On The Welfare Costs Of Unemployment Fluctuations *

Michael Reiter, Institute for Advanced Studies, Vienna

May 2012

Abstract

The paper investigates the welfare losses created by fluctuations in the labor markets. The framework is a model of heterogeneous workers who are subject to fluctuations in the probability of finding employment. They can self-insure against unemployment through saving in a riskless asset. It is explained how fluctuations increase the probability of long unemployment spells, which are hard to insure against for the workers.

Numerically, it turns out that the welfare gains from completely eliminating the fluctuations in the labor market are either small or tiny, depending on whether stabilization policy is supposed to keep the average job finding rate or the average unemployment rate unchanged. In either case, the welfare gain in a heterogeneous agent model is not much higher than in a representative agent model, or even smaller.

JEL classification: C68, D31, D61, E32
Keywords: cost of business cycles; heterogeneous agents; incomplete markets

Address of the author:

Michael Reiter
Institute for Advanced Studies
Stumpergasse 56
A-1060 Vienna
Austria

e-mail: michael.reiter@ihs.ac.at

*I am grateful to Christian Haefke, Tamas Papp and seminar participants in Vienna for helpful discussions. Support from the Austrian National Bank under grant Jubiläumsfonds 13810 is gratefully acknowledged.

1 Introduction

Business cycles are characterized by fluctuations in the unemployment rate which are large[1] and which are often considered to be inefficient.[2] If this is true, it is natural to assume that stabilization policy, by reducing inefficient fluctuations, can increase social welfare.

Lucas (1987) concluded that the potential welfare gains from stabilization policy are tiny for the representative agent of the economy. But the welfare costs of business cycles, in particular the burden of unemployment, are spread unevenly across households because of imperfect insurance. It is therefore natural to study welfare costs in the framework of heterogeneous agents and incomplete markets, which was first done by Imrohoruglu (1989). She still found that the welfare gains are very small, and this result was confirmed by most later studies, including Atkeson and Phelan (1994) and Krusell and Smith (1999). However, some more recent studies, notably Krebs (2007) and Krusell et al. (2009), find gains from stabilization of more than one percent.

This literature follows Lucas (1987) in investigating the consequences of a "miraculous elimination" of economic fluctuations. Government policy is not modeled explicitly, so that it is not clear how stabilization policy affects the individual income processes. Instead some general "integration principle" is used, which assumes that stabilization removes the portion of individual income that is correlated with the aggregate cycle. The details of how this principle is applied are important. They explain the difference in the results of Krusell and Smith (1999) and Krusell et al. (2009).

In the present paper, I stay in the tradition of considering the miraculous elimination of fluctuations, but I make a more specific assumption about the effects of stabilization on individual income. The dominant strand of the macroeconomic literature on unemployment fluctuations is the Mortensen/Pissarides search and matching framework. In its simplest version, it assumes a constant job separation rate, and a job finding rate that varies over the business cycle. In this framework, what stabilization policy can ideally achieve is to stabilize the job finding rate. There are two natural benchmarks to consider. In the first one, fluctuations in the job finding rate are eliminated, keeping average job finding rate constant. This leads to a reduction in the mean unemployment rate, and therefore an increase in average production. This is a consequence of the mechanics of a matching function, and was recently stressed by Jung and Kuester (2011). In the second benchmark, the job finding rate is stabilized at a lower mean value, so that the average unemployment rate is unchanged.

For any given average unemployment rate, however, the question is why *variations* in the unemployment rate should matter for the household. After all, the average time that a household spends in unemployment is determined by the average unemployment rate, not by

[1] The standard deviation of the detrended log unemployment rate is 0.19 Shimer (2005, Table 1), almost 8 times that of log GDP.

[2] For example (Hall 2009, p.284): "The employment rate of 94 or 95 percent seen in normal times may well be efficient, but a decline to 91 or 92 percent in a severe recession is almost certainly inefficiently high."

its variation over the cycle. The answer is that fluctuations in unemployment rates matter because they increase the probability of *long* unemployment spells. In an environment where households self-insure against unemployment through saving, the long spells are the ones that are difficult to insure against. The purpose of my paper is to analyze the welfare implications of this mechanism. I first give a qualitative analysis of the relationship between fluctuations, long unemployment spells, and welfare. I investigate how this depends on the size of asset market frictions. It turns out that the welfare gains from stabilization do not depend monotonically on asset market frictions, but are highest at an intermediate level of frictions. Then I provide a numerical analysis of the welfare gains, which turn out to be very small, using a calibration of the fluctuations in job finding probabilities which match US data. This finding is consistent with the existing literature, because the effect of stabilization on household income in my model is closer to Imrohoruglu (1989) and Krusell and Smith (1999) than to Krusell et al. (2009). The labor market matching framework leads to predictions of what stabilization policy can achieve that differ strongly from the integration principle as applied in Krusell et al. (2009). In Section 4.7 I will compare the two approaches in more detail.

In such a framework, the heterogeneity of agents does not play an important role for the aggregate gains from stabilization. If stabilization leaves the average unemployment rate constant, the welfare gain in the heterogeneous agent model is not higher than what Lucas (1987) found for the representative agent. It may even be smaller. If the average unemployment rate is reduced, the heterogeneous agent models can have higher welfare gains, because a higher unemployment rate requires more self insurance, and this is costly if financial frictions are large. Nevertheless, this additional effect on welfare is never more than half of the direct effect from increased average production. Moreover, the welfare gain does not depend very strongly on the asset level of workers. Of course, the limited role of household heterogeneity for stabilization policy does not diminish the importance of the heterogeneous agent framework for *long-term* policy analysis, such as the design of unemployment insurance.

The plan of the paper is as follows. Section 2 sets up the model. Before presenting numerical results in Section 4, I illustrate the logic of stabilization policy with a toy model in Section 3. Section 5 concludes.

2 Model

2.1 Partial Equilibrium

Compared to the literature cited in the introduction, I take a step back by looking at the problem in partial equilibrium, in the sense of keeping the wage and the interest rate fixed, and having workers' job finding rate fluctuate exogenously. I take the labor market dynamics as exogenous, because there is no consensus in the literature on why unemployment fluctuates so much, nor on what drives wage fluctuations and how sticky wages are. Moreover, by having job finding rates fluctuate exogenously rather than as a market response to changes in productivity, I give the model a better chance to generate sizable welfare gains from eliminating those fluctuations. The reason for keeping the interest rate fixed lies in the nature of the model, which only has workers. Considering the high concentration of financial wealth in the hand of capitalists and entrepreneurs, as well as the role of international capital markets, the implications of such a model for interest rate changes cannot be taken seriously.

2.2 The Household Problem

We assume there is a continuum of households (workers), who supply labor inelastically. This means, every worker looks for a job, and works a fixed amount of hours if she has one. A worker who has a job gets separated exogenously at the end of the period with constant probability σ.[3] A worker who gets separated will be short-term unemployed in the next period. A short-term unemployed worker fill find a job with probability p_t and start employment in period $t+1$. With probability λ she will become long-term unemployed next period. A long-term unemployed worker finds a job with probability $p_t \chi$ where $\chi < 1$. This means, the long-term unemployed have a lower job finding rate, but it is proportional to p. This seems to be a good description of the empirical findings in Elsby et al. (2010).

The worker chooses consumption so as to maximize expected utility, which is assumed to be of the CRRA form

$$\mathrm{E}_0 \sum_{t=1}^{\infty} \beta^t u(c_t) = \mathrm{E}_0 \sum_{t=1}^{\infty} \beta^t \frac{c_t^{1-\gamma} - 1}{1-\gamma} \tag{1}$$

subject to

1. the asset accumulation equation

   $$a_t = (1+r)a_{t-1} + y_t - c_t \tag{2}$$

   and the liquidity constraint

   $$a_t \geq \underline{a} \tag{3}$$

   Income is given by

   $$y_t = \begin{cases} w & \text{if employed in } t \\ b_s & \text{if short-term unemployed in } t \\ b_l & \text{if long-term unemployed in } t \end{cases} \tag{4}$$

2. The exogenous process of job finding probability, p_t and the constant separation probability σ.

[3] I follow most of the matching literature in keeping the separation rate fixed and explaining all the fluctuations in unemployment through variations of the job finding rate. As will become clear below, it is the variation in finding rates that generates a role for stabilization policy. Explaining unemployment partly through changes in the separation rate would probably reduce the gains from stabilization even further.

The decision problem of the worker can be characterized recursively by the following value functions. Denote by $V^E(a,x)$ the value of an employed worker in aggregate state x with beginning of period assets a. $V^U(a,x)$ and $V^L(a,x)$ are the corresponding value functions for short-term and long-term unemployed workers, respectively. Then the value of an employed worker satisfies

$$V^E(a,x) = \max_c \left\{ u(c) + \beta \sum_{x'} \pi(x,x') \left[(1-\sigma)V^E(a',x') + \sigma V^U(a',x') \right] \right\}$$
$$\text{with} \qquad a' \equiv (1+r)a + w - c \tag{5}$$

The value of a short-term unemployed satisfies

$$\begin{aligned} V^U(a,x) = \max_c \Big\{ u(c) + \beta \sum_{x'} \pi(x,x') \Big[& p(x)V^E(a',x') \\ & + \lambda(x)V^L(a',x') + (1-p(x)-\lambda(x))V^U(a',x') \Big] \Big\} \end{aligned}$$
$$\text{with} \qquad a' \equiv (1+r)a + b_s - c \tag{6}$$

The value of a long-term unemployed satisfies

$$V^L(a,x) = \max_c \left\{ u(c) + \beta \sum_{x'} \pi(x,x') \Big[p(x)\chi V^E(a',x') + (1-p(x)\chi)V^L(a',x') \Big] \right\}$$
$$\text{with} \qquad a' \equiv (1+r)a + b_l - c \tag{7}$$

2.3 The Dynamics of Aggregate Employment

The assumptions made in Section 2.2 imply the following dynamics of the aggregate employment rate E, short term unemployment rate U and long term unemployment rate U.

$$E_t = (1-\sigma)E_{t-1} + p_{t-1}U^S_{t-1} + p_{t-1}\chi U^L_{t-1} \tag{8a}$$
$$U^S_t = \sigma E_{t-1} + (1-p_{t-1}-\lambda)U^S_{t-1} \tag{8b}$$
$$U^L_t = \lambda U^S_{t-1} + (1-p_{t-1}\chi)U^L_{t-1} \tag{8c}$$
$$U_t = U^S_t + U^L_t \tag{8d}$$

Notice that the aggregate employment rates do not affect the household optimization problem. They are only used for aggregate bookkeeping.

2.4 Fluctuations and Stabilization

I assume there are two aggregate states, good and bad, denoted by the indices g and b, respectively. The probability of switching from state j to state j' is given by $\pi(j,j')$, where $j,j' \in \{g,b\}$. I will assume a symmetric transition law π, such that the two aggregate states are equally likely in the stationary distribution of (8).

The aggregate state determines the probability of finding a job, p. It is higher in the good state:

$$p_g > p_b \tag{9}$$

To eliminate fluctuations, or to stabilize the economy, simply means to eliminate the variation in p. I will do this in two different ways:

1. Keeping average p unchanged, by setting $p^* = (p_g + p_b)/2$.
2. Keeping the unemployment rate in the stationary distribution of (8) constant. This requires to set $p^* < (p_g + p_b)/2$.

Whether stabilization policy in the real world is better described by the first or the second type can only be answered in a full-fledged structural model of the labor market. The numerical results below will always be reported for both types of stabilization.

2.5 Definition of Aggregate Welfare Gain

The welfare gain (or loss) of a household is a function of the household's assets and employment status, and of the aggregate state at the time when stabilization is introduced. For an aggregate welfare measure, we integrate the individual welfare gain over the distribution of individual and aggregate state variables. We need a welfare measure that is not just a steady-state comparison, but includes the welfare gain along the transition path from the fluctuating to the stabilized economy. To obtain this, we integrate the household value function in each case over the same cross-sectional distribution, namely the stationary distribution of the fluctuating economy.

The details of the welfare computation are as follows. Given a stochastic process for the job finding probability, the solution to the household problem consists in a saving function

$$a_t = \mathcal{A}\left(a_{t-1}, x_t, e_t\right) \tag{10}$$

where $e_t \in \{e, u, l\}$ denotes the employment state of the household in period t, and $x_t \in \{b, g\}$ denotes the aggregate state of the economy. Here e stands for employed, u for (short-term) unemployed, l for long-term unemployed. b stands for bad aggregate state (low job finding probability) and g stands for good aggregate state (high job finding probability). In the stabilized economy, the job finding probability is the same in b and g, so that the two aggregate states are effectively identical.

The savings function (10) in the fluctuating economy, together with the exogenous stochastic processes for x_t and e_t imply a unique stationary distribution over (a, x, e), which we denote by $\mu^*(a, x, e)$. Denote by $V^e(a; x)$ the value function of the household with assets a and employment states $e \in \{E, U, L\}$ in the aggregate state $x \in g, b$ of the fluctuating economy.

Denote by $\tilde{V}^e(a;x)$ the value function in the corresponding stabilized economy. The welfare gain G from stabilization is then computed as

$$G \equiv \begin{cases} \left(\frac{\int \tilde{V}^e(a;x)\,d\mu^*(a,x,e)}{\int V^e(a;x)\,d\mu^*(a,x,e)}\right)^{1/(1-\gamma)} - 1 & \text{if } \gamma \neq 1 \\ \exp\left[(1-\beta)\left(\int \tilde{V}^e(a;x)\;d\mu^*(a,x,e) - \int V^e(a;x)\;d\mu^*(a,x,e)\right)\right] - 1 & \text{if } \gamma = 1 \end{cases} \tag{11}$$

To motivate this definition, define $\hat{V}^e(a;\kappa;x)$ as the value that a household achieves, starting from state (a,x,e), if it consumes $(1+\kappa)$ as much as in the non-stabilized economy, under each realization of the shocks in the indefinite future. Then

$$\hat{V}^e(a;\kappa;x) = \begin{cases} V^e(a;x)(1+\kappa)^{1-\gamma} & \text{if } \gamma \neq 0 \\ V^e(a;x) + \frac{1}{1-\beta}\log(1+\kappa) & \text{if } \gamma = 0 \end{cases} \tag{12}$$

The aggregate welfare gain is then defined as the value of κ such that

$$\int \hat{V}^e(a;\kappa;x)\;d\mu^*(a,x,e) = \int \tilde{V}^e(a;x)\;d\mu^*(a,x,e) \tag{13}$$

which using (12) gives (11).

2.6 Insurance

To assess the welfare gains from stabilization, it is interesting to compare them to the welfare gains obtained from an increase in unemployment insurance.

I implement insurance as actuarially fair in the following sense. Given the unemployment rate in steady state (which is not changed by insurance), the increase in benefits of the unemployed are exactly financed by the reduction in wages of the employed. I always consider an increase in insurance of 1 percent of labor productivity. Notice that the increase in insurance is modeled as an increase in the parameters b_s and b_l. While the increase in b_s and b_l is interpreted as redistribution that must be financed, the basic level of those parameters is treated like home production, not requiring redistribution. This is just for ease of exposition and could be changed without further consequences.

Notice that there is no moral hazard from insurance in the model, because labor supply is exogenous.

3 The Logic of Unemployment Stabilization

Given that the unemployment rate, averaged over time, is 6 percent, why does it matter whether it is always 6 percent, or sometimes 3 percent and sometimes 9 percent? In the long run, a household will be unemployed 6 percent of the time in either case. The answer is that if the unemployment rate varies (sometimes 3, sometimes 9 percent), it increases the probability of having a long unemployment spell, against which it is difficult to insure. This section presents a toy model that allows to formulate this argument analytically. Although this model makes some ad-hoc assumptions, it illustrates the main mechanisms that are at work in the full model of Section 2.

We assume a finite horizon of T time periods, and both the gross interest rate R and the time discount factor β equal 1, to simplify notation. The household maximizes

$$\mathrm{E}_o \left[\sum_{t=1}^{T} u(c_t) + \phi a_T \right] \tag{14}$$

where we assume that $u(c)$ is strictly concave in c. The household gets constant marginal utility ϕ from any assets left after T periods, which can be interpreted as bequests. Household income is $y_t = w = 2$ if employed, and $y_t = b = 1$ if unemployed. We set the job separation to $\sigma = 1$, so that in each period, the probability of being employed equals the probability of finding a job. As in the labor market matching literature, we assume that the job finding probability of each household is independent of its employment history, so that $\mathrm{E}_{t-1}\, p_t$ is the same for each household, and equals the average employment rate in t. The unemployment rate equals $1 - p_t$.

Let us first analyze the case where the household is excluded from the capital market. Then consumption equals income in all periods, and the objective function reduces to

$$\mathrm{E}_0 \sum_{t=1}^{T} u(c_t) = \mathrm{E}_0 \left[p_t u(2) + (1 - p_t) u(1) \right] = T u(1) + (u(2) - u(1)) \sum_{t=1}^{T} \mathrm{E}_0\, p_t \tag{15}$$

If the household cannot save, the only thing that matters is the unconditional probability of being employed, $\mathrm{E}_0\, p_t$, not whether this probability itself fluctuates. There is no case for stabilization policy. Notice that this is true not only in this simple model, but also in the model of Section 2. The argument only depends on the assumption that utility is separable across time and across states of the world.

Let us now study the problem where the household can save. To further simplify the setup, we assume that the set of feasible consumption levels is discrete, i.e., the household can only vary consumption in discrete steps, $C \in \{1, 2, 3, 4, \ldots\}$. We also assume

$$u(2) - u(1) > \phi > u(3) - u(2) \tag{16}$$

Then the optimal saving policy of the household is simple: if an agent is unemployed in period t and still has assets $a \geq 1$, it is optimal to reduce assets by 1 and consume 2 rather than 1 in period t. There is no point in keeping the assets for a later day, since there is no chance of getting higher marginal utility than in the current period. Similarly, it is never optimal to increase asset holdings over time. This would require an employed household to consume 1 rather than 2. But since there is no danger of consumption ever falling below 1, there is no point in reducing it to 1 today. However, it is not optimal to consume 3 rather than 2, because the additional utility does not even cover the foregone utility from the reduction in bequests, because of (16).

With this policy, we can write household utility as a function of the unemployment history. If $n_u \leq A$ (the household is unemployed not more than A times), then it can always maintain a consumption level of 2. Then utility is given by $Tu(2) + (A - n_u)\phi$. Otherwise, utility is given by $(T + A - n_u)u(2) + (n_u - A)u(1)$. (We assume, for simplicity, that A is integer.) Defining $\mathcal{I}_{n_u \geq A}$ is the indicator function with value 1 if $n_u \geq A$ and 0 otherwise, we can write expected utility as

$$\mathrm{E}_0 \left[Tu(2) + (A - n_u)\phi - \mathcal{I}_{n_u \geq A}(n_u - A)\left(u(2) - u(1) - \phi\right)\right] \tag{17}$$

As long as $n_u \leq A$, each additional period of unemployment causes a utility reduction of ϕ. If $n_u > A$, it costs $u(2) - u(1)$, which, from (16), is bigger than ϕ.

Let us now analyze (17) under different assumptions about initial assets A.

Case $A \geq T$

This is the case of perfect insurance. Unemployment risk is covered with probability 1, and utility in bequests is linear. Then (17) reduces to

$$Tu(2) + A\phi - \phi\,\mathrm{E}_0\, n_u$$

Employment risk enters expected utility only through $\mathrm{E}_0\, n_u = \sum_{t=1}^{T} \mathrm{E}_0(1 - p_t)$, the unconditional expectation of the average unemployment rate. If agents are perfectly insured, unemployment fluctuations do not matter.

Of course, this result is just a consequence of assuming linear utility in bequests, but this assumption represents well the case of agents with decreasing absolute risk aversion, where labor earnings risk is of little importance for households with high assets.

Case $A = 0$

Since the household does not have assets to begin with, and it is never optimal to save, this case is the same as the case without access to the capital market. Then (17) reduces to

$$Tu(2) - \mathrm{E}_0\, n_u\left[u(2) - u(1)\right]$$

Since the term in brackets is a constant, the only thing that matters is, again, $\mathrm{E}_0\, n_u$.

Case $A = T - 1$

This is the case of a high, but not perfect level of insurance. Among the cases we discuss here, it is the one that mimics the results of the full model most closely.

Then (17) reduces to

$$Tu(2) + (A - \mathrm{E}_0\, n_u)\phi - \mathrm{Prob}\,[n_u = T]\,(u(2) - u(1) - \phi)$$

For given average unemployment rate $\mathrm{E}_0\, n_u$, the objective is then to minimize $\mathrm{Prob}\,[n_u = T]$. From the assumption that the job finding probability of a household is independent from its employment history, this probability is equal[4] to $\mathrm{E}_0 \prod_{t=1}^{T}(1 - p_t)$.

With this objective function, is stabilization policy desirable? Stabilization means reducing the variance of p_t, keeping the unconditional expectation $\mathrm{E}_0\, p_t$ fixed. Consider first the case that the aggregate job finding rates p_t are independent over time. Then we get $\mathrm{Prob}\,[n_u = T] = \mathrm{E}_0 \prod_{t=1}^{T}(1 - p_t) = \prod_{t=1}^{T}(1 - \mathrm{E}_0\, p_t)$. This means that, for a fixed $\mathrm{E}_0\, p_t$, the stochastic process of the p_t does not matter.

Stabilization can only be important if we allow for the (realistic) case of positive correlation of the p_t over time. To illustrate this, let us look at the simple case that p is stochastic, but constant over time. In the world without stabilization, this probability is drawn randomly in period 0, and then is either $p_t = p_b$ for $t = 1, 2, \ldots, T$ with probability 0.5, or $p_t = p_g$ for $t = 1, 2, \ldots, T$ with probability 0.5. In the world with stabilization policy, uncertainty is removed, and $p_t = (p_b + p_g)/2$ for $t = 1, 2, \ldots, T$ with probability 1.

In that case, for given p, we have $\mathrm{Prob}\,[n_u = T] = (1 - p)^T$. Because of the convexity of the power function and assumption (9), we have

$$\frac{1}{2}\left[(1 - p_b)^T + (1 - p_g)^T\right] > \left[\frac{1}{2}\left((1 - p_b) + (1 - p_g)\right)\right]^T \tag{18}$$

Stabilization is therefore desirable, because it reduces $\mathrm{Prob}\,[n_u = T]$.

Case $A = 1$

This last case illustrates that stabilization can even be harmful. It is the case of a very low level of insurance. (17) now becomes

$$Tu(2) - \mathrm{E}_0(n_u - A)\,[u(2) - u(1)] + \mathrm{Prob}\,[n_u = 0]\,(u(2) - u(1) - \phi) \tag{19}$$

[4] To see this, note that the event that the household is unemployed, conditional on p_t, can be written as $\epsilon_t \geq p_t$, where ϵ_t is uniformly distributed and i.i.d. over time. Then

$$\begin{aligned}
\mathrm{Prob}\,[n_u = T] &= \mathrm{E}\left[\mathcal{I}_{\epsilon_1 \geq p_1 \text{ and } \epsilon_2 \geq p_2 \text{ and } \epsilon_3 \geq p_3 \text{ and } \ldots \text{ and } \epsilon_T \geq p_T}\right] = \mathrm{E}\left[\prod_{t=1}^{T} \mathcal{I}_{\epsilon_t \geq p_t}\right] \\
&= \mathrm{E}\left\{\mathrm{E}_{p_1, p_2, \ldots, p_T}\left[\prod_{t=1}^{T} \mathcal{I}_{\epsilon_t \geq p_t}\right]\right\} \\
&= \mathrm{E}\left\{\prod_{t=1}^{T} \mathrm{E}_{p_1, p_2, \ldots, p_T}\left[\mathcal{I}_{\epsilon_t \geq p_t}\right]\right\} = \mathrm{E}\left\{\prod_{t=1}^{T}(1 - p_t)\right\}
\end{aligned}$$

The before-last equality follows from the independence of the ϵ_t.

For given $\mathrm{E}_0\, n_u$, the aim is then to maximize $\mathrm{Prob}\,[n_u = 0] = \mathrm{E}_0 \prod_{t=1}^{T} p_t$. Analogous to the discussion of the case $A = T - 1$, this would reduce to $\prod_{t=1}^{T} \mathrm{E}_0\, p_t$ for uncorrelated p_t, so that stabilization does not matter. To illustrate what happens if the p_t are positively correlated over time, assume again that they are perfectly correlated, as in the case of $A = T - 1$ above. Then $\mathrm{Prob}\,[n_u = 0] = p^T$. Since

$$\frac{1}{2}\left[p_b^T + p_g^T\right] > \left[\frac{1}{2}\left(p_b + p_g\right)\right]^T \tag{20}$$

stabilization reduces expected utility, because it reduces $\mathrm{Prob}\,[n_u = 0]$.

Why do fluctuations in the unemployment rate improve welfare? The chance that the available resources are sufficient to ensure against unemployment are very small, but the best chance we have is to make one realization of p big; the convexity of the power function then helps to maximize this very small chance.

This result may appear to be an artifact of the overly simplistic toy model. But the numerical results of Section 4 will show that this can actually happen in the full model, if only in cases where asset holdings of households are extremely low.

4 Numerical Results

4.1 Parameter Values

The length of the model period is 1/48 of a year, to which we refer as a week. The parameter values of the benchmark parametrization are reported in Table 1. The value for β is standard.

Variable :	Meaning	*Value*
β:	discount factor	$0.96^{1/48}$
R:	interest factor	$1.02^{1/48}$
γ:	risk aversion	1
b_s:	replacement rate	0.4
b_l:	repl. rate long-term	0.4
σ:	separation rate	0.349/4
λ:	prob. becoming long-term	1/24
χ:	search efficiency long-term	0.5
ρ:	persistence unemployment rate	$0.73^{1/48}$
p_g:	job finding prob. good state	0.2287
p_b:	job finding prob. bad state	0.1263
$\underline{a}$:	minimum asset level	0

Table 1: Benchmark parameter values

The choice of R will be motivated in Section 4.2. The values for the replacement rate are the same as in Shimer (2005, Table 1). They are lower than in most other papers, the idea here is to give stabilization policy a change to be important. The separation rate is taken from Haefke and Reiter (2011), based on CPS data.

The probability of switching from long-term to short-term unemployed is 1/24 weekly, so that on average workers switch to long-term unemployment after 6 months, which is the usual definition of long-term unemployment. The search efficiency of the long-term relative to that of the short-term unemployed is set to $\chi = 0.5$, in line with the findings of Elsby et al. (2010).

The parameters ρ, p_g and p_b were chosen so as to match the mean (5.73 percent), standard deviation (1.20 percent) and first order autocorrelation (0.9837) of the unemployment rate (US Civilian Unemployment Rate, monthly, seasonally adjusted, 1948-2010). Second moments are computed after detrending monthly data by the Hodrick-Prescott filter with smoothing parameter $\lambda = 10^7$. I have chosen this high value of λ, because for smaller values, the trend would capture a large part of what is clearly classified as fluctuations by visual inspection.

I set the minimum asset level to zero, rather than allowing the household to go into debt. This choice should not matter too much, as Carroll (2001) observes that a shift in the borrowing limit tends to shift the whole asset distribution upwards, so as to leave the buffer stock of saving (assets above borrowing threshold) almost unaffected. By not allowing households to borrow, I avoid the problem of the spread between borrowing and lending rates. In any case, this makes self-insurance more difficult for households and should give greater welfare gains from stabilization.

Next to the benchmark values, I report results for many other parameter combinations. Typically, I keep benchmark values and vary one parameter systematically. Across different parametrizations, I keep the time discount parameter constant and vary the interest rate. The interpretation is that household face financial frictions that drive a wedge between their lending rate and the market interest rate that large investors can obtain.[5] The alternative would be to fix an interest rate and vary the impatience of households, but this has the disadvantage that it is difficult to compare value losses across households with different discount factors.

4.2 Interest Rates and Precautionary Saving

The upper panels of Figure 1 display results for the benchmark calibration, for varying interest rates. The upper left panel displays welfare gains from stabilization (which will be discussed in Section 4.3), and the upper right panel displays two aggregate statistics of the steady state in the fluctuating economy: the mean household asset levels,[6] expressed in percent of annual

[5] This is similar to Imrohoruglu (1989, p.1370), who assumes a time discount rate of 4 percent annually, and a zero percent real interest rate. I will study real interest rates varying between four and minus six percent.

[6] The figures report only mean, not median assets. Median assets are close to mean assets in the model, unlike in the data.

wages (solid line, refers to left y-axis), and the percentage reduction in consumption, at the moment when the household enters unemployment (dotted line, refers to right y-axis).

The results cover a wide interval of gross interest rates. At the maximum level $R = 1.04$, we have $\beta R = 0.9984$ in annual terms, quite close to unity. Nevertheless, average assets in the economy are less then 30 percent of annual wages. This is because the savings incentive is quite low in this economy: the only risk is unemployment, there are no persistent changes in individual productivity (wages). At the minimum level $R = 0.94$, the wide gap between the interest rate ($r = R - 1 = -0.06$) and the time discount factor ($1 - \beta = 0.04$) reduces average assets to about 5 percent of annual wages.

To pin down a benchmark value for the difference between the interest rate and the time discount rate, we can use recent empirical estimates of the amount of liquid assets accumulated by households, and of the estimated drop in consumption when a household enters unemployment. On the first point, Engen and Gruber (2001, Table 1) report that the median household has liquid assets of about 6 percent of annual income. Supposing a replacement rate of 40 percent, the results in Gruber (1997) imply that food consumption drops by 11.6 percent at the start of an unemployment spell. Browning and Crossley (2001) estimate a 14 percent drop of total consumption 6 months after the unemployment incidence.[7]

From this evidence, the empirically relevant range seems to be between $R = 1.00$ and $R = 1.02$. I will use $R = 1.02$ in most of the following exercises, also because this is the value that maximizes the welfare gain from stabilization.

4.3 Welfare Gains in the Benchmark Case

The upper left panel of Figure 1 displays three welfare measures for the benchmark calibration: the welfare gains from stabilization when keeping the job finding rate constant (solid line); the welfare gains from stabilization when keeping the unemployment rate constant (dashed line); the welfare gains from an actuarially fair increase in unemployment insurance of one percentage point (dotted line). All welfare measures are expressed as a percentage of permanent consumption.

The first, and perhaps surprising result, is that the welfare gains from stabilization are tiny if stabilization leaves the average unemployment rate unchanged. These gains reach their maximum at $R = 1.02$, where they equal 0.013 percent of consumption. This is even smaller than the welfare gains reported in Lucas (2003). That those gains reach a maximum at an interior value of R was expected from our discussion in Section 3. Figure 2 shows how those gains are spread between different workers. In a recession, workers in any employment status and with any level of assets would gain from stabilization. The higher expected job finding rates in the near future (as long as the boom lasts) outweighs the future gains from

[7]These results are based on Canadian data. Canada may have a higher replacement rate, but the paper also shows, consistent with my Figure 4, that the replacement rate has limited impact on the results.

stabilization. In a boom, all workers would lose from stabilization. Effects are strongest for workers with low levels of assets, but the differences between asset levels are not dramatic.

Figure 1 also shows that stabilization can achieve a welfare gain of around 0.2 percent of permanent consumption if it leaves the average job finding rate, rather than the unemployment rate, unchanged. The difference comes from the fact that the stabilization of job finding rates (keeping the average finding rate constant) reduces the mean unemployment rate by 0.254 percentage points, which leads to a 0.152 percent increase in output. This is almost identical to the difference between the solid and the dashed line in the case of very high self insurance ($R = 1.04$), which is 0.159. Interestingly, the welfare gain is monotonically decreasing in R. The reduction in mean unemployment is the more important the more costly it is to self-insure. While this welfare gain is not trivial, it is not much higher then the welfare gain from stabilization that would occur in a representative agent model. In the latter case, the gain would be at least 0.152 percent, corresponding to the increase in average output, and somewhat more because utility is concave in consumption and leisure.

To put those welfare gains into perspective, it is interesting to compare them to the welfare gains from increased unemployment insurance (cf. Section 2.6). Not surprisingly, the welfare gain decreases monotonically in R, because insurance is more valuable when self-insurance is more costly. Only at high levels of self-insurance, with $R = 1.03$, stabilization of unemployment rates brings the same gain as a 1 percentage point increase in unemployment insurance. One should mention that an increase in insurance is always welfare enhancing in this model, because labor supply is exogenous. Using cross-country panel data, Costain and Reiter (2008) estimate than a one percentage point increase in insurance raises steady state unemployment by 2 percent, i.e., about 0.12 percentage points. This would cause an output loss of more than 0.1 percent, which more than offsets the gain from increased insurance.

Notice that the stabilization of the job finding rate becomes more important, relative to insurance, when R increases. The higher R, the better households are insured against short unemployment spells, so that the dominant concern becomes long spells, and reducing the probability of long spells is exactly what stabilization does.

4.4 Comparison to Representative-Agent Model

A precise comparison between our heterogeneous agent model and a representative agent counterpart is hard to make. For example, with the benchmark interest rate of $R = 1.02$ and $\beta = 0.96$, an agent endowed with the aggregate income stream would not save anything.

So is appears better to look at the numbers cited in the literature. Lucas (2003, p. 4) presents estimates of 0.05 percent and 0.1 percent of steady state consumption for risk aversion parameters of one and two, respectively. These estimates are based on the fluctuations in consumption, to which one would have to add the welfare loss from fluctuations in labor. With the variability of employment in our data, fluctuations of labor supply create a utility loss of about 0.012 percent of steady state consumption if the Frisch labor supply elasticity is 0.5 (cf. Appendix B). This estimates slightly increase if variations in hours per worker are included.

In comparison, the welfare gain from unemployment stabilization is less than 0.02 percent in the benchmark case, and not much higher than 0.05 percent in any of the experiments considered. This means, for realistic calibrations, the gain tends to be even lower than in the representative agent model.

4.5 Varying Some Parameters

4.5.1 Lower Replacement Rates

Figure 3 show results for the case of lower unemployment benefits. For long term unemployed, it is now $b_l = 0.1$ rather than $b_l = 0.4$ as in the benchmark case. For the short-term unemployed, I keep $b_s = 0.4$ in the upper panels and use $b_s = 0.1$ in the lower panels. Benefits of ten percent of labor productivity should probably be seen as a lower bound of what is realistic, given that unemployment insurance must be interpreted as including home production and the value of leisure.

It turns out that the welfare gain (with average unemployment rate unchanged), although higher than in the benchmark case, is still only about 0.05 percent of permanent consumption. This is well in the range of the welfare gains obtained in representative-agent models. Qualitatively, the results are similar to the benchmark case. To show the effect of the unemployment insurance replacement rate more systematically, the upper panels of Figure 4, display results as a function of the replacement rate $b_s = b_l$, with the gross interest rate fixed at $R = 1.02$.

The reason why the gains from stabilization do not increase more is that self-insurance becomes more urgent when unemployment insurance is low, and therefore mean asset levels are much higher than in the benchmark case, for any given level of R. One might object that median liquid assets of workers are lower in the data, but the model already allows for a very large wedge between the discount and the interest rate factor.

4.5.2 Varying Degrees of Risk Aversion

The lower panels of Figure 4, display results as a function of the risk aversion parameter, with the gross interest rate fixed at $R = 1.02$. As one expects, the welfare gain from stabilization increases monotonically in the degree of risk aversion. However, this effect is mitigated by the change in the steady state level of assets: households buy more self-insurance if it is more urgently needed.

4.5.3 Disaster Risk

The lower panels of Figure 1 report on another robustness check. To capture the situation that households may find themselves without any assets after some rare event, I have made the asset return risky. With a chance of 5 percent annually, the worker faces some disaster so that all her wealth is lost, i.e., $R_t = 0$. The return in the no-disaster case was increased so that the expected gross interest rate $\mathrm{E}_{t-1} R_t$ is unchanged. The horizontal axis in the picture measures the expected gross rate.

The comparison to the benchmark case (upper panels of Figure 1) shows that this change in setup makes no substantial difference to the result. The reason is that, after such a shock, people accumulate more assets rather quickly, so that only a small portion of the work force is close to the borrowing constraint.

4.6 Welfare Loss in Case of Very Little Self-Insurance

The analysis of the simple model in Section 3 has suggested that stabilization might even decrease welfare, in the case of very low self-insurance. The results in the lower panel of Figure 5 show that this can indeed happen even in the full model, but only for very small degrees of self-insurance. The welfare gain is negative in a ranch of interest rates where average steady state asset levels are about 0.1 percent of annual wages; this does not appear to be a realistic scenario.

The results in Figure 5 were computed for a constant unemployment exit hazard, for the following reason. If we fix the job finding rate so as to keep the steady state unemployment rate constant, this leads to an increase in long-term versus short-term unemployment, which causes a utility loss if the long-term unemployed have lower exit hazards. This effect is very small and can be ignored for realistic values of self-insurance, but in the case of very low self-insurance, all effects are very small. To have a clear interpretation of the results, it is necessary to eliminate this effect, which can be done by giving long-term unemployed the same benefit and the same exit hazard as the short-term unemployed. Then long-term unemployment makes no difference.

4.7 Comparison to the Literature

Imrohoruglu (1989) finds very small welfare costs of fluctuations, in a setup where stabilization eliminates the fluctuations in employment probabilities, but where the stochastic processes of prices are exogenous and not affected by stabilization. Both Atkeson and Phelan (1994) and Krusell and Smith (1999) stress the welfare effects of stabilization through the changes in the stochastic processes of prices (asset prices and wages). They still only find small effects. Moreover, they attribute the fact that Imrohoruglu (1989) finds any nonzero costs of fluctuations, despite of unchanged price processes, to the way she constructs the individual income process in the stabilized economy. Concretely, she does not what is proposed by

Atkeson and Phelan (1994), namely to remove the correlation across individual income process by making the aggregate shock idiosyncratic.

In contrast to this earlier literature, Krusell et al. (2009) do find substantial welfare gains from stabilization. Compared to the present paper, Krusell et al. (2009) use a calibration that is more favorable to stabilization policy, because recessions are very hard times for the unemployed. They assume that in a recession, the long-term unemployed only have a 1.25 percent change per quarter to find a new job (cf. page 399), although their situation is mitigated by the fact that recessions only last 8 quarters on average. Furthermore, benefits for the long-term unemployed are only about 5 percent of wages. This means that in a recession of 8 quarters, most of the long-term unemployed live on 5 percent of the wage for two years. In contrast, I assume that the job finding probability of the long-term unemployed is half as big as those of the short-term unemployed, in line with the empirical findings of Elsby et al. (2010). I use unemployment replacement rates between 10 and 40 percent. Nevertheless, the calibration is not the main reason for the difference in results. This can be seen from the fact that Krusell and Smith (1999) use the same calibration as Krusell et al. (2009) but find much smaller welfare effects. What is key is their interpretation of the integration principle. It says the following (cf. their Section 2.3.2). Write individual income y_i as a function of an aggregate shock z and an idiosyncratic shock i, $y_i = h(i, z)$. Income in the stabilized economy is then defined as

$$y^{stab}(i) = \int h(i, z)\, dF_z(z) \tag{21}$$

where $F_z(z)$ is the distribution function of z. To illustrate what this means in our context, consider again the case where the job separation σ equals one. Denote the job finding probability at time t by p_t. Then income can be written as

$$y_t(i) = w \cdot \mathcal{I}_{i \le p_t} + b \cdot \mathcal{I}_{i > p_t} \tag{22}$$

where i is an idiosyncratic shock which is uniformly distributed on $[0, 1]$, $\mathcal{I}$ is again the indicator function, and $p_t \in \{p_b, p_g\}$. We assume again that good and bad states are equally likely.

In the stabilized economy, we integrate over the distribution of p_t, which gives

$$\begin{aligned} y_t^{stab}(i) &= \frac{1}{2}\left[w \cdot \mathcal{I}_{i \le p_g} + b \cdot \mathcal{I}_{i > p_g} + w \cdot \mathcal{I}_{i \le p_b} + b \cdot \mathcal{I}_{i > p_b}\right] \\ &= \begin{cases} b & \text{if } i > p_g \text{ (this happens with prob. } 1 - p_g) \\ \frac{w+b}{2} & \text{if } p_b < i \le p_g \text{ (prob. } p_g - p_b) \\ w & \text{if } i \le p_b \text{ (prob. } p_b) \end{cases} \end{aligned} \tag{23}$$

The mean income in the fluctuating as well as in the stabilized economy is given by

$$\mu \equiv pw + (1 - p)b \tag{24}$$

where p is defined as $p = (p_b + p_g)/2$. The variance in the fluctuating economy is

$$\frac{p_g + p_b}{2}(w - \mu)^2 + \frac{1 - p_g + 1 - p_b}{2}(b - \mu)^2 = (w - b)^2 p(1 - p) \tag{25}$$

The variance in the stabilized economy is

$$\begin{aligned}(w - \mu)^2 p_b + \left(\frac{w + b}{2} - \mu\right)^2 (p_g - p_b) + (b - \mu)^2 (1 - p_g) \\ = (w - b)^2 \left[p(1 - p) - \frac{p_g - p_b}{4}\right]\end{aligned} \tag{26}$$

Stabilization therefore reduces the variance of income by

$$\frac{1}{4}(w - b)^2(p_g - p_b) \tag{27}$$

This example highlights the difference between the integration principle and the approach of the matching model. Compared to the fluctuating economy, the integration principle reduces the probability of the extreme realizations b and w, and predicts a medium level of income $\frac{w+b}{2}$ with probability $(p_g - p_b)$. In the matching model, income is either b or w, both in the fluctuating and in the stabilized economy.

Another prominent paper that applies the integration principle is Krebs (2007). In this model, households face permanent shocks to income. Some of those shocks are idiosyncratic, but others are interpreted as job displacement shocks, whose probability and severity depends on business cycle conditions. In recessions, job displacements are both more frequent and imply a higher loss of lifetime income. In the stabilized economy, both the frequency and the income loss of displacements are constant. Depending on the degree of risk aversion, the gains from stabilization can reach several percent. This comes from the interaction of the integration principle and persistent income shocks.

While the calibration of the above models is carefully motivated by empirical studies, they are still in the tradition of Lucas (1987) and do not provide a structural explanation for the causes and effects of job displacements. Therefore I do not include some important features of those models, such as persistent changes in the worker's productivity, mainly because this would again raise the issue of how all this is affected by stabilization policy.

5 Conclusions

In this paper I have studied a heterogeneous agent model where economic fluctuations reduce welfare through variations in the unemployment risk of households. I have shown how stabilization raises welfare because it reduces the frequency of long unemployment spells. However, the welfare increase turns out to be very small; in particular, they are not substantially (if at all) higher than the welfare gains calculated from comparable representative agent models.

This is in line with most earlier findings in the literature, but in contrast to some recent contributions, for example Krebs (2007) and Krusell et al. (2009). The difference mainly comes from the assumptions on how aggregate stabilization affects idiosyncratic income processes. While Krusell et al. (2009) and most of the existing literature use an abstract integration principle, I assume that stabilization policy affects household income by stabilizing the job finding rate of the unemployed. This is a natural framework, given the recent literature on labor market matching.

These conflicting results highlight an important dilemma for applied theory. The empirical findings, as reported, for example, in Storesletten et al. (2004) and Krebs (2007), seem to indicate important gains from stabilization, but those findings are, to my knowledge, not yet satisfactorily explained by structural models. In the absence of a generally accepted model underpinning those empirical results, researchers are likely to analyze policy in the framework of relatively simple benchmark models, such as a Mortensen-Pissarides model. Those models, however, attribute little importance to stabilization policy.

A Computational Issues

The household problem is solved by dynamic programming on a discrete grid. Employment status and the aggregate state are discrete variables; household assets are discretized on a grid of 1000 points (changing the number of grid points to 2000 has a negligible effect on the results). Between grid points, the value function is interpolated by cubic splines.

Given a household savings function, one can compute a transition probability matrix for the aggregate and individual state variables. This serves to compute the stationary distribution of wealth. In the theoretical model, the stationary cross-sectional distribution of wealth has an infinite number of discrete mass points, because the distribution of idiosyncratic productivity is discrete such that households at the borrowing constraint $\underline{k} = 0$ return to the region of positive k in packages of positive mass. I approximate this complicated distribution by a finite number of mass points at a predefined grid $\underline{k} = \bar{k}_1, \bar{k}_2, \ldots, \bar{k}_{n_k} = \overline{k}$. The maximum level $\overline{k}$ must be chosen such that in equilibrium very few households are close to it.

The key element of the approximation is the following. If the mass p of households in period t saves the amount $\tilde{k}$ with $\bar{k}_i \leq \tilde{k} \leq \bar{k}_{i+1}$, I approximate this by assuming that $p \cdot \frac{\tilde{k}-\bar{k}_i}{k_{i+1}-k_i}$ households end up at grid point $\bar{k}_{i+1}$, while $p \cdot \frac{\bar{k}_{i+1}-\tilde{k}}{k_{i+1}-k_i}$ households end up at grid point $\bar{k}_i$. This random perturbation of capital leaves aggregate capital unaffected, because $\frac{\tilde{k}-\bar{k}_i}{k_{i+1}-k_i}\bar{k}_{i+1} + \frac{\bar{k}_{i+1}-\tilde{k}}{k_{i+1}-k_i}\bar{k}_i = \tilde{k}$.

The stationary distribution is then computed by sparse matrix methods as the eigenvector of the unit eigenvalue of the transition probability matrix.

B Utility Loss of Labor Fluctuations in the Representative Agent Model

With the per-period utility function

$$\log C_t - \eta \frac{L_t^{1+\nu}}{1+\nu} \tag{28}$$

we have the first order condition for labor supply

$$\eta L_t^{\nu} C_t = w \tag{29}$$

In steady state, with the approximation $\bar{C} = w\bar{L}$, this reduces to

$$\eta \bar{L}^{1+\nu} = 1 \tag{30}$$

Up to a second order approximation around the steady state, the effect of variations in L on utility is given by

$$\frac{\eta\nu}{2}\bar{L}^{\nu-1}\,\mathrm{E}(L_t - \bar{L})^2 = \frac{\eta\nu}{2}\bar{L}^{1+\nu}\,\mathrm{E}\left(\frac{L_t - \bar{L}}{\bar{L}}\right)^2 = \frac{\nu}{2}\,\mathrm{E}\left(\frac{L_t - \bar{L}}{\bar{L}}\right)^2 \tag{31}$$

The last term in parentheses is the variance of the proportional changes in L. In our sample, the standard deviation of the detrended unemployment rate is 1.20 percentage points.

With a mean unemployment rate of 5.73 percent, this translates into a standard deviation of the employment rate of $0.0120 \cdot \frac{1}{1+0.0573} \approx 0.0113$, and $\mathrm{E}\left(\frac{L_t-\bar{L}}{\bar{L}}\right)^2 \approx 0.0113^2 \approx 0.000128$. Since the per-period gain of a one-percent increase in C is 0.01, the welfare gain from reducing the variation in labor supply is equivalent to a $0.012808\frac{\nu}{2}$ percent increase in consumption. Since ν is the inverse of the Frisch-elasticity of labor supply, a labor supply elasticity of 0.5 gives $\nu = 2$ and a welfare gain from the elimination of labor fluctuations alone of about 0.013, which is about the same as the welfare gain from stabilization (unchanged average unemployment rate) in the benchmark calibration with $R = 0.2$. With higher labor supply elasticity, it is correspondingly lower.

References

Atkeson, A. and C. Phelan (1994). Reconsidering the cost of business cycles with incomplete markets. *NBER Macroeconomic Annual 9*, 187–207.

Browning, M. and T. F. Crossley (2001). Unemployment insurance benefit levels and consumption changes. *Journal of Public Economics 80*, 1–23.

Carroll, C. D. (2001). A theory of the consumption function, with and without liquidity constraints. *Journal of Economic Perspectives 15*(3), 23–45.

Costain, J. and M. Reiter (2008). Business cycles, unemployment insurance, and the calibration of matching models. *Journal of Economic Dynamics and Control 32*(4), 1120–55.

Elsby, M., B. Hobijn, and A. Sahin (2010). The labor market in the great recession. *Brookings Papers on Economic Activity Spring*, 1–48.

Engen, E. and J. Gruber (2001). Unemployment insurance and precautionary savings. *Journal of Monetary Economics 47*(3), 545–579.

Gruber, J. (1997). The consumption smoothing benefits of unemplyoment insurance. *American Economic Review 87*(1), 192–205.

Haefke, C. and M. Reiter (2011). What do participation fluctuations tell us about labor supply elasticities? IZA Discussion Paper No. 6039.

Hall, R. E. (2009). Reconciling cyclical movements in the marginal value of time and the marginal product of labor. *Journal of Political Economy 117*(2), 281–323.

Imrohoruglu, A. (1989). Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy 97*(6), 1364–83.

Jung, P. and K. Kuester (2011). The (un)importance of unemployment fluctuations for the welfare cost of business cycles. *Journal of Economic Dynamics and Control 35*(10), 1744–1768.

Krebs, T. (2007). Job displacement risk and the cost of business cycles. *American Economic Review 97*(3), 664–686.

Krusell, P., T. Mukoyama, A. Sahin, and A. A. Smith (2009). Revisiting the welfare effects of eliminating business cycles. *Review of Economic Dynamics 12*(3), 393–402.

Krusell, P. and A. A. Smith (1999). On the welfare effects of eliminating business cycles. *Review of Economic Dynamics 2*(1), 245–272.

Lucas, R. E. (1987). *Models of Business Cycles*. Oxford: Basil Blackwell.

Lucas, R. E. (2003). Macroeconomic priorities. *American Economic Review 93*(1), 1–14.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review 95*, 25–49.

Storesletten, K., C. I. Telmer, and A. Yaron (2004). Cyclical dynamics in idiosyncratic labor-market risk. *Journal of Political Economy 112*(3), 695–717.

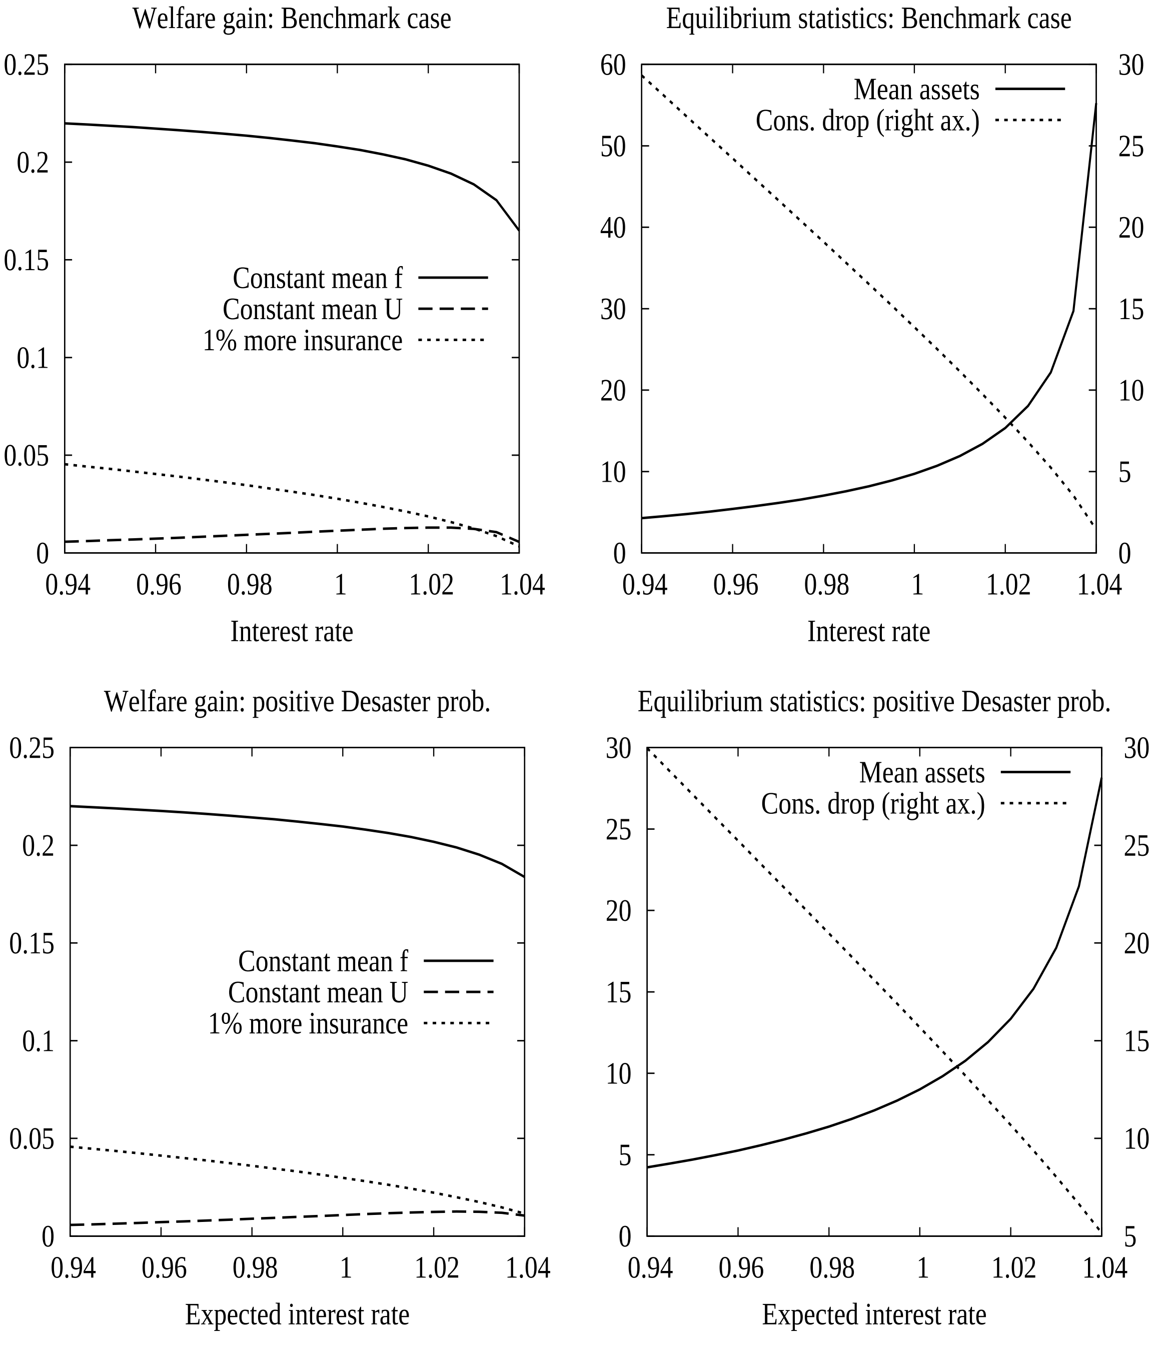

Figure 1: Results for benchmark model. Lower panels: disaster prob.=0.05

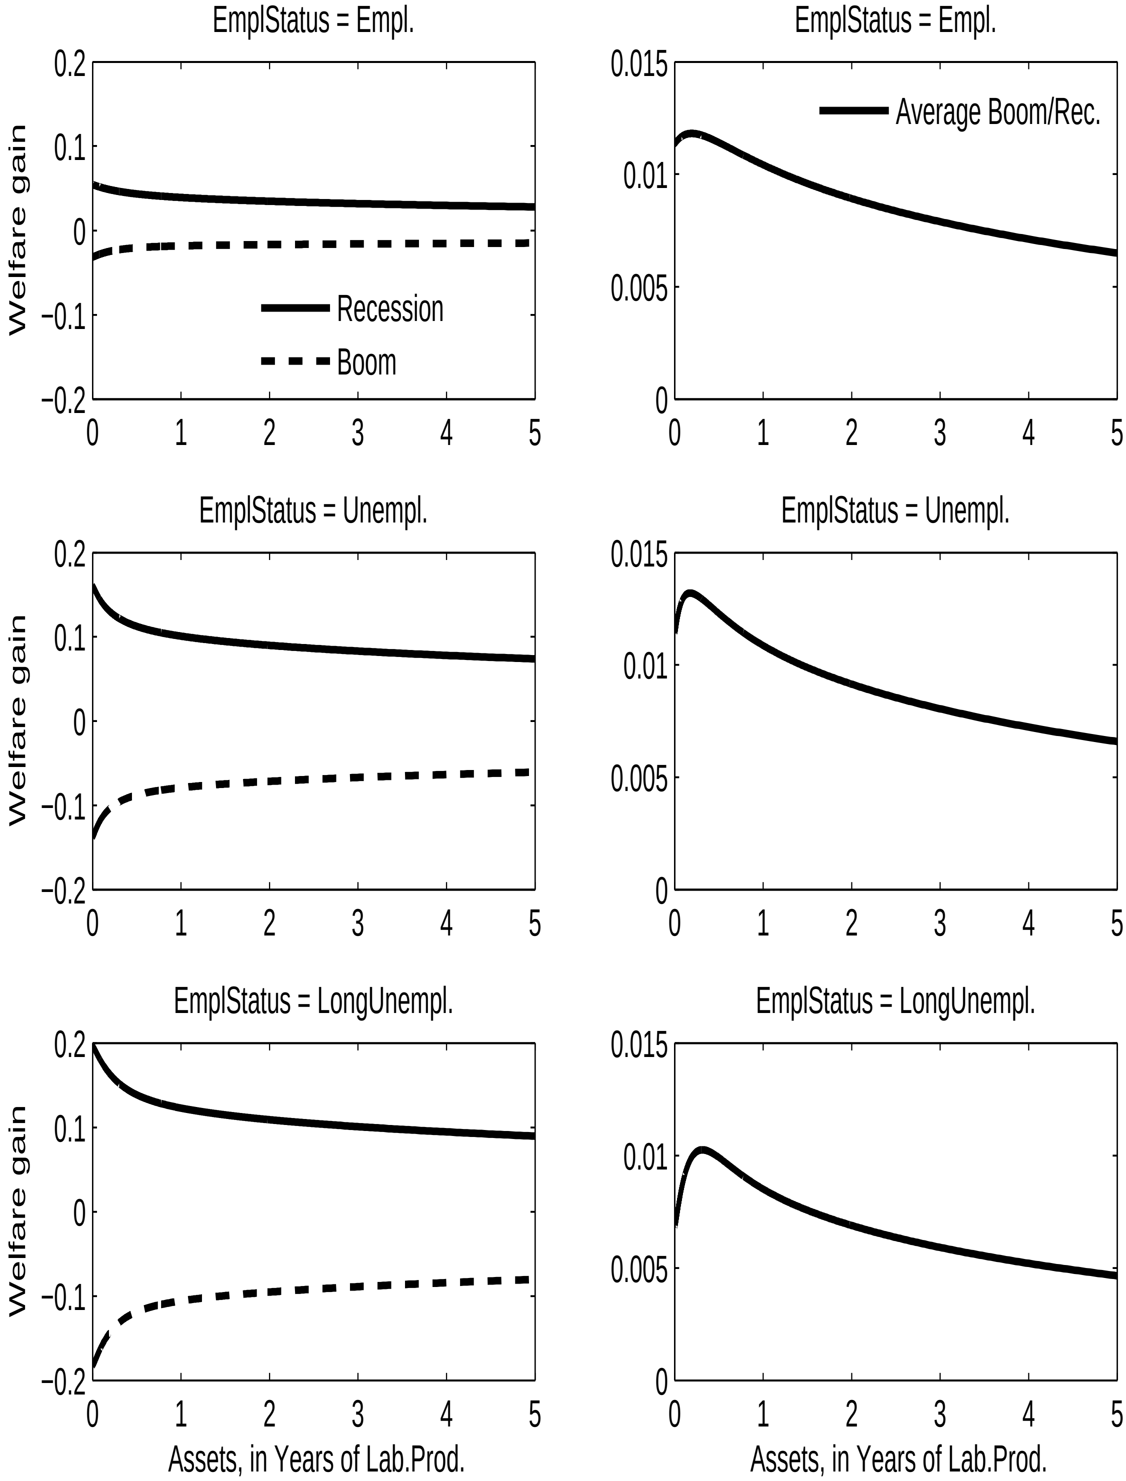

Figure 2: Welfare gains in benchmark case

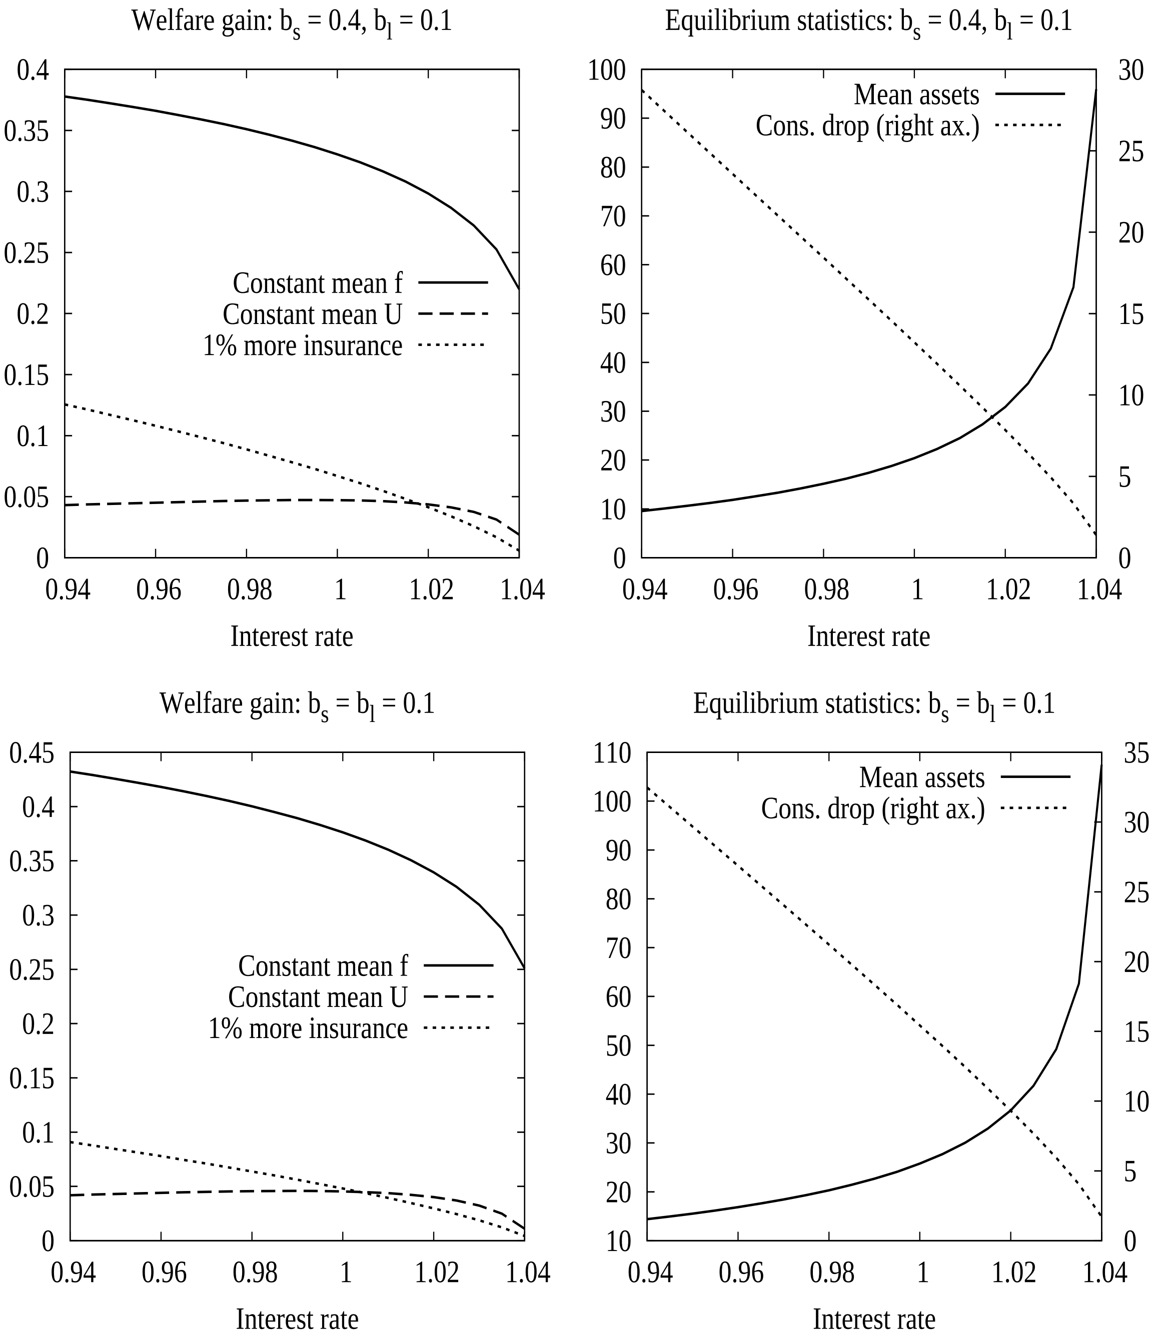

Figure 3: Results for lower unemployment insurance replacement rates

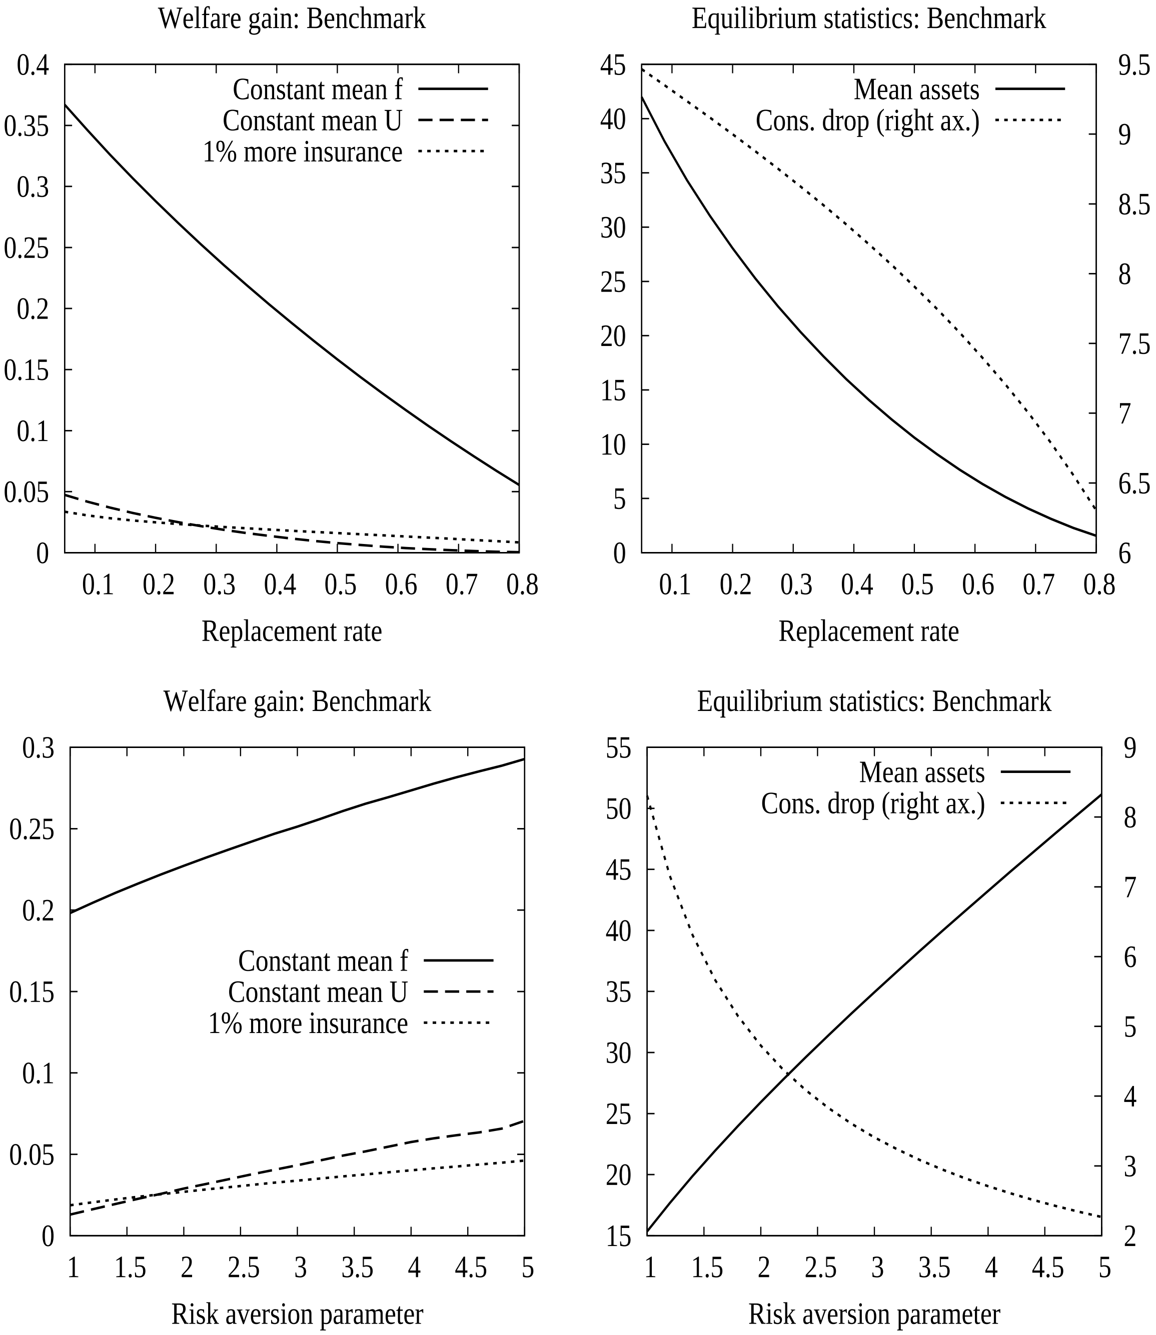

Figure 4: Results for varying replacement rates and risk aversion parameters; $R = 1.02$

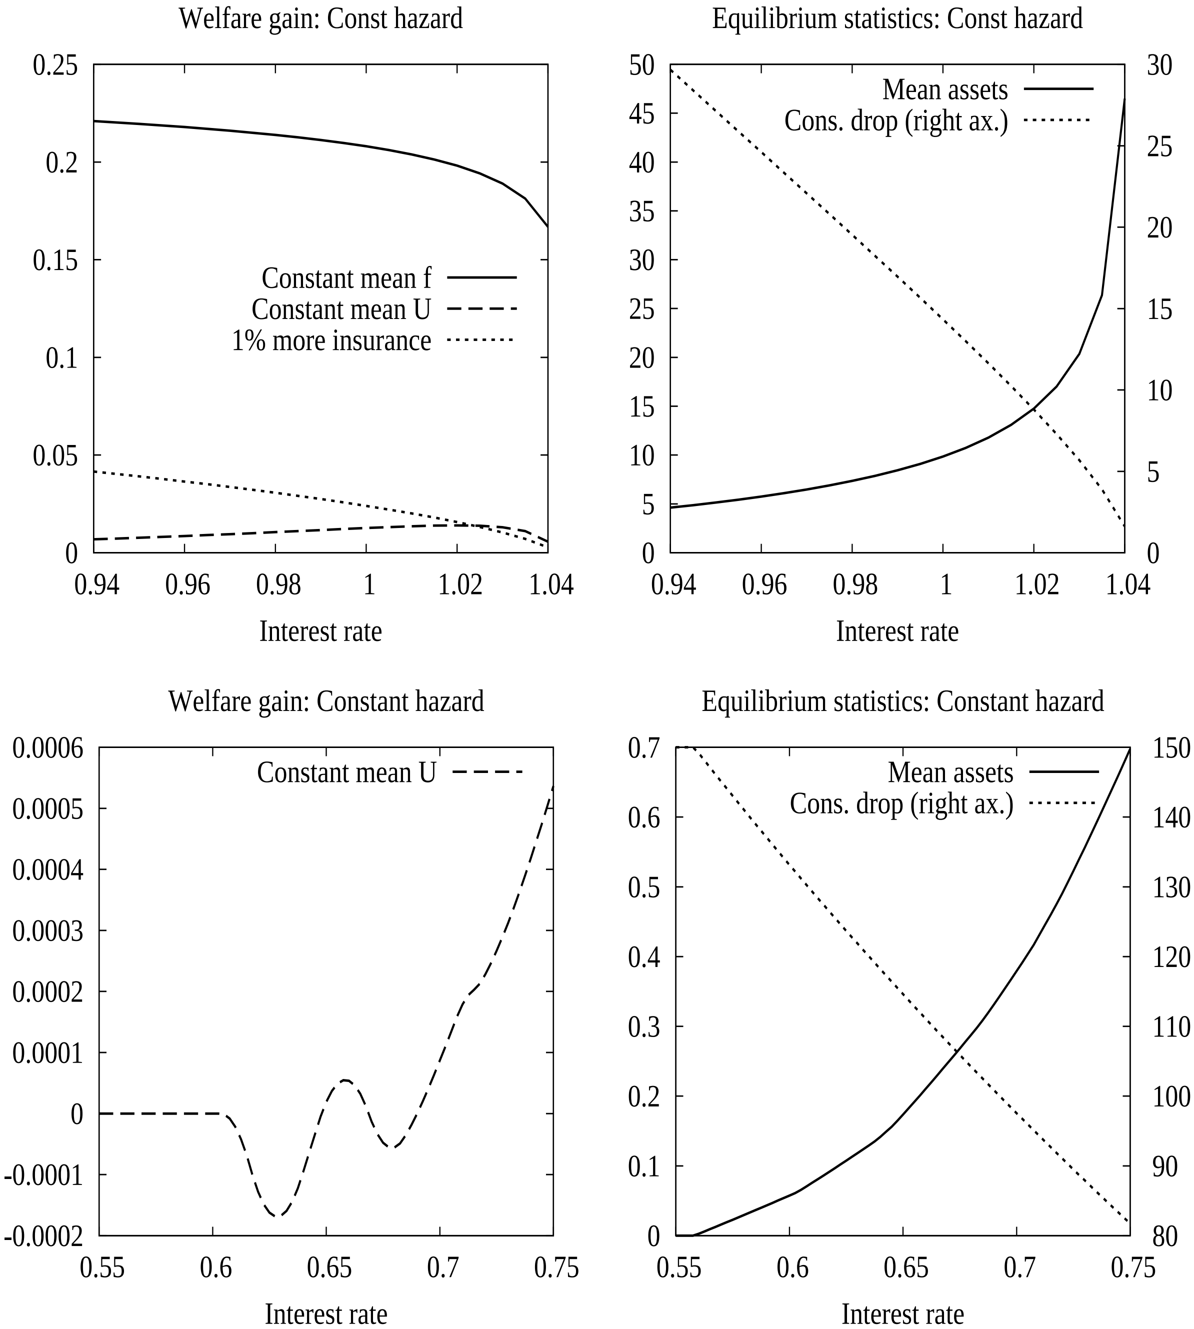

Figure 5: Results with constant unemployment exit hazard