

The Trade-Off between Insurance and Incentives in Differentiated Unemployment Policies*

JOHANNES SPINNEWIJN†

†*London School of Economics; Institute for Fiscal Studies*
(j.spinnewijn@lse.ac.uk)

Abstract

In this paper, I revisit the central trade-off between insurance and incentives in the design of unemployment insurance policies. The generosity of unemployment insurance benefits differs not only across countries, but also across workers within countries. After illustrating some important dimensions of heterogeneity in a cross-country analysis, I extend the standard Baily–Chetty formula to identify the key empirical moments and elasticities required to evaluate the differentiated unemployment policy within a country. I also review some prior work and aim to provide guidance for future work trying to inform the design of unemployment policies.

I. Introduction

One of the central adages of the Institute for Fiscal Studies (IFS) is to bring together theoretical models and empirical analysis to inform policy. Over the past 50 years, the research on the design of unemployment insurance (UI) has

* Submitted March 2019.

This article has benefitted immensely from my collaborations on unemployment insurance design with various co-authors and with my colleague Camille Landais in particular. I would also like to thank Gabriel Leite Mariante and Tatiana Pazem for excellent research assistance, and the editor Monica Costa Dias and two anonymous referees for valuable comments and suggestions. I acknowledge financial support from the ERC starting grant 716485.

Keywords: unemployment, policy design, heterogeneity.

JEL classification numbers: H20, J64, J65.

slowly adapted this adage, which has resulted in a rich literature that tightly integrates theory and empirics, and allows for comprehensive, evidence-based policy evaluations.

It was at the end of the first decade after the IFS was founded that Martin Baily characterised the central trade-off for UI design in his paper ‘Some aspects of optimal unemployment insurance’, published in the *Journal of Public Economics* in 1978. The trade-off is to provide insurance against the income loss due to unemployment, while maintaining incentives to reduce unemployment risk. In this classic paper, Baily emphasised the value that social insurance can provide, in response to emerging work that warned of the negative consequences of UI through lower employment incentives and increased unemployment rates.¹ He expressed the trade-off in a simple formula consisting of few moments, capturing both the consumption smoothing gains and the unemployment costs of UI. These moments can be readily taken to the data.

It took another 20 years until the early days of the empirical ‘credibility revolution’ when Jonathan Gruber was the first to implement the Baily formula empirically.² This spurred a significant amount of follow-up work, which took off further when Raj Chetty demonstrated the theoretical credibility of Baily’s simple formula ten years later.³ Chetty showed that the moments of the formula remain the same when enriching the underlying theoretical model. These moments are often referred to as ‘sufficient statistics’ for the evaluation of UI policy,⁴ akin to the elasticity of taxable income for tax policy.

Now, at the 50th anniversary of the IFS, the body of research on UI and the evaluation of its design, in particular, are impressive.⁵ Also, Baily’s formula has been extended and implemented to account for various other features that take us closer to the real world. This ranges from accounting for general equilibrium effects⁶ and behavioural frictions,⁷ to including informal sector work,⁸ allowing for duration-dependence⁹ and for dynamic selection.¹⁰ The work on UI relative to other social insurance programmes is arguably disproportional to the importance of UI policy as measured by its share of government expenditures. However, our improved understanding of the trade-off between insurance and incentives is just as applicable to other social

¹Feldstein, 1973, 1974.

²Gruber, 1997.

³Chetty, 2006.

⁴Chetty, 2009.

⁵See Chetty and Finkelstein (2013), Tatsiramos and Van Ours (2014) and Schmieider and Von Wachter (2016) for reviews.

⁶Landais, Michaillat and Saez, 2018a,b.

⁷Spinnewijn, 2015.

⁸Gerard and Gonzaga, 2016.

⁹Nekoei and Weber, 2017.

¹⁰Kolsrud et al., 2018.

insurance programmes, such as health insurance and social security. The focus on UI is partly explained by the presence of a tractable workhorse model (which is arguably lacking for other social insurance programmes) and the remarkable variation in UI policy and rules, allowing for the estimation of its causal effects.

Somewhat surprisingly, the evaluation of the specific policy rules used in the design of UI and resulting differentiation in UI generosity has been lagging behind. In fact, Martin Baily already acknowledged this omission in his original article in 1978, stating the following. ‘A compulsory government program prevents adverse selection from driving out the insurance coverage, but of course it is still true that when workers are not all alike, some of them have much more to gain from the program than others, and I am ignoring this.’ A number of studies have contributed to filling this gap: for example, Shavell and Weiss (1979), Hopenhayn and Nicolini (1997), Shimer and Werning (2008), Kolsrud et al. (2018) and Lindner and Reizer (2018) on the dynamic profile of benefits; Schmieder, Von Wachter and Bender (2012), Kroft, Lange and Notowidigdo (2013) and Landais, Michaillat and Saez (2018a,b) on benefits over the business cycle; and Michelacci and Ruffo (2015) on benefits over the life cycle. But, overall, understanding the vast differences in UI policies both across and within countries is an understudied area.

The question about how to optimally differentiate the unemployment policy – either based on exogenous conditions or on workers’ choices – closely relates to the question about which tags to use to target social transfers and to improve redistribution.¹¹ In this paper, I argue that the Baily formula and the sufficient statistics insights can be valuable to think about the design of differentiated UI. Thus, the main objective of this paper is to demonstrate the policy relevance of further research filling this gap in the UI literature, and to provide some guidance on how to get started.

The first part of the paper illustrates the variation in UI policy rules across countries and the within-country variation in UI generosity that the policy rules entail. I draw some key lessons from this institutional analysis. First, the main dimensions along which the UI generosity differs across different countries are pre-unemployment earnings and the length of the unemployment spell. Second, beyond these two dimensions, there are various other rules and conditions that countries apply. The desirability of these features seems untested, but the myriad of rules increases the complexity for potential beneficiaries, which seems undesirable in itself. Finally, as the variation in UI generosity within countries is substantial, a useful cross-country comparison of overall UI generosity is difficult. Moreover, a cross-country evaluation of UI generosity is complicated by the different context and objectives that apply in different

¹¹Akelof, 1978.

countries, which is less of a constraint for a within-country evaluation of the differentiated UI policies.

The second part of the paper revisits the trade-off between insurance and incentives as characterised by Baily (1978), but extending the framework to evaluate the optimal differentiation of UI policy. This extension builds on Kolsrud et al. (2018) who study how unemployment benefits should evolve over the unemployment spells, and it allows for some simple insights. First and foremost, the respective consumption smoothing gains and unemployment responses for the groups subject to different UI generosity are key (and sometimes sufficient) to evaluate the desirability of the differentiation itself. This highlights the value of further empirical work to estimate heterogeneity in the insurance gains and incentives, in particular along the particular dimensions used for the differentiation of UI benefits. The challenge is that the within-group policy variation required to estimate the relevant moments may not be readily available. Second, the evaluation is complicated when the eligibility rules are based on workers' choices, either directly or indirectly. Examples are when unemployment benefits are conditional on the time spent unemployed or the age at the onset of the unemployment spell. In general, it is necessary to estimate the response in the share of workers eligible to any benefit level to the specific change in the UI policy. Conveniently, these responses are of second-order importance when we start from a uniform policy and simply consider in which direction to differentiate the policy.

I conclude by discussing the empirical implementation of the insurance–incentive trade-off and I highlight some of the caveats. In particular, the recommendations on the optimal *differentiation* across groups can be more robust and/or less data-demanding than the evaluation of the optimal *level* for any given group. I also open the discussion on whether the analysis can be extended to start thinking beyond the traditional UI programmes that condition benefit receipt on unemployment so as to provide better protection against job loss in today's labour markets. Throughout my discussion, I review some recent contributions to the literature, but again I refer the interested reader to the aforementioned reviews for more comprehensive accounts of the relevant literature.

II. UI design in practice

In this section, I illustrate the substantial variation in UI generosity not only across, but also within countries. As countries differentiate their UI policies in very different ways, the protection from which workers can benefit when unemployed is very different depending on the country and the circumstances they are in. My focus is primarily on standard UI benefits, designed to (partially) replace pre-unemployment earnings while unemployed. In many countries, unemployed workers who are not entitled to (or have exhausted) standard UI

benefits can often rely on what the OECD defines as unemployment assistance (UA). These transfers are in general less generous and means-tested. I account for these transfers in the calculations below. Individuals who lose their job or are out of work can also benefit from other forms of income support or welfare (e.g., housing benefits, child benefits, food stamps, basic income, etc.), which are ignored in the calculations. The UI benefits, conditional on unemployment, can also be complemented with government-mandated severance payments, conditional on layoff. Severance payments have become more important over time and are the primary source of insurance against job loss in developing countries,¹² but these are also ignored in the calculations.

I have studied the UI policy rules for 18 countries, which include the United Kingdom, the United States and Canada, and a large number of European countries for which data availability and UI policy variation have been used in recent papers. This includes Austria,¹³ Germany,¹⁴ Hungary,¹⁵ Italy¹⁶ and Sweden.¹⁷ I also include Japan and South Korea. The only non-OECD country in the sample is Brazil, for which the UI policy has also been recently studied.¹⁸ The approach I take and the assumptions I have made for each country are explained in detail in the online appendix and the calculations are detailed in the spreadsheet provided online as supplementary material.

For these 18 different countries, Figure 1 plots the replacement rate at the start of the unemployment spell for a 35-year-old worker in a household with an employed partner and one child. The worker earned the respective country's average wage prior to unemployment. Large heterogeneity exists among the different countries, with replacement rates from as low as 0.1 in the United Kingdom to as high as 0.8 in Switzerland. No country offers full replacement of pre-unemployment wages for the average worker. There is a well-known divide in the generosity of social insurance programmes, with Anglo-Saxon and Eastern European countries on the less generous side of the benefit spectrum, and Western and Northern European countries on the more generous side. Asian countries tend to have replacement rates around the average of the sample, at least for this particular scenario.

A first important determinant of the unemployment benefit level to which workers are eligible is their earnings pre-unemployment. This is illustrated in Figure 2. I compare the replacement rates at the start of the unemployment spell for individuals who, prior to entering unemployment, earned their country's average wage (as in Figure 1) and for those who earned only half of that.

¹²Gerard and Naritomi, 2019.

¹³Card, Chetty and Weber, 2007; Nekoei and Weber, 2017.

¹⁴Schmieder, Von Wachter and Bender, 2016.


¹⁵Della Vigna et al., 2017; Lindner and Reizer, 2018.

¹⁶Citino, Russ and Scrutinio, 2019.

¹⁷Kolsrud et al., 2018; Landais et al., 2017; Landais and Spinnewijn, 2017.

¹⁸Gerard and Gonzaga, 2016; Gerard and Naritomi, 2019.

FIGURE 1
UI replacement rate across countries




Notes: The figure shows the replacement rate under the same baseline scenario in all 18 countries. I consider a 35-year-old worker, earning the average wage for the respective country, fully eligible and part of a household with a partner (also employed at the average wage) and one child. This scenario is changed in the following figures. Further details are provided in the online appendix.

Note that, in Sweden, the average pre-unemployment wage of unemployed individuals is almost exactly half the average national wage.¹⁹ Most countries lie below the 45-degree line in Figure 2, meaning that replacement rates tend to be higher for individuals with lower pre-unemployment wages. Typically, for those cases, benefits are calculated as a percentage of pre-unemployment earnings subject to an upper cap, which is binding at the average wage but not at half of the average wage. In Spain, for example, initial unemployment benefits are set at 70 per cent of the individual's pre-unemployment wage, capped at the monthly amount of 1254.96 euros, which is lower than 70 per cent of the average wage, but not lower than 70 per cent of half of the average wage. This results in a replacement rate of less than 70 per cent for individuals whose pre-unemployment wage was equal to the average wage. A few countries are shown precisely on the 45-degree line of the graph, meaning that the replacement rate is the same for the two considered cases of pre-unemployment wage. In those cases, the benefit cap is typically non-binding even at the average wage. This is the case for Portugal, where the 65 per cent replacement rate is applied both

¹⁹Kolsrud et al., 2018.


FIGURE 2

Cross-country UI replacement rates by pre-unemployment earnings

at the average wage and at half of the average wage. For the remainder of the analysis, I consider individuals who earned half of the national average wage.

A second important dimension along which unemployment benefits vary is the length of the ongoing unemployment spell. Figure 3 illustrates how unemployment benefits vary over the course of the unemployment spell, comparing the replacement rates at the start of the unemployment spell with those one year into the spell. I consider individuals who contributed to the UI system for one year, which in many countries is the minimum necessary in order to be eligible for benefits. In general, benefits decline over the course of the spell. In fact, in many countries, UI benefits are limited in time and workers exhaust their UI benefits before one year of unemployment. In the United States, for example, workers tend to receive UI benefits for only six months. In several countries, however, workers are eligible to further UA, typically characterised by smaller amounts, no contribution requirements and no duration limit. This is the case, for example, in Germany. At the start of the unemployment spell, a German worker with dependent children has the right to claim benefits worth 67 per cent of pre-unemployment wage (subject to a cap that is non-binding here), but after one year such eligibility has expired. The individual is then eligible for further UA, which comprises flat payments that amount to less than 67 per cent of pre-unemployment wage. The UA is often means-tested. For example, in the United Kingdom,

FIGURE 3
Cross-country UI replacement rates by unemployment duration



unemployed workers exhaust Job Seeker's Allowance after six months unless their household income and savings are below specific thresholds.

Figure 4 illustrates how contribution time and age, often interacted, affect benefits over the course of the unemployment spell. I compare the replacement rates for unemployed workers who are one year into the spell for two situations: younger workers who contributed to the UI system for only one year and older workers who contributed for five years. In certain countries, benefits one year into the spell are the same in both situations, either because all eligible payments are exhausted (e.g. Japan and Brazil) or because the worker has moved away from regular UI benefits into further UA, which does not depend on age and contribution (e.g. Greece and Austria). However, it is often the case that older individuals with a longer contribution history can claim higher benefits in the long run. For example, in Italy, older individuals can receive regular UI for longer than one year, while younger workers cannot. In Spain, it is the contribution history that determines the duration of regular UI benefits: an individual who contributed for five years can claim UI for longer than one year, while an individual who contributed for only one year moves to less generous UA before the end of the first year of the unemployment spell.

To further illustrate how the differentiation of the UI policy based on individual characteristics can vary across countries, I show the UI generosity in Sweden and Switzerland side-by-side under different scenarios in Figure 5.

FIGURE 4

Cross-country UI replacement rates for long-term unemployed with varying years of contribution and age





FIGURE 5

Within-country variation: comparing Switzerland and Sweden



In both countries, a worker earning half the national average wage, aged 35 with one child, is entitled to the same level of benefits at the start of the spell (i.e. 80 per cent of the pre-unemployment wage). However, different individual characteristics have different consequences for the level of benefits in both countries. For a worker with a pre-unemployment wage equal to the full national average wage, the replacement rate falls to 50 per cent in Sweden, but remains at 80 per cent in Switzerland. In both countries, benefits decrease after one year if the worker has only contributed for one year. However, for older workers with a longer history of UI contributions, benefits do not decrease after one year in Switzerland, but they do so in Sweden. In addition, Switzerland has a child premium, with the replacement rate for a childless worker falling to 70 per cent, while in Sweden the presence of a child does not affect the amount of benefit.

This brief summary has highlighted some important drivers of variation in insurance coverage and how their importance differs in different countries, but this does not do full justice to the often complex myriad of rules that countries apply and how these rules differ across countries. The UI generosity is determined not only by earnings and employment history before unemployment, but also by the type of job separation. In France, the benefits depend on whether the unemployment spell follows the end of a fixed-term contract, an individual layoff or an economic layoff. In most countries, only involuntary unemployment is covered by the UI programme. In Sweden, however, workers can receive unemployment benefits after voluntarily leaving a job, but only 50 days after unemployment. Waiting periods are also used for young workers who have just entered the labour market (e.g. in Belgium). In addition to the unemployment durations, countries such as Brazil and Portugal make the level of UI benefits dependent on the number of past unemployment spells. In the United States, the potential duration of benefits can be extended for workers when the unemployment rate is high. The effects of this countercyclical UI policy has been the subject of a recent debate.²⁰ In addition to extra eligibility conditions, there are other transfers than the unemployment benefits that the unemployed might be eligible for, specifically. One example is financial aid for housing and utilities, which can be sizeable. In Germany, for example, a couple with one child can receive up to 562 euros extra per month.

III. UI design in theory

In this section, I revisit the Baily–Chetty characterisation of optimal UI, but I allow the UI benefit level to depend on individual characteristics. In the spirit

²⁰Farber and Valletta, 2015; Chodorow-Reich and Karabarbounis, 2016; Hagedorn, Manovskii and Mitman, 2015.

of the sufficient statistics approach, the characterisation is a function of a few high-level moments that can, in principle, be estimated empirically, which is the topic of the next section. While I touch upon some of the issues here, I refer the interested reader to Chetty and Finkelstein (2013) and Kleven (2018) for a more in-depth discussion of the Baily–Chetty formula and the so-called sufficient statistics approach and its challenges. The extension and notation I use build on Kolsrud et al. (2018).

1. Set-up

We consider an unemployment policy designed to protect workers against unemployment risk. The unemployment policy $P = \{b_x\}_{x \in X}$ conditions the UI benefit level on a vector of characteristics x . This could include the workers' pre-unemployment earnings, the unemployment duration, the age at the onset of the spell, macro-economic indicators, etc. A key concern in the design of the policy is that the share of unemployed workers will depend on the generosity of the policy and thus will affect the expected cost of the policy. Denoting the total labour force by L and the share of workers on unemployment benefit b_x by S_x , the government's budget surplus can be written as

$$(1) \quad G(P) = [\bar{\tau} - \sum_{x \in X} S_x [b_x + \tau_x]] \times L,$$

where $\bar{\tau}$ is the average tax in the labour force and τ_x is the average tax for workers with characteristics vector x . The fiscal cost of unemployment depends not only on the UI benefits b_x , but also on the foregone tax revenues τ_x .

It is not only the share of unemployed workers that depends on the policy. The fact that the variables used to differentiate the policy can be endogenous to the policy itself matters for the evaluation of a differentiated unemployment policy. For example, when UI benefits are more generous for higher pre-unemployment wages, this can affect workers' incentives to find a high-wage job. Similarly, when UI benefits are less generous for the long-term unemployed, this changes workers' incentives to leave unemployment early in the spell. Even for seemingly exogenous characteristics, such as age, which is used to differentiate UI benefits, workers can still manipulate the age at which they start an unemployment spell.²¹ That is, when the potential duration of benefits is longer for older workers, this affects workers' incentives to delay the start of an unemployment spell. In general, the share of workers on b_x will depend not only on the efforts of workers' with characteristics x to avoid unemployment, but also on workers' incentives to be eligible for b_x . As the set-up is general, the insights extend beyond the differentiation of standard UI benefits and can account for other social assistance that unemployed workers

²¹Doornik, Schoenherr and Skrastins, 2018; Citino, Russ and Scrutinio, 2019.

can benefit from (e.g. housing benefits) or alternative policies that individuals who are out of work can consider applying for (e.g. disability benefits or early retirement benefits).

I assume that workers try to maximise their welfare, given the policy in place, potentially affecting their unemployment risk and UI benefit eligibility, and I denote the resulting utility by $V_i(P)$. Social welfare associated with an unemployment policy P can then be written as

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

where \bar{G} is an exogenous revenue constraint and λ equals the marginal cost of public funds. The policy's central trade-off is to provide insurance against unemployment while maintaining incentives to avoid unemployment. The characterisation naturally extends for general Pareto weights to account for redistributive motives.

2. Sufficient statics approach

The most powerful idea – but also its strongest limitation – underlying the sufficient statistics approach is to focus on local deviations from the current policy. The welfare effects of local deviations are easier to characterise as we only need to account for the externalities of individuals' responses to the policy change, as we discuss below. This can predominantly depend on the corresponding effect on the government's budget constraint, which workers do not internalise. As is well known, local deviations from the current policy can be considered to test the optimality of the policy in place, as no deviation from an optimal policy can increase welfare. However, local deviations also allow us to identify the direction in which a suboptimal policy should be changed to increase welfare, although an important caveat is that we are restricted to local recommendations. This can be particularly restrictive when evaluating multi-dimensional policies: the welfare effect of changing b_x may crucially depend on $b_{x'}$. We would need to embed the approach in a more structural framework to evaluate big policy reforms or to provide recommendations on the optimal policy. However, the workings of the structural framework will depend on model primitives that are harder to identify, leading to a trade-off between the two approaches.²² The frontier is to leverage the advantages of both approaches: for example, using the sufficient statistics approach to identify and estimate the welfare effect of local changes and using the structure of the model to gauge how this welfare effect would change as we move away from the current policy.²³

²²Chetty, 2009.

²³Kolsrud et al., 2018.

a) *Local deviations*

Consider now an increase in the benefit level b_x , keeping all other parts of the unemployment policy fixed. The total effect on welfare depends on how much the unemployed who benefit from this increase value it in comparison to its budgetary cost. This can be expressed as

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

As is well known now in the public finance literature, the welfare effect of a policy change depends on the agents' behavioural responses, but only to the extent that the agents' behaviour has consequences that they did not internalise themselves. Indeed, invoking the envelope theorem, agents' responses to a policy change will have only a second-order effect on their own welfare $V_i(P)$.²⁴ For example, agents can change their job-search strategy in response to an increase in UI benefits. However, as they were optimising before the policy change, this response will have a negligible effect on their own welfare. Similarly, agents can undertake action to become eligible for b_x , but because they were making this trade-off optimally before the policy change, the welfare effect will be of second order.²⁵ As a consequence, in the absence of other externalities, we only need to account for the effect of behavioural responses on the government's budget $G(P)$ (i.e. the fiscal externality) and the direct effect of the policy change on agents' welfare.

b) *Value of UI*

The direct effect of an increase in b_x depends on the agents' welfare gain from having the extra resources available when unemployed and eligible for b_x . This gain is fully captured by the marginal utility of consumption $v'_u(c_u)$ for each unemployed worker on b_x ,

$$(4) \quad \int \frac{\partial V_i(P)}{\partial b_x} di = S_x \times E[v'_u(c_u)|x].$$

Clearly, the value of UI will depend on the severity and persistence of the income shock that unemployed workers are exposed to, and the private means and social transfers they have access to in order to protect themselves against such a shock. Both will affect the extent to which workers can smooth consumption when unemployed. However, the utility gain from the extra dollar of consumption when unemployed, possibly due to the UI benefit increase, is

²⁴Chetty, 2006.

²⁵Changes in the choice variables might be discontinuous in response to small policy changes. In principle, one can allow for such discontinuous behavioural responses if they average out when integrating across heterogeneous individuals so that the social welfare function is differentiable.

all we need to know to capture the value of UI at the margin. The overall effect thus simplifies to the share of workers who are unemployed and eligible for b_x and their average marginal utility of consumption. Note that the income shock due to job loss may well extend beyond the unemployment spell as unemployed workers may become re-employed at lower wages and in more insecure jobs. Workers' consumption as they go in and out of unemployment will be interdependent. Nevertheless, as UI benefits condition on unemployment, the workers' consumption when unemployed is sufficient to determine the marginal value of these UI benefits.²⁶

c) Cost of UI

Let us now turn to the budgetary effect from an increase in b_x . The first effect from increasing the benefit level is mechanical and again depends on the share of unemployed workers receiving b_x . The second effect is behavioural and is determined by the corresponding budgetary cost. Conveniently, there is no need to know any individual responses, as only the aggregate effect on the government's budget matters. However, the budget will be affected not only through the increased unemployment risk in response to the more generous benefit, but also through the changed eligibility for the different benefit levels. Hence, we need to know the *full* vector of elasticities,

$$\varepsilon_{S_{x'}, b_x} = \frac{\partial S_{x'} / \partial b_x}{S_{x'} / b_x},$$

which captures the responses of the worker shares on unemployment benefit $b_{x'}$ when changing benefit level b_x :

$$\begin{aligned} (5) \quad \frac{\partial G(P)}{\partial b_x} &= -S_x - \sum_{x'} \frac{\partial S_{x'}}{\partial b_x} (b_{x'} + \tau) \\ &= -S_x \times \left[1 + \sum_{x'} \frac{S_{x'} (b_{x'} + \tau)}{S_x b_x} \varepsilon_{S_{x'}, b_x} \right]. \end{aligned}$$

While, in principle, the full set of responses determines the fiscal externality, the set of economically relevant elasticities can be smaller. Mapping out these interaction effects and estimating the ones that are relevant will be key to evaluate a differentiated UI programme.²⁷

²⁶In comparison, the marginal value of severance payments would be determined by the consumption at layoff, regardless of the unemployment spell that follows, or the marginal value of re-employment bonuses would be determined by the consumption on re-employment.

²⁷The elasticities are weighted by the relative share of the budget spent on different parts of the unemployment policy. The budgetary spillover effects of a change in b_i on other parts of the policy are 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.

d) *Uniform policy*

I first put together the gain and cost from increasing UI generosity for the uniform profile b_u , allowing the well-known formula in Baily (1978) and Chetty (2006) to be retrieved. The optimal uniform policy is characterised by

$$(6) \quad \frac{\partial W}{\partial b_u} = 0 \Leftrightarrow \frac{E[v'_u(c_u)] - \lambda}{\lambda} = \frac{b_u + \tau_u}{b_u} \varepsilon_{S_u, b_u}.$$

The left-hand side measures the premium that society *is willing* to pay, as expressed by the social welfare function, to transfer an extra dollar to an unemployed worker. This crucially depends on the difference in marginal utility of consumption when unemployed and when employed. The more workers can smooth their consumption when losing their job, the less valuable UI is. The right-hand side measures the premium society *ought to* pay when transferring an extra dollar to the unemployed due to the increase in unemployment. The higher the elasticity of the unemployment rate, the more costly UI is. At the optimum, the marginal value and cost of UI should be equalised.

3. Differentiated unemployment policy

We first generalise the characterisation of the optimal uniform policy to the optimal differentiated policy when workers have no control over the eligibility criteria. That is, $\varepsilon_{S_{x'}, x} = 0$ for $x' \neq x$. For each part of the policy, we need

$$(7) \quad \frac{\partial W}{\partial b_x} = 0 \Leftrightarrow \frac{E[v'_u(c_u)|x] - \lambda}{\lambda} = \frac{b_x + \tau_x}{b_x} \varepsilon_{S_x, b_x}.$$

The consumption smoothing gains and the unemployment response among the group of workers eligible for b_x are sufficient to evaluate whether the benefit level is optimally set. Put differently, for two different groups of workers, any differentiation of the UI policy is only justified if either the consumption smoothing gains or the unemployment response is different. The test for whether UI should be more generous for one group relative to another is simply whether the consumption smoothing gains relative to the unemployment cost are higher.

The difference in consumption smoothing gains and unemployment responses across different groups of unemployed workers remains key when the unemployment policy is differentiated based on endogenous conditions. To characterise the optimal differentiated policy, however, we do need to account for the effect of policy changes across different parts of the differentiated

policy:

$$(8) \quad \frac{\partial W}{\partial b_x} = 0 \Leftrightarrow \frac{E[v'_u(c_u)|x] - \lambda}{\lambda} = \frac{b_x + \tau_x}{b_x} \varepsilon_{S_x, b_x} + \sum_{x' \neq x} \frac{S_{x'}(b_{x'} + \tau_{x'})}{S_x b_x} \varepsilon_{S_{x'}, b_x}.$$

I provide two examples to shed light on the potential role of these cross-elasticities.

Consider, first, a two-part benefit policy differentiating the UI benefit level for two groups of workers (e.g. low-income versus high-income workers, $X = \{y < \bar{y}, y \geq \bar{y}\}$). The share of workers receiving $b_{<\bar{y}}$ depends on the share of low-income workers and the unemployment rate among low-income workers, $S_{<\bar{y}} = F_{<\bar{y}} U_{<\bar{y}}$. A change in UI for low-income workers has the following budgetary effect:

$$(9) \quad \frac{\partial G(P)}{\partial b_{<\bar{y}}} = -S_{<\bar{y}} - \frac{\partial S_{<\bar{y}}}{\partial b_{<\bar{y}}}(b_{<\bar{y}} + \tau_{<\bar{y}}) - \frac{\partial S_{\geq\bar{y}}}{\partial b_{<\bar{y}}}(b_{\geq\bar{y}} + \tau_{\geq\bar{y}})$$

$$(10) \quad \cong -S_{<\bar{y}} - F_{<\bar{y}} \frac{\partial U_{<\bar{y}}}{\partial b_{<\bar{y}}}(b_{<\bar{y}} + \tau_{<\bar{y}}) - \frac{\partial F_{<\bar{y}}}{\partial b_x} U_{\bar{y}}(b_{\geq\bar{y}} - b_{<\bar{y}}).$$

Here, the approximation relies on the effect of the low-income benefit level on the high-income unemployment rate being small ($\partial U_{\geq\bar{y}}/\partial b_{<\bar{y}} \approx 0$). We can obtain a similar expression for $b_{\geq\bar{y}}$. The endogeneity of the condition used to differentiate UI requires accounting for the share of workers who may lower their income to become eligible for the more generous unemployment benefit level for low-income workers. The fiscal externality depends on the unemployment rate among these workers at the margin, which I denote by $U_{\bar{y}}$, and the difference in UI generosity on the two parts of the policy. Importantly, starting from a uniform policy $b_{\geq\bar{y}} = b_{<\bar{y}}$, this fiscal externality becomes of second order, and so even with endogenous conditions, it is only the unemployment response of the directly affected group that matters when considering to differentiate the UI generosity of that group. By the same token, the more differentiated the UI policy, the more important the eligibility responses are when considering to further differentiate the policy.

Consider now a two-part benefit policy differentiating the UI benefit level for two parts of the unemployment spell: e.g., the first six months of unemployment and thereafter, $X = \{d < \bar{d}, d \geq \bar{d}\}$. This is the setting analysed in Kolsrud et al. (2018). The share of unemployed workers on the different parts of the UI policy are now directly related as the share of long-term unemployed depends on the share of short-term unemployed and their exit rate out of unemployment. More formally, we have $S_{d+1} = S_d(1 - h_d)$,

where h_d is the hazard rate at unemployment duration d . The corresponding cross-elasticities have been at the heart of the theoretical models used to study the optimal timing of UI benefits. On the one hand, generous unemployment benefits at the start of the spell ($b_{<\bar{d}}$) will discourage workers from finding employment and affect the share of workers moving on to the second part of the unemployment policy. On the other hand, generous unemployment benefits later in the spell ($b_{\geq\bar{d}}$) already reduce the incentives early in the spell for job seekers who are forward-looking and contemplate the risk of being long-term unemployed. In addition to the cross-elasticities, there are of course the direct effects. How responsive are short-term versus long-term unemployed workers to changes in unemployment benefits? How do these responses compare with consumption smoothing gains from the short-term and long-term unemployed? The power of the proposed approach is to capture all the forces together and turn these questions into empirical ones. In fact, when starting from a flat benefit profile b_u , the characterisation simplifies to

$$(11) \quad \frac{\partial W}{\partial b_{<\bar{d}}} = 0 \Leftrightarrow \frac{E[v'_u(c_u)|d < \bar{d}] - \lambda}{\lambda} = \frac{S_{<\bar{d}} + S_{\geq\bar{d}}}{S_{<\bar{d}}} \frac{b_u + \tau_u}{b_u} \varepsilon_{[S_{<\bar{d}} + S_{\geq\bar{d}}], b_{<\bar{d}}}$$

The benefit level at a given unemployment duration should be lower, when the consumption smoothing gains are lower and when the overall unemployment response to changes in that benefit level is higher, accounting for the share of workers who are still unemployed at that duration.

IV. From theory to practice

The key advantage of the identified trade-off between incentives and insurance, following the Baily–Chetty formula, is its potential to be evaluated empirically. The identified moments have clear empirical counterparts, which make the link from theory to recommendations very transparent. However, while the characterisation is general and robust to richer models, the empirical implementation naturally requires assumptions. These assumptions can be strong, so it is important, first of all, to be transparent about them. Moreover, it is valuable to use complementary approaches and data to try to relax the assumptions or to consider policy recommendations that do not depend on these assumptions. I discuss this briefly, in the context of differentiated unemployment policies.

1. Consumption smoothing

The value of higher unemployment benefits to a group of workers depends on the marginal utility of consumption for that group of workers when they are unemployed. It is standard in the literature to ignore the redistributive value of

UI, which would affect the wedge between $E[v'_u(c_u)|x]$ and the marginal cost of public funds λ . The focus is instead on the insurance value, comparing the marginal utility of consumption for a given worker when unemployed versus employed. At the margin, the insurance value of transferring an additional dollar from employment to unemployment depends on the drop in consumption that unemployed workers are exposed to and how much they care about the variation in consumption. The standard consumption-based approach in the literature is to estimate the drop in consumption when unemployed and then to scale the drop in consumption with a risk preference parameter. Indeed, assuming preferences with constant relative risk aversion $\gamma = -cv''(c)/v'(c)$, the relative difference in marginal utilities simplifies to

$$\frac{v'(c_u) - v'(c_e)}{v'(c_e)} \cong \gamma \frac{c_e - c_u}{c_e}.$$

Implementing the consumption-based approach is remarkably easy as it only requires linking data on consumption expenditures to data on unemployment status. In recent years, more registry-based data on consumption expenditures have become available, which can also be linked to employment registers.²⁸ A major advantage for the implementation is that, in principle, there is no need to know the means that workers use to smooth their consumption. However, this is conditional on knowing the preference parameter with which the consumption drop should be scaled. The relevant preference parameter would depend not only on whether workers' preferences over consumption are state-dependent (e.g. through complementarities with leisure), but also on whether observable consumption expenditures are state-dependent and affect welfare-relevant consumption differently when employed and unemployed. This is a specific concern in the context of unemployment because of differences in work-related or job-search-related expenditures, or because unemployed workers can substitute towards home production and shop at lower prices. More generally, the relevant preference parameter will depend on the type of consumption expenditures that unemployed job seekers can change. For example, being able to lower expenditures on durable goods will affect the marginal utility of consumption less,²⁹ while being forced to lower only non-committed expenditures will affect the marginal utility of consumption more.³⁰

The challenge of translating consumption wedges into marginal utility wedges does not disappear when evaluating the optimal differentiation of a policy, but the nature of the challenge changes. While differences in

²⁸Kolsrud et al., 2018; Gerard and Naritomi, 2019.

²⁹Browning and Leth-Petersen, 2003.

³⁰See Chetty and Szeidl (2007). Also, see Chetty and Finkelstein (2013) and Landais and Spinnewijn (2017) for a further discussion of the challenges for the consumption-based approach.

consumption drops across workers with different characteristics may seem indicative, it is necessary to know how preferences change as well. For example, to evaluate how high unemployment benefits should be for workers with low pre-unemployment earnings, we should consider the consumption drop in this group, but we would also need to know their respective preferences. In fact, differences in consumption may be more than offset by differences in preferences, as workers with different preferences will invest differentially in consumption smoothing.³¹ However, to evaluate the relative generosity of two parts of the unemployment policy, we no longer need to know the level of the preference scalars $\gamma_x, \gamma_{x'}$, but it is sufficient to know the relative preferences $\gamma_x/\gamma_{x'}$. In particular, when we know that preferences are comparable (i.e. $\gamma_x \approx \gamma_{x'}$), it becomes sufficient to know the relative drop in consumption. Kolsrud et al. (2018) find that the consumption drop is more pronounced for the long-term unemployed than for the short-term unemployed, consistent with the fact that workers run out of assets as they remain unemployed for longer. This indicates that the consumption smoothing gains are larger for the long-term unemployed, unless the workers who select into long-term unemployment are less averse to the reduced consumption.

It will always be challenging to convert observable responses or wedges into comparable objects for welfare analysis. The extra challenge in the context of UI is that coverage is mandated and workers do not directly reveal their valuation through their coverage choice.³² A number of recent papers in the UI literature further tackle this challenge by proposing alternative approaches. In the spirit of the consumption-based approach, it is possible to look at wedges in resources used when employed and unemployed, such as changes in spousal labour supply,³³ or to consider responses to changes in the anticipated unemployment risk.³⁴ These responses still need to be scaled by a preference parameter. A way to circumvent the scaling is to use differences in behavioural responses to different sources of income variation. Chetty (2008) and Landais (2015) show how the differential response in unemployment risk to changes in unemployment benefits, relative to other income changes, can capture consumption smoothing gains. Landais and Spinnewijn (2017) instead show how to use the differential marginal propensity to consume when unemployed and employed, in order to identify the revealed cost of smoothing consumption so as to bound the value of consumption smoothing.³⁵ Overall, the authors who

³¹Chetty, 2006; Landais and Spinnewijn, 2017.

³²An exception is the UI policy in Scandinavian countries where workers are offered to choose between basic and more comprehensive coverage (Landais et al., 2017).

³³Fadlon and Nielsen, 2019.

³⁴Hendren, 2017.

³⁵Landais and Spinnewijn (2017) also leverage the UI choices in Sweden to obtain revealed-preference estimates of the value of insurance, and to implement the various approaches in the same context on the same sample of workers.

use the alternative approaches suggest that the insurance value of UI is higher than the consumption-based approach suggests (using ‘standard’ values of risk aversion). The implied mark-ups that workers are willing to pay to transfer consumption from employment to unemployment are closer to 100 per cent and sometimes above, rather than around 10–50 per cent as suggested by the consumption-based approach. These differences are sizeable, but very little is known about how the estimated consumption gains differ for different groups.

2. Unemployment response

In comparison with studies of the insurance value, there is a vast body of literature estimating the labour supply effects of social insurance,³⁶ and the unemployment responses to changes in UI benefits in particular.³⁷ The predominant focus in the literature is on the unemployed themselves and how the unemployment duration changes when varying the benefit level or the potential duration of benefits. To estimate the effect of UI generosity, it is common to use exactly the variation that comes from the differentiation in unemployment benefits (e.g. by pre-unemployment earnings, by contribution years or by age at layoff).³⁸ Indeed, the differentiated schedules include jumps or kinks (e.g. at the cut-offs for earnings thresholds or age), providing plausibly exogenous variation. However, the variation does not allow us to evaluate the differentiated schedule itself.

Let me illustrate this for a two-part benefit policy differentiating the UI benefit level for low-income versus high-income workers, $X = \{y < \bar{y}, y \geq \bar{y}\}$. Assume there is a jump in the benefit level from $b_{<\bar{y}}$ to $b_{\geq\bar{y}}$ at income threshold $y = \bar{y}$. A standard approach in the literature³⁹ is to use a regression-discontinuity design linking the difference in unemployment outcomes just above and below the income threshold, $\lim_{y \rightarrow \bar{y}}^+ U_y - \lim_{y \rightarrow \bar{y}}^- U_y$, to the jump in the UI benefits at the threshold, $b_{\geq\bar{y}} - b_{<\bar{y}}$. To interpret this estimate as the causal effect of benefits on unemployment for workers at the income threshold, $\partial U_{y=\bar{y}} / \partial b_{y=\bar{y}}$, it is necessary that workers around the threshold are similar, except for the UI benefit for which they are eligible. To gauge this identifying assumption, a standard check is whether workers are comparable on observables and whether workers do not change their income in response to the benefit jump.⁴⁰

³⁶Krueger and Meyer, 2002.

³⁷Schmieder, Von Wachter and Bender, 2016.

³⁸Card et al. (2015) and Kolsrud et al. (2018) use the kink in the benefit schedule as a function of pre-unemployment earnings due to a cap on UI benefits. Card, Chetty and Weber (2007) and Schmieder, Von Wachter and Bender (2012) use jumps in the benefit schedule at the cut-offs for age and contribution years.

³⁹Card, Chetty and Weber, 2007; Schmieder, Von Wachter and Bender, 2012.

⁴⁰McCrary, 2008.

The object of the regression-discontinuity design is quite different from what is needed to evaluate the differentiated policy. First, the response in earnings to benefit differences below or above an earnings threshold is not a threat to identification, but one of a response in interest as the corresponding spillover effects ($\partial F_{<\bar{y}}/\partial b_{<\bar{y}}$ and $\partial F_{\geq\bar{y}}/\partial b_{\geq\bar{y}}$) need to be taken into account. If there is a non-zero response, it is the average unemployment rate for the workers who change their income in response to the benefit change that is relevant, as these responding workers increase the UI expenditures by $b_{\geq\bar{y}} - b_{<\bar{y}}$ when they become unemployed.⁴¹ If there is no response, the evaluation of the differentiated policy simplifies to the standard Baily–Chetty expression. Second, to evaluate the differentiated policy, the value of knowing workers behaviour at the threshold (i.e. $\partial U_{y=\bar{y}}/\partial b_{y=\bar{y}}$) is fairly limited. What we need to know instead is how the unemployment response to benefit differs for different income groups ($\partial U_{<\bar{y}}/\partial b_{<\bar{y}}$ versus $\partial U_{\geq\bar{y}}/\partial b_{\geq\bar{y}}$). This requires independent variation in benefits given to low-income workers $db_{<\bar{y}}$ and high-income workers $db_{\geq\bar{y}}$. The jump in the benefit schedule does not provide this variation.⁴²

In general, variation in the differentiated benefit schedule is hard to come by. A notable exception is Kolsrud et al. (2018), who exploit variation in UI benefits, both early in the spell and late in the spell, to evaluate the time profile of UI benefits. The variation is driven by a cap on unemployment benefits, which depends on the spell duration and changes over time, providing independent variation in $db_{<\bar{d}}$ and $db_{\geq\bar{d}}$. This allows the estimation of all relevant cross-elasticities, $\varepsilon_{S_{x'}, b_x}$ for $x, x' \in X = \{d < \bar{d}, d \geq \bar{d}\}$, to evaluate the optimal differentiation between the short-term and long-term unemployed.

The large body of empirical literature on the unemployment effects of UI benefits has provided a wide range of elasticity estimates, ranging from 0.5 to 1.5, by and large. Some of the differences in estimates can be simply reconciled by the potential duration over which the benefits are changes. The largest estimates are found when the potential benefit duration is long and the overall benefit level is increased.⁴³ In the United States, where benefits are exhausted after six months of unemployment, the responses to variation in unemployment benefit levels tend to be larger⁴⁴ than the responses to changes in the potential duration of benefits.⁴⁵ This is also consistent with the findings of Kolsrud et al. (2018), who estimate the incentive costs to be larger for

⁴¹An important concern in the presence of selection in a regression-discontinuity design is selection on treatment (Gerard, Rokkanen and Rothe, 2016), but this is only of second-order importance when evaluating the budgetary effect of small changes in b_x .

⁴²The identification issues here also relate to the distinction between the local average treatment effects, which the variation in UI policies allows us to estimate, and the heterogeneity in the marginal treatment effects, which we would need in order to evaluate the heterogeneous responses.

⁴³Card et al., 2015; Kolsrud et al., 2018.

⁴⁴Meyer, 1990; Landais, 2015.

⁴⁵Rothstein, 2011; Farber and Valletta, 2015.

UI benefits paid to the short-term unemployed compared with those paid to the long-term unemployed. This difference holds despite the forward-looking response by the short-term unemployed to changes in UI benefits when long-term unemployed.

a) Differentiated unemployment policy

Putting the estimates of the insurance value and incentive costs together, we should be able to conclude whether UI benefits are too high or too low, but reaching consensus can be challenging. Compared with the traditional consumption-based estimates of the insurance gains, even the low-range estimates of the incentive costs imply that UI would be too generous overall. However, the more recent estimates on the insurance gains indicate that even for high incentive costs, a further increase in generosity may be desirable.

Linking this back to the dimensions of policy differentiation discussed in Section II, it is also crucial to get a better understanding of the heterogeneity in consumption smoothing gains and unemployment responses. We currently have very little evidence on the corresponding heterogeneity by pre-unemployment earnings, by pre-unemployment history, by age and by available means, which are all common dimensions across which UI generosity differs.⁴⁶ Moreover, we would need to systematically compare the heterogeneity in effects with the eligibility responses to the differentiation in the policy itself.

The evidence in Kolsrud et al. (2018) allows us to evaluate the optimal differentiation by unemployment duration and illustrates that the challenges for drawing conclusions on the optimal differentiation are different from the challenges for pinning down the optimal level. In particular, as the incentive costs are larger for UI benefits paid to the short-term unemployed, while the consumption drops are more pronounced for the long-term unemployed, the robust conclusion is that benefits should be higher for the long-term unemployed, unless the long-term unemployed have very different preferences (or are assigned different social welfare weights). This is still a local recommendation, but it does not rely on whether UI benefits are too high or too low, on average. Hopefully, future research can identify other significant dimensions of heterogeneity and shed further light on the optimal differentiation of UI.

⁴⁶Landais and Spinnewijn (2017) document some observable heterogeneity in the consumption drops and the value of UI as revealed by workers' UI choices. The correlation in the heterogeneity in both measures is often insignificant or even negative, indicating the potential importance of heterogeneity in preferences or the confounding role played by frictions underlying choices.

V. Discussion

A central challenge in the design of UI is to protect workers against unemployment risk while providing incentives to reduce unemployment risk. In practice, countries provide different protection to workers with different characteristics or employment histories. This differentiation of UI could be evaluated according to the same trade-off, where the UI generosity should be higher for workers who value insurance more or for whom the incentive cost is lower. The benefit from differentiating the generosity will be higher when there is more heterogeneity in valuation relative to cost, assuming that workers have limited ability to change their eligibility status. More empirical work is needed to evaluate the current differentiation in UI benefits – in particular, the differentiation by unemployment duration and by pre-unemployment earnings. However, we should also go a step further to uncover key dimensions of heterogeneity in welfare effects that are currently ignored in UI policies. For example, unemployment benefits might be better targeted when conditioning on the unemployment history rather than just on the length of the ongoing unemployment spell.

While this discussion has focused on standard UI, the evaluation could be extended to other unemployment policies. As mentioned before, severance pay – only conditioning on layoff – is relatively more common in developing countries.⁴⁷ It is possible to compare the insurance loss with the incentive gain from such transfers – unconditional on being unemployed. Private UI savings accounts are a related policy, further reducing incentive costs, but still providing access to liquidity upon layoff. In general, the need for protection against the consequences of job loss may extend beyond the initial unemployment spell, as evidenced by the long-lasting effect of job loss on earnings⁴⁸ and consumption.⁴⁹ This may call for more targeted policy instruments such as wage insurance.

In principle, heterogeneity in valuations could also be accommodated by offering workers the choice over how much UI coverage to get. However, realising these gains is not obvious when workers also differ in costs due to adverse selection or when their choices are subject to behavioural frictions.⁵⁰ With the exception of some of the Scandinavian countries, in all countries, the UI coverage is pre-determined, leaving no choice to the worker.

Just as in most of the UI literature, I have restricted my analysis of the gains and costs of UI to the insurance value and the incentive cost, respectively. In practice, UI can be used for fiscal stabilisation⁵¹ and to

⁴⁷Gerard and Naritomi, 2019.

⁴⁸Jacobson, LaLonde and Sullivan, 1993.

⁴⁹Landaïs and Spinnewijn, 2017.

⁵⁰Landaïs et al., 2017; Landaïs and Spinnewijn, 2017.

⁵¹Kekre, 2017.

stimulate consumption.⁵² The general equilibrium effects may be different from the partial unemployment responses, which affects the characterisation of the optimal trade-off.⁵³ In order to evaluate the optimal differentiation of UI, the question becomes whether these general equilibrium effects are different for the unemployment benefits paid to different groups. My discussion has also overlooked the role of behavioural frictions, although they are shown to be important for both workers' job-search decisions and their consumption smoothing. Examples are reference-dependence,⁵⁴ impatience, excess sensitivity⁵⁵ and/or biased beliefs.⁵⁶ When evaluating different types of unemployment policies, it is crucial to understand how they affect workers' exposure to these biases. When comparing how desirable a policy is for different groups of workers, only the relative incidence of these biases will be important. So, even when accounting for these additional effects, a robust evaluation of the differential protection embedded in the differentiated UI policies seems valuable and feasible.

Supporting information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

- Online Appendix
- Supplementary Material

References

- Akerlof, G. A. (1978), 'The economics of "tagging" as applied to the optimal income tax, welfare programs, and manpower planning', *American Economic Review*, vol. 68, no. 1, pp. 8–19.
- Baily, M. N. (1978), 'Some aspects of optimal unemployment insurance', *Journal of Public Economics*, vol. 10, pp. 379–402.
- Browning, M. and Leth-Petersen, S. (2003), 'Imputing consumption from income and wealth information', *Economic Journal*, vol. 113, pp. F282–F301.
- Card, D., Chetty, R. and Weber, A. (2007), 'The spike at benefit exhaustion: leaving the unemployment system or starting a new job?', *American Economic Review*, vol. 97, no. 2, pp. 113–18.
- , Johnston, A., Leung, P., Mas, A. and Pei, Z. (2015), 'The effect of unemployment benefits on the duration of unemployment insurance receipt: new evidence from a regression kink design in Missouri, 2003–2013', *American Economic Review*, vol. 105, no. 5, pp. 126–30.
- Chetty, R. (2006), 'A general formula for the optimal level of social insurance', *Journal of Public Economics*, vol. 90, pp. 1879–901.

⁵²McKay and Reis, 2016.

⁵³Landais, Michaillat and Saez, 2018a,b.

⁵⁴DellaVigna et al., 2017.

⁵⁵Gerard and Naritomi, 2019; Ganong and Noel, 2019.

⁵⁶Spinnewijn, 2015; Mueller, Spinnewijn and Topa, 2018.

- (2008), ‘Moral hazard versus liquidity and optimal unemployment insurance’, *Journal of Political Economy*, vol. 116, pp. 173–234.
- (2009), ‘Sufficient statistics for welfare analysis: a bridge between structural and reduced-form methods’, *Annual Review of Economics*, vol. 1, pp. 451–88.
- and Finkelstein, A. (2013), ‘Social insurance: connecting theory to data’, in A. Auerbach, R. Chetty, M. Feldstein and E. Saez (eds), *Handbook of Public Economics*, vol. 5, pp. 111–93, Amsterdam: North-Holland.
- and Szeidl, A. (2007), ‘Consumption commitments and risk preferences’, *Quarterly Journal of Economics*, vol. 122, pp. 831–77.
- Chodorow-Reich, G. and Karabarbounis, L. (2016), ‘The cyclical nature of the opportunity cost of employment’, *Journal of Political Economy*, vol. 124, pp. 1563–618.
- Citino, L., Russ, K. and Scrutinio, V. (2019), ‘Happy birthday? Manipulation and selection in unemployment insurance’, Italian Social Security Institute (INPS), WorkINPS Paper no. 22.
- DellaVigna, S., Lindner, A., Reizer, B. and Schmieder, J. F. (2017), ‘Reference-dependent job search: evidence from Hungary’, *Quarterly Journal of Economics*, vol. 132, pp. 1969–2018.
- Doornik, B. V., Schoenher, D. and Skrastins, J. (2018), ‘Unemployment insurance, strategic unemployment and firm–worker collusion’, Central Bank of Brazil, Research Department, Working Paper Series no. 483.
- Fadlon, I. and Nielsen, T. H. (2019), ‘Household labor supply and the gains from social insurance’, *Journal of Public Economics*, vol. 171, pp. 18–28.
- Farber, H. S. and Valletta, R. G. (2015), ‘Do extended unemployment benefits lengthen unemployment spells? Evidence from recent cycles in the US labor market’, *Journal of Human Resources*, vol. 50, pp. 873–909.
- Feldstein, M. (1973), ‘The economics of the new unemployment’, *The Public Interest*, vol. 33, no. 3.
- (1974), ‘Unemployment compensation: adverse incentives and distributional anomalies’, *National Tax Journal*, vol. 27, pp. 231–44.
- Ganong, P. and Noel, P. (2019), ‘Consumer spending during unemployment: positive and normative implications’, *American Economic Review*, vol. 109, pp. 2383–424.
- Gerard, F. and Gonzaga, G. (2016), ‘Informal labor and the efficiency cost of social programs: evidence from the Brazilian unemployment insurance program’, National Bureau of Economic Research (NBER), Working Paper no. 22608.
- and Naritomi, J. (2019), ‘Job displacement insurance and consumption: evidence from Brazil’. National Bureau of Economic Research (NBER), Working Paper no. 25749.
- , Rokkanen, M. and Rothe, C. (2016), ‘Bounds on treatment effects in regression discontinuity designs with a manipulated running variable,’ National Bureau of Economic Research (NBER), Working Paper no. 22892.
- Gruber, J. (1997), ‘The consumption smoothing benefits of unemployment insurance’, *American Economic Review*, vol. 87, no. 1, pp. 192–205.
- Hagedorn, M., Manovskii, I. and Mitman, K. (2015), ‘The impact of unemployment benefit extensions on employment: the 2014 employment miracle?’, National Bureau of Economic Research (NBER), Working Paper no. 20884.
- Hendren, N. (2017), ‘Knowledge of future job loss and implications for unemployment insurance’, *American Economic Review*, vol. 107, no. 7, pp. 1778–823.
- Hopenhayn, H. A. and Nicolini, J. P. (1997), ‘Optimal unemployment insurance’, *Journal of Political Economy*, vol. 105, pp. 412–38.
- Jacobson, L. S., LaLonde, R. J. and Sullivan, D. G. (1993), ‘Earnings losses of displaced workers’, *American Economic Review*, vol. 83, no. 4, pp. 685–709.
- Kekre, R. (2017), ‘Unemployment insurance in macroeconomic stabilization’, Harvard University, Working Paper.

- Kleven, H. (2018), 'Sufficient statistics revisited', Princeton University, Working Paper.
- Kolsrud, J., Landais, C., Nilsson, P. and Spinnewijn, J. (2018), 'The optimal timing of unemployment benefits: theory and evidence from Sweden', *American Economic Review*, vol. 108, no. 4–5, pp. 985–1033.
- Kroft, K., Lange, F. and Notowidigdo, M. J. (2013), 'Duration dependence and labor market conditions: evidence from a field experiment', *Quarterly Journal of Economics*, vol. 128, pp. 1123–67.
- Krueger, A. B. and Meyer, B. D. (2002), 'Labor supply effects of social insurance', in A. J. Auerbach and M. Feldstein (eds), *Handbook of Public Economics*, vol. 4, pp. 2327–92, Amsterdam: North-Holland.
- Landais, C. (2015), 'Assessing the welfare effects of unemployment benefits using the regression kink design', *American Economic Journal: Economic Policy*, vol. 7, no. 4, pp. 243–78.
- , Michaillat, P. and Saez, E. (2018a), 'A macroeconomic approach to optimal unemployment insurance: theory', *American Economic Journal: Economic Policy*, vol. 10, no. 2, pp. 152–81.
- , Michaillat, P. and Saez, E. (2018b), 'A macroeconomic approach to optimal unemployment insurance: applications', *American Economic Journal: Economic Policy*, vol. 10, no. 2, pp. 182–216.
- , Nekoei, A., Nilsson, P., Seim, D. and Spinnewijn, J. (2017), 'Risk-based selection in unemployment insurance: evidence and implications', Institute for Fiscal Studies (IFS), Working Paper no. 17/22.
- and Spinnewijn, J. (2017), 'The value of unemployment insurance', London School of Economics (LSE), Working Paper.
- Lindner, A. and Reizer, B. (2018), 'Frontloading the unemployment benefit: an empirical assessment', *American Economic Journal: Applied Economics*, forthcoming (10.1257/app.20180138).
- McCrary, J. (2008), 'Manipulation of the running variable in the regression discontinuity design: a density test', *Journal of Econometrics*, vol. 142, pp. 698–714.
- McKay, A. and Reis, R. (2016), 'The role of automatic stabilizers in the U.S. business cycle', *Econometrica*, vol. 84, pp. 141–94.
- Meyer, B. D. (1990), 'Unemployment insurance and unemployment spells', *Econometrica*, vol. 58, pp. 757–82.
- Michelacci, C. and Ruffo, H. (2015), 'Optimal life cycle unemployment insurance', *American Economic Review*, vol. 105, no. 2, pp. 816–59.
- Mueller, A. I., Spinnewijn, J. and Topa, G. (2018), 'Job seekers' perceptions and employment prospects: heterogeneity, duration dependence and bias', National Bureau of Economic Research (NBER), Working Paper no. 25294.
- Nekoei, A. and Weber, A. (2017), 'Does extending unemployment benefits improve job quality?', *American Economic Review*, vol. 107, no. 2, pp. 527–61.
- Rothstein, J. (2011), 'Unemployment insurance and job search in the great recession', *Brookings Papers on Economic Activity*, vol. 43, pp. 143–213.
- Schmieder, J. F. and Von Wachter, T. (2016), 'The effects of unemployment insurance benefits: new evidence and interpretation', *Annual Review of Economics*, vol. 8, pp. 547–81.
- , Von Wachter, T. and Bender, S. (2012), 'The effects of extended unemployment insurance over the business cycle: evidence from regression discontinuity estimates over 20 years', *Quarterly Journal of Economics*, vol. 127, pp. 701–52.
- , Von Wachter, T. and Bender, S. (2016), 'The effect of unemployment benefits and nonemployment durations on wages', *American Economic Review*, vol. 106, no. 3, pp. 739–77.
- Shavell, S. and Weiss, L. (1979), 'The optimal payment of unemployment insurance benefits over time', *Journal of Political Economy*, vol. 87, pp. 1347–62.

- Shimer, R. and Werning, I. (2008), 'Liquidity and insurance for the unemployed', *American Economic Review*, vol. 98, no. 5, pp. 1922–42.
- Spinnewijn, J. (2015), 'Unemployed but optimistic: optimal insurance design with biased beliefs', *Journal of the European Economic Association*, vol. 13, pp. 130–67.
- Tatsiramos, K. and Van Ours, J. C. (2014), 'Labor market effects of unemployment insurance design', *Journal of Economic Surveys*, vol. 28, pp. 284–311.