



Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Optimal unemployment insurance in an equilibrium business-cycle model

Kurt Mitman^a, Stanislav Rabinovich^{b,*}^a Institute for International Economic Studies, Stockholm University, Sweden^b Department of Economics, Amherst College, United States

ARTICLE INFO

Article history:

Received 11 February 2013

Received in revised form

1 November 2014

Accepted 8 November 2014

Keywords:

Unemployment insurance

Business cycles

Optimal policy

Search and matching

ABSTRACT

The optimal cyclical behavior of unemployment insurance is characterized in an equilibrium search model with risk-averse workers. Contrary to the current US policy, the path of optimal unemployment benefits is pro-cyclical – positively correlated with productivity and employment. Furthermore, optimal unemployment benefits react non-monotonically to a productivity shock: in response to a fall in productivity, they rise on impact but then fall significantly below their pre-recession level during the recovery. As compared to the current US unemployment insurance policy, the optimal state-contingent unemployment benefits smooth cyclical fluctuations in unemployment and deliver substantial welfare gains.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

How should unemployment insurance (UI) respond to fluctuations in labor productivity and unemployment? This question has gained importance in light of the high and persistent unemployment rates following the 2007–2009 recession. In the United States, existing legislation automatically extends unemployment benefit duration in times of high unemployment. Nationwide benefit extensions have been enacted in every major recession since 1958, including the most recent one, in which the maximum duration of unemployment benefits reached an unprecedented 99 weeks. The desirability of such extensions is the subject of an active policy debate, which has only recently begun to receive attention in economic research. In this paper, the optimal cyclical behavior of unemployment insurance is characterized using an equilibrium search model.

The approach integrates risk-averse workers and endogenous worker search effort into the workhorse Diamond–Mortensen–Pissarides model, with business cycles driven by shocks to aggregate labor productivity. The key motivation for using the Diamond–Mortensen–Pissarides model is to explore the consequences of general equilibrium effects for the optimal design of UI policy over the business cycle. The equilibrium search approach is ideal for studying these effects: it accounts for the possibility that more generous unemployment benefits not only discourage unemployed workers from searching, but also raise the worker outside option in wage bargaining, thereby discouraging firms from posting vacancies. Although the framework is a classic one, commonly used to study labor market dynamics and policies, the normative implications of this framework – such as optimal UI – are still very much an open question and need to be more fully understood. Our paper is a step within this research agenda.

* Corresponding author. Tel.: +1 413 542 5537.

E-mail address: srabinovich@amherst.edu (S. Rabinovich).

The optimal state-contingent UI policy is the solution to the Ramsey problem of the government, taking the equilibrium conditions of the model as constraints. Specifically, the government chooses the generosity of unemployment benefits (level and duration) optimally over the business cycle, and can condition its policy choices on the past history of aggregate productivity shocks. The main result is that the optimal benefit schedule is *pro-cyclical* over long time horizons: when the model is simulated under the optimal policy, optimal UI benefits are positively correlated with labor productivity and negatively correlated with the unemployment rate. This overall pro-cyclicality of benefits, however, masks richer dynamics of the optimal policy. In particular, the optimal policy response to a one-time productivity drop is different in the short run and in the long run: optimal benefit levels and duration initially rise in response to a negative shock, but both subsequently fall below their pre-recession level. Thus, the behavior of optimal benefits in response to productivity is non-monotonic, and the fall in benefit generosity lags the fall in productivity. The intuition for these dynamics of the optimal policy is that the initial fall in productivity lowers the gains from creating additional jobs, hence the opportunity cost of raising the generosity of UI benefits is low. On the other hand, the subsequent rise in unemployment raises the social gains from posting vacancies but does not raise the private incentives for doing so. As a consequence, UI generosity optimally rises initially in response to a productivity drop, but then quickly falls in response to the subsequent rise in unemployment.

Compared to current US policy, in the short run (beginning of a recession) the optimal policy coincides with the practice of extending benefits. Where the policies diverge is that US policy typically extends benefits throughout the recession and for years well into the subsequent recovery. This results in a much slower recovery of employment as compared to the optimal policy.¹ The fact that the optimal policy accelerates the recovery of employment indicates that the distortionary effects of increasing benefit generosity in the short run are outweighed by the commitment from the government to lower them in the future.

Our paper contributes to the literature on optimal policy design within search and matching models, which emphasize that policy affects firm vacancy creation decisions. It is thus in the tradition of the general equilibrium approach to optimal unemployment insurance, exemplified by Cahuc and Lehmann (2000), Fredriksson and Holmlund (2001), Coles and Masters (2006), and Lehmann and van der Linden (2007). The novelty of our analysis is to determine how unemployment insurance should optimally respond to business cycle conditions, rather than analyzing optimal policy in steady state.

Our paper also contributes to the emerging literature on optimal unemployment insurance over the business cycle. Three recent papers in this literature are Kroft and Notowidigdo (2010), Landais et al. (2013), and Jung and Kuester (2014). Kroft and Notowidigdo (2010) examine optimal state-contingent UI in a principal-agent framework, extending the approach of Baily (1978), Shavell and Weiss (1979), Hopenhayn and Nicolini (1997), and Shimer and Werning (2008). This approach focuses on the tradeoff between insurance and incentive provision for an individual unemployed worker, but abstracts from the effects of policy on firm hiring decisions. Landais et al. (2013) incorporate firm hiring decisions into their model, but these decisions do not respond to UI policy because wages do not depend on the workers' outside option.² Finally, like our paper, Jung and Kuester (2014) examine optimal policy in a modified Diamond–Mortensen–Pissarides framework. There are differences between the two papers in terms of restrictions on policy instruments used. For example, Jung and Kuester (2014) allow both unemployment benefits and taxes to be business cycle dependent, but do not allow unemployment benefits to expire.³ In addition, Jung and Kuester (2014) depart from the standard framework by introducing shocks to worker bargaining power that are negatively correlated with productivity. As a result, similar to Landais et al. (2013), firm hiring decisions are less responsive to changes in unemployment benefit policy in recessions. Overall, our findings serve to illustrate that the choice of modeling framework, in particular the presence or the absence of general equilibrium effects, can have drastic implications for optimal policy. Our results on the pro-cyclicality of optimal unemployment benefits and their dynamics in response to shocks are new to this literature.

The paper is organized as follows. The model is presented in Section 2. Section 3 describes the calibration strategy. Section 4 defines the optimal policy and contains the main optimal policy results. In Section 5, we discuss our results and conduct sensitivity analysis.⁴ Finally, Section 6 concludes.

2. Model description

The model is a Diamond–Mortensen–Pissarides model with aggregate productivity shocks. Time is discrete and the time horizon is infinite. The economy is populated by a unit measure of workers and a larger continuum of firms.

¹ Recent research by Mitman and Rabinovich (2014) points to benefit extensions as a possible explanation for the “jobless recoveries” that occurred after the last three recessions in the US.

² In Section 5.2.2, we compare our model to the model in Landais et al. (2013) in terms of testable predictions and discuss a way to distinguish between the two empirically.

³ In Section 5.1, we examine the consequences of allowing taxes to be state-contingent, and show that this results in findings consistent with Jung and Kuester (2014).

⁴ Additional sensitivity analysis and all derivations can be found in the supplementary materials available online.

2.1. Economic environment

In any given period, a worker can be either employed (matched with a firm) or unemployed. Workers are risk-averse expected utility maximizers and have expected lifetime utility:

$$U = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [u(x_t) - c(s_t)],$$

where \mathbb{E}_0 is the period-0 expectation operator, $\beta \in (0, 1)$ is the discount factor, x_t denotes consumption in period t , and s_t denotes search effort exerted in period t if unemployed. Only unemployed workers can supply search effort: there is no on-the-job search. The within-period utility of consumption $u: \mathbb{R}_+ \rightarrow \mathbb{R}$ is twice differentiable, strictly increasing, strictly concave, and satisfies $u'(0) = \infty$. The cost of search effort for unemployed workers $c: [0, 1] \rightarrow \mathbb{R}$ is twice differentiable, strictly increasing, strictly convex, and satisfies $c'(0) = 0$, $c'(1) = \infty$. An unemployed worker produces h , which stands for the combined value of leisure and home production. There do not exist private insurance markets and workers cannot save or borrow.⁵

Firms are risk-neutral and maximize profits. Workers and firms have the same discount factor β . A firm can be either matched to a worker or a vacant. A firm posting a vacancy incurs a flow cost k .

Unemployed workers and vacancies match in pairs to produce output. The number of new matches in period t equals $M(S_t(1-L_{t-1}), v_t)$, where $1-L_{t-1}$ is the unemployment level in period $t-1$, S_t is the average search effort exerted by unemployed workers in period t , and v_t is the measure of vacancies posted in period t . The quantity $N_t = S_t(1-L_{t-1})$ represents the measure of efficiency units of worker search.

The matching function M exhibits constant returns to scale, is strictly increasing and strictly concave in both arguments, and has the property that the number of new matches cannot exceed the number of potential matches: $M(N, v) \leq \min\{N, v\} \forall N, v$. Letting $\theta_t = v_t/N_t$ be the market tightness in period t , denote the job-finding probability per efficiency unit of search by $f(\theta) = M(N, v)/N = M(1, \theta)$, and the probability of filling a vacancy by $q(\theta) = M(N, v)/v = M(1/\theta, 1)$. By the assumptions on M made above, the function $f(\theta)$ is increasing in θ and $q(\theta)$ is decreasing in θ . For an individual worker exerting search effort s , the probability of finding a job is $sf(\theta)$. When workers choose the amount of search effort s , they take as given the aggregate job-finding probability $f(\theta)$.

Existing matches are exogenously destroyed with a constant job separation probability δ . Thus, any of the L_{t-1} workers employed in period $t-1$ has a probability δ of becoming unemployed in period t .

All worker-firm matches are identical: the only shocks to labor productivity are aggregate shocks. Specifically, a matched worker-firm pair produces output z_t in period t , where z_t is aggregate labor productivity; $\ln z_t$ follows an AR(1) process:

$$\ln z_t = \rho \ln z_{t-1} + \sigma_\varepsilon \varepsilon_t, \quad (1)$$

where $0 \leq \rho < 1$, $\sigma_\varepsilon > 0$, and ε_t are independent and identically distributed standard normal random variables. Throughout, $z^t = \{z_0, z_1, \dots, z_t\}$ denotes the history of shocks up to period t .

2.2. Government policy

The government can insure against aggregate shocks by buying and selling claims contingent on the aggregate state and is required to balance its budget only in expectation.⁶ Further, it is assumed that the price of a claim to one unit of consumption in state z_{t+1} after a history z^t is equal to the probability of z_{t+1} conditional on z^t ; this would be the case, e.g., in the presence of a large number of out-of state risk-neutral investors with the same discount factor.

Government policies are restricted to take the following form. The government levies a constant lump sum tax τ on firm profits and uses its tax revenues to finance unemployment benefits. The government is allowed to choose both the level of benefits and the rate at which they expire. Benefit expiration is stochastic. This assumption is likewise made in Fredriksson and Holmlund (2001), Albrecht and Vroman (2005) and Faig and Zhang (2012) and will ensure the stationarity of the worker's optimization problem.⁷

A benefit policy at time t thus consists of a pair (b_t, e_t) , where $b_t \geq 0$ is the level of benefits provided to those workers who are eligible for benefits at time t , and $e_t \in [0, 1]$ is the probability that an unemployed worker eligible for benefits becomes ineligible in the following period. The eligibility status of a worker evolves as follows. A worker employed in period t is automatically eligible for benefits in case of job separation. An unemployed worker eligible for benefits in period t becomes ineligible in the following period with probability e_t , and an ineligible worker does not regain eligibility until he finds a job. All eligible workers receive the same benefits b_t ; ineligible workers receive no unemployment benefits.

The benefit policy can depend on the entire history of past aggregate shocks; thus the policy $b_t = b_t(z^t)$, $e_t = e_t(z^t)$ must be measurable with respect to z^t .⁸ Benefits are constrained to be non-negative: the government cannot tax h .

⁵ In Section 5.4 we discuss the possible consequences of relaxing this assumption.

⁶ The US UI system is financed by payroll taxes on firms and is administered at the state level. However, under the provisions of the Social Security Act, each state can borrow from a federal unemployment insurance trust fund, provided it meets certain federal requirements. This motivates the assumption that the government is not required to balance its budget within each state.

⁷ Our main results are robust to the possibility of benefit expiration. See Section 5.6 for a further discussion of benefit expiration.

⁸ Note, however, that b_t is not allowed to depend on an individual worker's history.

2.3. Timing

The government commits to a policy $(\tau, b_t(\cdot), e_t(\cdot))$ once and for all before the period-0 shock realizes. Within each period t , the timing is as follows. The economy enters period t with a level of employment L_{t-1} . Of the $1 - L_{t-1}$ unemployed workers, a measure $D_{t-1} \leq 1 - L_{t-1}$ is eligible for benefits, i.e. will receive benefits in period t if they do not find a job.

The aggregate shock z_t then realizes. Firms observe the aggregate shock and decide how many vacancies to post, at cost k per vacancy. At the same time, workers choose their search effort s_t at the cost of $c(s_t)$. Letting S_t^E and S_t^I be the search effort exerted by an eligible unemployed worker and an ineligible unemployed worker, respectively, the aggregate search effort is then equal to $S_t^E D_{t-1} + S_t^I (1 - L_{t-1} - D_{t-1})$, and the market tightness is therefore equal to

$$\theta_t = \frac{v_t}{S_t^E D_{t-1} + S_t^I (1 - L_{t-1} - D_{t-1})}. \quad (2)$$

A fraction δ of the existing L_{t-1} matches are exogenously destroyed. The number of workers who find jobs at the same time is $f(\theta_t) (S_t^E D_{t-1} + S_t^I (1 - L_{t-1} - D_{t-1}))$.

All the workers who are now employed produce z_t and receive a bargained wage w_t (below we describe wage determination in detail). Workers who (i) were employed and lost a job, or (ii) were eligible unemployed workers and did not find a job, consume h plus unemployment benefits, $h + b_t$ and lose their eligibility for the next period with probability e_t . Ineligible unemployed workers who have not found a job consume h , and remain ineligible for the following period.

This determines the law of motion for employment

$$L_t(z^t) = (1 - \delta)L_{t-1}(z^{t-1}) + f(\theta_t(z^t)) [S_t^E(z^t)D_{t-1}(z^{t-1}) + S_t^I(z^t)(1 - L_{t-1}(z^{t-1}) - D_{t-1}(z^{t-1}))], \quad (3)$$

and the law of motion for the measure of eligible unemployed workers

$$D_t(z^t) = (1 - e_t(z^t)) [\delta L_{t-1}(z^{t-1}) + (1 - S_t^E(z^t)f(\theta_t(z^t)))D_{t-1}(z^{t-1})]. \quad (4)$$

Thus, the measure of workers receiving benefits in period t is

$$\delta L_{t-1} + (1 - S_t^E f(\theta_t)) D_{t-1} = \frac{D_t}{1 - e_t}.$$

Since the government has access to financial markets in which a full set of state-contingent claims is traded, its budget constraint is a present-value budget constraint:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ L_t(z^t) \tau - \left(\frac{D_t(z^t)}{1 - e_t(z^t)} \right) b_t(z^t) \right\} \geq 0. \quad (5)$$

2.4. Worker value functions

A worker entering period t employed retains his job with probability $1 - \delta$ and loses it with probability δ . If he retains his job, he consumes his wage $w_t(z^t)$ and proceeds as employed to period $t + 1$. If he loses his job, he consumes $h + b_t(z^t)$ and proceeds as unemployed to period $t + 1$. With probability $1 - e_t(z^t)$ he then retains his eligibility for benefits in period $t + 1$, and with probability $e_t(z^t)$ he loses his eligibility. Denote by $W_t(z^t)$ the value after a history z^t for a worker who enters period t employed.

A worker entering period t unemployed and eligible for benefits chooses search effort s_t^E and suffers the disutility $c(s_t^E)$. He finds a job with probability $s_t^E f(\theta_t(z^t))$ and remains unemployed with the complementary probability. If he finds a job, he earns the wage $w_t(z^t)$ and proceeds as employed to period $t + 1$. If he remains unemployed, he consumes $h + b_t(z^t)$, and proceeds as unemployed to the next period. With probability $1 - e_t(z^t)$ he retains his eligibility for benefits in period $t + 1$, and with probability $e_t(z^t)$ he loses his eligibility. Denote by $U_t^E(z^t)$ the value after a history z^t for a worker who enters period t as eligible unemployed.

Finally, a worker entering period t unemployed and ineligible for benefits chooses search effort s_t^I and suffers the disutility $c(s_t^I)$. He finds a job with probability $s_t^I f(\theta_t(z^t))$ and remains unemployed with the complementary probability. If he finds a job, he earns the wage $w_t(z^t)$ and proceeds as employed to period $t + 1$. If he remains unemployed, he consumes h and proceeds as ineligible unemployed to the next period. Denote by $U_t^I(z^t)$ the value after a history z^t for a worker who enters period t as ineligible unemployed.

The Bellman equations for the three types of workers are then

$$W_t(z^t) = (1 - \delta) [u(w_t(z^t)) + \beta \mathbb{E} W_{t+1}(z^{t+1})] + \delta [u(h + b_t(z^t)) + \beta (1 - e_t) \mathbb{E} U_{t+1}^E(z^{t+1}) + \beta e_t \mathbb{E} U_{t+1}^I(z^{t+1})], \quad (6)$$

$$U_t^E(z^t) = \max_{s_t^E} -c(s_t^E) + s_t^E f(\theta_t(z^t)) [u(w_t(z^t)) + \beta \mathbb{E} W_{t+1}(z^{t+1})] + (1 - s_t^E f(\theta_t(z^t))) [u(h + b_t(z^t)) + \beta(1 - e_t(z^t)) \mathbb{E} U_{t+1}^E(z^{t+1}) + \beta e_t \mathbb{E} U_{t+1}^I(z^{t+1})], \quad (7)$$

$$U_t^I(z^t) = \max_{s_t^I} -c(s_t^I) + s_t^I f(\theta_t(z^t)) [u(w_t(z^t)) + \beta \mathbb{E} W_{t+1}(z^{t+1})] + (1 - s_t^I f(\theta_t(z^t))) [u(h) + \beta \mathbb{E} U_{t+1}^I(z^{t+1})]. \quad (8)$$

2.5. Firm value functions

A matched firm retains its worker with probability $1 - \delta$. In this case, the firm receives the output net of wages and taxes, $z_t - w_t(z^t) - \tau$, and then proceeds into the next period as a matched firm. If the firm loses its worker, it gains nothing in the current period and proceeds into the next period unmatched. A firm that posts a vacancy incurs a flow cost k and finds a worker with probability $q(\theta_t(z^t))$. If the firm finds a worker, it gets flow profits $z_t - w_t(z^t) - \tau$ and proceeds into the next period as a matched firm. Otherwise, it proceeds unmatched into the next period.

Denote by $J_t(z^t)$ the value of a firm that enters period t matched to a worker, and denote by $V_t(z^t)$ the value of an unmatched firm posting a vacancy. These value functions satisfy the following Bellman equations:

$$J_t(z^t) = (1 - \delta) [z_t - w_t(z^t) - \tau + \beta \mathbb{E} J_{t+1}(z^{t+1})] + \delta \beta \mathbb{E} V_{t+1}(z^{t+1}), \quad (9)$$

$$V_t(z^t) = -k + q(\theta_t(z^t)) [z_t - w_t(z^t) - \tau + \beta \mathbb{E} J_{t+1}(z^{t+1})] + (1 - q(\theta_t(z^t))) \beta \mathbb{E} V_{t+1}(z^{t+1}). \quad (10)$$

2.6. Wage bargaining

Wages are determined according to Nash bargaining⁹: the wage is chosen to maximize a weighted product of the worker's surplus and the firm's surplus. Further, the worker's outside option is being unemployed and eligible for benefits, since he becomes eligible upon locating an employer and retains eligibility if negotiations with the employer break down. The worker–firm pair therefore chooses the wage $w_t(z^t)$ to maximize

$$\Delta_t(z^t)^\xi \Gamma_t(z^t)^{1-\xi}, \quad (11)$$

where $\Delta_t(z^t)$ is the worker's surplus from being employed:

$$\Delta_t(z^t) = [u(w_t(z^t)) + \beta \mathbb{E} U_{t+1}^E(z^{t+1})] - [u(h + b_t(z^t)) + \beta(1 - e_t) \mathbb{E} U_{t+1}^E(z^{t+1}) + \beta e_t \mathbb{E} U_{t+1}^I(z^{t+1})], \quad (12)$$

and $\Gamma_t(z^t)$ is the firm's surplus from employing a worker:

$$\Gamma_t(z^t) = z_t - w_t(z^t) - \tau + \beta \mathbb{E} J_{t+1}(z^{t+1}) - \beta \mathbb{E} V_{t+1}(z^{t+1}), \quad (13)$$

and the parameter $\xi \in (0, 1)$ is the worker's bargaining weight.

2.7. Equilibrium

Given a policy $(\tau, b_t(\cdot), e_t(\cdot))$ and an initial condition (z_{-1}, L_{-1}) an equilibrium is a sequence of z^t -measurable functions for wages $w_t(z^t)$, search effort $S_t^E(z^t)$, $S_t^I(z^t)$, market tightness $\theta_t(z^t)$, employment $L_t(z^t)$, measures of eligible workers $D_t(z^t)$, and value functions:

$$\{W_t(z^t), U_t^E(z^t), U_t^I(z^t), J_t(z^t), V_t(z^t)\}$$

such that the value functions satisfy the worker and firm Bellman equations (6)–(10); the value $V_t(z^t)$ of a vacant firm is zero for all z^t ; the search effort S_t^E solves the maximization problem in (7) for s_t^E , and the search effort S_t^I solves the maximization problem in (8) for s_t^I ; the wage maximizes (11), with $\Delta(z^t)$ and $\Gamma(z^t)$ defined in (12) and (13); employment and the measure of eligible unemployed workers satisfy (3) and (4); and tax revenue and benefits satisfy (5).

The Online Appendix shows that the optimal choices of firms and workers can be characterized by four optimality conditions. The first is the free-entry condition for firms (dependence on z^t is understood throughout):

$$\frac{k}{q(\theta_t)} = z_t - w_t - \tau + \beta(1 - \delta) \mathbb{E} \frac{k}{q(\theta_{t+1})}, \quad (14)$$

Eq. (14) equates the marginal cost of creating a vacancy, weighted by the probability of filling that vacancy, to the profits from employing a worker. The next two conditions are the optimal search conditions for eligible and ineligible unemployed

⁹ This is standard in the literature. Alternative wage setting mechanisms are discussed in Section 5.2.

workers:

$$\begin{aligned} \frac{c'(S_t^E)}{f(\theta_t)} = & u(w_t) - u(h + b_t) + (1 - e_t)\beta\mathbb{E}_t \left(c(S_{t+1}^E) + (1 - \delta - S_{t+1}^E f(\theta_{t+1})) \frac{c'(S_{t+1}^E)}{f(\theta_{t+1})} \right) \\ & + e_t\beta\mathbb{E}_t \left(c(S_{t+1}^I) + (1 - S_{t+1}^I f(\theta_{t+1})) \frac{c'(S_{t+1}^I)}{f(\theta_{t+1})} - \delta \frac{c'(S_{t+1}^E)}{f(\theta_{t+1})} \right), \end{aligned} \quad (15)$$

$$\frac{c'(S_t^I)}{f(\theta_t)} = u(w_t) - u(h) + \beta\mathbb{E}_t \left(c(S_{t+1}^I) + (1 - S_{t+1}^I f(\theta_{t+1})) \frac{c'(S_{t+1}^I)}{f(\theta_{t+1})} - \delta \frac{c'(S_{t+1}^E)}{f(\theta_{t+1})} \right). \quad (16)$$

Eqs. (15) and (16) state that the marginal cost of increasing the job finding probability for the eligible and ineligible workers, respectively, equals the marginal benefit. The marginal cost (left-hand side of each equation) of increasing the job finding probability is the marginal disutility of search for that worker weighted by the aggregate job finding rate. The marginal benefit (right-hand side of each equation) equals the current consumption gain from becoming employed plus the benefit of economizing on search costs in the future.

Finally, Nash bargaining will imply a relationship between the worker's surplus from being employed and the firm's surplus from hiring the worker, which can be written compactly as

$$\xi u'(w_t)k\theta_t = (1 - \xi)c'(S_t^E). \quad (17)$$

It will be apparent in Section 4 that the conditions (14)–(17) will play the role of incentive constraints in the optimal policy problem, analogous to incentive constraints in principal-agent models of unemployment insurance, e.g., [Hopenhayn and Nicolini \(1997\)](#). In particular, these conditions illustrate the social planner's tradeoff between unemployment insurance and the job-finding rate. Consider the effect of an increase in unemployment benefits. All else equal, this reduces an eligible worker's surplus from becoming employed. The optimal search condition (15) then implies that the search effort of eligible workers falls. Furthermore, the Nash bargaining (17) condition implies that, since worker surplus falls, firm surplus from hiring must fall as well; the free-entry condition (14) then dictates that θ must fall. For both of these reasons – reduced aggregate job-finding rate and reduced search – the job-finding rate will fall.

Furthermore, conditions (14)–(17) provide intuition for how this insurance-employment tradeoff varies over the business cycle. Consider the effect of a decrease in z . All else equal, this will reduce the job-finding rate by lowering firm incentives to post vacancies, through Eq. (14). In addition, the Nash bargaining condition implies that wages fall in response to a fall in productivity. Since both wages and the aggregate job-finding rate fall, Eqs. (15) and (16) imply reduced incentives to search, which reduce the job-finding rate even further. This effect of a fall in z could potentially be mitigated by lowering unemployment benefits, which would, as argued above, push the job-finding rate in the opposite direction. A fall in productivity thus forces the social planner to choose between a lower job-finding rate and lower unemployment insurance. This is precisely the channel whose consequences are explored in Section 4 when the optimal policy is analyzed.

3. Calibration

The model is calibrated to match salient features of the US labor market. The model period is taken to be 1 week. Mean weekly productivity is normalized to one. The benefit level is set to $b=0.4$ to match the average replacement rate of unemployment insurance. A benefit duration scheme that mimics the benefit extension provisions currently in place within the US policy is used. The standard benefit duration is 26 weeks; local and federal employment conditions trigger automatic 20-week and 33-week extensions. In the model $e_t = 1/59$ when productivity is more than 3% below the mean, $e_t = 1/46$ when productivity is between 1.5% and 3% below the mean, and $e_t = 1/26$ otherwise. The tax rate $\tau = 0.023$ is set so that the government balances its budget if the unemployment rate is 5.5%.

The utility function is $u(x) = \log(x)$. The functional form of the cost of search is assumed to be $c(s) = (A/(1+\psi))[(1-s)^{-(1+\psi)} - 1] - As$. This functional form is chosen to ensure that the optimal search effort will always be strictly between 0 and 1. In particular, the functional form above guarantees that, for any $A > 0$, $c' > 0$, $c'' > 0$, as well as $c(0) = c'(0) = 0$, $c(1) = c'(1) = \infty$. The matching function is assumed to be $M(N, v) = Nv/[N^\chi + v^\chi]^{1/\chi}$, similar to [den Haan et al. \(2000\)](#). The choice of the matching technology is likewise driven by the requirement that the job-finding rate and the job-filling rate always be strictly less than 1.¹⁰

Following [Shimer \(2005\)](#), labor productivity z_t is taken to mean real output per person in the non-farm business sector. This measure of productivity is taken from the quarterly data constructed by the BLS for the time period 1951–2004. Unemployed is measured as the 1951–2004 seasonally adjusted unemployment series constructed by the BLS, and vacancies are measured using the seasonally adjusted help-wanted index constructed by the Conference Board.

¹⁰ The frequently used alternative is the Cobb–Douglas specification. However, commonly used local solution methods for the model do not guarantee that the job-finding rate is less than one under this specification.

The discount factor $\beta = 0.99^{1/12}$ implying a yearly discount rate of 4%. The parameters for the productivity shock process are estimated, at the weekly level, to be $\rho = 0.9895$ and $\sigma_\varepsilon = 0.0034$. The job separation parameter δ is set to 0.0081 to match the average weekly job separation rate.¹¹ The cost of posting a vacancy is set to $k=0.58$ following Hagedorn and Manovskii (2008), who estimate the combined capital and labor costs of vacancy creation to be 58% of weekly labor productivity.

This leaves five parameters to be calibrated: (1) the value h of non-market activity; (2) the worker bargaining weight ξ ; (3) the matching function parameter χ ; (4) the level coefficient of the search cost function A ; and (5) the curvature parameter of the search cost function ψ . These five parameters are calibrated jointly to simultaneously match five data targets: (1) the standard deviation of the vacancy–unemployment ratio; (2) the standard deviation of the unemployment rate; (3) the average weekly job-finding rate; (4) the elasticity of wages with respect to productivity; and (5) the elasticity of unemployment duration with respect to unemployment benefits. The first four of these targets are directly measured in the data. For the elasticity of unemployment duration with respect to benefits, $\varepsilon_{d,b}$, the micro-estimate of 0.9 reported by Meyer (1990) is targeted.¹² Note that the model counterpart of the measured elasticity is taken to be the *micro*- (partial-equilibrium) elasticity: the percentage change of unemployment duration due to decreased search effort alone, in response to a 1% increase in the benefit level, but keeping fixed the value of $f(\theta)$.¹³ Intuitively, the elasticity of wages with respect to productivity identifies the worker bargaining weight ξ ; the volatility of the market tightness and the unemployment rate jointly identify the value of non-market activity h and matching function parameter χ ; and the remaining two targets identify the parameters A and ψ , since these parameters govern the distortions in search behavior induced by unemployment benefits.

Table 1 reports the calibrated parameters and the matching of the calibration targets. Table 2 compares relevant labor market statistics in both the US economy and those from the calibrated model. The model does a good job of capturing the business cycle properties of the US labor market.¹⁴

4. Optimal policy

The government is assumed to be utilitarian: it chooses a policy to maximize the period-0 expected value of worker utility, taking the equilibrium conditions as constraints. A policy $\tau, b_t(z^t), e_t(z^t)$ is feasible if there exists a sequence of z^t -measurable functions $\{w_t(z^t), S_t^E(z^t), S_t^I(z^t), \theta_t(z^t), L_t(z^t), D_t(z^t)\}$ such that (3), (4), (14)–(17) hold for all z^t , and the government budget constraint (5) is satisfied. The government maximizes

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \begin{aligned} & L_t(z^t) u(w_t(z^t)) + \left(\frac{D_t(z^t)}{1 - e_t(z^t)} \right) u(h + b_t(z^t)) \\ & + \left(1 - L_t(z^t) - \frac{D_t(z^t)}{1 - e_t(z^t)} \right) u(h) - D_{t-1}(z^{t-1}) c(S_t^E(z^t)) \\ & - (1 - L_{t-1}(z^{t-1}) - D_{t-1}(z^{t-1})) c(S_t^I(z^t)) \end{aligned} \right\} \quad (18)$$

over the set of all feasible policies.

The government's problem can be written as one of choosing a policy $\tau, b_t(z^t), e_t(z^t)$ together with functions $\{w_t(z^t), S_t^E(z^t), S_t^I(z^t), \theta_t(z^t), L_t(z^t), D_t(z^t)\}$ to maximize (18) subject to (3), (4), (14)–(17) holding for all z^t , and subject to the government budget constraint (5). The optimal policy is found by solving the system of necessary first-order conditions for this problem. The period- t solution will naturally be state-dependent: in particular, it will depend on the current productivity z_t , as well as the current unemployment level $1 - L_{t-1}$, and current measure of benefit-eligible workers D_{t-1} with which the economy has entered period t . However, in general the triple $(z_t, 1 - L_{t-1}, D_{t-1})$ is not a sufficient state variable for pinning down the optimal policy, which may depend on the entire past history of aggregate shocks. In the Online Appendix, it is shown that the optimal period t solution is a function of $(z_t, 1 - L_{t-1}, D_{t-1})$ as well as $(e_{t-1}, \mu_{t-1}, \nu_{t-1}, \gamma_{t-1})$, where e_{t-1} is the previous period's benefit expiration rate and $\mu_{t-1}, \nu_{t-1}, \gamma_{t-1}$ are the Lagrange multipliers on the constraints (15), (16), (14), respectively, in the maximization problem (18). The tuple $(z_t, 1 - L_{t-1}, D_{t-1}, e_{t-1}, \mu_{t-1}, \nu_{t-1}, \gamma_{t-1})$ captures the dependence of the optimal b_t, e_t on the history z^t . The fact that the $z_t, 1 - L_{t-1}$ and D_{t-1} are not sufficient reflects the fact that the optimal policy is time-inconsistent: for example, the optimal benefits after two different histories of shocks may differ even though the two histories result in the same current productivity and the same current unemployment level. Intuitively, the government might want to induce firms to post vacancies – and workers to search – by promising low unemployment benefits, but has an

¹¹ The same procedure of adjusting for time aggregation as Hagedorn and Manovskii (2008) is used to obtain the weekly estimates for the job finding rate and the job separation rate from monthly data.

¹² There exist a range of estimates (e.g., Krueger and Meyer, 2002) in the literature for the elasticity of unemployment duration with respect to benefit level. However, we find that qualitatively our results are robust to calibrating to higher or lower values of the elasticity.

¹³ This is distinct from the macro elasticity, which would comprise the total effect of a 1% increase in UI benefits, and thus include the general equilibrium effect on θ . See Section 5.2.2 for a discussion.

¹⁴ Our calibration procedure implies a large value of h , which might seem surprising, considering that empirical studies (e.g., Gruber, 1997; Browning and Crossley, 2001; Aguiar and Hurst, 2005) report a consumption drop for workers upon becoming unemployed. However, Aguiar and Hurst (2005), who properly distinguish between consumption and expenditure, show that the consumption drop for the unemployed is only 5%, illustrating that most earlier estimates of the consumption drop are biased upward. Most importantly, we show in Section 5.8 that our optimal policy results are robust to the calibrated value of h : they hold even if we assume a substantially lower value for h .

Table 1
Internally calibrated parameters.

Parameter	Interpretation	Value	Target	Data	Model
h	Value of non-market activity	0.584	St. dev of $\log(v/(1-L))$	0.259	0.259
ξ	Bargaining power	0.078	$\varepsilon_{w,z}$	0.449	0.449
χ	Matching parameter	0.442	St. dev of $\log(1-L)$	0.125	0.125
A	Disutility of search	0.0053	Mean job finding rate	0.139	0.139
ψ	Search cost curvature	1.911	$\varepsilon_{d,b}$	0.9	0.9

Note: $\varepsilon_{d,b}$ is the elasticity of unemployment duration with respect to benefits. $\varepsilon_{w,z}$ is the elasticity of wages with respect to productivity.

Table 2
Summary statistics

Statistic		z	$1-L$	v	$v/(1-L)$
<i>Quarterly US Data, 1951:I–2004:IV</i>					
Standard deviation		0.013	0.125	0.139	0.259
Correlation matrix	z	1	–0.302	0.460	0.393
	$1-L$	–	1	–0.919	–0.977
	v	–	–	1	0.982
	$v/(1-L)$	–	–	–	1
<i>Calibrated model</i>					
Standard deviation		0.013	0.125	0.146	0.259
Correlation matrix	z	1	–0.851	0.869	0.912
	$1-L$	–	1	–0.768	–0.914
	v	–	–	1	0.950
	$v/(1-L)$	–	–	–	1

Note: Standard deviations and correlations are reported in logs as quarterly deviations from an HP-filtered trend with a smoothing parameter of 1600.

ex post incentive to provide higher benefits, so as to smooth worker consumption, after employment outcomes have realized. Including the variables $e_{t-1}, \mu_{t-1}, \nu_{t-1}, \gamma_{t-1}$ as state variables in the optimal policy captures exactly this trade-off. Note that it is assumed throughout that the government can fully commit to its policy.¹⁵

4.1. Optimal policy results

To investigate how the economy behaves over time under the optimal policy, the model is simulated both under the current benefit policy and under the optimal policy. Table 3 reports the summary statistics, under the optimal policy, for the behavior of unemployment benefit levels b and potential benefit duration $1/e$. Interestingly, the average values of unemployment benefit levels and duration under the optimal policy are very close to the current US policy.

The key observation is that, over a long period of time, the correlation of optimal benefits with productivity is positive: both benefit levels and potential benefit duration are pro-cyclical in the long run and, in particular, negatively correlated with the unemployment rate. Moreover, this result is not driven by any balanced budget requirement, since the government is allowed to run deficits in recessions.

In order to understand the mechanism behind this behavior of the optimal policy, Fig. 1(a) plots the optimal benefit policy function $b_t(z_t, 1-L_{t-1}, D_{t-1}, e_{t-1}, \mu_{t-1}, \nu_{t-1}, \gamma_{t-1})$ as a function of current z and last period's $1-L$ only, keeping $D_{t-1}, e_{t-1}, \mu_{t-1}, \nu_{t-1}$ and γ_{t-1} fixed at their average values. The optimal benefit level is decreasing in current productivity z and decreasing in unemployment $1-L$. The intuition for this result is that the optimal benefit is lower in states of the world when the marginal social benefit of job creation is higher, because lower benefits are used to encourage search effort by workers and vacancy creation by firms. The marginal social benefit of job creation is higher when z is higher, since the output of an additional worker–firm pair is then higher. The marginal social benefit is also higher when current employment is lower, because the expected output gain of increasing θ is proportional to the number of unemployed workers. Note, however, that although the social gains from creating jobs are high when unemployment is high, the private gains to firms of posting vacancies do not directly depend on unemployment. As a consequence, optimal benefits are lower, all else equal, when current unemployment is high. Fig. 1(b) illustrates the same result for the optimal duration of benefits: optimal benefit duration is lowest at times of high productivity and high unemployment. This shape of the policy function also implies that during a recession, there are two opposing forces at

¹⁵ The Online Appendix describes the method used to solve for the optimal policy.

Table 3
Optimal benefit behavior.

Statistic	Benefit level b	Potential duration $1/e$
Mean	0.417	26.74
Standard deviation	0.015	0.052
Correlation with z	0.756	0.679
Correlation with $1 - L$	-0.451	-0.361
Correlation with b	1	0.989

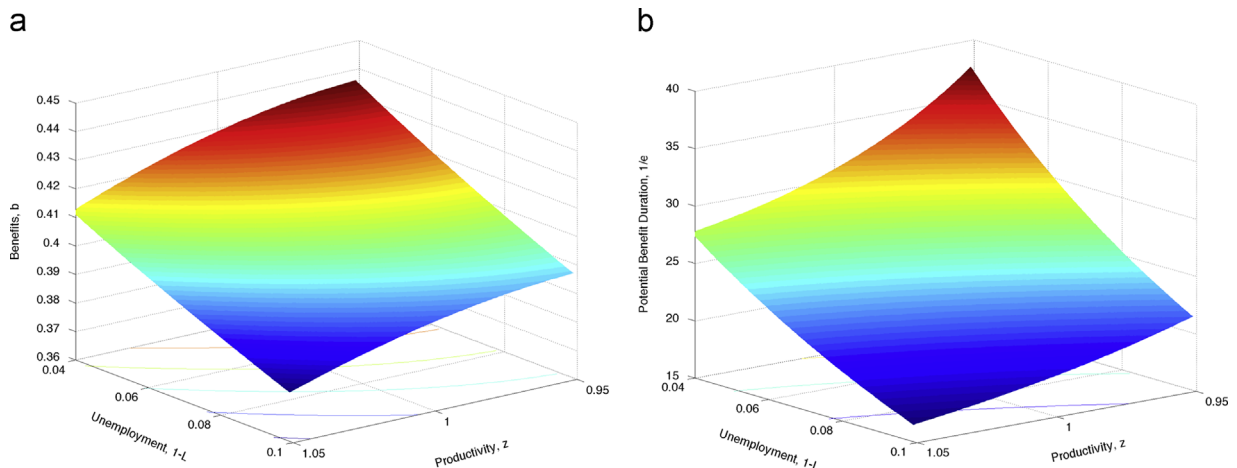


Fig. 1. Optimal policy. (a) Benefit level. (b) Benefit duration.

work – low productivity and high unemployment – which give opposite prescriptions for the response of optimal benefits. This gives an ambiguous prediction for the overall cyclical behavior of benefit levels and benefit duration.

In Fig. 2 the dynamic response of the economy to a negative productivity shock is presented under the optimal policy and the current policy. The impulse response is to a productivity drop of 1.5% below its mean. Note that under the current policy, benefit duration does not change in response to the shock, since automatic extensions are only activated when productivity is more than 1.5% below the mean. The optimal benefit level initially jumps up, but then falls for about two quarters following the shock, and slowly reverts to its pre-shock level. The same is true of optimal benefit duration. Unemployment rises in response to the drop in productivity and continues rising for about one-quarter before it starts to return to its pre-shock level. Note that the rise in unemployment is significantly lower than under the current benefit policy.

The intuition for this optimal policy response is that the government would like to provide immediate insurance against the negative shock and, expecting future productivity to rise, would like to induce a recovery in vacancy creation and search effort. This is consistent with the intuition explained in Section 2.7. Eqs. (14)–(17) imply that a fall in productivity leaves the social planner with a choice between lower unemployment insurance and lower employment. Early in the recession, however, productivity is low, and so is the social value of employment. As a result, the social planner tolerates the rise in unemployment. As unemployment increases, the social benefit of creating more vacancies increases relative to the benefit of providing insurance, and the planner therefore cuts unemployment benefits to reduce unemployment. Thus, benefit generosity responds positively to the initial drop in productivity but negatively to the subsequent rise in unemployment, precisely as implied in Fig. 1.

Turning to other key labor market variables, as compared to the current benefit policy, the optimal policy results in a faster recovery of the vacancy–unemployment ratio, the search intensity of unemployed workers eligible for benefits, and the job finding rate. Wages also fall less, in percent deviation terms, under the optimal policy than they do under the current policy. This is due to the fact that the initial rise in benefits smooths the fall in wages through an increase in the worker outside option. The fact that wages fall less in percentage terms indicates that firm profits fall more. Despite the fall in contemporaneous profits, there is not a large fall in market tightness. The reason for this is that firms expect future benefits to fall. The figure thus illustrates that the labor market response depends not only on the contemporaneous benefit policy but also on agents' expectations about future policy dynamics.

The fact that the optimal policy calls for an increase in benefit generosity in response to a negative shock, but a decrease as an economy recovers, begs the question of how benefits would respond to a more “typical” recession, where productivity falls for a protracted period of time and then recovers. The optimal policy is simulated under two alternative shock

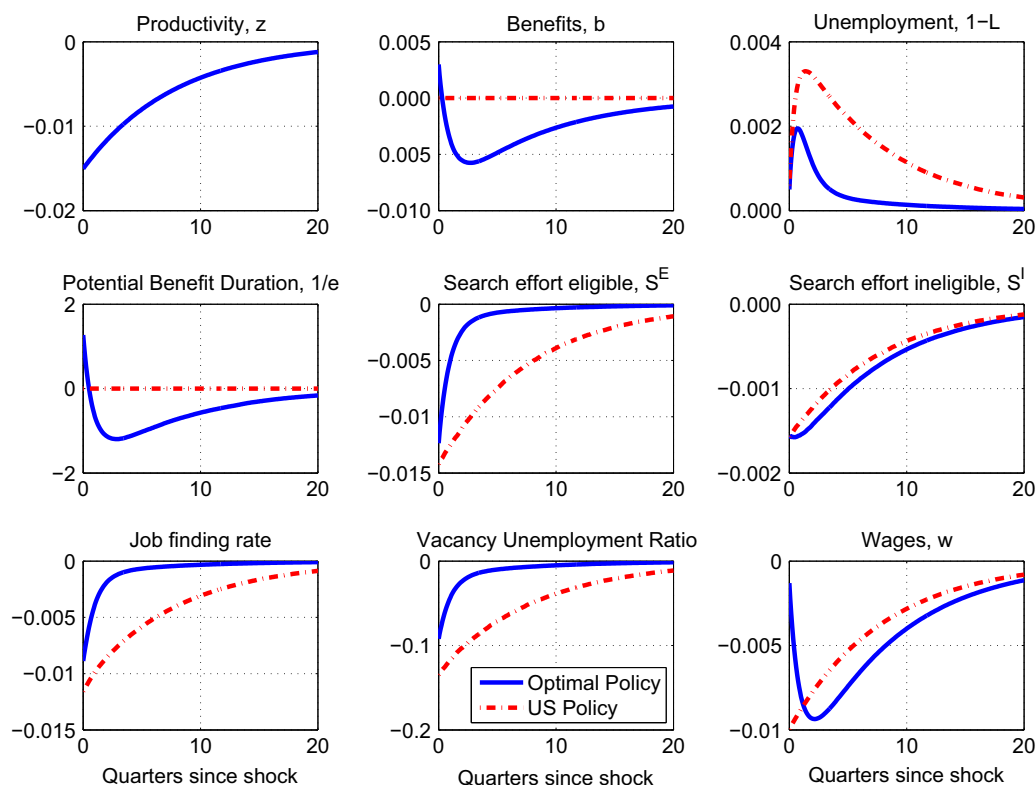


Fig. 2. Responses to 1.5% drop in productivity.

realizations in Fig. 3. The alternate shock paths have the economy contracting for two and eight quarters. In both cases benefit generosity increases for approximately the first quarter, but then falls below the pre-recession level. Increasing benefits initially is a robust feature, but it is important to note that the reversal to decrease benefits occurs before the economy begins to recover; this is driven by the rise in unemployment, whose effect eventually dominates the effect of the productivity drop.

Table 4 reports the moments of key labor market variables when the model is simulated under the current policy and the optimal policy. As compared to the optimal policy, the optimal policy results in lower average unemployment and lower unemployment volatility. These results show that the optimal benefit policy stabilizes cyclical fluctuations in unemployment.

Finally, the expected welfare gain from switching from the current policy to the optimal policy is computed. Implementing the optimal policy results in a significant welfare gain: 0.52% as measured in consumption equivalent variation terms.

5. Discussion of the results

In this section, the key assumptions made in the paper are discussed, as are the consequences of modifying or relaxing these assumptions.

5.1. Alternative policy instruments

Attention was deliberately focused on a particular restricted set of policy instruments: a constant lump-sum tax and state-dependent unemployment benefits. Allowing taxes to be history-dependent as well could potentially change the optimal behavior of unemployment benefits. This section investigates the consequences of allowing such history-dependent taxes. This will serve both to clarify the role of pro-cyclical unemployment benefits in our benchmark model, and to facilitate comparison of our paper to Jung and Kuester (2014).

First, the first-best allocation is characterized. Then it is shown that, in the absence of worker-side moral hazard, the first-best allocation can be implemented by using state-dependent taxes. When worker-side moral hazard is present, the first-best cannot be implemented, but allowing history-dependent taxes improves welfare, and changes the behavior of optimal unemployment insurance relative to the benchmark. The results are numerically characterized.

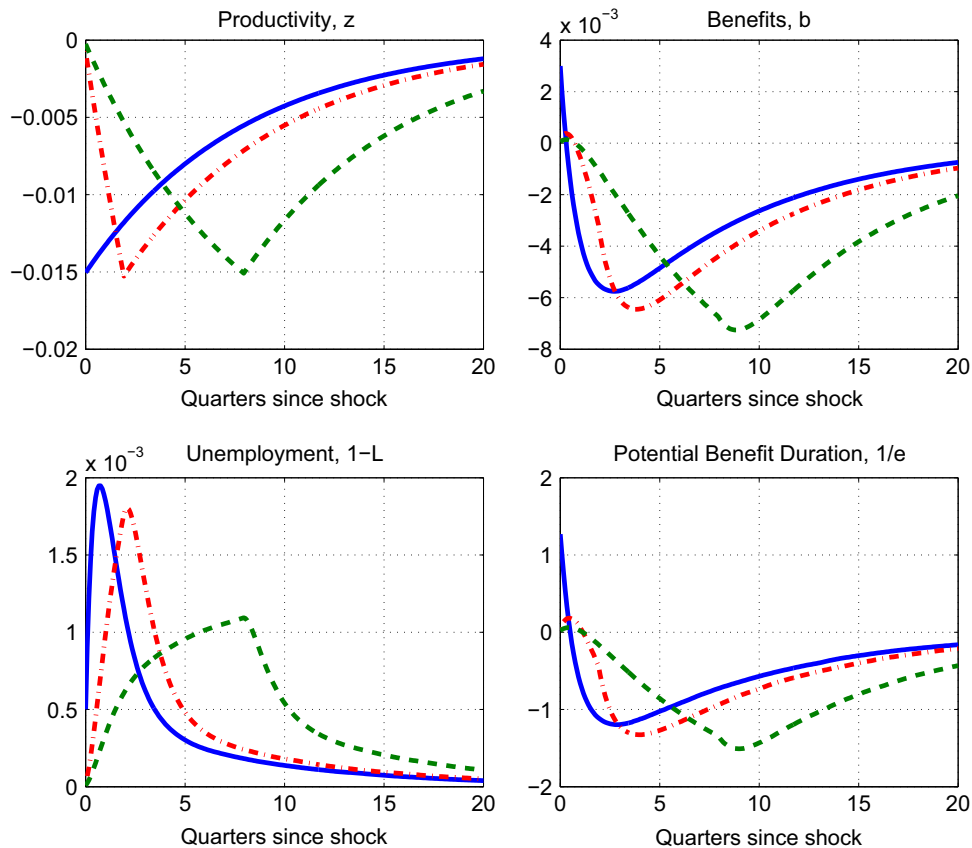


Fig. 3. Responses to 1.5% drop in productivity over different productivity realizations.

Table 4

Model statistics simulated under the current US and Optimal policies.

Statistic		z	$1-L$	$v/(1-L)$	\hat{f}	w	S^E	S^I
<i>Current US policy</i>								
Mean		1	0.057	0.974	0.139	0.945	0.505	0.720
Standard deviation		0.013	0.125	0.259	0.148	0.008	0.053	0.002
Correlation matrix	z	1	-0.851	0.912	0.892	0.772	0.884	0.951
	$1-L$	-	1	-0.913	-0.920	-0.457	-0.922	-0.861
	$v/(1-L)$	-	-	1	0.997	0.497	0.992	0.932
	\hat{f}	-	-	-	1	0.456	0.998	0.926
	w	-	-	-	-	1	0.444	0.696
	S^E	-	-	-	-	-	1	0.927
	S^I	-	-	-	-	-	-	1
<i>Optimal policy</i>								
Mean		1	0.057	0.793	0.135	0.948	0.502	0.720
Standard deviation		0.013	0.026	0.059	0.031	0.011	0.011	0.002
Correlation matrix	z	1	-0.889	0.828	0.7828	0.912	0.766	0.995
	$1-L$	-	1	-0.924	-0.905	-0.674	-0.897	-0.855
	$v/(1-L)$	-	-	1	0.997	0.525	0.995	0.774
	\hat{f}	-	-	-	1	0.459	0.999	0.724
	w	-	-	-	-	1	0.435	0.944
	S^E	-	-	-	-	-	1	0.706
	S^I	-	-	-	-	-	-	1

Note: Means are reported in levels, standard deviations and correlations are reported in logs as quarterly deviations from an HP-filtered trend with a smoothing parameter of 1600. \hat{f} denotes the weekly job finding rate.

5.1.1. The first-best allocation

The social planner is choosing sequences of consumption $x_t^E(z^t)$ and $x_t^U(z^t)$ for the employed and unemployed, respectively; sequences of search effort $S_t(z^t)$ of the unemployed; sequences of market tightness $\theta_t(z^t)$; and sequences of employment levels $L_t(z^t)$ and vacancies $v_t(z^t)$ to maximize¹⁶

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t \{L_t(z^t)u(x_t^E(z^t)) + (1-L_t(z^t))u(x_t^U(z^t)) - (1-L_{t-1}(z^{t-1}))c(S_t(z^t))\} \quad (19)$$

subject to the law of motion

$$L_t(z^t) = (1-\delta)L_{t-1}(z^{t-1}) + S_t(z^t)f(\theta_t(z^t))(1-L_{t-1}(z^{t-1})), \quad (20)$$

and the aggregate feasibility constraint

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t \{L_t(z^t)(z_t - x_t^E(z^t)) - (1-L_t(z^t))(x_t^U(z^t) - h) - kv_t(z^t)\} = 0. \quad (21)$$

Let η be the Lagrange multiplier on (21). The first-order conditions for $x_t^E(z^t)$ and $x_t^U(z^t)$ immediately imply that $x_t^E(z^t) = x_t^U(z^t) \equiv x^{FB} \forall z^t$, where x^{FB} is the solution to $u'(x^{FB}) = \eta$. Using this result, the first-best allocation denoted by x^{FB} , θ_t^{FB} , L_t^{FB} , S_t^{FB} can be characterized.¹⁷

Under the first-best allocation, workers receive full insurance: $x^E = x^U$ (consumption is constant both across employment states and over time). Therefore, in order for a policy to be able to implement the first-best in the decentralized equilibrium, two conditions are necessary: (1) there is no moral hazard on the worker side (otherwise, $x^E = x^U$ would imply $S=0$, due to the worker's optimal choice of search effort); and (2) the workers' surplus from being employed is zero.

This motivates us to consider the following special cases. First, the case in which there is no choice of search effort and the worker's bargaining power is zero. Second, the worker's bargaining power is allowed to be non-zero, but the assumption of no choice of search effort is maintained. Finally, the case when search effort is endogenous is examined as well.

5.1.2. No search effort choice, zero worker bargaining power

When there is no search effort choice,¹⁸ and zero worker bargaining power ($\xi = 0$), the first-best can be implemented by setting a time-invariant unemployment benefit b such that $h+b = x^{FB}$.¹⁹ As long as the unemployment benefit is set this way, the bargained wage will also equal $w = x^{FB}$, and so full insurance will be achieved. Of course, the policy needs to ensure that path of θ_t coincides with the first-best. But this can be achieved by setting time-varying lump-sum taxes

$$\tau_t(z^t) = z_t - x^{FB} + \beta(1-\delta)\mathbb{E} \frac{k}{q(\theta_{t+1}^{FB}(z^{t+1}))} - \frac{k}{q(\theta_t^{FB}(z^t))}. \quad (22)$$

In other words, the taxes are set such that the firm's free entry condition holds after every history of shocks when the market tightness follows the path θ_t^{FB} . It remains to verify that the proposed taxes, together with the unemployment benefit, indeed satisfy the government's budget constraint:

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t \{L_t(z^t)\tau_t(z^t) - (1-L_t(z^t))b\} = 0. \quad (23)$$

This follows from combining the free entry condition with the feasibility constraint. See the Online Appendix for details.

5.1.3. No search effort choice, non-zero worker bargaining power

If worker bargaining power is greater than zero, full insurance is not achieved with the above policy instruments alone, since workers receive a surplus from being employed. However, the government can ensure that the bargained wage always equals $w = x^{FB}$ by setting an appropriate nonlinear tax $T(w)$ on wages, in addition to the constant unemployment benefit b and the lump-sum taxes τ_t .²⁰ Thus, the fact that workers have non-zero bargaining power per se is inconsequential for implementing the first-best allocation, as long as a sufficiently rich (albeit unrealistic) set of policy instruments is allowed.

5.1.4. Endogenous search effort

When workers choose their search effort, there is now worker-side moral hazard, so $x^E = x^U$ clearly cannot be implemented. In this case, for the optimal policy is solved numerically when the government has access to history-dependent unemployment benefits (b_t, e_t) and history-dependent taxes τ_t . In other words, the maximization problem is identical to (18) except that constant tax τ is replaced by a state-contingent tax $\tau_t(z^t)$.

¹⁶ Note that the level of vacancies is $v_t = S_t\theta_t(1-L_{t-1})$.

¹⁷ The detailed conditions for the first-best allocation are shown in the Online Appendix.

¹⁸ For example, assume that $c(S)$ is identically zero, and exogenously set $S=1$.

¹⁹ Note that, throughout Section 5.1.2, x^{FB} refers to the first-best consumption allocation for the case with no search effort.

²⁰ For example, this can be achieved by $T(w) = \max\{0, w - x^{FB}\}$. The worker-firm pair then has no incentive to set a bargained wage above x^{FB} ; the wage also will not be set below x^{FB} if $h+b = x^{FB}$. In equilibrium, we will have $h+b = w = x^{FB}$ and so $T(w) = 0$. The rest of the policy will therefore be set as in the case of zero bargaining power.

Fig. 4 plots the impulse response to a 1.5% drop in productivity for three specifications: the benchmark, constant-tax model (the replication of Fig. 2), the first-best allocation with no search effort choice (the allocation considered in Section 5.1.2), and the model with time-varying taxes. The first-best allocation, as discussed above, features constant consumption for the employed and the unemployed. Note that achieving this constant consumption requires taxes (the taxes implementing the first-best, as in Eq. (22)) to be highly pro-cyclical. These pro-cyclical taxes, in turn, result in smoother unemployment than what we saw in the benchmark optimal policy case.²¹

Next, the model with time-varying taxes is examined; the Online Appendix provides more details on its properties. The first important observation is that, similar to the first-best, taxes are strongly procyclical. This is consistent with the results on subsidies to vacancy creation found in Jung and Kuester (2014). Remarkably, and similar to the first-best, the wages of employed workers are almost perfectly smoothed over the business cycle, as a result of the procyclical tax.

Comparison of the time-varying tax case to the benchmark optimal policy reveals several insights. First, an important function of the optimal policy is smoothing consumption of the employed. In the case with constant taxes, this was partly accomplished by allowing unemployment benefits (which affect wages) to rise in a recession before they fall. When time-varying taxes are used to smooth wages, this eliminates the need for a non-monotonic response of unemployment benefits. Second, unlike the benchmark case, unemployment benefit level and duration move in opposite directions: while unemployment benefit levels fall in response to the negative shocks, unemployment benefit duration rises. Since b and $1/e$ do not co-move positively, we no longer obtain the unambiguous result that pro-cyclicity of UI on both dimensions is optimal. One reasonable measure of the overall generosity of unemployment benefits is the expected present value of unemployment benefits, which is approximated by b/e . This quantity is mildly countercyclical, but it moves very little over the business cycle (as was expected, since the effects of b and $1/e$ almost cancel each other out), again consistent with the result in Jung and Kuester (2014) that unemployment benefit generosity is mildly countercyclical but almost constant. The intuition for this result is that the optimal policy acts to smooth unemployment over the business cycle. With constant taxes, this smoothing was accomplished through pro-cyclical unemployment benefits. In the presence of time-varying taxes, this is accomplished through pro-cyclical taxes, and this eliminates the need for pro-cyclical unemployment benefits.

5.2. Alternative wage-setting mechanisms

A key assumption made in the model is that wages are determined by Nash bargaining. This implies, in particular, that wages are sensitive not only to changes in productivity, but also to changes in unemployment insurance. This, in turn, creates the scope for unemployment benefits to affect vacancy creation. This section examines the consequences of shutting down this channel.

As explained in Section 3, the model is calibrated to match the observed elasticity of wages with respect to productivity. The model therefore replicates the wage rigidity observed in the data. However, under the assumption of Nash bargaining, low responsiveness of wages to productivity (captured by a low worker bargaining power, ξ) implies relatively high responsiveness of wages to the worker outside option, and therefore to unemployment benefits. In this section, wages are assumed not to be a function of the worker outside option at all. Specifically, similar, e.g., to Landais et al. (2013), the wage is assumed to be an exogenously specified function of productivity: $w = \bar{w}z^\gamma$, where \bar{w} is a constant and γ is a constant between 0 and 1. Notice that this wage specification is capable of generating exactly the same elasticity of wages with respect to productivity as the benchmark Nash-bargaining model, although the policy implications are, potentially, very different.

There are two natural choices for γ . The first is $\gamma = 0.449$, under which the model replicates the overall cyclicity of wages in the data (this is also the target used in the benchmark calibration). The second is $\gamma = 0.7$, consistent with the productivity–elasticity of wages of new hires only (this is also the number assumed by Landais et al., 2013). The optimal policy is solved under both parameterizations. The definition and solution for the optimal policy is similar to that in Section 4, with two modifications: the wage is no longer a choice variable for the government, and the Nash bargaining condition is replaced by the exogenous restriction $w = \bar{w}z^\gamma$.

Fig. 5 displays the impulse responses of the optimal policy for three specifications: $\gamma = 0.449$, $\gamma = 0.7$, and the benchmark (Nash bargaining).²² The figure provides two key results. First, under rigid wages, the impulse response generated by the optimal policy retains its non-monotonic shape, similar to the optimal policy under Nash bargaining. Second, all else equal, higher wage rigidity implies a more countercyclical response – a larger initial rise in unemployment benefits in a recession, and a smaller subsequent drop.²³

These observations motivate two questions. First, why is the response of the optimal policy to productivity still apparently ambiguous, even under rigid wages? Second, why does an increase in wage rigidity weaken the case for pro-cyclical unemployment benefits? Since rigid wages fully shut down the responsiveness of vacancy creation to policy, the

²¹ While consumption is perfectly smoothed under the first best, unemployment is not, since it is efficient for unemployment to vary in response to productivity.

²² Our numerical results indicate that, under this form of wage rigidity, unemployment benefits do not expire under the optimal policy, and therefore report the impulse response of the unemployment benefit level only.

²³ Fig. 5 shows this result for $\gamma = 0.449$ and $\gamma = 0.7$, but we obtained consistent predictions for other values of γ : we found that reducing γ increases the initial rise in benefits, and dampens the subsequent drop.

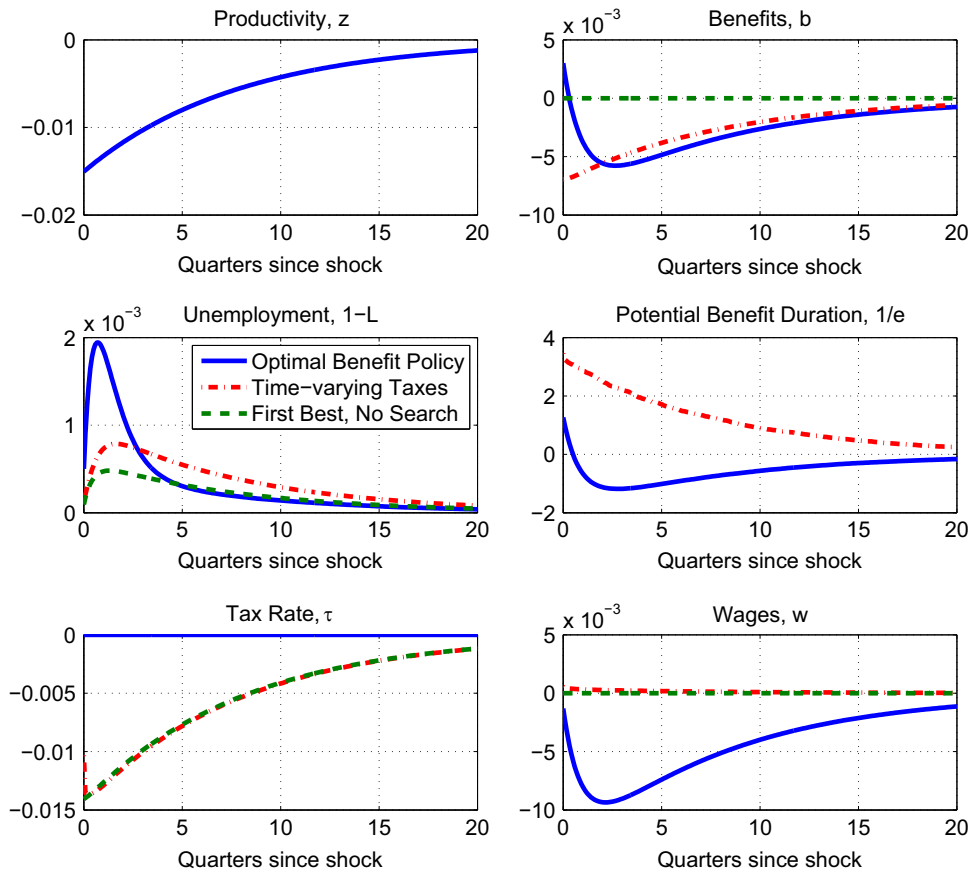


Fig. 4. Responses to 1.5% drop in productivity: benchmark model, first-best, and time-varying taxes.

entire distortionary effect of unemployment insurance is now coming from worker-side moral hazard. Below, the intuition for the observed numerical results is analyzed in a simple static model in which this moral hazard is the only friction.

Consider the following simplified environment: the economy lasts for one period, and all the workers are initially unemployed. A productivity shock z realizes at the beginning of the period; the wage w and the aggregate job-finding rate f are exogenously specified non-decreasing functions of z . The government chooses an unemployment benefit schedule $b(z)$. The workers, knowing the unemployment benefit schedule $b(z)$, choose search effort s in every possible state z . Under these assumptions, the optimal policy problem, for a fixed tax τ , is

$$\max \mathbb{E}\{sfu(w) + (1-sf)u(h+b) - c(s)\} \quad (24)$$

subject to the budget constraint,

$$\mathbb{E}\{sf\tau - (1-sf)b = 0\}, \quad (25)$$

and the worker's optimal search condition, which must hold for all z :

$$c'(s) = f \times (u(w) - u(h+b)), \quad (26)$$

where, throughout, the expectation is with respect to z , and w and f are understood to be exogenous functions of z .

Consider the effect of a decrease in z , assuming first that w is completely fixed. The decrease in z then affects the problem only by decreasing f , but this has two opposing effects on the optimal policy. On one hand, a decrease in f lowers the social return to search, and therefore lowers the level of s that the planner would like to induce. This reduces the social cost of providing unemployment insurance, since the only distortion of this insurance is reduced effort, and this forgone effort is unproductive. This effect of a lower z calls for higher unemployment benefits. On the other hand, the decrease in f also lowers the private returns to search, as evident from the worker's optimal search condition (26). This increases the punishment required to induce any given level of effort, thus calling for lower unemployment benefits. Thus, even in the simple static model in which the only distortion is on the worker side, the direction of the optimal policy is potentially ambiguous.

Next, suppose that w is also an increasing function of z . Now, a fall in z leads not only to a fall in f , but also to a fall in w . This does not affect the first channel described above – the social return to search – since wage affects the share of the

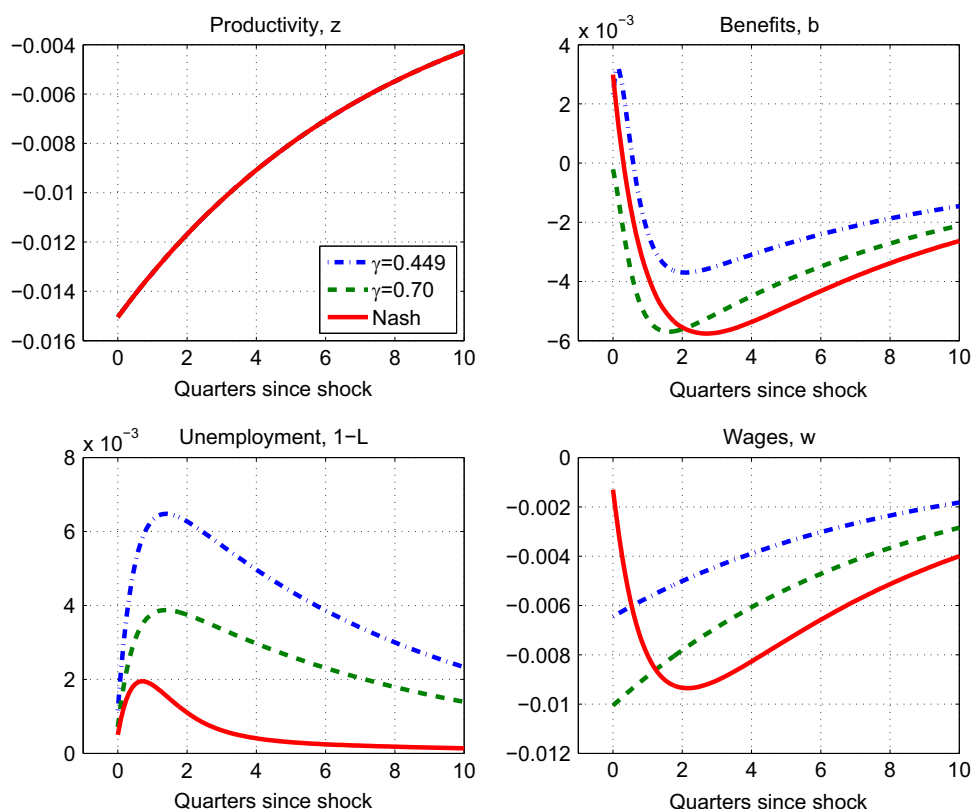


Fig. 5. Responses to 1.5% drop in productivity: Nash bargaining vs. rigid wages.

surplus going to the worker, not the total amount of this surplus. However, the response of the wage amplifies the second effect, since a fall in the wage even further reduces the workers' private incentives to search. Hence, the ability of the wage to respond to productivity exacerbates the moral hazard problem in recessions, strengthening the case for pro-cyclical unemployment benefits. This discussion shows that, although rigid wages do not eliminate the ambiguous predictions for optimal policy, stronger exogenous wage rigidity weakens the case for pro-cyclical unemployment benefits.

It is worth noting that this conclusion refers to the comparison between two rigid-wage regimes with different degrees of wage rigidity (captured by γ in our model). In comparing the rigid-wage model to Nash bargaining, additional complications arise. In particular, in the Nash-bargaining environment, the government can use time-varying unemployment benefits to smooth wages. Wage rigidity, by eliminating this ability, may reduce the case for the initial rise in unemployment benefits in a recession. This is most clearly apparent by comparing the impulse response of wages in Fig. 5 under the different specifications.

5.2.1. Comparison to Landais et al. (2013)

In Landais et al. (2013), there are two important assumptions relative to ours: wages are assumed to be an exogenous function of labor productivity, and there are aggregate decreasing returns in production. The results from our rigid wage model are thus not directly comparable to those in Landais et al. (2013), since the downward-sloping aggregate labor demand curve is not included, which, Landais et al. (2013) state, is important for generating their results.

Despite the differences, it is interesting that our rigid-wage model does not replicate the optimal policy results of Landais et al. (2013), who argue that optimal unemployment benefits are countercyclical. There are other differences between the two papers beyond the model environment itself. In particular, the government's choice variable in Landais et al. (2013) is neither the level of unemployment benefits nor the replacement ratio²⁴ (unemployment benefit level divided by wage), but the contemporaneous consumption ratio of the employed (the after-tax wage) and the unemployed (unemployment benefits). Furthermore, the analysis in Landais et al. (2013) rests on a comparison of steady states, whereas our results are obtained in a dynamic stochastic model. In fact, a key finding in our paper is the complicated dynamic behavior of optimal policy, which implies that policy prescriptions should go beyond a simple positive or negative relationship to current aggregate conditions.

²⁴ See Section 5.7 for a further discussion of replacement ratios versus levels.

5.2.2. The distortionary effect of unemployment insurance: micro- vs. macro-elasticities

Comparison of our results to Landais et al. (2013) illustrates that the choice of modelling framework can be crucial for optimal policy predictions. This raises the question of how to distinguish empirically between the various search models used in the literature on optimal unemployment insurance. In this section, the testable implications between our model and the model of Landais et al. (2013) are discussed, in particular their predictions for the distortionary effect of unemployment benefits.

A key implication of that model is that general equilibrium effects *dampen* the responsiveness of unemployment to UI policy.²⁵ Thus, Landais et al. (2013) predict that the sensitivity of unemployment to economy-wide changes in benefit policy should be smaller, in percentage terms, than its sensitivity to policy changes for a small group of workers: the macro elasticity of unemployment with respect to UI benefits is smaller than the micro-elasticity.

In contrast, in our model, wages are determined by bargaining and are therefore an increasing function of the workers' outside option. Unemployment benefits raise this outside option, thereby discouraging firms from posting vacancies. General equilibrium effects thus *amplify* the responsiveness of unemployment to UI policy. As a result, our model implies that the sensitivity of unemployment to large-scale policy changes should be greater than what would be measured in small-scale experiments: the macro elasticity of unemployment with respect to UI benefits is *larger* than the micro-elasticity. As stated above in Section 3, we calibrated our model parameters to match the empirical finding that the micro-elasticity is about 0.9. Consistent with the above intuition, our model predicts a macro elasticity of 2.4, substantially larger than the micro-elasticity.

Both our model and that of Landais et al. (2013) thus generate clear testable predictions regarding the micro- and the macro-elasticity of unemployment with respect to UI benefits. The relative size of these elasticities in the data is still an open empirical question. A large literature has estimated the micro-effect of UI: for example, the classic studies by Moffitt (1985) and Meyer (1990) estimate that the micro-elasticities of unemployment duration with respect to benefit duration and benefit level, respectively, are about 0.16 and 0.9. The Great Recession has motivated a resurgence of research on this topic. In particular, the recent studies by Farber and Valletta (2014) and Rothstein (2011) have both found that the unprecedented extensions of unemployment benefits in the Great Recession have had very modest effects.

Measuring the macro elasticity requires obtaining reliable estimates of general equilibrium effects of UI. These general equilibrium effects are difficult to measure, because large scale policy changes are typically endogenous to changing macroeconomic conditions. Several studies have attempted to measure the wedge between the micro- and the macro-elasticity by estimating the crowding-out effect of search in response to policy changes. Using French data on young long-term unemployed people, Crepon et al. (2013) evaluate the effect of a job placement counseling program, both on the workers who participated in the program and on those who did not. When they consider the entire sample of workers, they find no evidence that an increase in search by some workers crowds out the job finding probability of other job seekers. However, when they restrict their estimates to a selected group of males, they do find evidence of crowding out, suggesting – indirectly – that the macro elasticity of unemployment with respect to benefits may be smaller than the micro-elasticity for certain segments of the labor market.²⁶ Similarly, Lalive et al. (2013) find evidence of crowding-out effects of search by examining the effects of unemployment benefits extensions for Austrian workers. The results of both studies imply that the macro elasticity is smaller than the micro due to the existence of such crowding-out effects. Other studies instead attempt to directly estimate the macro elasticity itself, rather than the wedge. For the US, Hagedorn et al. (2013, 2014) use exogenous variation in benefit extensions across states and estimate an aggregate elasticity of unemployment with respect to benefit duration of about 0.55, significantly larger than the micro-estimate of 0.16 in Moffitt (1985).²⁷ Their result implies that the macro elasticity is large, and, in particular, larger than the existing estimates of the micro-elasticity. In a study of Sweden, Fredriksson and Soderstrom (2008) exploit cross-regional variation in unemployment benefit generosity and likewise find a very large macro elasticity. Our conclusion from these very recent studies is that the current evidence on the question is mixed, and that empirical work on measuring the macro elasticity is a promising research agenda that should stimulate further research. Our model's prediction provides a way of testing between the various models if further reliable estimates of the macro elasticity do become available in the future.

5.3. The effect of persistence

An important feature of the model environment is that agents' behavior depends not only on the current policy, but also on their expectations about future policy. Moreover, both the private and social incentives for job creation depend not only

²⁵ As Landais et al. (2013) explain, the combination of their two key assumptions – rigid wages and aggregate decreasing returns – results in a wedge between the micro- and the macro-elasticity of unemployment with respect to unemployment benefits. Specifically, since wages are exogenously fixed, the labor market does not adjust to equate labor supply and labor demand, and jobs are therefore rationed. A fall in unemployment benefits triggers an increase in search intensity by unemployed workers, but, because the number of jobs does not respond to this policy change, this increase in search intensity has a crowding-out effect that partially offsets the effect on unemployment.

²⁶ Note, however, that the policy evaluated in Crepon et al. (2013) is neither an unemployment benefit extension nor a change in the benefit level, and may therefore not have the same effect as UI on the workers' outside option; as such, it might not have delivered an amplifying general equilibrium effect through vacancy posting even if such an effect existed.

²⁷ Note that the measure of 0.9 from Meyer (1990) is the micro-elasticity with respect to level and not duration.

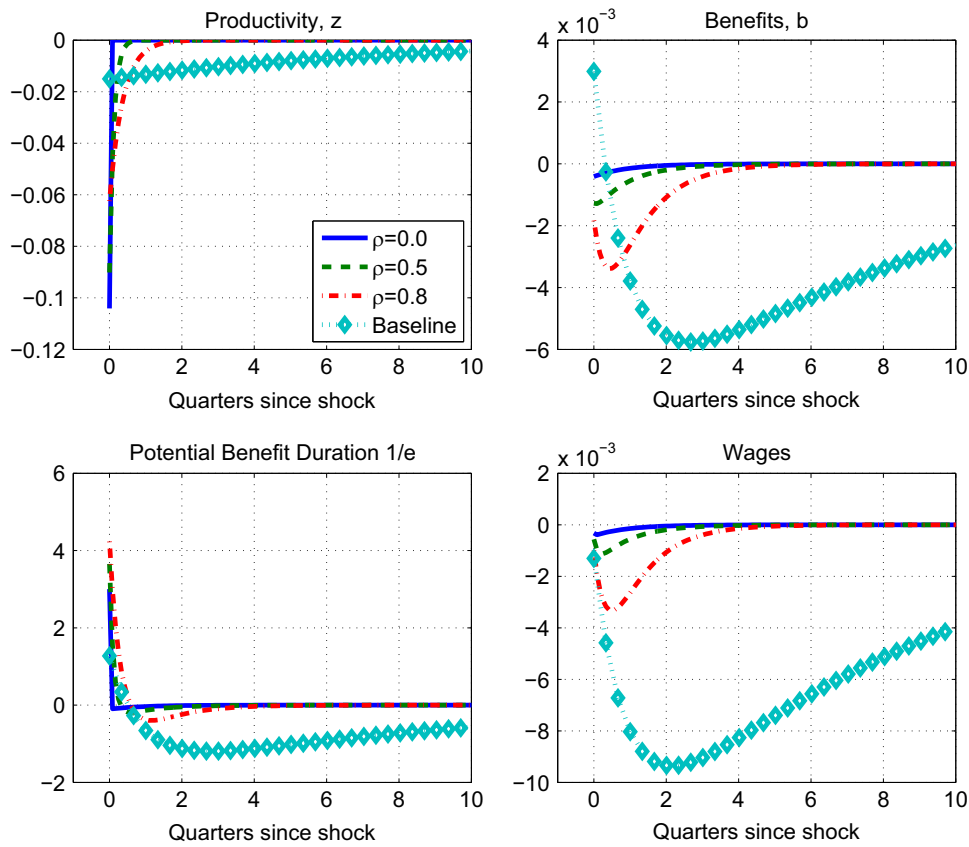


Fig. 6. Responses to 1.5% drop in productivity for different persistence parameters.

on current productivity, but also on expected future productivity. Both of these observations suggest that the persistence of productivity shocks may be important for the quantitative results. This section investigates the effect of changing the persistence, ρ , of the productivity process (1).

Fig. 6 reports the impulse responses of the optimal policy for the same drop in productivity, but for different values of ρ . Lower persistence of the shock dampens both the initial rise in benefit generosity and its subsequent fall.

The intuition for this observed pattern is that the optimal policy has a tendency to backload punishments in recessions. In particular, punishment for not finding a job (in the form of less generous unemployment benefits) can be imposed by either lowering current benefits or by lowering future benefits. As explained earlier in Section 4, there is less of a social benefit from job creation – and hence less of a reason to punish unemployed agents – when current productivity is low. As a result, the optimal policy responds to a fall in current productivity by raising the current generosity of unemployment insurance, but lowering the future generosity, expecting productivity to recover in the future. The latter depends crucially on how quickly the recovery is expected to occur, and hence on the persistence of the shock. When ρ is high (as it is in our baseline, $\rho = 0.9895$), low productivity is expected to persist, and hence the shock is expected to lead to significant rise in unemployment. The optimal policy thus prescribes a generous initial extension of current unemployment insurance, but a correspondingly large decline once the recovery is expected to occur. On the other hand, when the shock is closer to i.i.d., it leads neither to a large drop in the social gains from job creation (since these gains depend on the expected present value of productivity) nor to a large rise in unemployment. As a result, the optimal policy does not raise benefit generosity as much initially and does not subsequently need to lower it: the need to backload punishments almost disappears.

Moreover, the response of the benefit level b changes with the persistence of the shock: the initial rise in b disappears altogether, although the rise in benefit duration, $1/e$, does not. An important component of the optimal policy is smoothing the wages of employed workers, which respond much more strongly to the benefit level (which has a first-order effect on the outside option) than the benefit duration. Importantly, wages in a given period depend not only on current productivity but also on expected future productivity. Therefore, when productivity is not very persistent, wages respond quite weakly to a shock to current productivity, and hence the incentive to smooth wages disappears. This is most clearly visible in the lower-right panel of Fig. 6, which shows that the wage response tracks b quite closely, and more so for small ρ .

5.4. The assumption of no savings

An important assumption made for transparency in this paper is that workers cannot save or borrow. Relaxing this assumption may have several implications for the results.

On one hand, if workers are allowed to hold wealth, cyclical variations in this wealth will affect how government-provided insurance should vary over the business cycle. Periods in which unemployed workers' wealth is lower would warrant higher unemployment benefits. Intuitively, this effect is similar to the effect that would arise if h varied over the business cycle. In particular, if workers are more liquidity-constrained in recessions than in booms, this would provide a motive for raising unemployment benefits in recessions (or raising their duration), with the potential to reverse our optimal policy results.

On the other hand, the presence of savings reduces the responsiveness of the worker outside option to unemployment benefits. As a result, both worker search effort and firm vacancy posting will respond less to policy than they would in the absence of savings. Therefore, in the presence of savings, inducing any given behavioral response requires a larger change in benefits than it would have required otherwise. This effect would amplify the cyclical behavior of optimal benefits in our model, potentially making optimal benefits even more strongly pro-cyclical. The overall effect of introducing savings in the model on the pro-cyclicality of optimal benefits is thus ambiguous.

Finally, the presence of saving reduces the welfare gains from unemployment insurance, since it provides workers with a way to self-insure against unemployment shocks. Thus, regardless of how the presence of saving changes the direction of the optimal policy, the welfare gains from adopting the optimal policy are likely to be smaller if saving is allowed.

5.5. The Hosios condition and its relationship to our model

A concern in the Diamond–Mortensen–Pissarides model with Nash bargaining is that the laissez-faire equilibrium is not constrained efficient. Even with risk-neutral workers, the Hosios (1990) condition requires that the worker bargaining weight be equal to the elasticity of the matching function in order to attain efficiency. If the Hosios condition is violated, there is a role for government intervention – such as unemployment benefits – even in the absence of insurance considerations. The reason for this is that when an individual firm posts a vacancy, it reduces the matching probability of other firms and increases the matching probability for workers, thereby imposing an externality.

The Hosios condition is not applicable to our model: since workers are risk-averse in our model, output maximization is not equivalent to welfare maximization. Nevertheless, the question can still be posed to what extent our optimal policy results are driven by corrections for the externality that an individual firm imposes when entering. To answer this question, the value of the worker bargaining power ξ such that the optimal government intervention (UI benefit and tax) is zero in the steady state is found, keeping all the other parameters fixed at their calibrated values. This serves as the intuitive analogue of the Hosios condition in our model. The value is found to be $\xi = 0.72$. Next, the optimal policy is solved for this value of ξ , keeping all other parameters fixed at the benchmark calibration values. The result is displayed in the Online Appendix. A comparison of the impulse responses shows that the shape of the optimal policy is robust to raising ξ to 0.72. The same holds for the overall pro-cyclicality of optimal benefits. These qualitative results are also unchanged if a value of ξ higher than 0.72 is used. This indicates that our results are not driven by a search externality.

An alternative way to interpret the Hosios condition in our context would be to ask for what value of the bargaining power the *current* US policy is optimal in the steady state. Interestingly, as Table 3 indicates, our calibrated model implies that the current unemployment benefit level and duration in the US are close to optimal. In other words, under this interpretation, our calibrated value for the bargaining power is already at the equivalent of the Hosios condition. Our results can be interpreted as saying that, while the average generosity of unemployment insurance in the US is close to optimal, its cyclical properties are not.

5.6. Benefit level and duration

Stochastic benefit expiration has been assumed to jointly characterize the optimal behavior of benefit levels and duration. This assumption is made for tractability, since it renders the dynamic problem of the worker stationary. The optimal cyclical behavior of benefit levels and expected benefit duration is qualitatively similar: both are pro-cyclical and both exhibit the same dynamic response to a productivity shock. However, the presence of stochastic benefit expiration in the model is not important for our results. To illustrate this, the optimal policy is computed when the government is restricted to change only one of these two policy dimensions. Three alternative policy experiments are conducted. In the first, the benefit level is fixed at its current level: $b=0.4$, and only the duration to changes over the business cycle. The results, reported in the Online Appendix, show that the optimal policy response is similar qualitatively to our benchmark: in response to a negative productivity shock, potential duration of benefits should initially rise, and then fall considerably below its initial level. However, both the initial rise in the potential duration and its subsequent decline are greater than in the benchmark optimal policy result. In the second and third experiments, the benefit expiration rate is fixed at the levels $e = 1/26$ (the current level) and $e=0$ (benefits do not expired at all), respectively, and the optimal benefit policy is computed. The results are shown in the Online Appendix. The shape of the policy response is once again similar to the

benchmark: benefits initially rise and then fall. Thus, the main result is quite robust and, in particular, holds when expiration is shut down altogether and the only policy variable is the benefit level.

5.7. The replacement ratio of unemployment benefits

The actual UI system in the US indexes individual's unemployment benefits to his wage in the previous job. Because of this, the policy variable of interest in policy discussions is often not the benefit level, but the replacement ratio – the ratio of an unemployed worker's benefits to his previous wage. In our model, however, the benefit level, rather than the replacement ratio, was deliberately used as the government's choice variable. In order to realistically mimic the administration of the replacement ratio in the US, the model would need to assume that the replacement ratio is a function of wages received during the worker's previous employment spell – not the current aggregate wage. Computing welfare would require the government to keep track of the distribution of past employment histories,²⁸ making the model intractable.

On the other hand, assuming that the government's choice variable is b/w , where b is the current unemployment benefit and w is the current aggregate wage, could lead to misleading results.²⁹ Thus, our assumption that the government chooses b rather than b/w , while imperfect, appears to be a good compromise.

5.8. Sensitivity analysis

The Online Appendix illustrates the robustness of the results to the parameterization of the model. In particular, it shows that our main qualitative results are robust to changing the value of non-market activity h , the bargaining power ξ , and the risk aversion parameter σ . Similar robustness results hold when changing the values of other parameters.

6. Conclusion

The design of an optimal UI system over the business cycle is characterized in an equilibrium search and matching model. Optimal unemployment benefits respond non-monotonically to productivity shocks: while raising benefit generosity may be optimal at the onset of a recession, it becomes suboptimal as the recession progresses and inducing a recovery is desirable. Optimal unemployment benefits are pro-cyclical overall, counter to previous results in the literature and to the way UI policy is currently conducted. Our findings thus demonstrate that conventional wisdom guiding policymakers may be overturned in a quite standard equilibrium search model of the labor market.

The paper has focused on the optimal cyclical behavior of UI benefits and thus serves to inform the ongoing policy debate on the desirability of benefit extensions in recessions. UI benefits are a worker-side intervention, as they affect the economy by changing the workers' value of being unemployed. As discussed in Section 5.1, it may be important to consider the optimal behavior of UI benefits in conjunction with firm-side interventions, such as hiring subsidies. Increasing hiring subsidies in recessions can be desirable as another instrument for stimulating an employment recovery. A potential concern with hiring subsidies, frequently articulated in policy debates, is the firm-side moral hazard they generate: firms could, for example, fire existing employees only to hire them again in order to receive hiring subsidies. A thorough investigation of the tradeoffs involved with such policies seems to be a fruitful extension for future work.

Finally, an important direction for future research is investigating the role of government commitment. The ability of the government to commit matters because the behavior of agents in our model depends not only on the current policy, but also on their expectations about future policy. Throughout the paper, we have assumed that the government can fully commit to its policy. A government without commitment power might be tempted not to lower benefits when there are a lot of unemployed workers. It will therefore be interesting to characterize the time-consistent policy and compare it to the optimal policy in the presence of aggregate shocks.

Acknowledgments

We thank the editor and an anonymous referee for constructive comments that improved the paper. We are grateful to Bjoern Bruegemann, Harold Cole, Hanming Fang, Greg Kaplan, Dirk Krueger, Iouri Manovskii, Guido Menzio, Pascal Michaillat, B. Ravikumar and Randall Wright for valuable advice and support. We would like to thank seminar and

²⁸ At any point in time, unemployed workers differ in the past wages they received while employed, and would thus differ in their benefit levels if the replacement ratio were used. This would also imply that workers would differ in their outside option during wage negotiations, leading to a distribution of wages at any point in time.

²⁹ For example, consider a worker who had been employed in a boom and gets fired at the beginning of a recession. The US unemployment insurance system would assign this worker an unemployment benefit based on his previous wages, which are likely to have been high. A system that conditions b on the current aggregate w would assign this worker a considerably lower unemployment benefit level. Furthermore, unlike the US system, a policy that varies b based on the aggregate w , rather than the worker's own history, would result in an unemployment benefit level that fluctuates throughout the worker's unemployment spell, whereas it is constant in the data. These features make this alternative problematic.

conference participants at McGill University, Indiana University, Amherst College, University of California – Davis, the Federal Deposit Insurance Corporation, the Guanghua School of Management at Peking University, the Macro Seminar and the Empirical Micro Lunch at the University of Pennsylvania, the Search and Matching Workshop at the Federal Reserve Bank of Philadelphia, the 2011 Midwest Macro Meetings, the 2011 Society for Economic Dynamics Meetings, the IZA Conference on Unemployment Insurance, the 2011 Cologne Workshop for Macroeconomics, and the 2012 Philadelphia Workshop on Macroeconomics, as well as colleagues at numerous institutions, for helpful comments and discussions.

Appendix A. Supplementary data

Supplementary data associated with this paper can be found in the online version at <http://dx.doi.org/10.1016/j.jmoneco.2014.11.009>.

References

- Aguilar, M., Hurst, E., 2005. Consumption versus expenditure. *J. Polit. Econ.* 113 (5), 919–948.
- Albrecht, J., Vroman, S., 2005. Equilibrium search with time-varying unemployment benefits. *Econ. J.* 115 (505), 631–648.
- Baily, M.N., 1978. Some aspects of optimal unemployment insurance. *J. Public Econ.* 10 (3), 379–402.
- Browning, M., Crossley, T.F., 2001. Unemployment insurance benefit levels and consumption changes. *J. Public Econ.* 80 (1), 1–23.
- Cahuc, P., Lehmann, E., 2000. Should unemployment benefits decrease with the unemployment spell? *J. Public Econ.* 77 (1), 135–153.
- Coles, M., Masters, A., 2006. Optimal unemployment insurance in a matching equilibrium. *J. Labor Econ.* 24 (1), 109–138.
- Crepon, B., Duflo, E., Gurgand, M., Rathelot, R., Zamora, P., 2013. Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *Q. J. Econ.* 128 (December (2)), 531–580.
- den Haan, W.J., Ramey, G., Watson, J., 2000. Job destruction and propagation of shocks. *Am. Econ. Rev.* 90 (3), 482–498.
- Faig, M., Zhang, M., 2012. Labor market cycles, unemployment insurance eligibility, and moral hazard. *Rev. Econ. Dyn.* 15 (1), 41–56.
- Farber, H.S., Valletta, R.G., 2014. Do Extended Unemployment Benefits Lengthen Unemployment Spells? Evidence From Recent Cycles in the U.S. Labor Market. Working Paper 19048, National Bureau of Economic Research.
- Fredriksson, P., Holmlund, B., 2001. Optimal unemployment insurance in search equilibrium. *J. Labor Econ.* 19 (2), 370–399.
- Fredriksson, P., Soderstrom, M., 2008. Do Unemployment Benefits Increase Unemployment? New Evidence on an Old Question. IZA Discussion Papers 3570, Institute for the Study of Labor.
- Gruber, J., 1997. The consumption smoothing benefits of unemployment insurance. *Am. Econ. Rev.* 87 (1), 192–205.
- Hagedorn, M., Karahan, F., Mitman, K., Manovskii, I., 2013. Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects. Working Paper 19499, National Bureau of Economic Research.
- Hagedorn, M., Manovskii, I., 2008. The cyclical behavior of equilibrium unemployment and vacancies revisited. *Am. Econ. Rev.* 98 (4), 1692–1706.
- Hagedorn, M., Manovskii, I., Mitman, K., 2014. Quasi-Experimental Evidence on the Effects of Unemployment Benefit Extensions. Working Paper, University of Pennsylvania.
- Hopenhayn, H.A., Nicolini, J.P., 1997. Optimal unemployment insurance. *J. Polit. Econ.* 105 (2), 412–438.
- Hosios, A.J., 1990. On the efficiency of matching and related models of search and unemployment. *Rev. Econ. Stud.* 57 (2), 279–298.
- Jung, P., Kuester, K., 2014. Optimal Labor Market Policy in Recessions. Working Paper 11–48, Federal Reserve Bank of Philadelphia.
- Kroft, K., Notowidigdo, M.J., 2010. Does the Moral Hazard Cost of Unemployment Insurance Vary with the Local Unemployment Rate? Theory and Evidence. Economics Working Paper, University of Chicago Booth Graduate School of Business.
- Krueger, A.B., Meyer, B.D., 2002. Labor supply effects of social insurance. In: Auerbach, A.J., Feldstein, M. (Eds.), *Handbook of Public Economics*, vol. 4. Elsevier, pp. 2327–2392. Ch. 33, <http://www.elsevier.com/books/handbook-of-public-economics/feldstein/978-0-444-87612-6>.
- Lalive, R., Landais, C., Zweimüller, J., 2013. Market Externalities of Large Unemployment Insurance Extension Programs. Discussion Paper 9673, Center for Economic Policy Research.
- Landais, C., Michailat, P., Saez, E., 2013. Optimal Unemployment Insurance Over the Business Cycle. Working Paper 16526, National Bureau of Economic Research.
- Lehmann, E., van der Linden, B., 2007. On the optimality of search matching equilibrium when workers are risk averse. *J. Public Econ. Theory* 9 (5), 867–884.
- Meyer, B.D., 1990. Unemployment insurance and unemployment spells. *Econometrica* 58 (4), 757–782.
- Mitman, K., Rabinovich, S., 2014. Do Unemployment Benefit Extensions Explain the Emergence of Jobless Recoveries? Working Paper 14–013, Pennsylvania Institute for Economic Research.
- Moffitt, R., 1985. Unemployment insurance and the distribution of unemployment spells. *J. Econom.* 28 (1), 85–101.
- Rothstein, J., 2011. Unemployment Insurance and Job Search in the Great Recession. Working Paper 17534, National Bureau of Economic Research.
- Shavell, S., Weiss, L., 1979. The optimal payment of unemployment insurance benefits over time. *J. Polit. Econ.* 87 (6), 1347–1362.
- Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. *Am. Econ. Rev.* 95 (1), 25–49.
- Shimer, R., Werning, I., 2008. Liquidity and insurance for the unemployed. *Am. Econ. Rev.* 98 (5), 1922–1942.