American Economic Review 2018, 108(4-5): 985–1033
https://doi.org/10.1257/aer.20160816

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

By Jonas Kolsrud, Camille Landais, Peter Nilsson, and Johannes Spinnewijn*

This paper provides a simple, yet robust framework to evaluate the time profile of benefits paid during an unemployment spell. We derive sufficient-statistics formulae capturing the marginal insurance value and incentive costs of unemployment benefits paid at different times during a spell. Our approach allows us to revisit separate arguments for inclining or declining profiles put forward in the theoretical literature and to identify welfare-improving changes in the benefit profile that account for all relevant arguments jointly. For the empirical implementation, we use administrative data on unemployment, linked to data on consumption, income, and wealth in Sweden. First, we exploit duration-dependent kinks in the replacement rate and find that, if anything, the moral hazard cost of benefits is larger when paid earlier in the spell. Second, we find that the drop in consumption affecting the insurance value of benefits is large from the start of the spell, but further increases throughout the spell. In trading off insurance and incentives, our analysis suggests that the flat benefit profile in Sweden has been too generous overall. However, both from the insurance and the incentives side, we find no evidence to support the introduction of a declining tilt in the profile.(JEL D82, E21, E24, J64, J65)

The key objective of social insurance programs is to provide insurance against adverse events while maintaining incentives. The impact of these adverse events is dynamic and so are the insurance value and incentive cost of social protection against

* Kolsrud: National Institute of Economic Research, Box 3116, SE-103 62 Stockholm, Sweden (email: jonas@kolsrud.se); Landais: London School of Economics, Houghton Street, London WC21 2AE, United Kingdom (email: c.landais@lse.ac.uk); Nilsson: Institute for International Economic Studies, Stockholm University, SE-106 91, Stockholm, Sweden (email: peter.nilsson@iies.su.se); Spinnewijn: London School of Economics, Houghton Street, London WC21 2AE, United Kingdom (email: j.spinnewijn@lse.ac.uk). This paper was accepted to the *AER* under the guidance of Hilary Hoynes, Coeditor. We thank Tony Atkinson, Richard Blundell, Raj Chetty, Liran Einav, Hugo Hopenhayn, Philipp Kircher, Henrik Kleven, Alan Manning, Arash Nekoei, Nicola Pavoni, Torsten Persson, Jean-Marc Robin, Emmanuel Saez, Florian Scheuer, Robert Shimer, Frans Spinnewyn, Iván Werning, Gabriel Zucman, and seminar participants at the NBER PF Spring Meeting, SED Warsaw, EEA Mannheim, DIW Berlin, Kiel, Zurich, Helsinki, Stanford, Leuven, Uppsala, IIES, Yale, IFS, Wharton, Sciences Po, UCLA, Sussex, Berkeley, MIT, Columbia, LMU, and LSE for helpful discussions and suggestions. We also thank Iain Bamford, Albert Brue-Perez, Jack Fisher, Benjamin Hartung, Panos Mavrokonstantis, and Yannick Schindler for excellent research assistance. We acknowledge financial support from the ERC (grant 716485 and 716485), FORTE (2015-00663) the Sloan foundation (NBER grant 22-2382-15-1-33-003), STICERD, and the CEP.
† Go to https://doi.org/10.1257/aer.20160816 to visit the article page for additional materials and author disclosure statement(s).

these events. As a consequence, the design of social insurance policies tends to be dynamic as well, specifying a schedule of benefits and taxes that are time-dependent. In the context of unemployment insurance (UI), the UI policy specifies a full benefit profile designed to balance incentives and insurance throughout the unemployment spell. Solving this dynamic problem can prove daunting, especially when adding important features of unemployment dynamics involving selection and nonstationarities. Indeed, there seems to be little consensus in practice on the optimal profile of UI benefits. Unemployment policies vary substantially across countries in the time profile of benefits paid during an unemployment spell, above and beyond differences in the overall generosity. In the United States, benefits are paid only during the first six months of unemployment. In other countries, like Belgium and Sweden, the unemployed could receive the same benefit level forever. Recent policy reforms, however, reduced the benefits for the long-term unemployed relative to the short-term unemployed.

This paper proposes and implements an evidence-based framework to characterize the optimal time profile of UI benefits and evaluate the welfare consequences of changes in the profile of existing UI policies. In doing so, this paper aims to bridge three different strands of the literature. There is an influential theoretical literature on optimal dynamic policies, but derived in stylized models that are often difficult to connect to the data (e.g., Shavell and Weiss 1979; Hopenhayn and Nicolini 1997; Werning 2002). An important empirical literature has analyzed the structural dynamics of unemployment, but without drawing the consequences for dynamic policies (e.g., Van den Berg 1990; Eckstein and Van den Berg 2007). Finally, a recent but growing empirical literature started evaluating social insurance design using the so-called *sufficient statistics* approach, but this literature has been mostly silent about the dynamic features of social insurance programs (e.g., Chetty 2008, Schmieder, von Wachter, and Bender 2012).

In the spirit of the sufficient-statistics approach, we derive a characterization of the optimal profile of unemployment benefits based on a limited set of high-level statistics. This simple, yet robust characterization provides new and transparent insights on the forces affecting the optimal trade-off between insurance and incentives costs throughout the unemployment spell. Our approach also identifies the relevant behavioral responses in this dynamic context to evaluate the welfare consequences of (local) changes in the policy. Our analysis therefore provides a clear guide for dynamic policy design and in particular for analyzing how insurance value and incentive cost of unemployment benefits evolve over the unemployment spell. We implement this approach empirically, using Swedish administrative data on unemployment, linked with survey data on consumption and tax register data on income and wealth.

We start by setting up a rich, dynamic model of unemployment that incorporates job search and consumption decisions and which allows for unobservable heterogeneity and duration dependence in job-finding rates in addition to unobservable heterogeneity in assets and preferences. Using dynamic envelope conditions, we show that the Baily-Chetty intuition (Baily 1978, Chetty 2006) generalizes for a dynamic unemployment policy: the UI benefits paid at time t of the unemployment spell should balance the corresponding insurance value with the implied moral hazard (or incentive) cost at the margin. At the optimal policy, the marginal value and cost are equalized for any part of the benefit profile. If they are not, one can identify (local) policy changes that increase welfare.

Like in the original Baily-Chetty formula, the insurance value and moral hazard cost of the dynamic policy can be expressed as a function of identifiable and estimable statistics. The incentive cost of benefits paid at time t of the unemployment spell depends only on the behavioral revenue effect, i.e., the effect of this benefit level on the government expenditures through agents' unemployment responses. This fiscal externality is fully captured by the responses of the survival rate throughout the unemployment spell, weighted by the benefit levels paid. In other words, regardless of the primitives underlying the dynamics of the agents' search behavior (e.g., heterogeneity versus true duration dependence in exit rates), these survival rate responses are sufficient to evaluate the incentive cost of changes in the benefit profile. From the insurance perspective, the marginal value of benefits paid at time t of the unemployment spell depends only on the average marginal utility of consumption for agents unemployed at time t. To capture this insurance value, we explore the robustness of the so-called consumption implementation approach, which consists in evaluating the marginal utility of consumption using observed consumption patterns over the unemployment spell and calibrated values of risk aversion. We demonstrate how the nature of selection into longer unemployment spells can affect the relative consumption smoothing gain from benefits paid at different time t of the unemployment spell.

The empirical part of this paper provides novel insights on the incentive costs and insurance value of UI benefits over the unemployment spell. We use a unique administrative dataset in Sweden based on unemployment and tax and asset registers for the universe of Swedish individuals from 1999 to 2007, combined with surveys on household consumption for a subset of the population. 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 nonparametrically 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 leverage the comprehensive information on income, transfers, and wealth from Swedish registers to construct a residual measure of household expenditures and, linking this measure to unemployment records, we identify how consumption expenditures change with unemployment and the duration of an unemployment spell in particular. We provide complementary and robustness analysis using survey data on consumption expenditures linked to unemployment records.

Our empirical analysis provides the following main results.

First, unemployment durations respond significantly to changes in benefit levels, whether these benefits are paid early or later in the spell. Furthermore, we find that the response to changes in benefits paid earlier in the spell is larger than the response to benefits paid later in the spell. 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 to comparable changes in the policy later in the spell. Importantly, such nonstationary forces, which may be driven by duration dependence or dynamic selection on returns to search effort over the unemployment spell, are large enough to offset the significant effect of forward-looking incentives.

Second, consumption expenditures drop substantially and early in the spell. We find that expenditures drop on average by 4.4 percent in the first 20 weeks of unemployment, compared to their pre-unemployment level. This drop deepens to 9.1 percent on average for those who are unemployed for longer. We also leverage the richness of the data to document the mechanisms underlying the observed patterns of consumption, and how they translate into consumption smoothing gains of UI benefits over the spell. We show that the role of selection effects in explaining the observed consumption patterns is rather limited. We document the role of assets and liquidity constraints in explaining the drop in consumption over the spell, and show the limited role of the added-worker effect in smoothing the unemployment shock, even for long-term unemployed. The consumption surveys also shed light on the types of consumption goods that individuals adjust over the spell, including substitution toward home production and away from durable goods. Taken together, our evidence consistently indicates that the consumption smoothing value of UI is higher for the long-term unemployed.

Finally, our empirical estimates can be mapped into the sufficient statistics derived in our theoretical analysis, allowing for a transparent local evaluation of the benefit levels in a two-part benefit profile. Our baseline implementation assumes preference homogeneity, separability between consumption and leisure, and the absence of other externalities (besides the described fiscal externality). Since we find that in Sweden the incentive costs are high relative to the drop in consumption throughout the unemployment spell, our implementation suggests that reducing the generosity of either part of the benefit profile increases welfare for reasonable values of risk aversion. The incentive cost, however, decreases over the unemployment spell as do the consumption expenditures of the unemployed. In the absence of offsetting selection on preferences, our estimates suggest a welfare gain of decreasing the marginal krona spent on the short-term unemployed that is more than twice as high as decreasing the marginal krona spent on the long-term unemployed. As the benefit profile was flat during our period of study, this suggests that the introduction of an inclining benefit profile could have increased welfare. We provide a complementary welfare analysis based on a structural model. We use our empirical analysis of mechanisms to inform the choice of primitives of the model that we then calibrate to match the sufficient statistics underlying our local policy recommendations. The structural analysis allows us to go beyond these local policy recommendations, but relies on the structure of the calibrated model. The calibration exercise indicates that an inclining tilt remains welfare improving when lowering the overall generosity of the policy.

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), generalized by Chetty (2006), and recently reviewed in Chetty and Finkelstein (2013). 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

[1] Recent dynamic extensions of the Baily-Chetty formula can be found in Schmieder, von Wachter, and Bender (2012), analyzing the potential benefit duration for a given benefit level, and in Spinnewijn (2015), providing a formula for the optimal intercept and slope of a linear benefit profile.

are hard to take to the data. Our analysis shows how the previously identified forces (e.g., 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 relates to a large literature on labor supply effects of social insurance. This literature has focused on exploiting isolated sources of variation in one part of the benefit profile.[2] We contribute by explicitly using duration-dependent variation in benefits and identifying the welfare-relevant unemployment responses for *multiple parts* of the benefit profile. Our analysis 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 drops as a response of income shocks and unemployment in particular (e.g., Gruber 1997). We provide novel insights on the evolution of consumption as a function of time spent unemployed. We do this using administrative data on income and wealth to construct a residual, registry-based measure of consumption, which allows us to identify moral hazard costs and consumption responses for the very same sample of unemployed.[3]

The remainder of the paper proceeds as follows. Section I analyzes the characterization and implementation of sufficient-statistics formulae for the evaluation of local policy changes in a dynamic model of unemployment. Section II describes our data and the policy context in Sweden. Section III describes our regression kink design and provides estimates of the policy-relevant unemployment elasticities. Section IV analyzes how consumption evolves during the unemployment spell and how this translates to the consumption smoothing gains of UI. Section V analyzes welfare complementing the implementation of the sufficient statistics with a calibration exercise of our structural model. Section VI concludes.

I. Model

This section sets up a dynamic model of unemployment and identifies the key trade-offs in designing the time profile of the unemployment benefits. We provide a characterization of the optimal profile in a nonstationary environment with heterogeneous agents. In the spirit of the "sufficient-statistics" literature, our approach consists in identifying the minimal level of information necessary for this characterization. Our focus goes beyond the primitives of the environment and the assumptions on agents' behavior in our specific model. Instead, we aim to identify the observable variables that are relevant for policy in a broad class of models and can be estimated empirically.

A. *Setup*

We first describe the setup of our dynamic unemployment model, building on the dynamic model in Chetty (2006), the agents' choices, and the unemployment

[2] See Krueger and Meyer (2002) for a review on the labor supply effects of social insurance. Recent examples analyzing variation in UI are Rothstein (2011); Farber and Valletta (2011); Landais (2015); Card et al. (2015); and Johnston and Mas (2015).

[3] See, for instance, Kostøl and Mogstad (2015), Kreiner, Lassen, and Leth-Petersen (2014), and Pistaferri (2015) for a survey of recent developments of consumption analysis using registry data.

policy. We try to save on notation in the main text, but provide more details in the online technical Appendix A. We consider a partial equilibrium framework with a continuum of agents with mass 1. The model is in discrete time t, starts at $t = 1$, and ends at $t = T$.

Each agent i starts unemployed and remains unemployed until she finds work. Once an agent has found work, she remains employed until the end. When employed, the agent earns w, when unemployed she earns 0. Before the start of the model, the government commits to an unemployment policy P providing insurance against the unemployment risk: the policy specifies an unemployment benefit profile depending on the duration of the ongoing unemployment spell (i.e., a benefit level b_t for each time t if the unemployment spell is still ongoing) and a uniform tax τ paid when employed.

Job Search.—Each agent i decides at each time t how much search effort $s_{i,t}$ to exert as long as she is unemployed. This effort level determines the agent's exit probability at time t. We denote the agent's exit rate out of unemployment at time t by $h_{i,t}(s_{i,t})$. We allow this mapping to depend on the type of agent i, capturing heterogeneity in employability across agents, and the time t she has spent unemployed, capturing differences in employment prospects due to the time spent unemployed.[4] The agent's probability to be unemployed after t periods equals the survival probability $S_{i,t} \equiv \prod_{t'=1}^{t-1} (1 - h_{i,t'}(s_{i,t'}))$ with $S_{i,1} = 1$. While we cannot observe an agent's specific survival probability, we can observe the population average of survival probabilities $S_t \equiv \int S_{i,t} di$.

Intertemporal Consumption.—Each agent i decides at each time t how much to borrow or save (at interest rate r). An agent starts the unemployment spell with asset level $a_{i,1}$, but borrowing constraints prevent her from running down her asset below $\bar{a}_i$ at any time. The agent's savings decisions determine her consumption level throughout the unemployment spell and when reemployed. We denote these levels by $c_{i,t}^u$ and $c_{i,t}^e$ for when unemployed and employed respectively.[5] While we cannot observe an agent's contingent consumption plan, we can observe average levels of consumption, for example at different spell lengths, $\bar{c}_t^u = \int (S_{i,t}/S_t) c_{i,t}^u di$.

Preferences.—Each agent i has time-separable preferences (with discount factor β) with per-period utility increasing in consumption but decreasing in search efforts exerted when unemployed. We denote agent i's per-period utility by $v_i^u(c_{i,t}^u, s_{i,t})$ and $v_i^e(c_{i,t}^e)$. We allow for preference heterogeneity and nonseparable preferences in the characterization of the optimal policy, but will assume preference homogeneity and separability between consumption and efforts in our baseline implementation. Each agent chooses how much to search and how much to consume in order to maximize her expected utility, taking the unemployment policy P as given. The

[4] Potential reasons for *true* duration-dependence in exit rates are human capital depreciation (see Acemoglu 1995 and Ljungqvist and Sargent 1998) and stock-flow sampling (see Coles and Smith 1998). We assume exogenous exit rate functions that only depend on the agent's search, but do not directly depend on other job seekers' search like in rationing models (e.g., Michaillat 2012) or on the unemployment policy like in employer screening models (e.g., Lockwood 1991). We discuss this further in Section ID and online Appendix Section A.2.

[5] The agent's consumption choice at time t will depend on her unemployment history. In particular, even when employed, the agent's unemployment history will have affected her asset accumulation and thus her optimal consumption at time t. We introduce formal notation to denote the relevant state variables in the online technical Appendix.

dynamics of the agent's behavior depend on her assets and the time spent unemployed in addition to the unemployment policy. To reduce notation, we will drop the arguments of the agent's behavior. We denote the agent's value function of her maximization problem by $V_i(P)$, accounting for her optimal consumption and search choices and potentially binding borrowing constraints.

Unemployment Policy.—We characterize the welfare impact of local deviations from the unemployment policy P. In particular, we consider changes in the benefit b_t paid at time t of the unemployment spell, starting from any benefit profile $\{b_t\}_{t=1}^{T}$. Our expressions naturally generalize for changes in step-wise policies, paying benefit level b_k for part k of the unemployment spell (from time B_{k-1} until B_k). Flat benefit profiles with no or few steps are very common in practice. In Sweden, the unemployment policy in Sweden is entirely flat for some workers and exists of two parts for others, with the benefit dropping to a lower (positive) level at 20 weeks of unemployment.[6] Our implementation evaluates the benefit levels of this two-tier benefit profile.

The government's budget depends on the expected benefit payments paid to the unemployed and the expected tax revenues received from the employed. Note that the average *unemployment duration D* simply equals the sum of the survival rates at each duration $\sum_{t=1}^{T} S_t$. Similarly, $D_k = \sum_{B_{k-1}+1}^{B_k} S_t$ denotes the expected time spent unemployed while receiving benefit b_k, which we refer to as the average *benefit duration*. We ignore time discounting in our characterization of the optimal policy (i.e., $1 + r = \beta = 1$), but generalize this in online technical Appendix A.

The government's budget simplifies to

$$G(P) = [T - D]\tau - \sum_{t=1}^{T} S_t b_t. \tag{1}$$

Social welfare associated with an unemployment policy P can be written as the Lagrangian

$$W(P) = \int V_i(P)\, di + \lambda\left[G(P) - \bar{G}\right], \tag{2}$$

where λ equals the Lagrange multiplier on the government's budget constraint and $\bar{G}$ is an exogenous revenue constraint. We assume that the social welfare function is differentiable.

B. *Dynamic Sufficient Statistics*

We consider the welfare impact of local policy deviations, which we decompose into the corresponding consumption smoothing gains and moral hazard cost. Our approach does not provide an explicit characterization of the optimal policy, but is sufficient to test for the (local) optimality of the policy in place. Evaluating local policy changes, away from the optimal policy, is of interest to identify how welfare can be increased and how the policy can be changed toward the optimal policy (if welfare is concave in the policy variables).

[6] We discuss the details of the Swedish unemployment policy in Section IIA.

Consider now an increase in the benefit level b_t in period t of the unemployment spell. The total impact on welfare depends on how much the unemployed value this increase in benefits b_t relative to its budgetary cost,

$$(3) \qquad \frac{\partial W(P)}{\partial b_t} = \int \frac{\partial V_i(P)}{\partial b_t} di + \lambda \frac{\partial G(P)}{\partial b_t}.$$

This welfare effect depends on the agents' behavioral responses to the policy, but only to the extent that the agents' behavior has consequences that they did not internalize themselves. Indeed, an agent's response to a policy change will have only a second-order impact on her own welfare $V_i(P)$. Assuming differentiability, this follows from the envelope conditions $\partial V_i / \partial x^z_{i,t'} = 0$, which hold for any behavior $x^z_{i,t'}$ the agent optimizes over, at any time t', when employed $(z = e)$ or unemployed $(z = u)$ and when the borrowing constraint is binding or not (see Chetty 2006).[7] So we only need to account for the impact of behavioral responses on the government's budget $G(P)$ and the direct impact of the policy change on agents' welfare, which proves particularly powerful in this dynamic context.

Moral Hazard.—Consider first the budgetary impact from an increase in b_t. The first effect from increasing the benefit level is mechanical and depends on the share of workers still unemployed after t periods, S_t. The second effect is behavioral and is determined by the budgetary cost of the agents' reduced search in response to the more generous benefit. This depends on the induced change in the average survival rates *throughout* the unemployment spell,

$$(4) \qquad \frac{\partial G(P)}{\partial b_t} = -S_t - \sum_{t'=1}^{T} \frac{\partial S_{t'}}{\partial b_t}(b_{t'} + \tau) = -S_t \times \left[1 + \sum_{t'=1}^{T} \frac{S_{t'}(b_{t'} + \tau)}{S_t b_t} \varepsilon_{t',t}\right]$$

$$(5) \qquad \equiv -S_t \times [1 + MH_t].$$

The moral hazard cost MH_t of an increase in b_t simply equals the weighted sum of the elasticities $\varepsilon_{t',t} = (\partial S_{t'} / \partial b_t) / (S_{t'} / b_t)$ of the average survival rate $S_{t'}$ with respect to the benefit level b_t. The elasticities are weighted by the relative share of the budget spent at different times during the unemployment spell. The budgetary spillover effects of a change in b_t on other parts of the policy is less relevant the less generous these other parts are. There is, however, a correction for the tax rate because more time spent unemployed also reduces the taxes received from employment.

Evaluated at a flat profile ($b_t = \bar{b}$ for all t), the moral hazard cost of an increase at time t is fully determined by the response in the average duration D, scaled by the survival rate at t,

$$(6) \qquad MH_t = \frac{\partial D / \partial b_t}{S_t}(\bar{b} + \tau) = \frac{D(\bar{b} + \tau)}{S_t \bar{b}} \varepsilon_{D, b_t}.$$

[7] Changes in the choice variables might be discontinuous in response to small policy changes. In principle we can allow for such discontinuous behavioral responses if they average out when integrating across heterogeneous individuals so that the social welfare function is differentiable.

This average duration response combines the potentially heterogeneous responses by unemployed workers throughout the unemployment spell, including responses earlier in the spell in anticipation of the increase in b_t and selection effects later in the spell due to the increase in b_t.

Consumption Smoothing.—Let us now turn to the insurance value of an increase in the benefit b_t. Due to the envelope conditions, the welfare increase is completely captured by the marginal utility of consumption at this time of the spell for the agents who are still unemployed,

$$\int \frac{\partial V_i(P)}{\partial b_t} di = \int S_{i,t} \frac{\partial v_i^u(c_{i,t}^u, s_{i,t})}{\partial c_{i,t}^u} di,$$
$$\equiv S_t \times E_t^u \left[\frac{\partial v_i^u(c_{i,t}^u, s_{i,t})}{\partial c_{i,t}^u} \right].$$

As defined above, the expectation operator E_t^u takes the weighted average over all individuals' marginal utility of consumption in the tth period of the unemployment spell (with weights $S_{i,t}/S_t$). By analogy to the budgetary cost, we can write

$$(7) \qquad \int \frac{\partial V_i(P)}{\partial b_t} di / \lambda = S_t \times [1 + CS_t],$$

where the consumption smoothing gain $CS_t \equiv \left\{ E_t^u \left[\frac{\partial v_i^u(c_{i,t}^u, s_{i,t})}{\partial c_{i,t}^u} \right] - \lambda \right\} / \lambda$. Since the Lagrange multiplier λ equals the shadow cost of the government's budget constraint, the consumption smoothing gains can be interpreted as the return of a government dollar spent to the unemployed in period t of the unemployment spell relative to the value of an unconditional transfer.[8] Importantly, in spite of potential heterogeneity across agents and in their responses to the policy change, the welfare gain is fully captured by the average marginal utility of consumption at time t of the unemployment spell.

Welfare Impact.—An optimal unemployment policy balances consumption smoothing gains and moral hazard costs. A dynamic benefit profile allows solving this trade-off at each point during the unemployment spell. Combining expressions (3), (5), and (7), we find

$$\frac{\partial W(P)}{\partial b_t} = \lambda S_t \times [CS_t - MH_t].$$

[8] Note that when the government can provide such lump sum transfer, it would be optimally set such that λ equals the average marginal value of resources at the start of this model. More generally, the consumption smoothing gain CS_t corresponds to the net social marginal welfare weight assigned to the unemployed at time t.

An increase (decrease) in benefit b_t increases welfare as long as the consumption smoothing gains are larger (smaller) than the moral hazard cost. This implies a natural characterization of the optimal policy.

PROPOSITION 1: *Consider an unemployment policy P, charging tax τ to the employed and paying a dynamic benefit profile $\{b_t\}_{t=1}^T$ to the unemployed. Assuming differentiability, an interior, optimal policy needs to satisfy*

$$\frac{E_t^u\left[\frac{\partial v_i^u(c_{i,t}^u, s_{i,t})}{\partial c_{i,t}^u}\right] - \lambda}{\lambda} = \sum_{t'=1}^{T} \frac{S_{t'}(b_{t'} + \tau)}{S_t b_t} \times \varepsilon_{t',t} \quad \textit{for each } t, \tag{8}$$

$$\frac{\lambda - E^e\left[\frac{\partial v_i^e(c_{i,t}^e)}{\partial c_{i,t}^e}\right]}{\lambda} = \sum_{t'=1}^{T} \frac{S_{t'}(b_{t'} + \tau)}{(T-D)\,\tau} \times \varepsilon_{t',\tau}, \tag{9}$$

and the budget constraint $G(P) = \bar{G}$.

PROOF:

See online Appendix A.

The expectation operator E_t^u is defined as before and takes the weighted average over all individuals unemployed at time t (with weights $S_{i,t}/S_t$). Similarly, E^e takes the weighted average over all individuals and times in employment (with weights $(1 - S_{i,t})/[T-D]$). The conditions for all benefit levels and the tax level in Proposition 1 can be combined to recover the well-known Baily-Chetty formula (Baily 1978 and Chetty 2006) for a flat benefit profile ($b_t = \bar{b}$),[9]

$$\frac{E^u\left[\frac{\partial v_i^u(c_{i,t}^u, s_{i,t})}{\partial c_{i,t}^u}\right] - E^e\left[\frac{\partial v_i^e(c_{i,t}^e)}{\partial c_{i,t}^e}\right]}{E^e\left[\frac{\partial v_i^e(c_{i,t}^e)}{\partial c_{i,t}^e}\right]} \cong \frac{\bar{b} + \tau}{\bar{b}} \varepsilon_{D,\bar{b}}. \tag{10}$$

Our analysis extends the "sufficient statistics" approach to the dynamics of the unemployment policy, aiming to provide a simple, yet robust guide for dynamic policy design. As we argue below, the characterization depends on a limited set of empirically implementable moments and is robust to the primitives and specific assumptions of the underlying model. Our dynamic extension also overcomes challenges that have constrained empirical and theoretical work in identifying the key dynamic forces. Empirically, identifying the role of different, nonstationary forces underlying a job seeker's environment, including the role of unobserved heterogeneity, proves daunting. Several studies have tried to estimate or calibrate

[9] The expectation operator E^u in condition (10) takes the weighted average over all unemployment periods (with weights $S_{i,t}/D$). The approximation relies on the unemployment response to taxes to be small. The exact expression for the right-hand side is $[1 + \frac{\bar{b}+\tau}{\bar{b}}\varepsilon_{D,\bar{b}}]/[1 + \frac{\bar{b}+\tau}{\tau}\varepsilon_{T-D,\tau}] - 1$. Note that the standard Baily-Chetty formulation uses the elasticity with respect to a budget-balanced increase in the benefit level, joint with an increase in the tax level, and ignores other tax distortions in the economy. Our model allows the tax to cover general expenditures $\bar{G}$ and our expressions are in terms of partial elasticities, which are more transparent for multidimensional policies and correspond more directly to the policy variation we exploit in the empirical analysis.

the contribution to the negative duration-dependence of exit rates from dynamic selection effects, *true* duration-dependence in the search environment (e.g., skill-depreciation or stock-flow sampling of vacancies), or an interaction of the two (e.g., duration-based employer screening).[10] Theoretically, it has also proven difficult to derive the optimal benefit profile and, in particular, the impact of nonstationary forces and heterogeneity (see Shimer and Werning 2006 and Pavoni 2009). In contrast, Proposition 1 provides a robust mapping from a nonstationary model with heterogeneous agents into a set of implementable moments to evaluate the benefit profile.[11]

C. *Implementation*

We first consider the implementability of our characterization, which guides our empirical analysis in Sections III and IV. Our focus is on the benefit levels of a two-part policy (b_1, b_2, B) paying benefit b_1 until time B and b_2 thereafter, like in place in Sweden (see Section IIA) and illustrated in panels A.I and B.I of Figure 1. This section clarifies additional assumptions and the policy variation required for the implementation we propose.

Moral Hazard Cost.—An extensive literature has analyzed unemployment responses to changes in the unemployment policy. Our analysis indicates that it is essential to have variation in unemployment benefits at different times during the unemployment spell.

For a two-part profile, the benefit duration D_1 (D_2), which denotes the expected time spent receiving benefit b_1 (b_2), corresponds to the area under the survival function before (after) B, as illustrated in panels A.II and B.II of Figure 1. The moral hazard cost of changing the benefit level b_k during part k of the policy fully depends on the response in both benefit durations,

$$MH_1 = \frac{b_1 + \tau}{b_1} \varepsilon_{D_1, b_1} + \frac{D_2(b_2 + \tau)}{D_1 b_1} \varepsilon_{D_2, b_1}, \tag{11}$$

$$MH_2 = \frac{D_1(b_1 + \tau)}{D_2 b_2} \varepsilon_{D_1, b_2} + \frac{b_2 + \tau}{b_2} \varepsilon_{D_2, b_2}. \tag{12}$$

Panels A and B of Figure 1 illustrate how estimating the moral hazard costs requires duration-dependent policy variation. Rather than having benefits change throughout the spell, which would be sufficient to evaluate a flat policy, we need changes in benefits paid only to the short-term db_1 or to the long-term unemployed db_2.

Our evaluation of the benefit profile is conditional on B, which determines the potential duration of the two parts. Still, our expressions can be used to approximate the moral hazard cost of changing the potential benefit durations, as analyzed in

[10] See Ljungqvist and Sargent (1998) and Machin and Manning (1999) for reviews on the negative duration dependence of exit rates out of unemployment. See Kroft, Lange, and Notowidigdo (2013) and Alvarez, Borovicková, and Shimer (2016) for recent examples.

[11] In Section VB, we also show how this mapping can be useful to uncover the role of stationary versus nonstationary forces from the estimated moments and understand their respective role for the optimal timing of benefits.

Schmieder, von Wachter, and Bender (2012). Indeed, an increase in potential benefit duration from B to $B+1$ corresponds to a (discrete) change of benefit b_{B+1} at $B+1$ (from level b_2 to level b_1), where the moral hazard cost of a (marginal) change in b_{B+1} equals

$$MH_{b_{B+1}} = \frac{\partial D_1 / \partial b_{B+1}}{S_{B+1}}(b_1 + \tau) + \frac{\partial D_2 / \partial b_{B+1}}{S_{B+1}}(b_2 + \tau).$$

While there is no such policy variation in the Swedish context, Schmieder and von Wachter (2016) review recent empirical work that analyzes either duration responses in unemployment or UI benefit receipt to changes in potential benefit duration. The expression above shows that the responses in both the benefit duration D_1 and the average duration D (since $D_2 = D - D_1$) affect the fiscal externality, unless the tax is small ($\tau \cong 0$) and no benefits are paid after time B ($b_2 = 0$). Importantly, any evaluation of the potential benefit duration would be conditional on the benefit levels and does not allow one to evaluate the tilt of the profile itself.

Consumption Smoothing.—Attempts at quantifying the consumption smoothing gains of UI policies have been more scarce as the estimation of differences in marginal utility levels proves difficult in practice. We follow the "consumption implementation" approach (Gruber 1997; Chetty 2006), relating the difference in marginal utilities to the difference in consumption levels, and extend this approach to our dynamic setting. Using consumption wedges to actually quantify the relevant consumption smoothing gains of UI requires the following assumptions.[12]

First, we rely on approximations of the marginal utility of consumption using Taylor expansions, assuming that third- and higher-order derivatives of the utility function are small. That is,

$$\frac{\partial v_i^u}{\partial c}(c_{i,t}^u, s_{i,t}) \cong \frac{\partial v_i^u}{\partial c}(\tilde{c}, s_{i,t}) \times \left[1 - \tilde{\gamma}_{i,t} \times \frac{\tilde{c} - c_{i,t}^u}{\tilde{c}}\right],$$

where $\tilde{\gamma}_{i,t} \equiv \tilde{c}\frac{\partial^2 v_i^u}{\partial c^2}(\tilde{c}, s_{i,t}) / \frac{\partial v_i^u}{\partial c}(\tilde{c}, s_{i,t})$ equals the relative risk aversion.[13]

Second, we assume that preferences over consumption are separable from leisure, i.e., $\partial v_i^u(c,s)/\partial c = \partial v_i^e(c)/\partial c = v_i'(c)$, so that consumption smoothing benefits do not depend on other behavior or the employment status itself, but only on the consumption wedges. This excludes potentially important complementarities between consumption and leisure during unemployment.[14]

[12] Note that the limitations of the consumption-based implementation have inspired alternative approaches relating the marginal utility gap to observable behavioral responses: Chetty (2008) decomposes unemployment responses in liquidity and substitution effects, Shimer and Werning (2007) analyze reservation wage responses. The extension of these alternative approaches to a dynamic setting seems promising, but requires that the policy variation used for the static implementation also changes over the unemployment spell.

[13] If the third-order derivative of the utility function is nonnegligible, the consumption smoothing gains depend on an additional term that depends on the coefficient of relative prudence, corresponding to precautionary saving motives (see Chetty 2006). We calculate the magnitude of this approximation error in Section VC.

[14] In Section IVB, we discuss the issues related to this assumption in more detail. One example is the substitution toward home production when no longer employed (e.g., Aguiar and Hurst 2005).

Third, we express the consumption smoothing gains from an increase in unemployment benefits relative to an increase in resources just before the onset of the unemployment spell (denoted by $v_i'(c_{i,0})$). This normalization emphasizes the insurance value of the policy.[15]

Fourth, we assume that preferences are homogeneous (i.e., $v_i(c) = v(c)$), but we consider the implications of preference-based selection over the unemployment spell in Section VB.

Under these four assumptions, we can approximate the consumption smoothing gains by[16]

$$CS_k = \frac{\frac{1}{D_k}\int \sum_{B_{k-1}+1}^{B_k} S_{i,t} v'\left(c_{i,t}^u\right) di - \int v'\left(c_{i,0}\right) di}{\int v'\left(c_{i,0}\right) di} \tag{13}$$

$$\cong \frac{v'\left(\bar{c}_k^u\right) - v'\left(\bar{c}_0\right)}{v'\left(\bar{c}_0\right)} \cong -\frac{v''\left(\bar{c}_0\right)\bar{c}_0}{v'\left(\bar{c}_0\right)} \times \frac{\bar{c}_0 - \bar{c}_k^u}{\bar{c}_0}, \tag{14}$$

where $\bar{c}_0$ and $\bar{c}_k^u$ denote the average consumption level before the onset of the spell and during part k of the spell. The resulting expression directly relates to the original approximation in Baily (1978) and highlights the role of the profile of the average consumption level over the unemployment spell to evaluate the unemployment benefit profile. If the unemployed consume less the longer they are unemployed, ceteris paribus, unemployment benefits are more valuable later in the spell.

For a two-part profile, the implementation thus comes down to calculating the average wedge in consumption for the short-term unemployed and the long-term unemployed, as illustrated in panel C of Figure 1. Importantly, no policy variation is needed to estimate these consumption wedges.

D. *Robustness*

We briefly consider the robustness of our sufficient-statistics characterization in a dynamic context. As argued before by Chetty (2006), the set of moments in Proposition 1 are sufficient to provide local evaluations of the unemployment policy, independently of the underlying primitives. That is, when different values for our models' parameters map into the same values for the identified moments for a given policy, the local policy recommendations remain the same. In particular, our

[15] In our stylized model, all individuals start unemployed, so the Lagrange multiplier (evaluated at the optimal policy) equals $\lambda = \int \partial V/\partial a_{i,1} di = \int \partial V/\partial b_1 di$. In online Appendix Section A.2.4, we consider a more general model where agents can start employed or unemployed and may experience multiple spells. The Lagrange multiplier then corresponds to the average marginal utility of consumption at the start of the model across all individuals. By considering the marginal utility of consumption before the onset of the unemployment spell in our implementation, we are capturing the insurance value of the unemployment policy, while ignoring the value of redistributing between different types of workers who face different layoff risks. Importantly, the evaluation of (budget-balanced) changes in the benefit profile is independent of this normalization.

[16] The first approximation relies on a Taylor expansion of $v'\left(c_{i,t}^u\right)$ for each individual i and time t around the average consumption level during the corresponding part k of the unemployment spell, $\bar{c}_k^u \equiv \frac{1}{D_k}\int \sum_{B_{k-1}+1}^{B_k} c_{i,t}^u S_{i,t} di$, using

$$\frac{1}{D_k}\int \sum_{B_{k-1}}^{B_k} S_{i,t} v''\left(\bar{c}_k^u\right)\left[c_{i,t}^u - \bar{c}_t^u\right] di = 0.$$

This also applies at time $t = 0$, just before the onset of the unemployment spell. The second approximation simply uses a Taylor expansion of $v'\left(\bar{c}_k^u\right)$ around $\bar{c}_0$.

dynamic model explicitly allows for (exogenous) heterogeneity in exit rate functions across agents $(h_{i,t}(\cdot) \neq h_{j,t}(\cdot))$ and variation in exit rates over the unemployment spell $(h_{i,t}(\cdot) \neq h_{i,t'}(\cdot))$. While separating unobserved heterogeneity and *true* duration-dependence in exit rates is hard, our approach shows that this is unnecessary for estimating the moral hazard cost and its evolution over the unemployment spell. The intuition is that only the survival responses need to be known, as they fully determine the fiscal externality of job seekers' behavior. As is well known, this result critically relies on the application of the envelope conditions for the job seekers' behavior (see Chetty 2006, Chetty and Finkelstein 2013). The result also indicates that the foundations of our dynamic model of search and consumption can be further extended without (substantially) changing the characterization of the optimal benefit profile. That is, the same set of moments will continue to determine the local evaluation of the unemployment policy.[17] Our baseline setup illustrates this robustness explicitly for exogenous heterogeneity across agents and across durations, but the intuition generalizes to other models of search and self-insurance.[18]

While our implementation is robust to heterogeneity in employment prospects and assets and the corresponding dynamic selection, our baseline implementation assumes preferences that are homogeneous and separable in consumption and search efforts. Under this assumption the relative consumption smoothing gains, for example for short-term and long-term unemployed, simplify to the corresponding relative consumption drops, which thus become sufficient to recommend welfare-improving changes in the tilt of the benefit profile. This avoids the well-known challenge for the consumption-based implementation of translating consumption wedges to welfare (see Chetty and Finkelstein 2013).[19] We assess this challenging conversion from consumption into welfare in more depth and also gauge the potential for dynamic selection on preferences in Sections IVB and VB.[20]

Finally, as our characterization critically relies on the application of the envelope theorem, by the same token, the presence of other externalities (not internalized by agents, but relevant for welfare) would affect the optimal policy characterization. Recent work has analyzed the impact of different types of externalities on the

[17] Chetty (2006) has shown how the simple formula (10) characterizing the flat benefit profile continues to apply with leisure benefits from non-employment, alternative means of self-insurance, spousal labor supply, human capital decisions, etc. See the review chapter by Chetty and Finkelstein (2013) for a more detailed discussion of different advantages and challenges for the sufficient-statistics approach.

[18] In online Appendix Section A.2.2, we show how our model can be indeed extended to multiple unemployment spells, allowing for moral hazard on the job and different means of self-insurance, while the same formulae continue to apply. The relevant variables for evaluating the benefit profile are the overall unemployment rate and the survival rates S_t at different unemployment durations, averaged over multiple spells. Layoff responses to UI policy affect the magnitude of the policy-relevant elasticities, but only if moral hazard on-the-job were to be important. In online Appendices A and B, we provide and discuss evidence based on the probability distribution function (PDF) of pre-unemployment wages around a kink in the unemployment policy that indicates that layoff rates do not respond strongly to the unemployment policy in our empirical context.

[19] See Chetty (2008) or Landais (2015) for the development of alternative methods, exploiting comparative statics of effort choices, in order to evaluate consumption smoothing gains. These methods could circumvent the issue of having to make assumptions regarding dynamic selection on risk preferences.

[20] Note also that our model assumes a utilitarian social welfare function. See Andrews and Miller (2013) for a discussion on the aggregation of individual welfare gains under preference heterogeneity in the context of the Baily-Chetty formula. With heterogeneous Pareto weights, the dynamic selection based on these weights will matter as well.

characterization of a static unemployment policy.[21,22] These insights generalize in our dynamic setting, but are of particular relevance for our analysis when the externality depends on the timing of the unemployment benefits. In online Appendix Section A.2, we demonstrate this in the context of employer screening (e.g., Lockwood 1991), which gives rise to negative duration-dependence when employers use unemployment spell length as a negative signal of unobserved productivity. In such a setting, job seekers do not internalize their impact on the hiring probability for other job seekers, which happens through the relative survival rates of different types. As shown by Lehr (2017) for a flat benefit profile, the unemployment policy can affect this hiring externality, but only if the relative survival rate of different productivity types depends on the unemployment policy. We show in the online Appendix that for a dynamic benefit profile, the externality-adjusted moral hazard cost equals

$$MH_t^x \equiv \sum_{t'=1}^{T} \frac{S_{t'}}{S_t}\left[\frac{b_{t'}+\tau}{b_t}\varepsilon_{t',t} - E_{t'}^u\left(\omega_{t'}^h \frac{\partial h_{t'}}{\partial b_t}\right)\right], \tag{15}$$

where $\omega_{t'}^h$ corresponds to an agent's welfare gain of finding a job at time t' and $\partial h_{t'}/\partial b_t$ equals the change in the job finding rate due to the employer's hiring response. If the hiring response depends on the timing of the benefits, the externality-adjustment could vary over the spell. This requires the relative survival rate of different productivity types to change with the timing of benefits. In the Appendix, we demonstrate that types with higher returns to search are more responsive to changes in benefits early on, but due to their low survival into longer unemployment spells, can be less responsive to changes in benefits later on. Hence, when types with higher returns to search are also more productive, the hiring externality can be positive for benefits paid early in the spell, but negative for benefits paid late in the spell.

II. Context and Data

To implement our formulae and evaluate the profile of UI benefits, two 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 merging unemployment registers, tax registers (with exhaustive information on income and wealth), and household consumption surveys. We present here the institutional background and data used in our empirical implementation.

[21] For example, Nekoei and Weber (2015) account for the fiscal impact of reservation wage responses, conditional on unemployment duration, which tend to be small relative to the duration responses themselves. Spinnewijn (2015) accounts for internalities due to biased beliefs about employment prospects. Landais, Michaillat, and Saez (2010) adjust the characterization to account for frictions in the labor market and general equilibrium effects.

[22] In online Appendix Section A.2.4, we show how our framework can account for the fiscal externality created by the presence of an income tax used to fund other government expenditures, which we use to analyze the sensitivity of our welfare recommendations.

A. *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 percent 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 percent 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 counseling activities and, potentially, active labor market programs set up by the Public Employment Service. Like 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. Over the period 1999 to 2007, more than 85 percent of all workers were contributing to an unemployment fund. Our sample focuses on workers with more than six months of employment history prior to being laid-off and who contribute to UI funds.

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.[23] Full-time workers would get daily benefits of 80 percent of their pre-unemployment daily wage throughout the spell (i.e., for as long as they remain unemployed), with daily benefits capped at 580SEK a day ($\approx$ US\$63 a day, or US\$320 a week).[24] The cap thus applies for daily wages above 725SEK ($\approx$ US\$399 a week).[25] 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 $\approx$ US\$467.5 a week) while the cap for benefits 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 $\approx$ US\$500 a week) and the cap for benefits received after the first 20 weeks was increased to 680SEK.[26]

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 III how the 2001 and 2002 variations in the duration-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.

[23] The potential duration of benefits is theoretically infinite in Sweden during our period of interest, and there is no exhaustion point of UI benefits.

[24] We use the following exchange rate: 1SEK $\approx$ US\$0.11.

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

[26] Some unions have launched their own complementary UI schemes which further increased the cap (by up to three 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.

B. *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 Arbetslöshetsförsäkringen) in Sweden. The data contain information from 1999 to 2007 on the date the unemployed registered with the PES (which is a prerequisite 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. We define unemployment as a spell of non-employment following an involuntary job loss, and during which an individual has zero earnings, receives unemployment benefits, and reports searching for a full-time job.[27] To define the start date of an unemployment spell, we use the registration date at the PES. 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.).[28]

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 sociodemographic 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 come from the wealth tax register (Förmögenhetsregistret), which covers the asset portfolios for the universe of Swedish individuals from 1999 to 2007. The register contains detailed information on the stock of all financial assets (including debt) and real assets as of December of each year.[29] For the 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.[30] Data on asset balances as of December are complemented with data on financial asset transactions (KURU) and real estate transactions from the housing registries. The comprehensiveness and detailed nature of both the income and wealth data in Sweden is exceptional, providing a unique opportunity to investigate what means individuals use to smooth consumption (transfers, asset rebalancing, increase in debt, etc.). We take advantage of the richness of the income and wealth data to construct a registry-based measure of annual household consumption expenditures as of December of each year. Our approach builds on previous attempts

[27] Involuntary job loss is defined as a layoff or a quit following a "valid reason." Valid reasons for quitting a job are defined as being sick or injured from working, being bullied at work, or not being paid out one's wage by one's employer. Quits are reviewed by the Public Employment Service at the moment an individual registers a new spell and if the quit is made because of a "valid reason," the individual is eligible for UI, and a notification is made in the PES data, allowing us to observe such "quits under valid reasons." Note, however, that quits are a small fractions of spells in our sample: 95.0 percent of job separations observed in our data are due to layoffs.

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

[29] All 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.

[30] All asset holdings are reported to the tax administration at the individual level. We aggregated assets at the household level using household identifiers from the registry data.

to measure consumption from registry data (e.g., Koijen, Van Nieuwerburgh, and Vestman 2015; Browning and Leth-Petersen 2003) and is closely related to Eika, Mogstad, and Vestad (2017) in exploiting additional information on asset portfolio choices and returns to reduce measurement error and excess dispersion of consumption measures based solely on first-differencing asset stocks. All details regarding the construction of the registry-based measure of consumption are given in online Appendix C.1.

The richness of the Swedish administrative data, its universal coverage and panel structure, enables us to identify duration responses and within-household consumption responses to unemployment shocks for the very same sample of individuals, and with a unique degree of precision. An important characteristic of the registry-based measure of consumption, though, is that it captures the sum of annual consumption expenditures as of December of each year. While this peculiarity does not represent a serious limitation to identify higher frequency flows of consumption throughout the unemployment spell, as shown in Section IV, we complement our data with information on consumption available through the yearly household budget survey (HUT, Hushållens utgifter), which provides direct measures of biweekly consumption expenditures at the moment the household is surveyed. From 2003 to 2009, individuals sampled in the HUT can be matched to the registry data, which allows us to reconstruct the full employment history of individuals whose household is surveyed in the HUT. While the HUT has a small sample size, and does not have a panel structure, it provides a flow measure of consumption at the time of the survey, and allows us to explore patterns of consumption responses for different categories of expenditures.

In Table 1, we provide summary statistics on unemployment, demographics, income, and wealth for our sample of unemployed individuals used in the duration analysis of Section III and in the consumption analysis of Section IV. The average unemployment spell of unemployed in this sample is 26.8 weeks. The average time spent unemployed during the first 20 weeks of the spell equals $D_1 = 13.5$ weeks. The average time spent unemployed in the second part of the benefit profile (after the first 20 weeks of the spell) equals $D_2 = 13.2$ weeks. The average replacement rate is 72 percent. Sociodemographic characteristics, reported in panel II of Table 1, show that the unemployed are relatively young (35 years old on average) and a minority of them are married or cohabiting (39 percent on average).

Prior to the onset of a spell, the average unemployed in the sample has yearly gross earnings (before any tax or payroll contribution) of 190,300 SEK ($\approx$ US$21,000).[31] A large fraction of unemployed in the sample starts unemployment with no or negative net wealth. Most wealth is held in the form of real estate. Liquid assets such as bank holdings represent less than 30 percent of yearly earnings at the start of the spell. Total debt, which mostly comprises mortgage, student loans, and credit card debt, is fairly large and represents on average almost 200 percent of the yearly earnings of an unemployed at the onset of her spell.

[31] All figures are expressed in constant SEK2003. 1SEK2003 $\approx$ 2003 US$0.11.

III. Duration Responses

This section analyzes 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.

A. *Regression Kink Design: Strategy and 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 between January 1999 and July 2001 (panel A.I), for spells starting between July 2001 and July 2002 (panel B.I), and for spells starting after July 2002 (panel C.I). 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, generating variation in the policy that allows us to identify the effect of an overall change in the benefit level (i.e., 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 durations 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, $\alpha_k = \partial Y / \partial b_k$. 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 distribution of unobserved heterogeneity μ is assumed to be evolving smoothly around the kink point. This means that the conditional density of the wage $(f_{w|\mu}(\cdot))$ and its partial derivative with respect to the wage $(\partial f_{w|\mu}(\cdot)/\partial w)$ are assumed to be continuous in the neighborhood of $\bar{w}_k$. This second assumption implies imperfect sorting around the threshold $\bar{w}_k$, i.e., individuals cannot have perfect control over their assignment in the schedule. We provide in online Appendix B various tests to assess the robustness of these identifying assumptions and the validity of our RK design.

Under these two identifying assumptions, α_k can be identified as

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

In practice, we provide estimates $\hat{\alpha}_k = \hat{\delta}_k / \nu_k$ 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:

$$(16) \qquad E[Y|w] = \beta_0 + \beta_1(w - \bar{w}_k) + \delta_k(w - \bar{w}_k) \cdot \mathbf{1}[w \geq \bar{w}_k].$$

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.I). 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 3, where we report for each policy period the point estimate and 95 percent robust confidence interval of the change in slope $\hat{\delta}_k$ for the regression model in (16). We use a linear specification and a bandwidth $h = 90$SEK around the 725SEK threshold. Because 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 Tobit models on the right-censored data.[32] In line with the evidence presented in Figure 2, the estimated change in slope is large and significant for spells starting before July 2001. In Figure 3 we also report the implied benefit elasticity of unemployment duration, computed as $\varepsilon_{D,b_k} = \hat{\delta}_k \cdot \frac{\bar{w}_k}{D^{cap}}$, where $\hat{\delta}_k$ is the estimated marginal slope change, and D^{cap} is the observed average duration at the kink.[33] Standard errors on the elasticities are obtained from bootstrapping using 50 replications. The implied elasticity of unemployment duration with respect to an overall change in the benefit level (both b_1 and b_2) is $\varepsilon_{D,b} = 1.53$ (0.13). 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 with respect to b_2 of $\varepsilon_{D,b_2} = 0.68$ (0.13).

[32] An alternative solution is to get rid of observations who have an ongoing spell at the moment the schedule changes. In online Appendix Figure B-9, we provide evidence showing that the two techniques deliver identical results.

[33] To get this formula for the elasticity $\varepsilon_{D,b_k} = \hat{\delta}_k \cdot \frac{\bar{w}_k}{D^{cap}}$, we simply use the fact that, at the kink, $b_k = 0.8.\bar{w}_k$, and the fact that the deterministic change in slope in the benefit schedule at the kink is $\nu_k = 0.8$

The same approach can be used to estimate the effect of benefits on the survival rate in unemployment at any spell duration. Table 2 shows the RK estimates for the effect of the benefit changes on the benefit durations D, 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.

We provide a comparison of these duration responses to estimates of labor supply responses to UI benefits available in the literature in online Appendix Section B.6. For these comparisons to be meaningful, we focus on elasticities that use similar variations in benefits. When benchmarked against conceptually similar elasticities, our duration responses prove quite similar to moral hazard estimates in the literature, although probably in the high end of the range of existing estimates.

The most common threat to identification and inference in the RK design is the presence of nonlinearity 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, we use spells starting after July 2002, for which the UI schedule is linear at the threshold, as a placebo and reject the presence of nonlinearity 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 nonlinear functional dependence between wages and unemployment duration around the 725SEK threshold.[34]

We provide several additional tests to assess the validity of the RK design and the robustness of the RK estimates in online Appendix B. We start by testing for manipulation of the assignment variable around the kink point $\bar{w}$, which would constitute a clear violation of the second identifying assumption of the RK design. Such manipulation could arise due to selection into job loss, but also due to selection into UI, as UI is voluntary in Sweden. Online Appendix Figure B-4 displays the probability density function of wages and reports tests in the spirit of McCrary (2008) that confirm continuity of the PDF and of its first derivative around the 725SEK threshold. Online Appendix Table B-1 further indicates that the removal of the kinks did not significantly affect the distribution of daily wages above and below the kink, suggesting that the presence of kinks in the UI schedule does not significantly affect the distribution of daily wages around the kink. This rules out clear manipulation of the assignment variable in response to the kinked schedule.

We also find smoothness in the relationship between observable characteristics of unemployed workers and wages at the 725SEK threshold plotted in online Appendix Figure B-5. This is reassuring, as non-smoothness in the distribution of observable heterogeneity would have cast doubt on the validity of the assumption of smoothness in the distribution of unobservable heterogeneity around the kink.

[34] In online Appendix Figure B-3, we also plot the evolution of the estimates of the change in slope year by year from 1999 to 2007. All estimates for the placebo years 2002 to 2007 are close to zero and insignificant. The figure provides clear evidence that our estimated responses are indeed due to the policy changes, and not due to time trends in the distribution of durations around the kink.

Online Appendix B finally provides additional tests for the presence of confounding nonlinearities in the relationship between the assignment variable and the outcomes. The sensitivity of the RK estimates to the size of the bandwidth is explored in Figure B-6. The stability of the RK estimates across bandwidth sizes further alleviates the concern that the RK estimates pick up some underlying nonlinearity in the relationship between wages and unemployment duration. In Figure B-10 we perform tests aimed at detecting nonparametrically 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. We explore the sensitivity of our inference to alternative strategies in Table B-2. In particular, we report 95 percent confidence interval based on permutation tests as in Ganong and Jäger (2014).[35] Interestingly, due to the linearity in the relationship between unemployment duration and wages across the whole support of the assignment variable, these confidence intervals are much tighter than those based on bootstrapped or robust standard errors. We finally explore in Table B-3 sensitivity of the RK estimates to variation in the order of the polynomial used to fit the data. The table displays estimates of the change in slope at the kink for linear, quadratic, and cubic specifications and assesses the model fit for these different specifications. For spells starting between 1999 and July 2001, the estimates are very similar across polynomial orders. For spells starting between July 2001 and July 2002, estimates from the quadratic model are, although imprecisely estimated, somewhat larger in magnitude than estimates using a linear or cubic specification. Yet, the linear specification is having similar root mean squared errors (RMSE) and minimizes the Aikake information criterion (AIC) compared to the quadratic and cubic specification, suggesting that linear estimates should be preferred.

B. *Implications for Moral Hazard Costs*

Our results carry important implications for the moral hazard costs of modifying the time profile of UI benefits. Table 2 displays our elasticity estimates and the implied moral hazard costs with bootstrapped standard errors using 50 replications.

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) equals $MH_{\bar{b}} = \frac{\bar{b}+\tau}{\bar{b}}\varepsilon_{D,b} = 1.64$ (0.14). This means that because of behavioral responses when increasing all benefits by 1 percent, the planner would need to levy from the employed 1.64 times more resources to balance its budget than the implied static cost absent behavioral responses. For these moral hazard computations, we use $\bar{b} = 0.72$, which corresponds to the observed average replacement rate in the sample, and we assume a tax $\tau = 0.05$, implying $\frac{\bar{b}+\tau}{\bar{b}} = 1.07$. The tax rate corresponds to the rate required to balance the government's UI budget for the average unemployment rate

[35] Ganong and Jäger (2014) propose to evaluate 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 counterfactual relationship between the assignment variable and the outcome (i.e., in the absence of the kink in the budget set) is nonlinear, then the curvature in this relationship will result in many of the placebo estimates being large and statistically significant.

during the period 1999–2007 (i.e., assuming no other expenditures $\bar{G}$ in equation (1)). Note that when $\bar{G} > 0$, the required tax level to balance the government's total budget will be higher. Moreover, when accounting for the fiscal externality through the rest of the tax system, the implied moral hazard cost would be even higher. In online Appendix Section A.2.4, we show that in the presence of an income tax τ^y, levied on both employed and unemployed, the additional fiscal externality equals $\tau^y \frac{w - \bar{b}}{\bar{b}} \varepsilon_{D,\bar{b}}$, which would increase the total moral hazard cost to $MH_{\bar{b}} = 1.76$ (resp. 1.70) for an effective income tax rate of $\tau^y = 0.20$ (resp. 0.10).[36]

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 expression in equation (6). The moral hazard cost of increasing benefits b_2 after 20 weeks of unemployment is therefore $MH_2 = \frac{\bar{b} + \tau}{\bar{b}} \frac{D}{D_2} \varepsilon_{D,b_2} = 1.44\ (0.28)$. Compared to $MH_{\bar{b}} = 1.64\ (0.14)$, this implies that the incentive cost 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 somewhat 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.[37] This then implies that $MH_1 = 1.82\ (0.40)$.

Our estimates thus indicate that the point estimate of the moral hazard cost of increasing benefits for the first 20 weeks is 26 percent larger than that of increasing benefits after 20 weeks of unemployment. Formal z-tests of equality of MH_1 and MH_2 provided in Table 2 do not allow to reject that MH_1 is equal to MH_2. We also conduct inference using permutation tests instead of z-tests. We draw placebo kinks and obtain a placebo distribution of estimates of MH_1 and MH_2. The procedure is described in online Appendix B.5.4, and we report in Table 2 the corresponding p-value (5.98 percent) from a test of equality of MH_1 and MH_2, which is much tighter than the p-value from our z-test.

To further assess the robustness of our findings regarding the relative magnitude of MH_1 and MH_2, we also exploit additional sources of variations in b_1 and b_2 stemming from variation in the location of the kink in b_1 and b_2 during the period 1999 to 2007. As explained in detail in online Appendix Section B.5, these sources of variations provide two broad strategies to identify the relative moral hazard costs of b_1 versus b_2, corresponding to four different potential estimates of the ratio MH_1/MH_2, summarized in online Appendix Table B-4. The first strategy consists in comparing estimates at the same "kink" over time. This approach has the advantage of comparing similar individuals over time at the same level of income. One drawback may be that behavioral elasticities are time varying due to business-cycle fluctuations, for instance. The second strategy consists of comparing estimates at different "kinks" within the same time period. This second approach has the advantage of comparing individuals within the same time period, therefore controlling for the fact that behavioral elasticities are time varying due to business-cycle fluctuations, for instance. A potential drawback of this approach is that individuals at

[36] In Sweden, UI benefits are fully included in individuals' income subject to the personal income tax.

[37] We simply use $\varepsilon_{D,b} = \varepsilon_{D,b_1} \frac{b}{b_1} + \varepsilon_{D,b_2} \frac{b}{b_2} = \varepsilon_{D,b_1} + \varepsilon_{D,b_2}$ for $b_1 = b_2$.

different kinks may differ in their responsiveness to the policy. Results from online Appendix Table B-4 show that estimates are very robust to the sources of variation used for identification, and suggest that the larger magnitude of MH_1 relative to MH_2 is a very robust finding, across all identification strategies.[38]

IV. Consumption Smoothing over the Unemployment Spell

In this section we provide empirical evidence on the evolution of consumption at job loss and throughout the unemployment spell. We then discuss how this evidence relates to the evolution of consumption smoothing gains of UI.

A. *Registry-Based Measure of Consumption: Estimating Consumption Drops*

Annual Consumption Drop.—We start by providing evidence that annual household consumption drops significantly at unemployment, following an approach similar to Gruber (1997). We implement an event study approach using the panel dimension of the data. We define event year $n = 1$ as the year an unemployment spell starts. Because our annual measure of consumption is observed as of December, we focus on individuals who are observed unemployed in December of year $n = 1$, (and who were employed in December of year $n = 0$). We estimate the following event study specification:

$$(17) \qquad \dot{\mathcal{C}}_{i,n} = \sum_j \gamma_j \cdot \mathbf{1}[n = j] + \sum_l \eta_l \cdot \mathbf{1}[\text{Calendar Year}_{i,n} = l] + \mu_{i,n},$$

where $\dot{\mathcal{C}}_{i,n} = \ln \mathcal{C}_n - \ln \mathcal{C}_{n-1}$, is the change in the log annual consumption of individual i between event years $n-1$ and n. Specification (17) only controls for a set of calendar year fixed effects η_l. In Figure 5, we report estimates of the event year dummies γ_j, where the omitted category is year $n = 0$. Our estimates provide compelling evidence of a sharp consumption drop at unemployment. First, the results show that annual consumption growth is remarkably similar in the five years prior to the start of a spell, which suggest that consumption profiles show limited signs of anticipation effects prior to unemployment. Second, annual consumption drops sharply and significantly in year $n = 1$, when individuals start an unemployment spell. Consumption growth is still negative in year $n = 2$, reflecting the fact that some individuals may still be unemployed through part, or all, of year $n = 2$. Consumption growth then picks up slightly as individuals find new jobs. Yet, as consumption *growth* remains similar to pre-unemployment levels after year $n = 3$, this implies that consumption *levels* remain significantly lower than their pre-unemployment levels, even five years after the start of a spell. This evidence of persistent consumption effects of unemployment is in line with the abundant literature documenting the long-lasting earnings consequences of job displacement (e.g., Couch and Placzek 2010).

[38] In online Appendix Figure B-13, we also provide inference using the permutation-based approach (taking random draws of placebo kinks) for all four estimates, which confirms the robustness of our conclusion.

Consumption Drop over the Spell.—To evaluate the consumption smoothing benefits of benefits given in different parts of the UI profile, we need to recover higher frequency consumption measures. Here, we wish to retrieve weekly consumption flows in order to estimate $\Delta C_1 = \bar{c}_1^u - \bar{c}_0$, the average consumption drop in the first part of the profile (first 20 weeks of an unemployment spell) and $\Delta C_2 = \bar{c}_2^u - \bar{c}_0$, the average consumption drop in the second part of the profile (after 20 weeks of unemployment). Our data enable us to estimate nonparametrically ΔC_1 and ΔC_2 with minimal and easily testable assumptions. We leverage two things. First, annual consumption is the sum of higher frequency expenditure flows during the year. Second, at the moment we measure annual consumption (December of each year), individuals differ in the time they have been unemployed, which provides variation to identify how consumption evolves over the unemployment spell. Based on these two simple observations, and defining event time t as the number of weeks since the start of an unemployment spell, the annual household consumption of an individual who is, in December, observed in the tth week of her unemployment spell is: $\mathcal{C}_t = \sum_{j=t-51}^{t} c_j$, where c_j is the weekly consumption flow in event week j. We also define $\Delta\mathcal{C}_t = \mathcal{C}_t - \mathcal{C}_0$, the drop in annual consumption of an individual observed in her tth week of a spell, compared to pre-unemployment annual consumption $\mathcal{C}_0$.[39] Under the assumption that consumption profiles are constant prior to unemployment, i.e., $c_j = \bar{c}_0, \forall j < 0$, we have $\Delta\mathcal{C}_t = \sum_{j=t-51, j>0}^{t} (c_j^u - \bar{c}_0)$. The expression above shows that the comparison of drops in annual consumption of individuals observed at different weeks t of their unemployment spell directly reveals information about the drops in flow consumption $c_j^u - \bar{c}_0$ throughout the unemployment spell. Compare for instance the drop in annual consumption of individuals observed in their first week of unemployment $\Delta\mathcal{C}_1 = c_1^u - \bar{c}_0$ and the drop in annual consumption of individuals observed in their second week of unemployment $\Delta\mathcal{C}_2 = c_2^u - \bar{c}_0 + c_1^u - \bar{c}_0$. The difference $\Delta\mathcal{C}_2 - \Delta\mathcal{C}_1 = c_2^u - \bar{c}_0$ reveals the drop in flow consumption in the second week of unemployment.[40]

It follows that we can easily estimate nonparametrically our two statistics of interest ΔC_1 and ΔC_2 from the observed drops in annual consumption of individuals observed at different points in their unemployment spell. In practice, we have

$$\Delta C_1 = \sum_{j=1}^{20} \frac{S_j}{D_1} (\Delta\mathcal{C}_j - \Delta\mathcal{C}_{j-1}), \tag{18}$$

$$\Delta C_2 = \sum_{j=21}^{104} \frac{S_j}{D_2} \big((\Delta\mathcal{C}_j - \Delta\mathcal{C}_{j-1}) + \mathbf{1}[j > 52] \cdot (\Delta\mathcal{C}_{j-52} - \Delta\mathcal{C}_{j-53}) \big), \tag{19}$$

[39] Pre-unemployment annual consumption $\mathcal{C}_0$ is measured as the annual consumption in event year $n = 0$, i.e., the last pre-unemployment year during which the individual is observed being employed. Formally, $\mathcal{C}_0 = \mathcal{C}_{t-51}$ for individuals who are in their first year of unemployment as of December ($t < 52$), and $\mathcal{C}_0 = \mathcal{C}_{t-103}$ for individuals who are in their second year of unemployment as of December ($t \geq 52$).

[40] More generally, for individuals in their first year of unemployment, ($t < 52$), the difference in annual consumption drops between t and $t-1$, $\Delta\mathcal{C}_t - \Delta\mathcal{C}_{t-1} = c_t^u - \bar{c}_0$ directly reveals the flow drop in consumption in week t. For individuals in their second year of unemployment, ($t \geq 52$), the difference in annual consumption drops between t and $t-1$, $\Delta\mathcal{C}_t - \Delta\mathcal{C}_{t-1} = \sum_{j=t-51}^{t} (c_j^u - \bar{c}_0) - \sum_{j=t-52}^{t-1} (c_j^u - \bar{c}_0) = (c_t^u - \bar{c}_0) - (c_{t-52}^u - \bar{c}_0)$ identifies the sum of the flow drop in consumption in week t and the flow drop in consumption in week $t-52$. The drop in flow consumption $c_t^u - \bar{c}_0$ for $t > 52$ is therefore identified by $(\Delta\mathcal{C}_t - \Delta\mathcal{C}_{t-1}) + (\Delta\mathcal{C}_{t-52} - \Delta\mathcal{C}_{t-53})$.

where S_j is the survival rate after j weeks of unemployment, and D_1 and D_2 are the average time spent in the first and the second part of the benefit schedule.[41]

To implement this strategy, we start by estimating, on the sample of individuals who are unemployed in December of any year between 2000 and 2007, the average annual consumption drops as a function of the time t that individuals have spent unemployed as of December, using the following specification:

$$(20) \qquad \Delta\mathcal{C}_{i,t} = \sum_j \beta_j \cdot \mathbf{1}[t = j] + \sum_l \eta_l \cdot \mathbf{1}[\text{Calendar Year}_{i,t} = l] + \varepsilon_{i,t}.$$

We display in panel A of Figure 6 our baseline estimates of the drop in annual consumption where, in practice, we group individuals in 10-week intervals of unemployment duration for estimation.[42] Panel A reports $\hat{\beta}_t/\mathcal{C}_0 = \widehat{\Delta\mathcal{C}_t}/\mathcal{C}_0$, i.e., we scale our estimates by the average annual consumption in the last year prior to unemployment, to express consumption drops relative to pre-unemployment levels. The figure shows that the drop in annual consumption increases steadily with time spent unemployed as of December for individuals observed throughout the first year of unemployment, and then stabilizes around 10 percent for individuals observed in their second year of unemployment. The decrease in the first year of unemployment is partly mechanical, as individuals who are observed at a later time t in the spell have spent more time of the year unemployed and thus experienced a larger drop in annual consumption. Still, the shape of the slope of these estimates reveals all the necessary information about the evolution of flow consumption throughout the spell.[43]

Panel A of Figure 6 also reports our nonparametric estimates of the average drops in consumption for both parts of the benefit profile $\widehat{\Delta C_1} = \sum_{j=1}^{20} \frac{\hat{S}_j}{D_1}(\hat{\beta}_j - \hat{\beta}_{j-1})$ and $\widehat{\Delta C_2} = \sum_{j=21}^{104} \frac{\hat{S}_j}{D_2}\Big((\hat{\beta}_j - \hat{\beta}_{j-1}) + \mathbf{1}[j > 52] \cdot (\hat{\beta}_{j-51} - \hat{\beta}_{j-52})\Big)$. Standard errors are computed using the Delta method. Our results suggest that the average drop in (flow) household consumption for individuals in the first part of the profile is 4.4 percent, while it is 9.1 percent on average for individuals in the second part of the profile. We also report the p-value from a test of equality of $\widehat{\Delta C_1}$ and $\widehat{\Delta C_2}$ which strongly rejects that consumption drops are the same in the first and second part of the benefit profile. Our evidence thus indicates that consumption expenditures do significantly decline over the unemployment spell.[44]

Robustness.—Our implementation so far has rested on two assumptions, which we will investigate next. First, for simplicity we assumed that $c_j = \bar{c}_0, \forall j < 0$, or in other

[41] Note that we censor durations at two years to deal with a few observations without any end date. Hence, our summation in ΔC_2 only goes up to week 104.

[42] The left-hand-side variable in specification (20) is a change in levels (and not a log change in consumption). The reason is that we need the change in levels (and not the log changes) to retrieve higher frequency consumption from the annual changes.

[43] Intuitively, a linear slope in the first year would imply that $c_j^u - c_0 = c_k^u - c_0, \forall k, j < 52$, which means that flow consumption drops right at job loss and remains constant after. Concavity in the estimates for the first year would indicate that $c_j^u - c_0 < c_k^u - c_0, \forall k < j < 52$, which means that flow consumption continues to decline throughout the unemployment spell.

[44] Our approach allows to estimate nonparametrically how consumption evolves at, potentially, *any* frequency over the unemployment spell. We report, for instance in online Appendix Figure C-6, the estimated flow consumption drops over the first year of unemployment at a 10-week frequency.

words, that the consumption profile is flat prior to the start of an unemployment spell. To investigate how consumption profiles evolve prior to the unemployment spell, we report in online Appendix Figure C-1 the annual log household consumption changes $\mathcal{C}_{i,t} - \mathcal{C}_{i,t-52}$ as a function of time t since the onset of a spell (in five-weeks bins), relative to the last five weeks prior to the onset of a spell. We go as far back as four years prior to the onset of a spell, to detect longer anticipation effects. The figure shows two interesting patterns. First, the consumption patterns of households are extremely stable in the four years prior to displacement, suggesting no long-term anticipation behaviors. Second, there are no sharp consumption changes immediately preceding displacement, suggesting no significant short-term anticipation behaviors. Overall, this evidence strongly supports the assumption that flow consumption profiles are flat prior to the onset of an unemployment spell.

The second assumption used in our baseline implementation is that average flow consumption profiles are similar for individuals observed at different time t in their spell. In practice, the sample of individuals who survive until t (i.e., individuals with a completed spell length $L \geq t$) may differ from the sample of individuals who survive until $t-1$ (i.e., individuals with a completed spell length $L \geq t-1$), and this may bias our estimates if individuals who select into different spell lengths L differ in their underlying consumption profiles (dynamic selection on consumption profiles). Note, however, that we can easily relax the assumption that average consumption profiles are independent of actual spell length. To ensure that we compare the same sample of individuals at different time t, we estimate separate flow profiles by completed spell length L using the following specification:

$$\Delta \mathcal{C}_{i,t} = \sum_k \sum_j \beta_{j,k} \cdot \mathbf{1}[t = j] \cdot \mathbf{1}[L = k] + \sum_l \eta_l \cdot \mathbf{1}[\text{Calendar Year}_{i,t} = l] + \varepsilon_{i,t},$$

where we break down spell length into 4 brackets: 1 to 20 weeks, 21 to 40 weeks, 41 to 60 weeks, and longer than 61 weeks. In panel B of Figure 6 we report our estimates of $\hat{\beta}_{t,L}$, representing the drops in annual consumption as a function of time spent unemployed for the 4 categories of completed spell lengths. Results show that the profile of drops in annual consumption is very similar across spell length, pointing to a limited role for dynamic selection on consumption profiles. We also report the implied estimates of ΔC_1 and ΔC_2, allowing for heterogeneity in profiles across spell length, $\widehat{\Delta C_1} = \sum_k \sum_{j=1}^{20} \frac{\hat{S}_{j,k}}{D_1}(\hat{\beta}_{j,k} - \hat{\beta}_{j-1,k})$ and $\widehat{\Delta C_2} = \sum_k \sum_{j=21}^{104} \frac{\hat{S}_{j,k}}{D_2}\big((\hat{\beta}_{j,k} - \hat{\beta}_{j-1,k}) + \mathbf{1}[j > 52] \cdot (\hat{\beta}_{j-52,k} - \hat{\beta}_{j-53,k})\big)$. Results are very similar to our baseline estimates, and confirm that flow consumption significantly decreases over the unemployment spell. The average drop in (flow) household consumption for individuals in the first part of the profile is 4.8 percent, while it is 9.6 percent on average for individuals in the second part of the profile.

Additional Survey Evidence.—We further investigate the robustness of our results using the household consumption surveys (HUT) in online Appendix C.4. While the HUT sample is much smaller, and does not have a panel structure, it directly measures flow consumption at the time of the survey interview. Results from our preferred specification of column 4 of online Appendix Table C-3 suggest that the drop in the first part of the profile $\widehat{\Delta C_1}$ is 4.6 percent and the drop in the second part of the profile $\widehat{\Delta C_2}$ is 10.8 percent. Reassuringly, despite much larger standard errors, these point estimates are remarkably similar to our baseline estimates using the registry-based measure of consumption.

B. *Implications for Consumption Smoothing Gains*

The empirical evidence presented above indicates that consumption declines significantly over the unemployment spell. This decline in consumption suggests that the marginal value of unemployment benefits increases over the spell, evaluated for the flat policy in place in Sweden.

An important outstanding issue is that our empirical analysis considers changes in expenditures. The question is to what extent expenditures translate into consumption and thus capture the potential welfare value of unemployment benefits.[45] In particular, unemployed workers may try to reallocate certain categories of expenditures to smooth the shocks in their consumption. A first example is the substitution toward home production, which has been analyzed extensively in the context of retirement (e.g., Aguiar and Hurst 2005). A second example is the substitution away from expenditures on durable goods that provide a consumption flow for future periods as well.

The HUT offers insights into the type of consumption goods that households adjust over the spell. In online Appendix Table C-4, we investigate how various categories of expenditures evolve over the unemployment spell. Consumption of nondurable, uncommitted goods, such as food, recreation, transportation, or restaurants (columns 2, 6, 7, and 8) drops significantly early on in the spell and further decreases over the spell.[46] More committed expenditures like housing rents paid by renters (column 3) do not seem to decline significantly, neither early nor later in the unemployment spell.[47] Interestingly, we find a larger drop in restaurant expenditures than in food expenditures, consistent with substitution toward home production. Substitution toward home production may affect the level of the consumption smoothing gains of UI benefits.[48] However, its effect on the evolution of the consumption smoothing gains of UI over the spell is ambiguous, and will depend on the relative availability of substitution toward home production for households that select into long versus short spells.[49]

[45] See Campos and Reggio (2016).

[46] Gruber (1997) considers only food consumption using US data from the PSID for the period 1968–1987 and finds an average drop of 6.8 percent in the first year of unemployment, which is very similar to our estimates.

[47] Fixed commitments reduce the ability to smooth consumption and can increase the relevant value of γ to translate consumption drops into welfare (Chetty and Szeidl 2007).

[48] The sign of the impact of such substitution patterns on the consumption smoothing gains of UI benefits is nevertheless ambiguous. The availability of consumption insurance through home production means that the drop in consumption will be smaller than the actual drop in expenditures, which decreases the marginal utility of unemployment benefits. At the same time, complementarity between expenditures and household production increases the marginal utility of benefits. The relative magnitude of these two effects will determine the sign of the impact on the consumption smoothing gains of UI.

[49] Note that results displayed in online Appendix Tables C-5 and C-6 indicate that the expenditure profiles over food, recreation, transportation, or restaurants are not significantly different for households that select into long versus short spells. This suggests that there is no significant dynamic selection over the spell based on the availability of substitution toward home production.

Online Appendix Table C-4 also indicates that expenditures on durable goods such as the purchase of new vehicles or the purchase of furniture and home appliances (columns 4 and 5) decline strongly early during the spell but increase later during the spell, yet remain largely below their pre-unemployment level. These results suggest that the unemployed can initially smooth the marginal utility of consumption services by shifting spending away from durables, but they lose the capacity to do so after some time.[50] This in turn will tend to further increase the inclining profile of the consumption smoothing value of UI benefits over the unemployment spell.

To sum up, empirical evidence seems to substantiate that, in the Swedish context, and for a flat benefit profile, consumption declines over the spell, and that the consumption smoothing gains of UI benefits are larger after 20 weeks of unemployment than during the first 20 weeks of unemployment. This conclusion seems robust to the presence of dynamic substitution across categories of expenditures.

V. Welfare Analysis

This concluding section brings our theoretical and empirical analysis together to provide an evidence-based assessment of the UI benefit profile in Sweden. We use our empirical estimates to implement our sufficient-statistics formulae and to shed light on forces underlying how the sufficient statistics evolve over the unemployment spell. While our formulae allow for a transparent evaluation of local changes in the policy profile, we also illustrate how a calibrated structural model allows us to go beyond the local recommendations.

A. *Sufficient-Statistics Implementation*

The welfare consequences of a marginal increase in UI benefits are reported in Table 3. The first row examines the consequences of a marginal increase in the benefit level $\bar{b}$ throughout the unemployment spell; the second row examines the consequences of a marginal increase in the benefit b_1 during the first 20 weeks of unemployment and the third row examines the consequences of a marginal increase in the benefit b_2 after 20 weeks of unemployment. Different columns show the different components of the welfare impact.

The first column repeats our estimates of the moral hazard costs (panel IV of Table 2). As noted before, the estimated moral hazard cost is high overall (all estimates exceed 1), but the cost is 26 percent higher for benefits paid during the first 20 weeks of unemployment compared to those paid after 20 weeks of unemployment. The second column repeats our estimates of the drops in average consumption during the first 20 weeks of unemployment and after 20 weeks of unemployment, but also averaged over the full unemployment spell. We convert the respective consumption drops into estimates of the consumption smoothing gains CS_k, following the implementation in equation (14), which relies on a Taylor approximation and

[50] Note again that, although results are relatively imprecise due to the small sample size, online Appendix Table C-5 indicates that dynamic selection based on durable expenditure profiles is not significant. The availability of consumption smoothing through shifting expenditures away from durables is not significantly different for households that select into long versus short spells.

homogeneous preferences. As the appropriate value of risk aversion in this context is uncertain, the consumption smoothing gains are reported for a range of plausible values (see Chetty 2009, Chetty and Finkelstein 2013) and range between 0.05 and 0.5. Regardless of the level of risk aversion, when risk preferences are homogeneous, the estimated consumption smoothing gains are twice as high for benefits paid later in the spell.

Putting the estimates of the CS gains and the MH costs together indicates that the MH costs are substantially larger than the CS gains, even for high risk aversion.[51] This is true both for benefits paid to short-term and long-term unemployed and thus suggests that UI benefits are too generous overall. Indeed, expressing the welfare effects in terms of "benefit-cost" ratios CS_k/MH_k, we find that the return to the marginal krona spent on the unemployed is substantially lower than 1. These estimates are, however, sensitive to our implementation assumptions including the risk preference parameter γ, the tax distortion τ, and our normalization relative to the value of krona spent pre-unemployment.[52]

Comparing the benefit-costs ratios for different parts of the policy allows us to formulate recommendations on the tilt of the benefits profile b_1/b_2.[53] We find that the relative moral hazard cost MH_1/MH_2 $(= 1.26)$ is more than twice as high as the relative consumption smoothing gain CS_1/CS_2 $(= 0.50)$. The implied return to the marginal krona spent on the short-term unemployed is thus at least twice as low compared to the long-term unemployed, indicating that welfare could be increased by introducing an inclining tilt $(b_1 < b_2)$ in the flat Swedish benefit profile. Importantly, this evaluation of the tilt is less sensitive to the assumptions above. In particular, CS_1/CS_2 does not depend on the preference parameter γ, while MH_1/MH_2 does not depend on the tax distortion τ. More generally, our recommendations on the tilt of the benefit profile are robust to implementation errors that are uncorrelated with the unemployment duration.

In sum, our implementation suggests that on average UI benefits should be decreased and especially so for the short-term unemployed. The 2001 reform in Sweden did the exact opposite by increasing the benefit cap for the first 20 weeks of unemployment.

B. *From Sufficient Statistics to Mechanisms*

While the previous local welfare implementation does not need to identify the forces that shape the sufficient statistics, it is in practice interesting to explore and analyze these forces, for at least two reasons. First, from a descriptive point of view,

[51] We note again that, although our duration elasticities and consumption drops are estimated from different samples, individuals in the two samples are almost identical in terms of observable characteristics and unemployment durations as discussed in Section IIB.

[52] Note that consumption drops are experienced at the household level. Our analysis only accounts for the consumption smoothing gains for the unemployed individual, hereby underestimating the overall consumption smoothing gain of UI.

[53] Expressing the welfare effects in terms of "cost-benefit" ratios (see Hendren 2013) is a common approach to formulate policy recommendations. Whenever $CS_1/MH_1 > CS_2/MH_2$, welfare can be increased by increasing the tilt b_1/b_2. Moreover, a budget-balanced increase in the tilt b_1/b_2 increases welfare if and only if $\frac{1+CS_1}{1+MH_1} > \frac{1+CS_2}{1+MH_2}$, where the ratios correspond the so-called "marginal value of public funds." We show this in Corollary 1 in online Appendix Section A.2.

it is interesting to understand what mechanisms can explain the observed patterns of these statistics. Second, from a welfare perspective, in order to extend our recommendations to nonlocal variations, one needs to impose more structure on the data. Investigating what mechanisms drive the estimated sufficient statistics is a useful way to inform the choice of structural assumptions to impose on the underlying model used to simulate the welfare consequences of nonlocal variations, as we do in Section VC.

In order to shed light on mechanisms, we can leverage the fact that different assumptions regarding the underlying structure of the model map into different dynamic behaviors of our sufficient statistics.

Dynamic Profile of Sufficient Statistics.—Our sufficient statistics characterization indicates that, everything else equal, unemployment benefits should decline over the spell when the consumption smoothing gains are lower for benefits paid later in the spell or when their moral hazard cost is higher. The optimal benefit profile has been the topic of a series of seminal papers, mostly studying single-agent, stationary environments (see Shavell and Weiss 1979; Hopenhayn and Nicolini 1997; and Shimer and Werning 2008). The relevant forces in a stationary setting and how they affect the optimal benefit profile are well understood from this literature. We can re-express the impact of these forces on the unemployment policy through their impact on the gradient of the moral hazard costs and consumption smoothing gains.[54]

PROPOSITION 2: *Consider a flat benefit profile* $(b_t = \bar{b} < w - \tau \ \forall t)$ *in a single-type, stationary environment* $(h_{i,t}(\cdot) = \bar{h}(\cdot) \ \forall i, t)$ *with* $\beta(1+r) = 1$, $T = \infty$ *and assuming differentiability:*

(*i*) *when agents are borrowing-constrained throughout the unemployment spell* $(a_{i,t} = 0 \ \forall i, t)$,

$$MH_t \leq MH_{t'} \quad \text{and} \quad CS_t = CS_{t'} \quad \text{for } t < t';$$

(*ii*) *when agents are not borrowing-constrained, but preferences are separable* $(\partial v_i^u(c,s)/\partial c = \partial v_i^e(c)/\partial c = v'(c))$,

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

The force underlying the increasing moral hazard cost is the forward-looking behavior of job seekers. Increasing unemployment benefits later in the spell discourages forward-looking job seekers already early in the spell and is therefore always more costly than increasing benefits earlier in the spell. This causes the optimal benefit profile to be declining, as shown before in Shavell and Weiss (1979) and Hopenhayn and Nicolini (1997). The force underlying the increasing consumption

[54] Online Appendix Section A.3 provides the proof (closely following Shimer and Werning 2008) and gauges the robustness of Proposition 2 within a stationary environment.

smoothing gain is that, when available, unemployed job seekers use liquid assets to smooth their marginal utility of consumption. As assets are depleted while unemployed, the marginal utility of consumption and thus the value of unemployment benefits is always higher later in the spell. This causes the optimal benefit profile to be inclining, as shown before in Shavell and Weiss (1979) and Shimer and Werning (2008).[55]

In practice, however, unemployment dynamics are nonstationary. Nonstationary forces, and dynamic selection in particular, can mitigate the forces underlying Proposition 2 and, when strong enough, actually reverse the gradient of the sufficient statistics. We discuss this below in relation to our empirical estimates, but also develop this argument formally in online Appendix Section A.4.[56] While characterizing the impact of nonstationarities on the optimal policy is challenging and previous work has mostly relied on calibrated models (e.g., Shimer and Werning 2006 and Pavoni 2009), our framework provides an alternative approach, which is to analyze how nonstationary forces affect the gradient of the moral hazard costs and consumption smoothing gains respectively.

In the end, how both the consumption smoothing gains and moral hazard costs evolve over the unemployment spell remains an empirical question. Yet, because different assumptions regarding the underlying structure of dynamic models of unemployment map into different dynamic behaviors of our sufficient statistics, we can use our estimated sufficient statistics to shed interesting light on the relative importance of different mechanisms in shaping consumption smoothing and moral hazard over the spell.

Moral Hazard Cost.—Our empirical analysis indicates that the moral hazard costs increase over the unemployment spell in Sweden, suggesting the importance of nonstationary forces dominating the forward-looking incentives. To illustrate how these forces shape the dynamics of moral hazard costs over the spell, we can use the earlier expression (11) to rewrite the moral hazard cost of an increase in benefit level b_2 after 20 weeks of unemployment, as the sum of two parts,

$$MH_2 = \frac{\bar{b}+\tau}{\bar{b}} \times \varepsilon_{\tilde{D}_B,b_2} + \frac{\bar{b}+\tau}{\bar{b}} \times \left[\frac{D_1}{D_2}\varepsilon_{D_1,b_2} + \varepsilon_{S_B,b_2}\right].$$

The first part captures the response in the *remaining duration* of unemployment, conditional on still being unemployed at time $B = 20$ weeks $(\tilde{D}_B \equiv \sum_{s=0}^{T-B} S_{B+s}/S_B)$. In a stationary, single-agent environment, this first part is the same for any time B and equal to the moral hazard cost of an overall increase in the benefit level $MH_{\bar{b}}$.

[55] Due to the opposing forces coming from moral hazard and consumption smoothing, even in a stationary environment, deriving the optimal timing of benefits is difficult. Werning (2002) and Shimer and Werning (2008) analyze these two opposing forces in models with both search and savings being endogenous and show that they exactly cancel out in case of CARA preferences; a flat benefit profile is optimal conditional on the unemployed having access to liquidity.

[56] Online Appendix Section A.4 starts from a specific stationary search environment with borrowing-constrained agents (Section A.4.2). We then consider depreciation in the returns to search over the unemployment spell (Section A.4.3) and unobservable heterogeneity in the returns to search (Section A.4.4). We finally allow for unobservable savings and consider heterogeneity in assets or preferences. We show how the extensions can mitigate and/or reverse the difference in moral hazard costs or consumption smoothing gains at different times during unemployment spell.

The second part captures the forward-looking incentives by which an increase in benefits in the second part of the policy reduces the exit out of unemployment in the first part of the spell (D_1) and increases the survival rate into the second part of the spell (S_B). This second part converges to 0 as B goes to 1. Following the argument in Shavell and Weiss (1979), the forward-looking incentives drive the result that in a stationary environment the moral hazard cost of increasing unemployment benefits is higher when timed later in the spell (see Proposition 2). We find the exact opposite in the Swedish context, which necessitates the presence of significant nonstationary forces that reduce the responsiveness of job seekers to changes in the policy later in the spell.

Evidence of Forward-Looking Behavior.—We first note that our results unambiguously show that unemployed individuals are indeed forward-looking. We already reported the estimated elasticity of D_1, the time spent on the first part of the profile, with respect to benefits b_2 received in the second part of the profile is positive and significant ($\varepsilon_{D_1,b_2} = 0.60$ (0.11); see Table 2). We can also study the effect of UI benefits on the hazard rates out of unemployment. Online Appendix Figure B-1 reports the corresponding RKD estimates and clearly shows that benefits b_2 received after 20 weeks do have a negative effect on the hazard rate in the first 20 weeks. Unemployed individuals are thus not fully myopic: they react early in the spell to variation in benefits paid later in the spell.

Evidence on Nonstationary Forces.—To highlight the nonstationary forces underlying the reversed pattern of moral hazard costs in the Swedish context, we compare at different spell lengths t, the elasticity of the remaining duration $\tilde{D}_t \equiv \sum_{s=0}^{T-t} S_{t+s}/S_t$ with respect to the overall benefit level $\bar{b}$. In a stationary, single-agent environment, these elasticities would remain constant and be equal to $\varepsilon_{D,\bar{b}}$.[57] Instead, our estimates, reported in Figure 4, show that the elasticity of the remaining duration strongly declines as a function of t. Five months into the spell, the elasticity is only one-third of the elasticity at the onset of the spell. This evidence of strong nonstationarities is corroborated in online Appendix Figure B-1, which shows that the effect of UI benefits $\bar{b}$ on hazard rates is mostly concentrated in the first two to three months. After three months, the effect of UI benefits on the hazard rate is small and almost always insignificant.

Both depreciation and heterogeneity in the returns to search can explain why the responsiveness of exit rates decreases substantially over the unemployment spell. The reduced responsiveness to changes in the benefit profile, either through depreciation or through selection, makes it less costly to increase benefits later in the spell. While depreciation by itself cannot, selection, when strong enough, can in fact reverse the decline of moral hazard costs over the spell, as shown in online Appendix Section A.4. Separating true depreciation from selection on heterogeneity in returns to search is notoriously difficult and beyond the scope of this paper, but

[57] Because exit rates only depend on the continuation policy in a stationary environment and there are no selection effects in a single-agent environment in response to policy variation earlier in the spell, the elasticity of the remaining duration $\tilde{D}_t$ with respect to an overall change in the benefit level is independent of t.

our analysis nevertheless suggests the importance of embedding such nonstationary features in dynamic models of unemployment.

Consumption Responses.—We now turn to the dynamic profile of the consumption smoothing gains. We find that the average consumption drop is more pronounced later in the spell. We therefore revisit mechanisms that can explain this pattern and whether they necessarily translate in consumption smoothing gains that increase over the spell.

Assets and Liquidity Constraints: As stated above, a single-agent model of intertemporal consumption predicts that the marginal utility of consumption is weakly increasing during the unemployment spell. Intuitively, long-term unemployment implies a larger shock in resources, so we expect the unemployed to run down their assets and consume less the longer they are unemployed. Or they become liquidity constrained and start consuming "hand-to-mouth."

We find clear evidence of the role of assets and liquidity constraints in determining the consumption smoothing over the spell. First, following the same methodology as in Figure 6, we report the average annual consumption drops by time spent unemployed as of December, but now breaking down the sample by the level of net wealth of the household at the start of the spell. As shown in online Appendix Figure C-2, individuals with higher net wealth experience a lower drop in consumption. Second, we provide direct evidence that individuals do use their liquid assets to smooth consumption over the spell. Figure C-3 displays the relative change in liquid bank asset holdings by time spent unemployed. The figure shows that individuals use their liquid assets to smooth consumption, especially earlier on in the spell, by depleting their bank accounts or reducing their savings. Finally, evidence from registry data also indicates that debt does not offer much help in smoothing consumption over the unemployment spell, hinting at the presence of liquidity constraints. In Figure C-4, we provide evidence of a reduction in the use of non-mortgage-related credit over the unemployment spell. This suggests that as the duration of the spell increases, access to credit becomes harder and consumption out of debt falls significantly.

Overall, this analysis confirms that assets and debt may offer means to smooth consumption over short spells. But as the unemployment spell continues, these means get quickly exhausted. Households in the second part of the profile seem therefore closer to being hand-to-mouth.

Spousal Labor Supply: Within a household, the labor supply of other members of the household may help reduce the drop in household consumption over the spell. In online Appendix Figure C-5 we investigate how the labor supply of other members of the household affects the drop in household consumption over the spell. Following the same methodology as in Figure 6, we report the average change in total earnings and in total disposable income of all other members of the household as a function of time spent unemployed, scaled by the annual household consumption level prior to unemployment. Results show that within-household changes in the earnings and disposable income of all other members of the household are small and almost never significantly different from zero throughout the unemployment spell. This suggests that, in our context, labor supply responses of other household members are not playing a significant role in increasing household consumption in response to an unemployment shock, even for long-term unemployed.

Dynamic Selection: In practice, agents are heterogeneous and selection into longer unemployment spells may affect consumption responses over the spell and the gradient of consumption smoothing gains. As shown in online Appendix Section A.4.5, if those selecting into longer unemployment spells suffer less from a drop in consumption, the positive gradient of consumption smoothing benefits due to the falling consumption is reduced and can potentially be reversed.

In theory, as long as preferences are homogeneous, we can use the approximation in condition (14) to translate the average consumption wedge into the consumption smoothing gains of increasing benefits. When preferences are heterogeneous, the consumption smoothing gains are approximated (using Taylor expansions of $v'(c^u_{i,t})$ around $c_{i,0}$ for each individual) by

$$CS_t \cong \frac{E^u_t\left[v'_i(c_{i,0})\right] - E_0\left[v'_i(c_{i,0})\right]}{E_0\left[v'_i(c_{i,0})\right]} - \frac{E_t\left[v''_i(c_{i,0})\,(c_{i,0} - c^u_{i,t})\right]}{E_0\left[v'_i(c_{i,0})\right]}, \tag{21}$$

where E^u_t takes the weighted average among the unemployed at time t (with weights $S_{i,t}/S_t$) and E_0 takes the average before the onset of the unemployment spell. The first part of this expression indicates that the consumption smoothing gains of long-term benefits will be reduced if individuals with lower marginal utility of consumption remain unemployed for longer. The second part of this expression indicates that the consumption smoothing gains are reduced if households with lower risk aversion and or lower consumer drops remain unemployed for longer.[58]

Our analysis of the consumption profiles by completed spell length already suggested that the dynamic selection on consumption profiles was limited (see Figure 6; Section IVA). To further assess the potential magnitude of dynamic selection, we investigate in online Appendix Table C-1 how various observable characteristics that have been shown to correlate with consumption and risk preferences are distributed across short-term and long-term unemployed. The patterns of selection are generally small in magnitude and often ambiguous in sign. Income levels and net wealth levels have quantitatively small and nonmonotonic effects on the probability to select into longer spells. Compared to households with no or negative net wealth, households with some small wealth ($<$500,000SEK $\approx$ US$55,000) have a slightly lower probability to be long-term unemployed. But individuals with high net wealth have a slightly higher probability to select into long spells. This result suggests that in our context, individuals with better means to smooth consumption do not unambiguously select into longer spells, which corroborates the evidence of no clear patterns of selection on consumption profiles. Also, portfolio characteristics (i.e., the fraction of portfolio wealth invested in stocks, and leverage defined as total debt divided by gross assets), which have been shown to be correlated with risk

[58] Note that if preferences are heterogeneous within the group of unemployed at a given time during the unemployment spell, any negative correlation between the drop in consumption and risk aversion would further reduce the consumption smoothing gains (see Andrews and Miller 2013).

preferences, have small and nonmonotonic impacts on the probability to experience a long unemployment spell. In contrast, the probability of experiencing long unemployment spells ($\mathbf{1}[L > 20wks]$) is significantly and monotonically correlated with age. However, existing evidence from the literature suggests a U-shape relationship between age and risk aversion (Cohen and Einav 2007), so the dynamic selection on age has an ambiguous effect on the evolution of risk preferences over the spell.

Overall, the observed dynamic selection patterns on consumption and risk preferences do not suggest that relative consumption smoothing gains would be significantly different than our estimates based on the average drops in consumption.

C. *Beyond Local Recommendations*

Our local recommendation regarding the policy and the tilt in particular may not be informative about how the optimal policy should look.[59] To go beyond our local recommendations, we need to know how the sufficient statistics vary with the policy parameters. We briefly illustrate how this can be explored using the structure of a model. That is, we calibrate a rich, nonstationary structural model to match the sufficient statistics for local policy evaluation and simulate how these sufficient statistics change when moving away from the local policy.

We describe the details of the calibration and simulation in online Appendix D. We consider a model with constant relative risk aversion (CRRA) preferences and additive search costs like in Hopenhayn and Nicolini (1997), Lentz and Tranæs (2005), and Chetty (2008), but motivated by the evidence discussed in the previous subsection, we allow for an exit rate function with heterogeneous returns to search that depreciate exponentially and for a heterogeneous asset distribution at the start of the unemployment spell. As our calibration targets our sufficient statistics estimates, the structural model delivers the same local recommendations around the existing policy. That is, the flat benefit $\bar{b} = 0.72$ is too generous overall and especially so for the short-term unemployed.

Figure 7 illustrates how the moral hazard costs and consumption smoothing gains change when reducing the overall generosity of the policy in our calibrated model. Two findings emerge. First, the consumption smoothing gains increase while the moral hazard costs decrease as we reduce the benefit level.[60] In particular, $CS_{\bar{b}}$ and $MH_{\bar{b}}$ are equalized for a flat replacement of $\bar{b} = 0.58$. Second, as we decrease the overall generosity, the consumption smoothing gains remain higher for benefits paid to the long-term unemployed ($CS_2 > CS_1$), while the moral hazard costs remain lower ($MH_2 < MH_1$). The introduction of an inclining tilt ($b_2 > b_1$) thus remains welfare-improving for lower replacement rates. In particular, our model predicts that welfare would be maximized by setting $b_1 = 0.48$ for the short-term unemployed and $b_2 = 0.68$ for the long-term unemployed (while keeping the tax rate unchanged).

[59] Of course, this issue may also arise for a one-dimensional change in benefit level b_k. Only if the policy problem is strictly concave in all benefit levels, can we be certain that a welfare-increasing change in the benefit level b_k actually moves the policy closer to the optimal policy.

[60] Note that we keep the asset distribution at the start of unemployment fixed, while this is likely to respond to the generosity of the unemployment policy.

In online Appendix D, we further illustrate the complementary value of the structural approach and sufficient-statistics approach. In particular, we explore how our findings depend on the assumptions of the nonstationary features of the search environment and the assumptions made for our consumption-based implementation of the consumption smoothing gains.[61]

VI. 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 behavior of job seekers, important nonstationary forces can make the moral hazard costs of benefits offered to the long-term unemployed higher than the costs 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 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. Most important, 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.

[61] In line with Chetty (2006), we find that for CRRA preferences the Taylor approximation of the marginal utilities substantially underestimates the consumption smoothing gains (by more than 20 percent for $\gamma = 2$). The approximation error on the relative consumption smoothing gains, however, is much smaller (equal to 5 percent for $\gamma = 2$), leaving the recommendation on the tilt basically unaffected.

Appendix

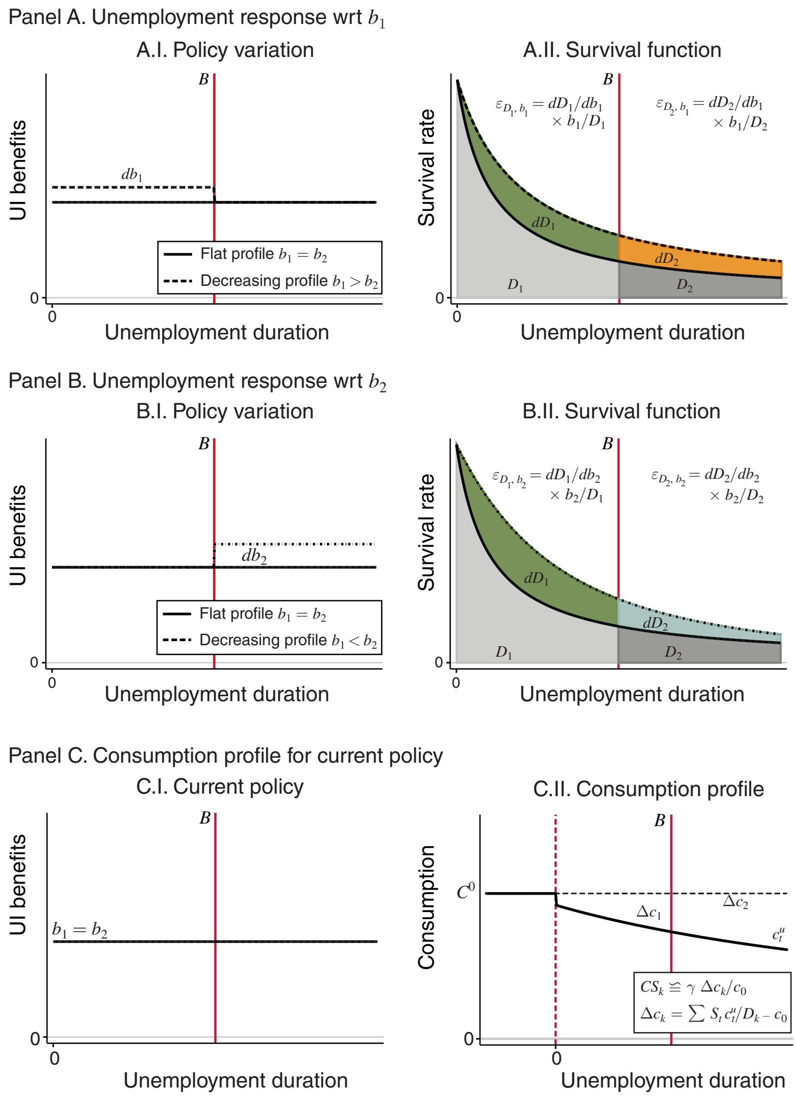

Figure 1. Sufficient Statistics for Welfare Analysis of Two-Part Policy

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 afterward. In panel A, we display policy variation db_1 in benefits given for the first B weeks that allows evaluating the moral hazard cost of a change in b_1, MH_1. The moral hazard cost depends on 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, as exemplified in the right panel. These responses enable the identification of the cross-duration elasticities ε_{D_1,b_1} and ε_{D_2,b_1} entering the RHS of the corresponding dynamic Baily-Chetty formula (11). Since we start from a flat profile, as is the case in our empirical application in Sweden, the response in total unemployment duration D is in principle sufficient, following equation (6). In panel B, we display policy variation that allows to evaluate the moral hazard cost of a change in b_2. Like in panel A, we start from a flat profile and display variation db_2 in benefits given after B weeks, which enables the identification of the cross-duration elasticities ε_{D_1,b_2}, ε_{D_2,b_2} entering the RHS of the corresponding dynamic Baily-Chetty formula. To evaluate the consumption smoothing gains of the two-part policy, following the implementation in (14), 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 C. Note that these consumption statistics need to be evaluated at the current profile, and do not require any policy variation.

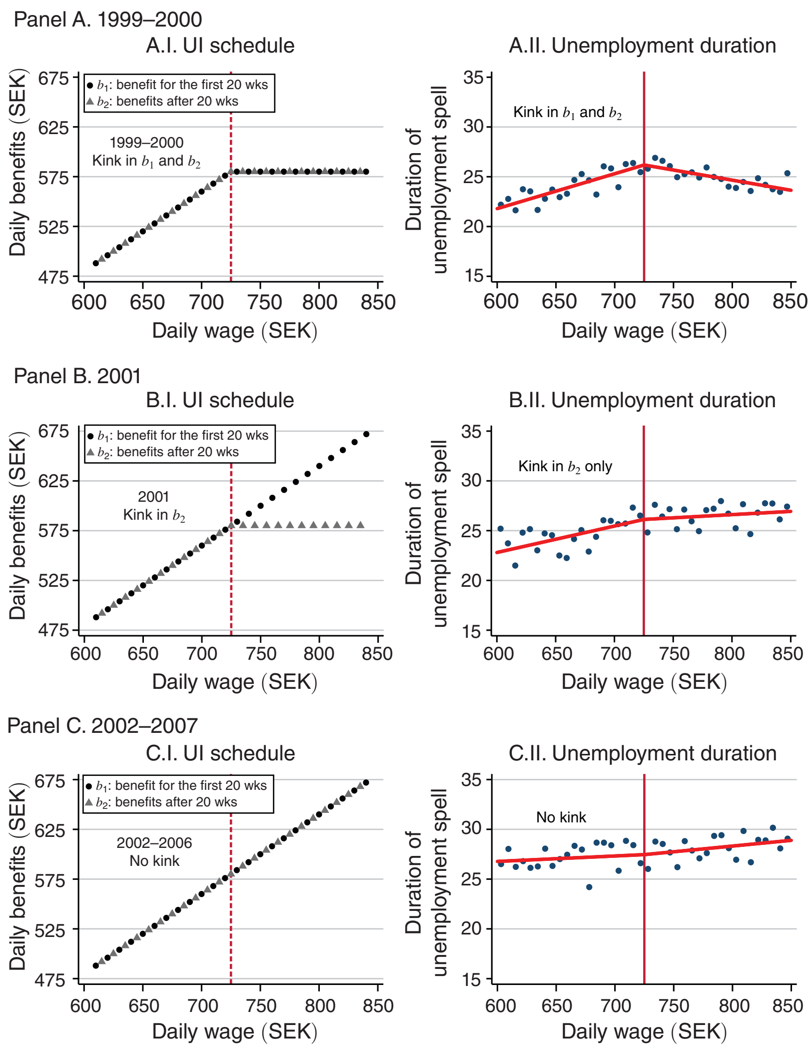

Figure 2. UI Benefits and Unemployment Duration as a Function of Daily Wage around the 725SEK Threshold

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 (panel 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 (panel 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 (panel 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 linear specifications of the form of equation (16) are displayed in Table 2. The vertical lines display predicted values of the regressions in the linear case.

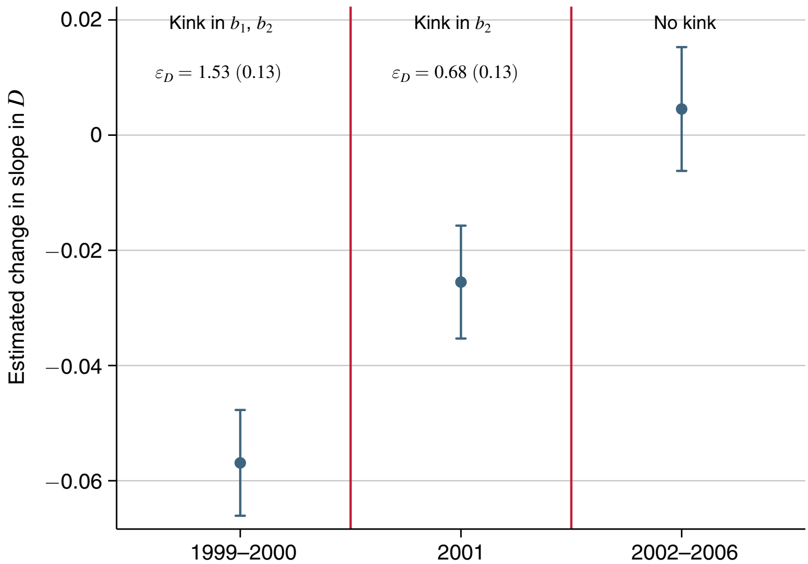

Figure 3. RKD Estimates at the 725SEK Threshold

Notes: The figure reports estimates of the change in slope with 95 percent robust confidence interval in the relationship between unemployment duration and daily wage at the 725SEK threshold using linear regressions of the form of equation (16) with a bandwidth size $h = 90$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). Bootstrapped standard errors computed using 50 replications are in parentheses. Formal estimates of the change in slope using linear specifications are displayed in Table 2.

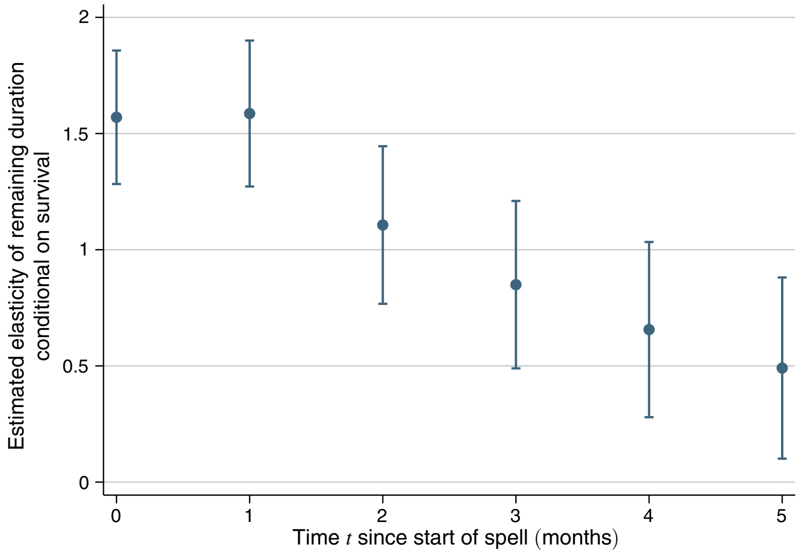

FIGURE 4. TESTING FOR STATIONARITY: ELASTICITY OF THE REMAINING DURATION OF UNEMPLOYMENT, CONDITIONAL ON SURVIVING UNTIL t, WITH RESPECT TO CHANGES IN THE FLAT BENEFIT LEVEL

Notes: The figure reports RKD estimates (with 95 percent 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 $\bar{b}$. Estimates use the presence, for spells starting before July 2001, of a kink in the benefit schedule of the flat benefit $\bar{b}$ at the 725SEK wage threshold. We use polynomial regressions of the form of equation (16) 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 $\bar{b}$ 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.).

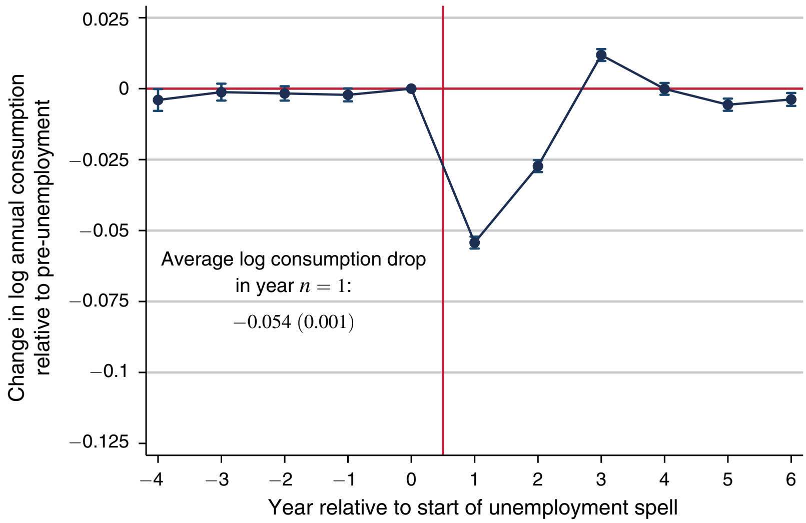

Figure 5. Estimated Change in log Annual Consumption as a Function of Time since Start of Unemployment Spell

Notes: The figure analyzes changes in log annual household consumption from our registry-based measure of consumption expenditures at unemployment, following the standard approach in the literature (e.g., Gruber 1997). We define event year $n = 1$ as the year an unemployment spell starts, and focus on individuals who are observed unemployed in December of year $n = 1$, and who were employed in December of year $n = 0$. We report coefficient estimates from the event year dummies γ_j from event study specification (17). Coefficients are relative to event year $n = 0$. See text for details.

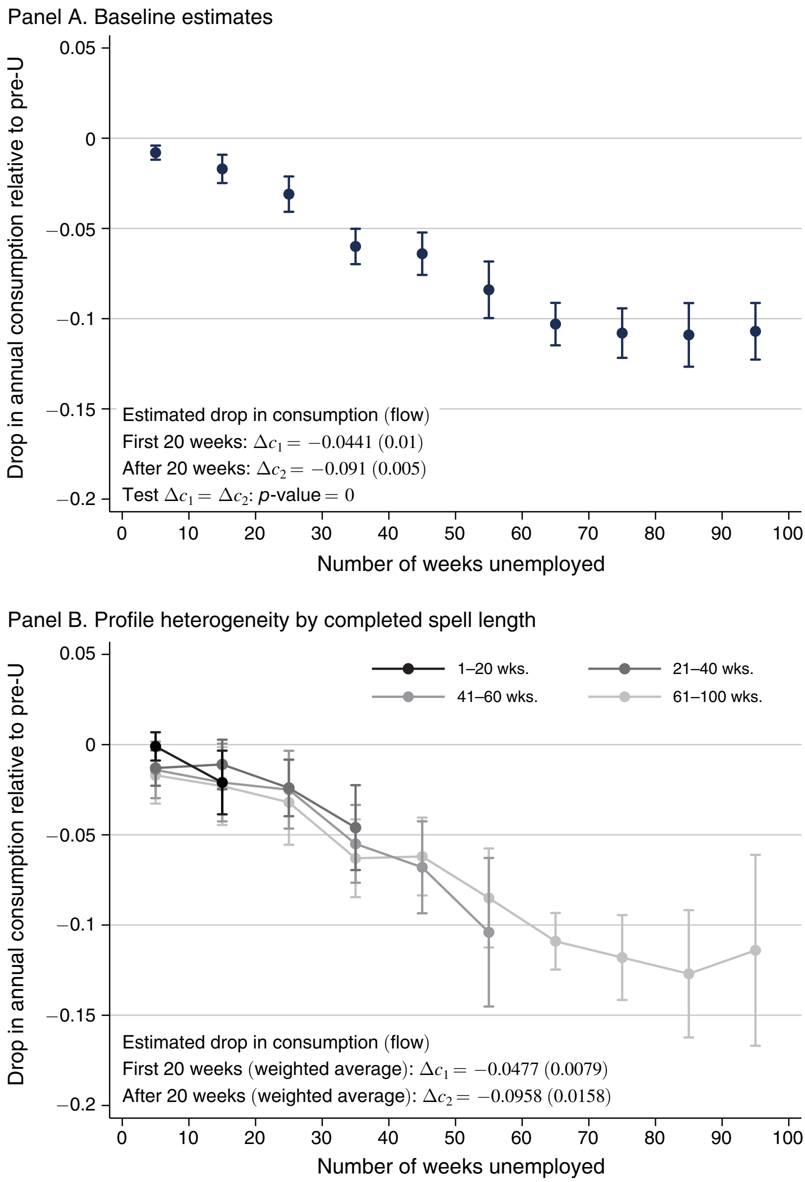

Figure 6. Estimated Drop in Annual Consumption Relative to Pre-Unemployment as a Function of Time Spent Unemployed

Notes: This figure shows average annual consumption drops compared to pre-unemployment $\Delta\mathcal{C}_t$ by time t spent unemployed as of December (when annual consumption is observed in the registry data), and uses this variation to non parametrically estimate the evolution of higher frequency household consumption throughout the unemployment spell. Panel A reports $\hat{\beta}_t/\mathcal{C}_0 = \widehat{\Delta\mathcal{C}_t}/\mathcal{C}_0$, i.e., estimates from equation (20) scaled by the average annual consumption in the last year prior to unemployment, so that all consumption drops are expressed relative to pre-unemployment levels. The figure also reports non parametric estimates of the average drops in consumption in each parts of the benefit profile $\widehat{\Delta C_1}$ and $\widehat{\Delta C_2}$ following the methodology explained in Section IVA. Standard errors are computed using the Delta method. We also report the p-value from a test of equality of $\widehat{\Delta C_1}$ and $\widehat{\Delta C_2}$. In panel B, we estimate separate profiles $\Delta\mathcal{C}_{t,k=1,\ldots,4}$ for 4 categories of total completed spell length L, and report estimates of $\widehat{\Delta C_1}$ and $\widehat{\Delta C_2}$ allowing for profile heterogeneity, following the methodology explained in Section IVA. Both panels provide compelling evidence of a significant drop in average flow consumption over the unemployment spell.

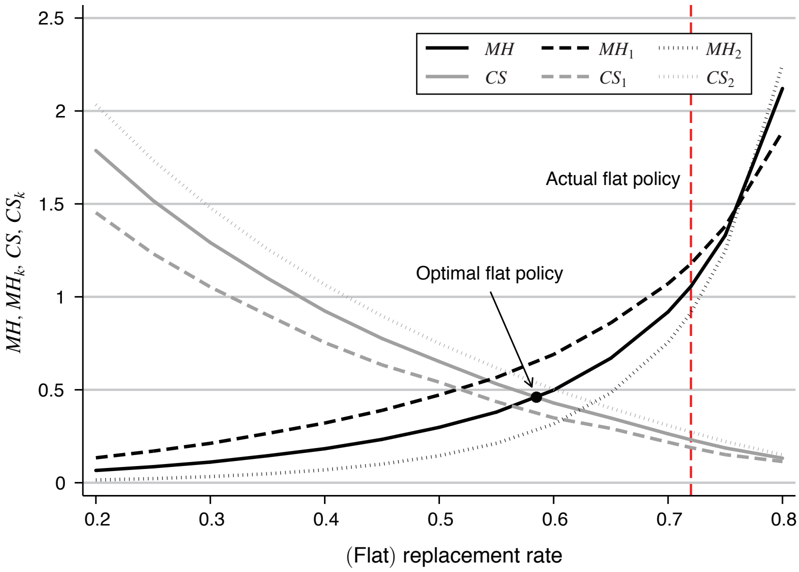

Figure 7. Structural Model: Welfare Effects for Different Benefit Levels

Notes: The figure illustrates how the moral hazard costs and consumption smoothing gains change for different levels of the flat benefit profile in our structural model. The model is calibrated to match our sufficient-statistics estimates, evaluated at the policy in place, which corresponds to a flat profile with average replacement rate of 0.72 as indicated by the vertical dashed line. We report the simulated moral hazard costs and consumption smoothing gains for an overall change in the flat benefit profile $\bar{b}$, for an increase in the benefit level in the first 20 weeks of unemployment, and for an increase in the benefit level after 20 weeks of unemployment. The simulated values of $CS_{\bar{b}}$ and $MH_{\bar{b}}$ (labeled CS and MH in the figure) are equalized for $\bar{b} = 0.58$. The consumption smoothing gains remain higher for benefits paid after 20 weeks ($CS_2 > CS_1$), while the moral hazard costs remain lower ($MH_2 < MH_1$), indicating that the introduction of an inclining tilt ($b_2 > b_1$) remains welfare improving for lower replacement rates.

TABLE 1—SUMMARY STATISTICS: DURATION RESPONSE AND CONSUMPTION RESPONSE SAMPLE

	Mean	P10	P50	P90
Panel I. Unemployment				
Duration of spell (wks)	26.78	3	16	73.43
Duration on b_1 (wks)	13.52	3	16	20
Duration on b_2 (wks)	13.26	0	0	53.43
Replacement rate	0.72	0.64	0.78	0.8
Panel II. Demographics				
Age	34.88	23	32	52
Fraction men	0.48	0	0	1
Fraction married	0.39	0	0	1
Fraction with higher education	0.25	0	0	1
Panel III. Income and wealth, SEK 2003(K)				
Gross earnings (individual)	190.3	171.3	191.3	2,07.52
Household net wealth	481.2	−222.4	20.6	1,318.1
Household bank holdings	52.9	0	0	139.1
Household real estate	631.2	0	163.9	1,605.6
Household debt	385.1	0	176.2	935.9

Notes: The table provides summary statistics for our main sample of unemployed individuals used in the RKD analysis of Section III and in the consumption response analysis of Section IV. The sample is drawn from the universe of unemployed individuals from the unemployment registers (PES–IAF) in Sweden from 1999 to 2007. We kept in the sample all individuals with daily wage in a bandwidth of 200SEK around the kink point in the benefit schedule. All earnings, income, and asset level measures are from wealth and income registers, and are yearly measures aggregated at the household level in constant k2003SEK for the last calendar year of full employment prior to the start of the unemployment spell. 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. Note that 1SEK2003 ≈ 2003 US$0.11

Table 2—RKD Estimates at the 725SEK Threshold

	Unemployment Duration D (1)	Duration D_1 ($<$ 20 weeks) (2)	Duration D_2 ($\geq$ 20 weeks) (3)
Panel I. 1999–2000: kink in b_1 and b_2			
δ_k	−0.0569	−0.0246	−0.0299
	(0.0050)	(0.0012)	(0.0039)
$\varepsilon_{D_k,\bar{b}}$	1.530	1.319	1.615
	(0.1300)	(0.0645)	(0.1986)
Observations	187,518	187,518	187,518
Panel II. 2001: kink in b_2 only			
δ_k	−0.0255	−0.0115	−0.0105
	(0.0049)	(0.0020)	(0.0030)
ε_{D_k,b_2}	0.6765	0.6015	0.5921
	(0.1312)	(0.1061)	(0.1642)
Observations	65,545	65,545	65,545
Panel III. 2002: placebo			
δ_k	0.0045	−0.0016	0.006
	(0.0055)	(0.0011)	(0.0049)
Observations	172,645	172,645	172,645

Panel IV. Moral hazard estimates		
$MH = \frac{\bar{b}+\tau}{\bar{b}} \cdot \varepsilon_{D,\bar{b}}$	1.637	(0.1391)
$MH_2 = \frac{\bar{b}+\tau}{\bar{b}} \cdot \frac{D}{D_2} \cdot \varepsilon_{D,b_2}$	1.445	(0.2829)
$MH_1 = \frac{\bar{b}+\tau}{\bar{b}} \cdot \frac{D}{D_1} \cdot (\varepsilon_{D,\bar{b}} - \varepsilon_{D,b_2})$	1.819	(0.4032)
Hypotheses testing: $MH_1 = MH_2$	z-stat	p-value
Z-test	−0.57	0.569
Permutation test		0.059

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 (column 1), the time D_1 spent on the first part of the Swedish UI profile (column 2), and the time D_2 spent on the second part of the Swedish UI profile (column 3). Estimates are obtained from linear regressions of the form of equation (16) with a bandwidth size $h = 90$SEK. 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 implied elasticities, computed as $\varepsilon_{D_k,b_k} = \hat{\delta}_k \cdot \frac{\bar{w}_k}{D_b^{cap}}$, where $\hat{\delta}_k$ is the estimated marginal slope change, and D_b^{cap} is the observed average duration at the kink. In Panel IV we also report implied moral hazard costs estimates defined in equation (6). *MH* is the moral hazard cost of increasing benefits throughout the unemployment spell. MH_1 is the moral hazard cost of increasing benefits given for the first 20 weeks of the spell. MH_2 is the moral hazard cost of increasing benefits given after the first 20 weeks of the spell. Computations assume $\tau = 0.05$ which balances the UI budget on average during the period 1999–2007. It follows that $\frac{\bar{b}+\tau}{\bar{b}} = 1 + 0.05/0.72 = 1.07$. See text for details. Standard errors are obtained from bootstrapping using 50 replications.

TABLE 3—WELFARE EVALUATION OF ACTUAL PROFILE USING ESTIMATED SUFFICIENT STATISTICS

	Moral hazard hazard costs	Average consumption drop	Consumption smoothing gains CS_k			Benefit-cost ratio
	MH_k	Δc_k	$\gamma = 1$	$\gamma = 2$	$\gamma = 5$	CS_k/MH_k
	(1)	(2)	(3)	(4)	(5)	(6)
Benefits given throughout the spell: $\bar{b}$	1.64	0.08	0.08	0.16	0.40	$\gamma \times 0.049$
	(0.14)	(0.01)	(0.01)	(0.02)	(0.04)	
Benefits given for the first 20 weeks: b_1	1.82	0.05	0.05	0.10	0.24	$\gamma \times 0.026$
	(0.40)	(0.01)	(0.01)	(0.02)	(0.04)	
Benefits given after first 20 weeks: b_2	1.45	0.10	0.10	0.19	0.48	$\gamma \times 0.066$
	(0.28)	(0.02)	(0.02)	(0.03)	(0.08)	

Notes: The table reports estimates of the sufficient statistics needed to evaluate the benefit profile in place in Sweden. The first row analyzes the welfare consequences of an increase in benefits $\bar{b}$ 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 repeats our estimates of the moral hazard cost of an increase in benefits (panel IV of Table 2). The second column repeats our estimates of the average consumption drop (column 1 of online Appendix Table C-3). To estimate the average consumption drop over the entire spell, we run the same regression as in (8) but with only one dummy $\mathbf{1}[t > 0]$ that indicates being observed while unemployed. We convert the respective consumption drops into estimates of the consumption smoothing gains CS_k, following the implementation in (14), which relies on a Taylor approximation and homogeneous preferences. The consumption smoothing gains are reported for a range of plausible values of the relative risk aversion γ in columns 3 to 5. Column 6 shows the ratio of consumption smoothing gains to moral hazard costs, CS_k/MH_k, depending on the uniform relative risk aversion γ. This ratio corresponds to the marginal value of a (tax-funded) kroner spent on unemployment benefits, accounting for the unemployment responses. Bootstrapped standard errors in parentheses.

REFERENCES

Acemoglu, Daron. 1995. "Public Policy in a Model of Long-Term Unemployment." *Economica* 62 (246): 161–78.

Aguiar, Mark, and Erik Hurst. 2005. "Consumption versus Expenditure." *Journal of Political Economy* 113 (5): 919–48.

Alvarez, Fernando E., Katarína Borovičková, and Robert Shimer. 2016. "Decomposing Duration Dependence in a Stopping Time Model." National Bureau of Economic Research Working Paper 22188.

Andrews, Isaiah, and Conrad Miller. 2013. "Optimal Social Insurance with Heterogeneity." Unpublished.

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

Browning, Martin, and Søren Leth-Petersen. 2003. "Imputing Consumption from Income and Wealth Information." *Economic Journal* 113 (488): F282–F301.

Campos, Rodolfo G., and Iliana Reggio. 2016. "Optimal Unemployment Insurance: Consumption versus Expenditure." *Labour Economics* 38: 81–89.

Card, David, David S. Lee, Zhuan Pei, and Andrea Weber. 2015. "Inference on Causal Effects in a Generalized Regression Kink Design." *Econometrica* 83 (6): 2453–83.

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

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

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods." *Annual Review of Economics* 1: 451–88.

Chetty, Raj, and Amy Finkelstein. 2013. "Chapter 3: Social Insurance: Connecting Theory to Data." In *Handbook of Public Economics*, Vol. 5, edited by Alan J. Auerbach et al., 111–93. New York: Elsevier.

Chetty, Raj, and Adam Szeidl. 2007. "Consumption Commitments and Risk Preferences." *Quarterly Journal of Economics* 122 (2): 831–77.

Cohen, Alma, and Liran Einav. 2007. "Estimating Risk Preferences from Deductible Choice." *American Economic Review* 97 (3): 745–88.

Coles, Melvyn G., and Eric Smith. 1998. "Marketplaces and Matching." *International Economic Review* 39 (1): 239–54.

Couch, Kenneth A., and Dana W. Placzek. 2010. "Earnings Losses of Displaced Workers Revisited." *American Economic Review* 100 (1): 572–89.

Eckstein, Zvi, and Gerard J. Van den Berg. 2007. "Empirical Labor Search: A Survey." *Journal of Econometrics* 136 (2): 531–64.

Eika, Lasse, Magne Mogstad, and Ola Vestad. 2017. "What Can We Learn about Household Consumption from Information on Income and Wealth?" Unpublished.

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

Ganong, Peter, and Simon Jäger. 2014. "A Permutation Test and Estimation Alternatives for the Regression Kink Design." Unpublished.

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

Hendren, Nathaniel. 2013. "The Policy Elasticity." National Bureau of Economic Research Working Paper 19177.

Hopenhayn, Hugo A., and Juan Pablo Nicolini. 1997. "Optimal Unemployment Insurance." *Journal of Political Economy* 105 (2): 412–38.

Johnston, Andrew C., and Alexandre Mas. 2015. "Potential Unemployment Insurance Duration and Labor Supply: The Individual and Market-Level Response to a Benefit Cut." Unpublished.

Koijen, Ralph, Stijn Van Nieuwerburgh, and Roine Vestman. 2015. "Judging the Quality of Survey Data by Comparison with 'Truth' as Measured by Administrative Records: Evidence From Sweden." In *Improving the Measurement of Consumer Expenditures,* edited by Christopher D. Carroll, Thomas F. Crossley, and John Sabelhaus, 308–46. Chicago: University of Chicago Press.

Kolsrud, Jonas, Camille Landais, Peter Nilsson, and Johannes Spinnewijn. 2018. "The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden: Dataset." *American Economic Review.* https://doi.org/10.1257/aer.20160816.

Kostøl, Andreas Ravndal, and Magne Mogstad. 2015. "Earnings, Disposable Income, and Consumption of Allowed and Rejected Disability Insurance Applicants." *American Economic Review* 105 (5): 137–41.

Kreiner, Claus Thustrup, David Dreyer Lassen, and Søren Leth-Petersen. 2014. "Liquidity Constraint Tightness and Consumer Responses to Fiscal Stimulus Policy." Unpublished.

Kroft, Kory, Fabian Lange, and Matthew J. Notowidigdo. 2013. "Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment." *Quarterly Journal of Economics* 128 (3): 1123–67.

Krueger, Alan B., and Bruce D. Meyer. 2002. "Labor Supply Effects of Social Insurance." In *Handbook of Public Economics*, Vol. 4, edited by Alan J. Auerbach and Martin Feldstein, 2327–92. Amsterdam: Elsevier.

Landais, Camille. 2015. "Assessing the Welfare Effects of Unemployment Benefits Using the Regression Kink Design." *American Economic Journal: Economic Policy* 7 (4): 243–78.

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

Lehr, Brandon. 2017. "Optimal Unemployment Insurance with Endogenous Negative Duration Dependence." *Public Finance Review* 45 (3): 395–422.

Lentz, Rasmus, and Torben Tranæs. 2005. "Job Search and Savings: Wealth Effects and Duration Dependence." *Journal of Labor Economics* 23 (3): 467–89.

Ljungqvist, Lars, and Thomas J. Sargent. 1998. "The European Unemployment Dilemma." *Journal of Political Economy* 106 (3): 514–50.

Lockwood, Ben. 1991. "Information Externalities in the Labour Market and the Duration of Unemployment." *Review of Economic Studies* 58 (4): 733–53.

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

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

Michaillat, Pascal. 2012. "Do Matching Frictions Explain Unemployment? Not in Bad Times." *American Economic Review* 102 (4): 1721–50.

Nekoei, Arash, and Andrea Weber. 2015. "Does Extending Unemployment Benefits Improve Job Quality?" Institute for the Study of Labor Discussion Papers 9034.

Pavoni, Nicola. 2009. "Optimal Unemployment Insurance, With Human Capital Depreciation, and Duration Dependence." *International Economic Review* 50 (2): 323–62.

Pistaferri, Luigi. 2015. "Household Consumption: Research Questions, Measurement Issues, and Data Collection Strategies." Unpublished.

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

Schmieder, Johannes F., and Till von Wachter. 2016. "The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation." *Annual Review of Economics* 8: 547–81.

Schmieder, Johannes F., Till von Wachter, and Stefan Bender. 2012. "The Effects of Extended Unemployment Insurance over the Business Cycle: Evidence from Regression Discontinuity Estimates over 20 Years." *Quarterly Journal of Economics* 127 (2): 701–52.

Shavell, Steven, and Laurence Weiss. 1979. "The Optimal Payment of Unemployment Insurance Benefits over Time." *Journal of Political Economy* 87 (6): 1347–62.

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

Shimer, Robert, and Iván Werning. 2007. "Reservation Wages and Unemployment Insurance." *Quarterly Journal of Economics* 122 (3): 1145–85.

Shimer, Robert, and Iván Werning. 2008. "Liquidity and Insurance for the Unemployed." *American Economic Review* 98 (5): 1922–42.

Spinnewijn, Johannes. 2015. "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs." *Journal of the European Economic Association* 13 (1): 130–67.

van den Berg, Gerard J. 1990. "Nonstationarity in Job Search Theory." *Review of Economic Studies* 57 (2): 255–77.

Werning, Iván. 2002. "Optimal Unemployment Insurance with Unobservable Savings." Unpublished.