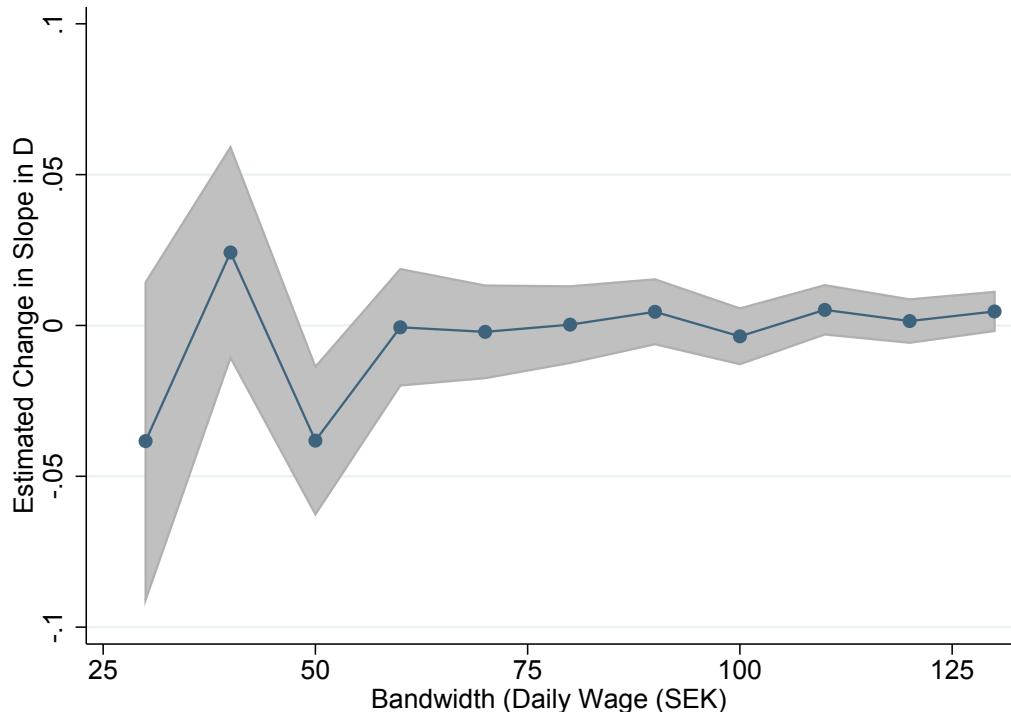


Figure 17: RKD ESTIMATES AS A FUNCTION OF BANDWIDTH SIZE (*continued*)

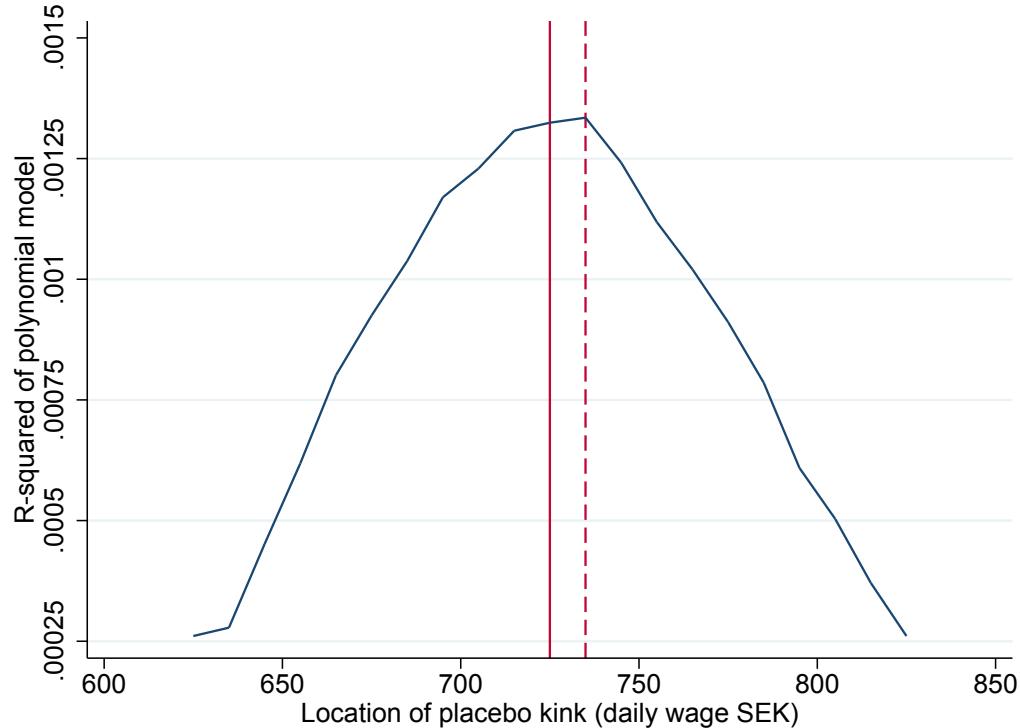
A. 2002-2005



Notes: The figure reports estimates of the change in slope with 95% robust confidence interval in the relationship between unemployment duration and daily wage at the 725SEK threshold using polynomial regressions of the form of equation (7) as a function of bandwidth size h . These estimates are reported for three periods of interest: 1999-2000 (i.e., spells starting before July 2001), 2001 (i.e., spells starting after July 2001 and before July 2002) and 2002- (i.e., spells starting after July 2002). Unemployment duration is defined as the number of weeks between registration at the PES and exiting the PES or finding any employment (part-time or full-time employment, entering a PES program with subsidized work or training, etc.). Unemployment duration is capped at two years. Sample is restricted to unemployed individuals with no earnings who report being searching for full-time employment.

Figure 18: NON-PARAMETRIC DETECTION OF KINK LOCATION

A. 1999 - 2000



B. 2001

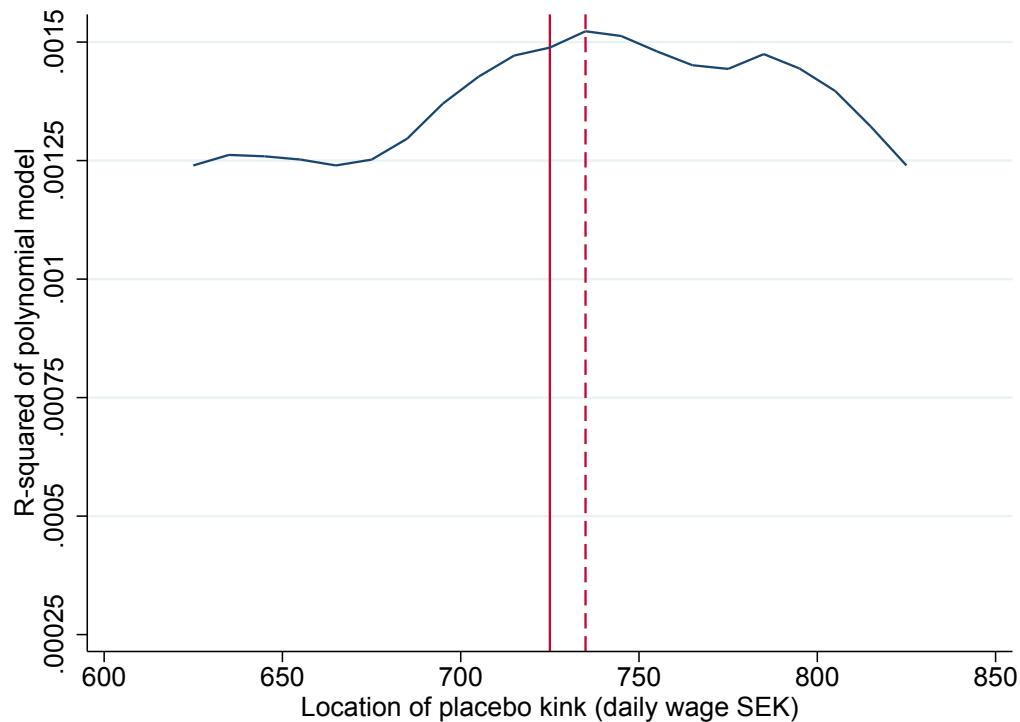
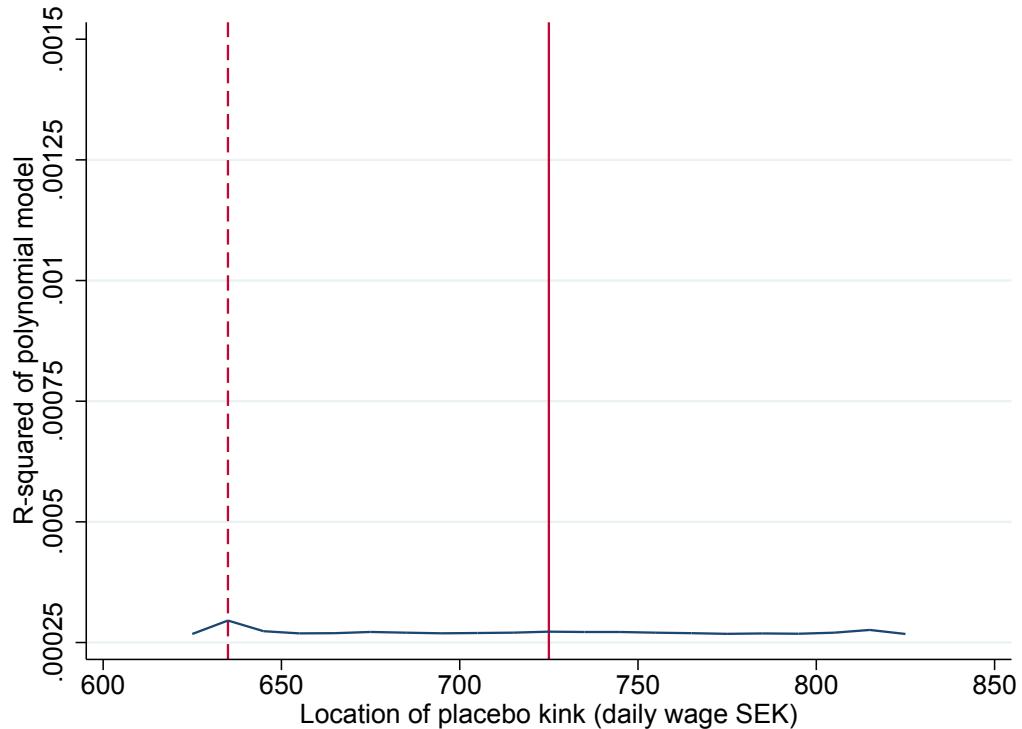


Figure 19: NON-PARAMETRIC DETECTION OF KINK LOCATION (*continued*)

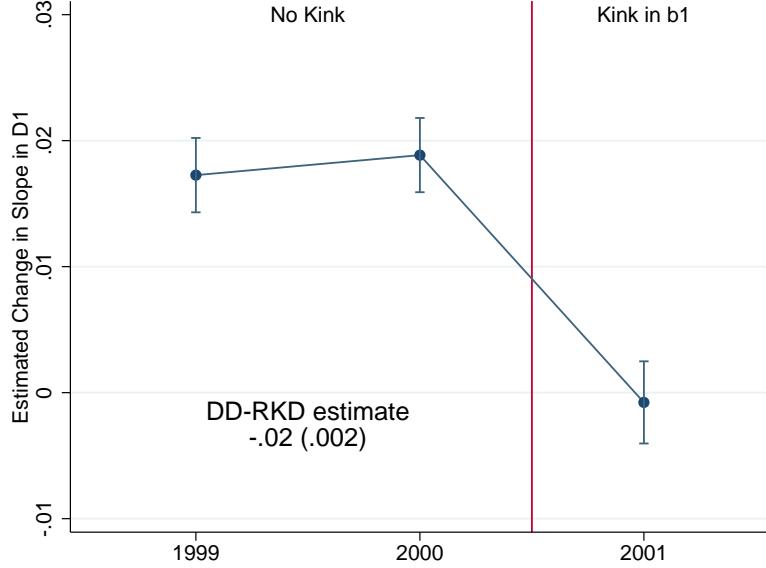
A. 2002-2005



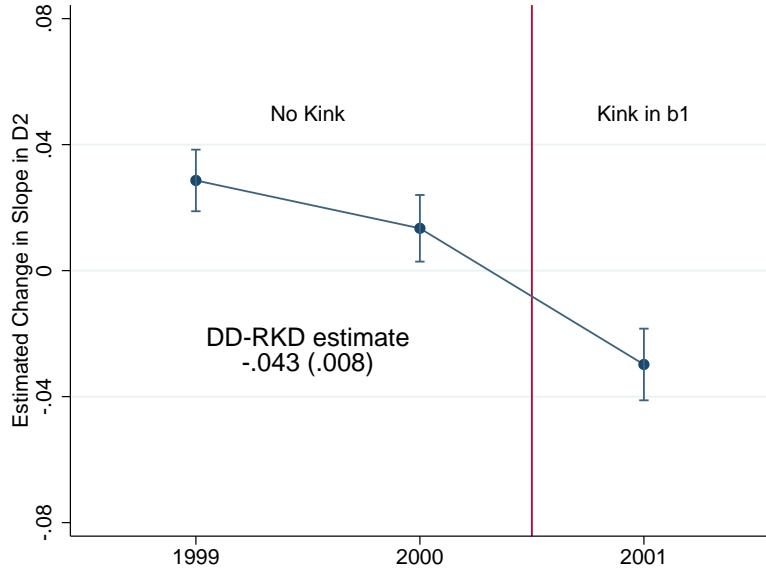
Notes: The figure reports the R-squared of polynomial regressions of the form of equation (7) for alternative (placebo) locations of the kink point \bar{w}_k for all observations with wages between 625SEK and 825SEK. The red line indicates the location of the true kink in the schedule. The dashed red line indicates the preferred location of the kink non-parametrically detected in the data, maximizing the R-squared of the model. These estimates are reported for three periods of interest: 1999-2000 (i.e., spells starting before July 2001), 2001 (i.e., spells starting after July 2001 and before July 2002) and 2002- (i.e., spells starting after July 2002). In both panels A and B, the preferred location of the kink is extremely close to the true kink and the relationship between the placebo kink location and R-squared of the model exhibits a clear concave shape. In panel C, when there is no true kink at 725SEK, this relationship is perfectly flat.

Figure 20: RKD ESTIMATES AT THE 850SEK THRESHOLD

A. Time spent on first part of the profile (D_1):
unemployment duration up to 20 weeks

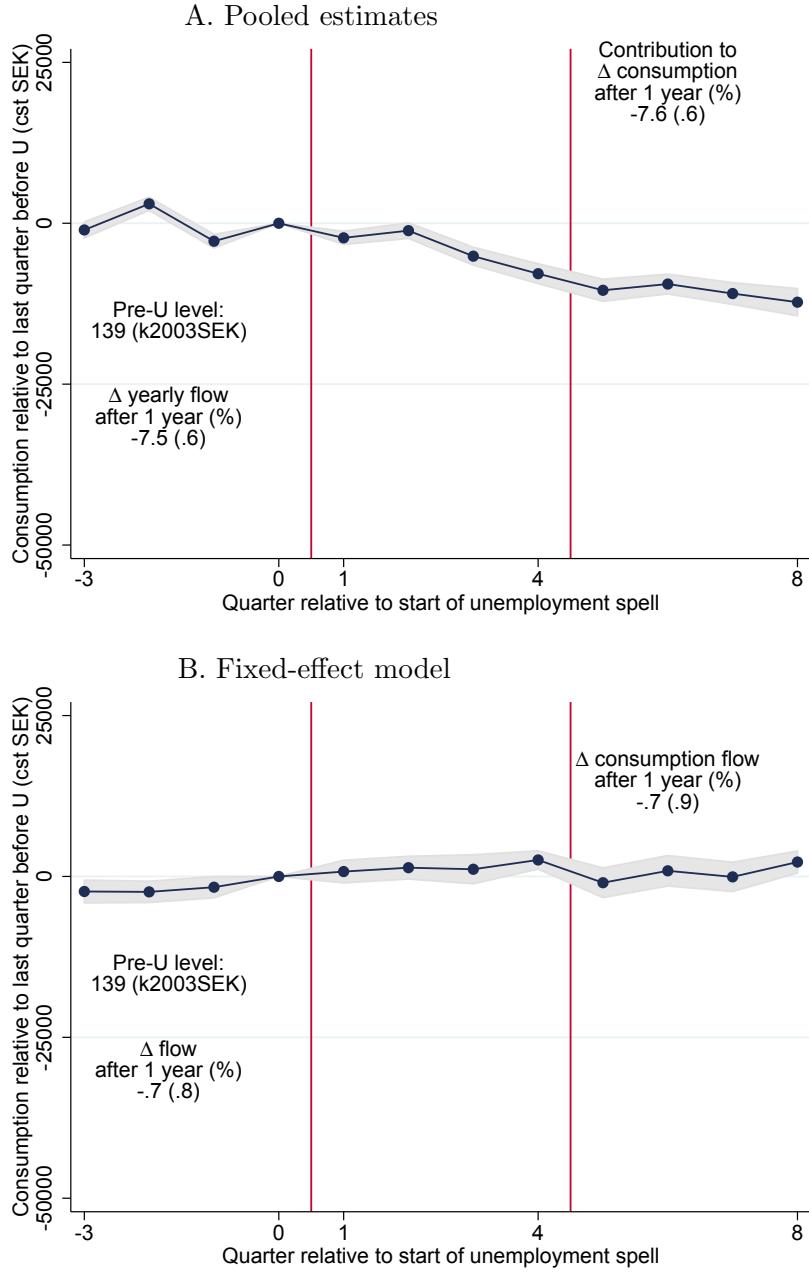


B. Time spent on second part of the profile (D_2):
unemployment duration after 20 weeks



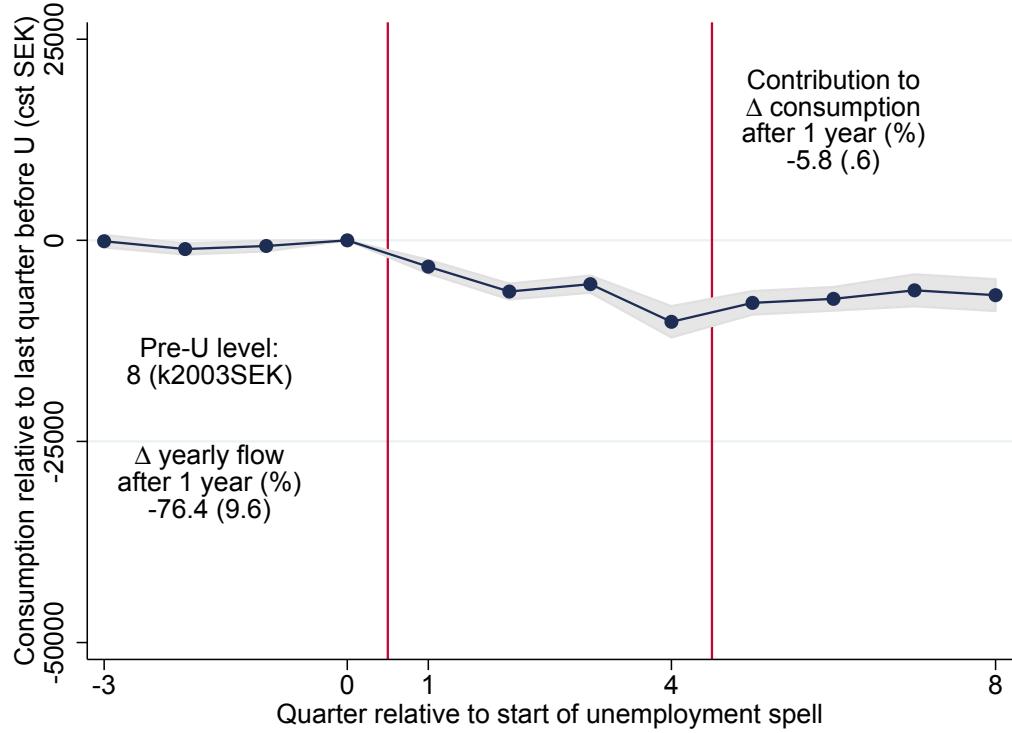
Notes: The figure reports estimates of the change in slope with 95% robust confidence interval in the relationship between daily wage and time spent in the first part of the benefit profile D_1 and time spent in the second part of the profile D_2 , at the 850SEK wage threshold, using polynomial regressions of the form of equation (7) with a bandwidth size $h = 90\text{SEK}$. These estimates are reported for years 1999-2000 (when the schedule is linear in both b_1 and b_2 at the 850SEK threshold), and for 2001 (i.e., spells starting after July 2001 and before July 2002 when the schedule exhibits a kink in b_1 only at the 850SEK threshold). For the period 1999-2000, when the schedule of both b_1 and b_2 is linear at the 850SEK threshold, the placebo test does not pass at the 850SEK kink. We therefore implement a strict DD-RKD, where the effect of b_1 is identified from the change between 2000 and 2001 in the change in slope of unemployment duration at the 850SEK wage threshold following the introduction of the kink in the schedule of b_1 . Formally, this implies $\alpha^{DD} = \frac{\Delta \hat{\delta}_k}{2000, 2001} / \frac{\Delta \nu_k}{2000, 2001}$.

Figure 21: YEARLY INCOME OF ALL OTHER MEMBERS OF THE HOUSEHOLD AS A FUNCTION OF UNEMPLOYMENT DURATION



Notes: In this figure, we explore how total yearly disposable income of all members of the household other than the unemployed individual ($Y_{-i,t}$) correlates with the time t (in quarters) since (or until) the onset of the unemployment spell. In Panel A, we report estimates from regression model (9) and plot the estimated coefficients β_t for the set of indicators $\mathbb{1}[T = t]$ for being observed, as of December of year n , in the t -th quarter since the onset of one's spell. Additionally, we report the average level of $Y_{-i,t}$ prior to becoming unemployed and the percentage change in $Y_{-i,t}$ after a full year of unemployment ($t = 5$). We finally report the contribution of this change in $Y_{-i,t}$ to the overall change in household consumption after a full year of unemployment. In Panel B, we take advantage of the panel structure of the data, and estimate a model with individual fixed effects using a within estimator. The figure follows from regression model (10) and plots the estimated coefficients $\tilde{\beta}_t$ for the same set of indicators $\mathbb{1}[T = t]$. Results show that after a year, the within-household change in the sum of disposable income of all other members of the household is not significantly different from zero. This suggests that in our context, the added-worker effect is not playing any significant role in increasing household consumption in response to an unemployment shock.

Figure 22: YEARLY CHANGE IN DEBT OF OF INDIVIDUALS WITH NO REAL ESTATE ASSETS AT THE START OF THE SPELL AS A FUNCTION OF UNEMPLOYMENT DURATION



Notes: In this figure, we explore how the yearly consumption flow from changes in debt of individuals with no real estate wealth correlates with the time t (in quarters) since (or until) the onset of the unemployment spell. Because we cannot precisely separate mortgage debt from other types of credit in the data, restricting the sample to individuals with no real estate is a direct and simple way to identify how non-mortgage related debt evolves over the unemployment spell. We report estimates from regression model (9) and plot the estimated coefficients β_t for the set of indicators $\mathbb{1}[T = t]$ for being observed, as of December of year n , in the t -th quarter since the onset of one's spell. Prior to becoming unemployed, increases in non-mortgage related debt contributes positively to yearly consumption by 8kSEK on average. After a year of unemployment ($t = 5$), unemployed individuals have decreased their consumption from changes in debt by 76.4%, which amounts to 5.8% reduction in total consumption.

Table 5: PRE-UNEMPLOYMENT CHARACTERISTICS OF INDIVIDUALS WITH SPELLS LONGER THAN 20 WEEKS. LINEAR PROBABILITY MODEL ESTIMATES

	(1)	(2)	(3)	(4)	(5)
Duration of future spell \geq 20 weeks					
Age: 30 to 39	0.129*** (0.00237)	0.118*** (0.00250)	0.116*** (0.00251)	0.119*** (0.00305)	0.120*** (0.00311)
Age: 40 to 49	0.164*** (0.00277)	0.153*** (0.00293)	0.153*** (0.00295)	0.162*** (0.00357)	0.163*** (0.00363)
Age: 50+	0.272*** (0.00288)	0.261*** (0.00307)	0.265*** (0.00319)	0.281*** (0.00367)	0.282*** (0.00371)
Married	0.0289*** (0.00243)	0.0283*** (0.00243)	0.0287*** (0.00244)	0.0185*** (0.00280)	0.0190*** (0.00281)
Gender: Female	-0.00226 (0.00192)	-0.00209 (0.00193)	-0.00279 (0.00193)	-0.0146*** (0.00230)	-0.0135*** (0.00230)
# of children	-0.0329*** (0.00193)	-0.0315*** (0.00193)	-0.0288*** (0.00193)	-0.0318*** (0.00231)	-0.0311*** (0.00233)
2nd quintile of income		0.0412*** (0.00319)	0.0436*** (0.00319)	0.0321*** (0.00409)	0.0319*** (0.00412)
3rd quintile of income		0.0842*** (0.00319)	0.0885*** (0.00319)	0.0850*** (0.00403)	0.0849*** (0.00415)
4th quintile of income		0.0471*** (0.00328)	0.0532*** (0.00329)	0.0532*** (0.00404)	0.0545*** (0.00421)
5th quintile of income		0.0453*** (0.00341)	0.0518*** (0.00345)	0.0589*** (0.00411)	0.0635*** (0.00431)
0<Net wealth \leq 200k		-0.0503*** (0.00234)	-0.0116*** (0.00271)	-0.0122*** (0.00315)	
200k<Net wealth \leq 500k		-0.0466*** (0.00324)	-0.0146*** (0.00350)	-0.0114*** (0.00425)	
500k<Net wealth \leq 5M		-0.0186*** (0.00300)	0.00576* (0.00336)	0.00774* (0.00418)	
Net wealth>5M		0.0731*** (0.0173)	0.0852*** (0.0172)	0.0866*** (0.0174)	
Fraction of portfolio in stocks					
3rd quartile			-0.000542 (0.00787)		
4th quartile			0.0303*** (0.00259)		
Leverage: debt / assets					
2nd quartile				0.0153*** (0.00390)	
3rd quartile				-0.0120*** (0.00322)	
4th quartile				-0.00629* (0.00361)	
<i>R</i> ²	0.0465	0.0490	0.0511	0.0624	0.0620
N	269931	269931	269931	190176	190176

Notes: The Table assesses the robustness of our welfare conclusions to dynamic selection on risk preferences and/or Pareto weights over the unemployment spell. We investigate how various observable characteristics correlate with the probability to experience a long unemployment spell. To do so, we restrict the sample to all individuals about to become unemployed in the next quarter and estimate a linear probability model where the outcome is an indicator variable for experiencing a future spell longer than 20 weeks. The default age category is 18 to 30 years old. Income refers to individual disposable income and results are relative to the first quintile. Net wealth results are relative to individuals with zero or negative net wealth at the start of the spell. We also investigate the effect of two portfolio characteristics, that, conditional on net wealth, are traditionally correlated with risk preferences. First, we look at the fraction of total wealth invested in stocks, and results are relative to the first two quartile of this distribution (50% of the sample have zero stocks prior to becoming unemployed). Second, we look at leverage defined as total debt divided by gross assets, and results are relative to the first quartile of leverage.