Discount Rates, Learning by Doing and Employment Fluctuations *

Patrick Kehoe[†] Virgiliu Midrigan[‡] Elena Pastorino[§]

October 2015

Abstract

We revisit the Shimer (2005) puzzle in a search and matching model with on-the-job human capital accumulation in which households exhibit preference for consumption smoothing. We parameterize the model so that it accords with the micro-evidence on returns to tenure and experience as well as individual life-cycle earning profiles. We find that employment fluctuations in response to productivity shocks are greatly amplified in this environment.

Keywords: Search and Matching, Employment, Human Capital, Discount Rates.

JEL classifications: E21, E24, E32, J21, J64.

*We thank Eugenia Gonzalez Aguado, Sonia Gilbukh and Sergio Salgado Ibanez for excellent research assistance. Kehoe thanks the National Science Foundation for financial support.

[†]Stanford University, University of Minnesota and University College London, patrickjameskehoe@gmail.com.

[‡]New York University, virgiliu.midrigan@nyu.edu.

[§]Stanford University, University of Minnesota and University College London, elenapastorino1@gmail.com.

1 Introduction

As Shimer (2005) points out, the textbook Mortensen-Pissarides search and matching model predicts fluctuations in labor market variables that are much smaller than those observed in the data. The textbook model misses, however, two ingredients that are empirically important. First, the textbook model assumes that consumers have linear preferences over consumption so that discount rates are constant over time. In the data, however, discount rates fluctuate greatly and are strongly correlated with employment, as Hall (2014) documents. Second, the textbook model assumes no on-the-job human capital accumulation. The model cannot therefore reproduce wage-tenure profiles of the type documented by Buchinsky et al. (2010) or life-cycle earnings profiles of the type documented by Elsby and Shapiro (2012).

This paper shows that versions of the search and matching model in which discount rates fluctuate over the cycle and in which workers accumulate human capital on the job predict much larger fluctuations in employment, up to five times more volatile than in the textbook search and matching model. As in Shimer (2012), the primitive driving force in our model is changes in aggregate productivity. Yet with preferences for consumption smoothing, a drop in aggregate productivity increases the rate at which agents discount the future. The surplus from forming a match thus falls, both because of the drop in productivity, as well as because of the increase in discount rates. As we show, the latter channel is much stronger and accounts for the bulk of the drop in match surplus in our model in response to a productivity shock.

The model we study is a parsimonious extension of the textbook search and matching model. We follow Merz (1995) and Andolfatto (1996) in assuming that households have preferences for consumption smoothing. As these researchers do, we assume that individual workers are organized in families that pool idiosyncratic risks. This is a closed economy so changes in aggregate productivity generate changes in aggregate consumption and thus discount rates. We follow Kehoe et al. (n.d.) who build on the work of Ljungqvist and Sargent (1998) in assuming that an individual worker's productivity changes over time according to laws of motion that depend on whether the worker is employed or not. Employed workers experience increases in productivity on average, while non-employed workers experience declines in productivity on average. We choose parameters characterizing these laws of motion to ensure that the model replicates the Buchinsky et al. (2010) evidence on returns to work as well as the Elsby and Shapiro (2012) evidence on how average earnings change over the workers' lifecycle.[1]

[1]See also Altonji and Shakotko (1987) and Topel (1991), Altonji and Williams (2005), Kambourov and Manovskii (2009), Jeong et al. (2014), Lagakos et al. (2012) for evidence on the returns to tenure and

Allowing for on-the-job human capital accumulation (and out-of-the job human capital depreciation) is an important ingredient for our model's prediction that changes in discount rates greatly reduce the surplus from a match. To see why this is the case, notice that in the textbook search and matching model the surplus from a match is fairly transitory. A non-employed agent has a high likelihood of becoming employed in any given period so being employed confers an advantage for only a short amount of time. The rate at which agents discount the future thus has little effect on the surplus from a match. In our model with human capital accumulation, in contrast, the returns to forming a match are *backloaded.* Since human capital acquired in any given match can be used in future matches, working is a long-lived investment the returns to which are sensitive to changes in discount rates. Similarly, the loss in human capital during a non-employment spell has long-term consequences on an agent's wages throughout its life-time. How much a worker values these losses depends greatly on the discount rate. In our baseline experiment, on-the-job human capital accumulation increases the standard deviation of employment by about 2.5 times.

In our baseline model agents differ in their productivity on the job but are assumed to posses identical technologies for producing goods at home. The least productive workers may thus choose to quit a match and become non-employed. Similarly, some non-employed workers turn down job offers. We show, however, that these endogenous separations and acceptance decisions are not critical in driving our results. Indeed, a version of the model in which an agent can produce at home a constant fraction of what it produces on the job implies that all offers are accepted and no endogenous separations occur. Yet fluctuations in employment in this model are nearly as large as in our baseline model. In this alternative version of the model fluctuations in employment are driven solely by changes in the rate at which firms post vacancies. Since after an aggregate productivity shock the cost of posting a vacancy is unchanged, but the value of forming a match falls due to discounting and the lower productivity, firms post fewer vacancies in equilibrium and employment falls.

Our work is related to a number of papers that address the Shimer puzzle. Hagedorn and Manovskii (2008) suggest an alternative calibration. Hall and Milgrom (2008) propose replacing the Nash bargaining protocol with an alternating offer bargaining protocol. Gertler and Trigari (2009) assume that wages are bargained infrequently. Ljungqvist and Sargent (2015) survey these and other approaches and argue that these alternative models generate larger employment responses to movements in productivity by reducing the fundamental surplus of a match. Our analysis is complementary to these alternative approaches.

experience, as well as average life-cycle earnings profiles.

2 Model

We consider a version of the Mortenssen-Pissarides model with risk-averse consumers. A continuum of such consumers face idiosyncratic shocks to their own productivity as well as aggregate productivity shocks. We let s^t denote the history of aggregate and idiosyncratic shocks. The probability of any history is $\pi(s^t)$. Let $y(s^t)$ denote the income of an individual consumer. This income evolves over time both because of the on-the-job human capital accumulation, the out-of-the-job human capital depreciation, as well as the idiosyncratic productivity shocks. We consider a stochastic OLG environment in which each consumer survives from one period to another with a constant hazard ϕ. A measure $1-\phi$ of consumers replaces those that die each period.

We assume that consumers are organized into representative families that pool all of their members' idiosyncratic risks. Though risk-sharing is perfect within the family, this is a closed economy that cannot trade with the rest of the world, so changes in aggregate productivity affect the family's discount factor. Each member of the family maximizes the present value of the income it receives using the family's discount factor. We first describe the family's problem and then the problem of its individual members.

2.1 Family Problem

The preferences of the family are

$$\max_{c_t} \sum_{t=0} \beta^t u(c_t) \tag{1}$$

where c_t is the consumption of any of its members. Because of full risk-sharing all members enjoy the same amount of consumption, regardless of the idiosyncratic shocks they experience. This type of risk sharing arrangement is familiar from the work of Merz (1995) and Andolfatto (1996).

The family faces a budget constraint

$$c_t + q_t \omega_{t+1} = \omega_t + \int y_{it} \mathrm{d}i + T_t$$

where ω_{t+1} are the family's holdings of a one-period-ahead security that trades at a price q_t. The term $\int y_{it}\mathrm{d}i$ represents the total income of all the members of the family, which includes wages for those members that are employed and home production for those that are not employed. Finally, T_t are the profits net of vacancy posting costs the family receives from its ownership of firms in the economy. As we show below, since firms are owned by families, they discount future flows using the same discount rate that workers use to discount future wages.

Here, for ease of notation, we have represented the insurance arrangements using a family construct. It should be clear that we can also represent these insurance arrangements by having appropriately defined contingent claims markets.

2.2 Individual Employment Problem

An individual member of the family, henceforth worker, chooses whether to work or not, $e(s^t) \in \{0, 1\}$, in order to maximize the present value of income

$$y(s^t) = e(s^t)w(s^t) + (1 - e(s^t))b,$$

which consists of wages $w(s^t)$ when employed, and home production b when non-employed. (Here we do not formally distinguish between unemployment and non-participation but when quantifying the model, we think of non-employment as capturing both states.)

We assume here that the amount non-employed workers earn is constant and thus independent of the worker's characteristics. We relax this assumption in a robustness section below. Since all members have a constant survival hazard ϕ, the objective of the individual member is

$$\max_{e(s^t)} \sum_{t=0} \sum_{s^t} \phi^t Q_t \pi(s^t) y(s^t), \tag{2}$$

where

$$Q_t = \beta^t \frac{u'(c_t)}{u'(c_0)}$$

is the family's marginal valuation of date t output. The one-period-ahead discount factor is thus equal to

$$\frac{Q_{t+1}}{Q_t} = \beta \frac{u'(c_{t+1})}{u'(c_t)}$$

The constraints on this problem are the search and matching frictions we describe below.

2.3 Human Capital and Output

A worker is characterized by a productivity level z which gives the amount of output the worker produces if matched with a firm. We think of z as representing the individual worker's human capital. We let $w_t(z)$ denote the wage received by a worker with productivity z. The wage is the solution to a Nash bargaining problem described below. An employed worker's human capital evolves according to:

$$\log z' = (1 - \rho_e)\mu_e + \rho_e \log z + \sigma_z \varepsilon' \tag{3}$$

where ρ_e determines the persistence, σ_z the volatility and μ_e the mean of the process, while ε' is a Gaussian disturbance with zero mean and unit variance. A non-employed worker's human capital evolves according to a similar AR(1) process with the same volatility but with a different mean (which we normalize to 0) and persistence ρ_u:

$$\log z' = \rho_u \log z + \sigma_z \varepsilon'. \tag{4}$$

We represent below the Markov processes in (3) and (4) as $F_e(z'|z)$ and $F_u(z'|z)$. Finally, newborn workers enter the economy endowed with a draw of z from a log-normal distribution with mean μ_0 and variance equal to the unconditional variance of productivity process for employed.

$$\log(z) \sim N\left(\mu_0, \sigma_z^2/(1-\rho_e^2)\right).$$

To see what these laws of motion imply for an individual's life-cycle earnings profiles and returns to work, suppose that $0 < \mu_0 < \mu_e$ as is the case in our calibration. Then a newborn worker's productivity increases on average over time since its productivity converges to $\exp(\mu_e)$ from below. The speed at which productivity increases is determined by ρ_e. Similarly, if ρ_u is relatively low, non-employed workers experience a reduction in their human capital since their productivity converges to $\exp(0)$ from above. The model thus produces saw-toothed individual productivity profiles: growth while employed and decline while non-employed. If ρ_u is sufficiently low, much of the increase in human capital during an employment spell is eroded during non-employment, implying flat average life-cycle earnings profiles. With a suitable choice of μ_0, μ_e, ρ_e and ρ_u the model can thus replicate simultaneously the Buchinsky et al. (2010) facts on returns to work as well as the Elsby and Shapiro (2012) average life-cycle earnings profiles.

2.4 Matching Technology

We assume a standard aggregate matching function $M(u_t, v_t)$, which represents the measure of matches in a period with u_t non-employed workers and v_t vacancies. We assume that M is given by

$$M(u, v) = Bu^\eta v^{1-\eta},$$

where B is the efficiency of the matching technology. Let $\theta_t = v_t/u_t$ denote the market tightness. The probability that a vacant job is filled in the current period is

$$\lambda_{f,t} = \frac{M(u_t, v_t)}{v_t} = B\left(\frac{u_t}{v_t}\right)^\eta = B\theta_t^{-\eta}.$$

Similarly, the probability that a non-employed worker finds a match is

$$\lambda_{w,t} = \frac{M(u_t, v_t)}{u_t} = B\left(\frac{v_t}{u_t}\right)^{1-\eta} = B\theta_t^{1-\eta}.$$

We assume that individual matches are exogenously destroyed with probability σ. In addition, since home production b is constant while z evolves over time, some matches, those that no longer have a positive surplus, are endogenously destroyed as well.

2.5 Worker Values

The value of an employed worker is given by its current wage as well as the discounted sum of future output:

$$W_t(z) = w_t(z) + \mathbb{E}_t \phi \frac{Q_{t+1}}{Q_t} (1-\sigma) \int_{z'} \max\left[W_{t+1}(z'), U_{t+1}(z')\right] \mathrm{d}F_e(z'|z)$$

$$+\mathbb{E}_t \phi \frac{Q_{t+1}}{Q_t} \sigma \int_{z'} U_{t+1}(z') \mathrm{d}F_e(z'|z),$$

where $\mathbb{E}_t$ is the mathematical expectation over aggregate productivity conditioned on the date t information. Notice the max operator which reflects the decision of whether to continue a relationship, as well as the family's discount factor Q_{t+1}/Q_t with which the workers discount future output. Similarly, the value of a non-employed worker is given by the amount it produces at home, b as well as the discounted sum of future output:

$$U_t(z) = b + \mathbb{E}_t \phi \frac{Q_{t+1}}{Q_t} \lambda_{w,t} \int_{z'} \max\left[W_{t+1}(z'), U_{t+1}(z')\right] \mathrm{d}F_u(z'|z)$$

$$+\mathbb{E}_t \phi \frac{Q_{t+1}}{Q_t} (1-\lambda_{w,t}) \int_{z'} U_{t+1}(z') \mathrm{d}F_u(z'|z)$$

Notice here that the worker may choose to turn down an offer even in the event it matches with a firm which occurs with probability $\lambda_{w,t}$. Finally, notice that the laws of motion for productivity differ depending on whether an agent is currently employed or non-employed.

2.6 Firm Values

Each firm pays a vacancy cost to form a match, produces output when matched, and pays dividends to the family. The objective of each firm is to maximize the discounted value of dividends using the discount factor Q_t of the family it belongs to. A firm that is matched with a worker with productivity z produces $A_t z$ units of output, where A_t is common to all firms and evolves over time according to

$$\log A_{t+1} = \rho_A \log A_t + \nu_{t+1},$$

where ν_{t+1} is an i.i.d. $\mathrm{N}(0, \sigma_\nu^2)$ random variable, the only source of aggregate uncertainty in this economy.

The value of a vacancy is given by the firm's current profits, output net of wages, as well as the present discounted value of future profits:

$$J_t(z) = A_t z - w_t(z) + \mathbb{E}_t \phi \frac{Q_{t+1}}{Q_t} (1-\sigma) \int_{z'} \max\left[J_{t+1}(z'), 0\right] \mathrm{d}F_e(z'|z).$$

Notice again that this formulation captures the possibility that a match no longer yields a positive value to the firm and is thus destroyed.

There is a large number of potential entrants that can post a vacancy cost κ to create a new vacancy. Let $n_t^u(z)$ denote the measure of non-employed workers at the beginning of period t. Let $\tilde{n}_t^u(z) = n_t^u / \int dn_t^u(z)$ denote the distribution of productivity among the non-employed. Free entry drives the expected value of posting a vacancy to 0:

$$0 = -\kappa + \mathbb{E}_t \phi \frac{Q_{t+1}}{Q_t} \lambda_{f,t} \int_{z'} \max\left[J_{t+1}(z'), 0\right] \mathrm{d}F_u(z'|z)\, \mathrm{d}\tilde{n}_t^u(z).$$

The free entry condition pins down the vacancy-unemployment ratio θ_t and thus the flows out of non-employment.

2.7 Wage Bargaining

We assume that wages are renegotiated period by period and are set by a generalized Nash bargaining protocol and thus solve

$$\max_{w_t(z)} \left[W_t(z) - U_t(z)\right]^{\gamma} J_t(z)^{1-\gamma}$$

or

$$\frac{\gamma}{W_t(z) - U_t(z)} = \frac{1-\gamma}{J_t(z)}. \tag{5}$$

Here γ represents the worker's bargaining weight.

2.8 Match Surplus

The dynamics of the match surplus is critical in determining the firm's incentives to post vacancies and thus the dynamics of labor market variables. Let

$$S_t(z) = W_t(z) - U_t(z) + J_t(z)$$

denote the surplus from a match. The Nash Bargain equation (5) implies that workers and firms split this surplus according to the bargaining weights:

$$W_t(z) - U_t = \gamma S_t(z), \qquad J_t(z) = (1-\gamma) S_t(z)$$

Some algebra shows that the surplus from the match satisfies the following functional equation:

$$
\begin{aligned}
S_t(z) \quad &= \quad A_t z - b + \mathbb{E}_t \phi \left[(1-\sigma) - \lambda_{w,t}\gamma\right] \frac{Q_{t+1}}{Q_t} \int \max\left[S_{t+1}(z'), 0\right] \mathrm{d}F_e(z'|z) \\
&\quad + \mathbb{E}_t \phi \frac{Q_{t+1}}{Q_t} \left(\int \left(\lambda_{w,t} \max\left[W_{t+1}(z'), U_{t+1}(z')\right] + (1-\lambda_{w,t}) U_{t+1}(z')\right) \left[\mathrm{d}F_e(z'|z) - \mathrm{d}F_u(z'|z)\right] \right)
\end{aligned}
$$

The first line on the right hand side of this expression is familiar from the textbook model in which the surplus from a match is simply the discounted sum of the difference between what a worker produces when matched and what it produces at home. Notice that the discount factor applied here, $\phi\left[(1-\sigma) - \lambda_{w,t}\gamma\right] \frac{Q_{t+1}}{Q_t}$, is much lower than the subjective discount factor of the family, $\phi \frac{Q_{t+1}}{Q_t}$. This reflects the transitory nature of the match surplus. If the separation probability σ or worker matching probability $\lambda_{w,t}$ is high, an employed worker's advantage relative to a non-employed work is likely to disappear next period, thus reducing the effective discount factor.

Allowing for differences in the rate at which workers accumulate human capital when employed or non-employed gives rise to an additional benefit from being matched, captured by the second line of the above expression. If a worker's productivity grows faster when employed than when non-employed, the surplus from a match also reflects the difference between the two laws of motion, $\mathrm{d}F_e(z'|z) - \mathrm{d}F_u(z'|z)$. This difference is scaled by a weighted average of the value of a matched worker, $\max\left[W_{t+1}(z'), U_{t+1}(z')\right]$ and a non-employed worker, $U_{t+1}(z')$, with weights given by the probabilities that a non-employed worker matches or not.

In contrast to the textbook term, this second component capturing differences in the rate of human capital accumulation for employed and non-employed workers is discounted at a much lower rate given by the rate of time-preference of the worker, $\phi Q_{t+1}/Q_t$. The lower discount rate reflects the fact that the worker's productivity z is transferrable across jobs and a higher productivity acquired when working will be used in all future periods.

2.9 Equilibrium

This is a closed economy so household assets are in zero net supply:

$$\omega_{t+1} = 0$$

Let $n_t^w(z)$ and $n_t^u(z)$ be the measure of employed and non-employed agents at the beginning of period t. The aggregate resource constraint is then

$$c_t + \kappa\theta_t u_t = A_t \int z \mathrm{d}n_t^w(z) + b u_t \tag{6}$$

where

$$u_t = \int \mathrm{d}n_t^u(z)$$

is the measure of non-employed. To understand the resource constraint (6), notice that each employed worker produces $A_t z$ units of output while each non-employed worker produces b units of output. The right hand side thus gives the total output produced in the economy – the sum of market and home production of all workers. The left hand side adds total consumption and investment in new vacancies. The total cost of posting vacancies is equal to κv_t where

$$v_t = \theta_t u_t$$

is the number of vacancies and θ_t is the market tightness.

The expression for match surplus derived above, together with the Nash surplus sharing condition, allows us to re-write the free entry condition as

$$\kappa = \mathbb{E}_t \phi \lambda_{f,t} (1-\gamma) \frac{Q_{t+1}}{Q_t} \int \max\left[S_{t+1}(z'), 0\right] \mathrm{d}F_u(z'|z) \mathrm{d}\tilde{n}_t^u(z)$$

A transitory decline in aggregate productivity reduces the surplus from a match both due to the direct effect on the amount produced by employed workers as well as due to the decline in the discount factor Q_{t+1}/Q_t. The discount factor falls because consumption falls in period t leading to an increase in the marginal utility of consumption. In response firms find it optimal to create fewer vacancies. Market tightness θ_t thus falls, leading to an increase in the firm's matching probability $\lambda_{f,t}$, and a decline in the worker matching probability $\lambda_{w,t}$, thus reducing flows out of non-employment and reducing the employment rate.

An alternative way of seeing why the preference for consumption smoothing reduces employment is to re-write the free entry condition in terms of date 0 prices:

$$Q_t \kappa = \mathbb{E}_t \phi \lambda_{f,t} (1-\gamma) Q_{t+1} \int \max\left[S_{t+1}(z'), 0\right] \mathrm{d}F_u(z'|z) \mathrm{d}\tilde{n}_t^u(z)$$

A reduction in consumption in period t increases its marginal valuation $Q_t \sim u'(c_t)$, thus increasing the firm's cost of investing in new vacancies. As we show below, this force is more potent in our economy with on-the-job human capital accumulation because returns to investment in period t are long-lived – human capital acquired by workers in any given match will be used in all future matches.

3 Quantification and Steady-State Implications

We next discuss how we have chosen parameters for our model, the model's steady state implications, and compare the model's predictions for returns to work and life-cycle earnings profiles with those in the data.

3.1 Parameterization

The model period is one quarter. We set the discount factor equal to $\beta = 0.98^{1/4}$ which corresponds to a 2% annual discount rate in the steady state. We set the survival rate ϕ equal to $1 - 1/160$ implying that workers are in the market for 40 years on average. The probability of separation $\sigma = 0.10$ is set as in Shimer (2005) implying that the average employment spell is equal to 2.5 years. The bargaining weight γ is set equal to 1/2, as is the elasticity of the matching function η.

The utility function is

$$u(c_t) = \frac{c_t^{1-\alpha}}{1-\alpha}$$

A large literature (see Guvenen (2006) and the references therein) finds that the elasticity of intertemporal substitution (EIS) is very low when estimated using data on households, on the order of 0.1 to 0.2. We follow this literature and set this elasticity, $1/\alpha$, equal to 0.2. We report results based on alternative values of α below.

We follow Shimer (2005) in choosing the fixed cost of posting vacancies, κ, so as to ensure that the vacancy-unemployment ratio in the steady state of the model is equal to 1. This is simply a normalization that reflects a choice of units and is otherwise inconsequential.

We have seven additional parameters that are jointly chosen using a moment matching approach in which we minimize the distance between a number of moments in the data and the model. B, the efficiency of the matching function; ρ_e, the persistence of productivity when employed; ρ_u, the persistence of productivity when non-employed; σ_z, the standard deviation of productivity shocks; μ_e, the parameter governing the mean productivity of employed worker and thus the returns to employment; μ_0, the mean of productivity draws for newly born workers, and b, the home production parameter.

At this point we use the numbers from the"Debt Constraints" paper. We restrict $\rho_u = \rho_e$, $\mu_0 = 0$ and choose the remaining five parameters to match i) an employment-population ratio of 0.8, ii) a ratio of home production to average wage of 0.4, iii) a standard deviation of wage changes of 0.21, iv) a standard deviation of wages in new employment spells of 0.94, and v) an average wage growth of continuously employed workers of 5.2%. We are currently exploring an alternative parameterization in which we not only match the Buchinsky et al. (2010) evidence in v), but also the Elsby and Shapiro (2012) estimates of the average lifecycle earnings profiles.

Table 1 summarizes our original parameterization strategy and shows the moments used in our calibration, the parameters that we assigned and those that we endogenously chose. As the table shows, the model matches all 5 parameters exactly given that the model is exactly identified.

4 Aggregate Implications

Next we study the responses of the economy to transitory productivity shocks. We show that employment falls sizably in response to a drop in productivity compared to the textbook model. We then illustrate that the household's preference for smoothing consumption and differences in the rate of human capital accumulation for employed and non-employed workers are critical in explaining the employment responses in our model. Finally, we study a robustness check on our model in which we assume that home production is proportional to a worker's productivity.

4.1 Response to Transitory Productivity Drop

We will eventually solve the model by explicitly introducing aggregate uncertainty and using the Krusell and Smith (1998) approach to aggregate individual decision rules. For now we shut down aggregate uncertainty and study the following experiment. We suppose that the economy is in its ergodic steady state. We then suppose that there is an unanticipated 1% drop in aggregate productivity A_t in period one. This drop is unexpected prior to the first period. Agents have perfect foresight afterwards. We set $\rho_A = 0.9$ so that the model matches the speed of postwar consumption recoveries in the data.

We compute the transitions dynamics by guessing a path for market tightness θ_t and consumption c_t from period one, the first period of the aggregate productivity drop, to some large date T at which we guess and verify that the economy converges to the steady state. Given this guess, we solve for the match surplus, the value functions of workers, non-employed and firms by backward induction. These value functions allow us to calculate the endogenous acceptance and separation thresholds for each period along the transition. Given these thresholds, we calculate the evolution of the measures of employed and non-employed, $n_t^w(z)$ and $n_t^u(z)$, by forward iteration, starting from the steady state in period one. Finally, we use the equilibrium conditions to update our initial guess for market tightness θ_t and consumption c_t and repeat until convergence.

Figure 1a displays the resulting paths of productivity, consumption, output, and the one-period ahead discount rate. The 1% drop in output is accompanied by a 0.95% drop in total (market and home) production and a 0.7% drop in consumption. The drop in consumption leads to a 1% increase in the annualized discount rate from about 6.3% to about 7.3%.

Figure 1b contrasts the response of consumption to that of labor market variables. Employment falls gradually, to a trough of about 0.38% below the steady state. As we show below, the bulk of this employment drop is accounted for by the 3% drop in market tightness, θ_t. The average wage falls by about 0.8%, thus about twice more than employment does.

A comparison of Figures 1a and 1b shows that the employment drop is much more persistent compared to the drop in productivity. The half-life of the employment response occurs about 12 quarters after it reaches its trough, while the half-life of the productivity response is about 6.5 quarters after the trough.

Table 2 summarizes this discussion by reporting the standard deviation and autocorrelation of several variables in our model. We simulate time-series for variables in our model using

$$x_t = \sum_{i=0}^{T} \beta_i^x \nu_{t-i}$$

where ν_t is the innovation to aggregate productivity in period t and β_i^x is the response of variable x to a productivity innovation i periods ago which we retrieve from the transitions to a one-time productivity shock reported in Figures 1a and 1b. The implicit assumption underlying these calculations is that the responses to productivity shocks in our model are constant over time, an assumption that we will dispense with when we explicitly introduce aggregate uncertainty and solve for the equilibrium of our model non-linearly.

Table 2 shows that employment is 0.57 as volatile as productivity, market tightness is about 4 times more volatile, while wages are about 1.1 times more volatile than productivity. Both employment and market tightness are much more persistent than productivity is: their autocorrelation is equal to 0.98 and 0.97, respectively, much greater than the 0.9 assumed autocorrelation of productivity.

We next argue that the bulk of the employment drop in our model is accounted for by a drop in the market tightness θ_t and thus implicitly by the drop in the number of vacancies firms create. To see this, consider a Shimer (2012) − type decomposition of changes in employment. In our model employment evolves over time according to:

$$e_{t+1} = (1 - s_t)e_t + f_t(1 - e_{t-1})$$

where s_t is the separation rate and f_t is the rate at which non-employed workers find a new job. The job finding rate f_t is equal to the product of the probability with which a non-employed worker matches with a firm, $\lambda_{w,t} = B\theta_t^{1-\eta}$ and the probability a_t that a match is accepted:

$$f_t = a_t B\theta_t^{1-\eta}$$

We next construct a counterfactual employment series by setting s_t and a_t equal to their steady state values and only allowing θ_t to change:

$$e^c_{t+1} = (1 - \bar{s})e^c_t + \bar{a}B\theta_t^{1-\eta}(1 - e^c_{t-1}).$$

Figure 3 compare this counterfactual employment series with the actual employment series in the model. The figure shows that the employment drop would have been about 1/3 smaller absent changes in the separation and acceptance rates. Consistent with the data, the bulk the employment drop is thus accounted for by changes in the market tightness.

4.2 Role of IES

We next argue that the household's preference for consumption smoothing is an important ingredient that amplifies the response of employment in our model. Figure 3 shows the response of employment to a 1% productivity drop in variations of our model in which we change the inverse elasticity of intertemporal substitution, α, from 0 (as in Shimer (2005)) to 10. Clearly, employment responses are much greater when the EIS is lower. The maximal drop in employment in the aftermath of a 1% productivity drop is equal to .07% when $\alpha = 0$, thus about one-fifth as large as in our Benchmark model with $\alpha = 5$. The maximal drop in employment is equal to .21% when $\alpha = 2$ and increases to as much as 0.53 when $\alpha = 10$. Intuitively, the increase in discount rates following a productivity shock is larger when the EIS is lower, reducing the returns to creating vacancies and thus the number of vacancies firms post.

Table 2 shows that employment is about 6 times more volatile in our Benchmark model with $\alpha = 5$ compared to an economy with linear preferences (the standard deviation falls from 0.57 vs. 0.09). Reducing α to 2 cuts the volatility of employment by about one-half (from 0.57 to 0.29), while increasing it to 10 makes employment roughly as volatile as productivity. The volatility of wages increases with the inverse EIS, from 1.04 when $\alpha = 0$ to 1.17 when $\alpha = 10$.

4.3 Role of human capital accumulation

We next gauge the role of on-the-job human capital accumulation in driving the employment responses in our Benchmark model. We shut down the Benchmark model's returns to tenure by setting μ_e equal to 0 so that there is not productivity growth on average on the job. We leave the parameter governing the productivity of the matching function, B, unchanged and choose a value of home production, b, so as to ensure that the employment-population ratio is equal to 0.8, as in our Benchmark value. The implied value of b is equal now to 20% of a worker's median wage, half its value in our Benchmark experiment, reflecting the lower returns to work in this version of the model.

Figure 4 contrasts the employment responses in our Benchmark model with that in the economy without returns to work. We set $\alpha = 5$ as in our original experiment. The maximal drop in employment in the economy without on-the-job human capital accumulation is equal to about 0.19%, thus about half of that in our Benchmark model. The employment drop is also much less persistent: employment is about 0.15% below its steady state 40 quarters after the productivity drop in the Benchmark model and essentially at the steady state in the economy without returns to work.

Table 2 shows that the standard deviation of employment is now about 40% as large as in the Benchmark model (0.23 v.s. 0.57) and that employment is much less persistent (the autocorrelation is 0.96 vs. 0.99). Wages are somewhat more volatile (1.35 vs. 1.12 in the Benchmark model), suggesting that learning by doing mitigates the drop in wages following a productivity shock.

In an alternative robustness check, we have considered an alternative specification of the model without returns to tenure. Here we have assumed away all heterogeneity in idiosyncratic productivity, z and have chosen the level of home production to reproduce a ratio of home production to wages of 0.4. In this alternative specification the maximal employment drop is equal to 0.22%, thus very close to that in the economy described above.

4.4 Economy with Proportional Benefits

In our Benchmark economy the amount non-employed workers produced is a constant number, independent of productivity. This implies that a portion of fluctuations in employment is due to changes in the rate at which unproductive matches are dissolved as well as the acceptance rate of new matches. We next study an economy in which unemployed workers' productivity is equal to a constant fraction of what they would have produced on the job:

$$b(z) = \lambda z,$$

where we choose λ, the replacement ratio, to ensure that the median non-employed worker produces 40% of what an employed worker earns in wage income.

Figure 5 shows that the employment response in this version of the model is about two-thirds as large as in our Benchmark model. This suggests that fluctuations in the acceptance and separate rate in our Benchmark setup play an important, but not critical role. Table 2 shows that the standard deviation of employment falls by about one-third (0.38 vs. 0.57), even though market tightness is almost equally volatile in the two economies.

5 Conclusions

We have revisited the Shimer (2005) puzzle in a search and matching model with on-the-job human capital accumulation in which households exhibit preference for consumption smoothing. We parameterize the model so that it accords with the micro-evidence on returns to tenure and experience. We find that employment responds much more in our environment to productivity shocks. In this setup, a shock to productivity reduces the flow surplus from a match but also changes the rate at which workers and firms discount future returns to employment. With sufficiently high returns to work, of the type documented by Buchinsky et al. (2010) using micro-level data on wages, the model implies that this discounting effect is quite potent, substantially magnifying employment responses.

References

Altonji, Joseph G and Nicolas Williams, "Do Wages Rise With Job Seniority? A Reassessment," *Industrial & Labor Relations Review*, 2005, *58* (3), 370–397.

_ **and Robert A Shakotko**, "Do Wages Rise With Job Seniority?," *The Review of Economic Studies*, 1987, *54* (3), 437–459.

Andolfatto, David, "Business Cycles and Labor Market Search," *The American Economic Review*, 1996, pp. 112–132.

Buchinsky, Moshe, Denis Fougère, Francis Kramarz, and Rusty Tchernis, "Interfirm Mobility, Wages and the Returns to Seniority and Experience in the United States," *The Review of Economic Studies*, 2010, *77* (3), 972–1001.

Elsby, Michael WL and Matthew D Shapiro, "Why Does Trend Growth Affect Equilibrium Employment? A New Explanation of an Old Puzzle," *The American Economic Review*, 2012, pp. 1378–1413.

Gertler, Mark and Antonella Trigari, "Unemployment Fluctuations with Staggered Nash Wage Bargaining," *Journal of Political Economy*, 2009, *117* (1), 38–86.

Guvenen, Fatih, "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective," *Journal of Monetary Economics*, 2006, *53* (7), 1451–1472.

Hagedorn, Marcus and Iourii Manovskii, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," 2008.

Hall, Robert E, "High Discounts and High Unemployment," Technical Report, National Bureau of Economic Research 2014.

_ **and Paul R Milgrom**, "The Limited Influence of Unemployment on the Wage Bargain," *The American Economic Review*, 2008, pp. 1653–1674.

Jeong, Hyeok, Yong Kim, and Iourii Manovskii, "The Price of Experience," Technical Report, National Bureau of Economic Research 2014.

Kambourov, Gueorgui and Iourii Manovskii, "Occupational Specificity of Human Capital," *International Economic Review*, 2009, *50* (1), 63–115.

Kehoe, Patrick, Virgiliu Midrigan, and Elena Pastorino, "Debt Constraints and Employment."

Krusell, Per and Anthony A Smith Jr, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 1998, *106* (5), 867–896.

Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman, "Experience Matters: Human Capital and Development Accounting," Technical Report, National Bureau of Economic Research 2012.

Ljungqvist, Lars and Thomas J Sargent, "The European Unemployment Dilemma," *Journal of political Economy*, 1998, *106* (3), 514–550.

_ **and** _ , "The Fundamental Surplus in Matching Models," *Working Paper*, 2015.

Merz, Monika, "Search in the Labor Market and the Real Business Cycle," *Journal of Monetary Economics*, 1995, *36* (2), 269–300.

Shimer, Robert, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American economic review*, 2005, pp. 25–49.

_ , "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 2012, *15* (2), 127–148.

Topel, Robert, "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority," *Journal of Political Economy*, 1991, pp. 145–176.

Table 1: Parameterization

Panel A: *Moments*

	Data	Model
Moments used in calibration		
Fraction employed	0.80	0.80
Home production / mean wage	0.40	0.40
Std. dev. wage changes	0.21	0.21
Std. dev. initial wages	0.94	0.94
Mean wage growth of employed	0.052	0.052
Additional model predictions		
Fraction workers with $w < b$		0.181
Fraction voluntary separations		0.002
Probability worker matches		0.595
Fraction rejected matches		0.278
Wage drop after non-employment spell		1.9%
Wage drop if non-employed 1 year		6.1%
Wage drop if non-employed 2 years		8.8%

Panel B: *Parameters*

Endogenously chosen	
B, efficiency matching function	0.595
ρ_z, persistence	$0.952^{1/4}$
σ_z, volatility	0.112
μ_z, mean efficiency employed	2.82
b, home production (rel. mean output)	0.27
Assigned	
period length	1 quarter
β, discount factor	$0.94^{1/4}$
$1 + r$, interest rate	$0.96^{-1/4}$
ϕ, survival probability	1-1/160
σ, probability of separation	0.10
α, inverse IES	5
η, elasticity matching function	0.50
γ, worker's bargaining share	0.50
κ, vacancy cost (rel. mean output)	0.90

Table 2: Business Cycle Implications

	Benchmark	Vary EIS			No returns to work	Proport. benefits
		$\alpha = 0$	$\alpha = 2$	$\alpha = 10$		
Std. dev., rel. productivity						
employment, e	0.57	0.09	0.29	0.98	0.23	0.38
market tightness, θ	4.15	0.82	2.31	6.51	2.28	3.85
average wage	1.12	1.04	1.10	1.17	1.35	1.19
market production, y	1.39	1.06	1.19	1.76	1.16	1.39
consumption, c	1.05	1.09	1.04	1.24	1.07	1.08
Autocorrelation						
productivity, a	0.90	0.90	0.90	0.90	0.90	0.90
employment, e	0.98	0.95	0.97	0.99	0.96	0.98
market tightness, θ	0.97	0.91	0.96	0.98	0.96	0.97
average wage, e	0.95	0.90	0.93	0.98	0.93	0.95
market production, y	0.95	0.91	0.93	0.97	0.92	0.95
consumption, c	0.95	0.92	0.94	0.98	0.93	0.96

Figure 1a: Effect of Transitory Productivity Drop

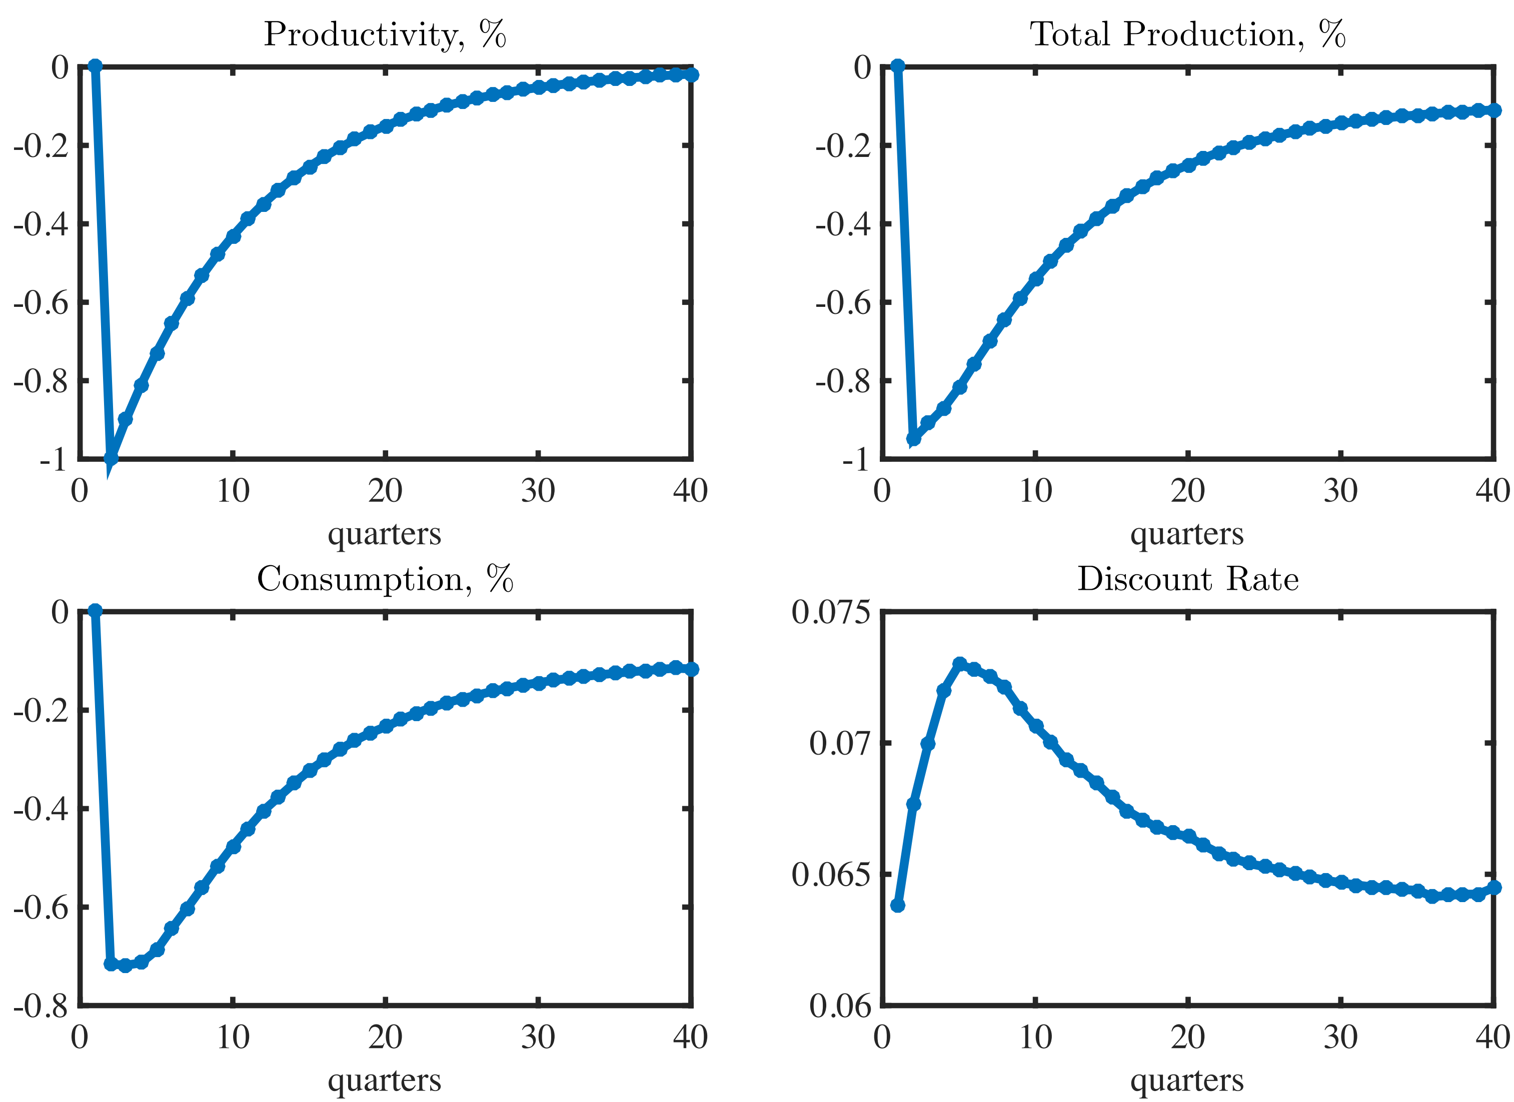

Figure 1b: Effect of Transitory Productivity Drop

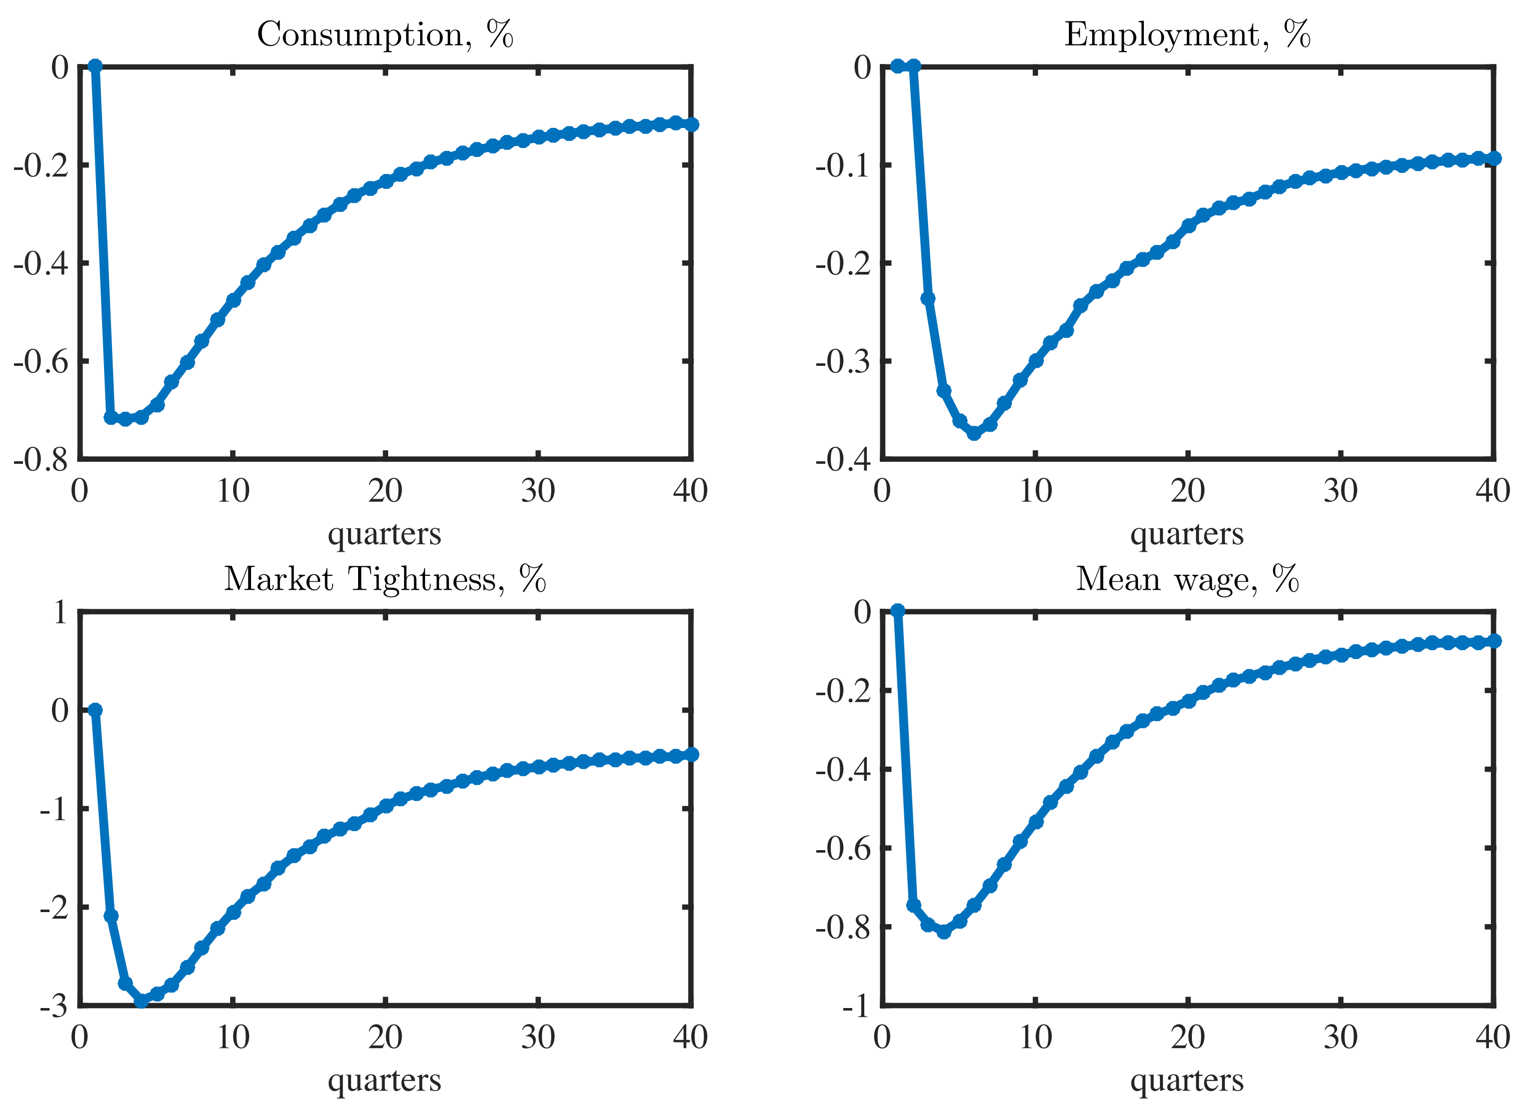

Figure 2: Shimer (2012) decomposition

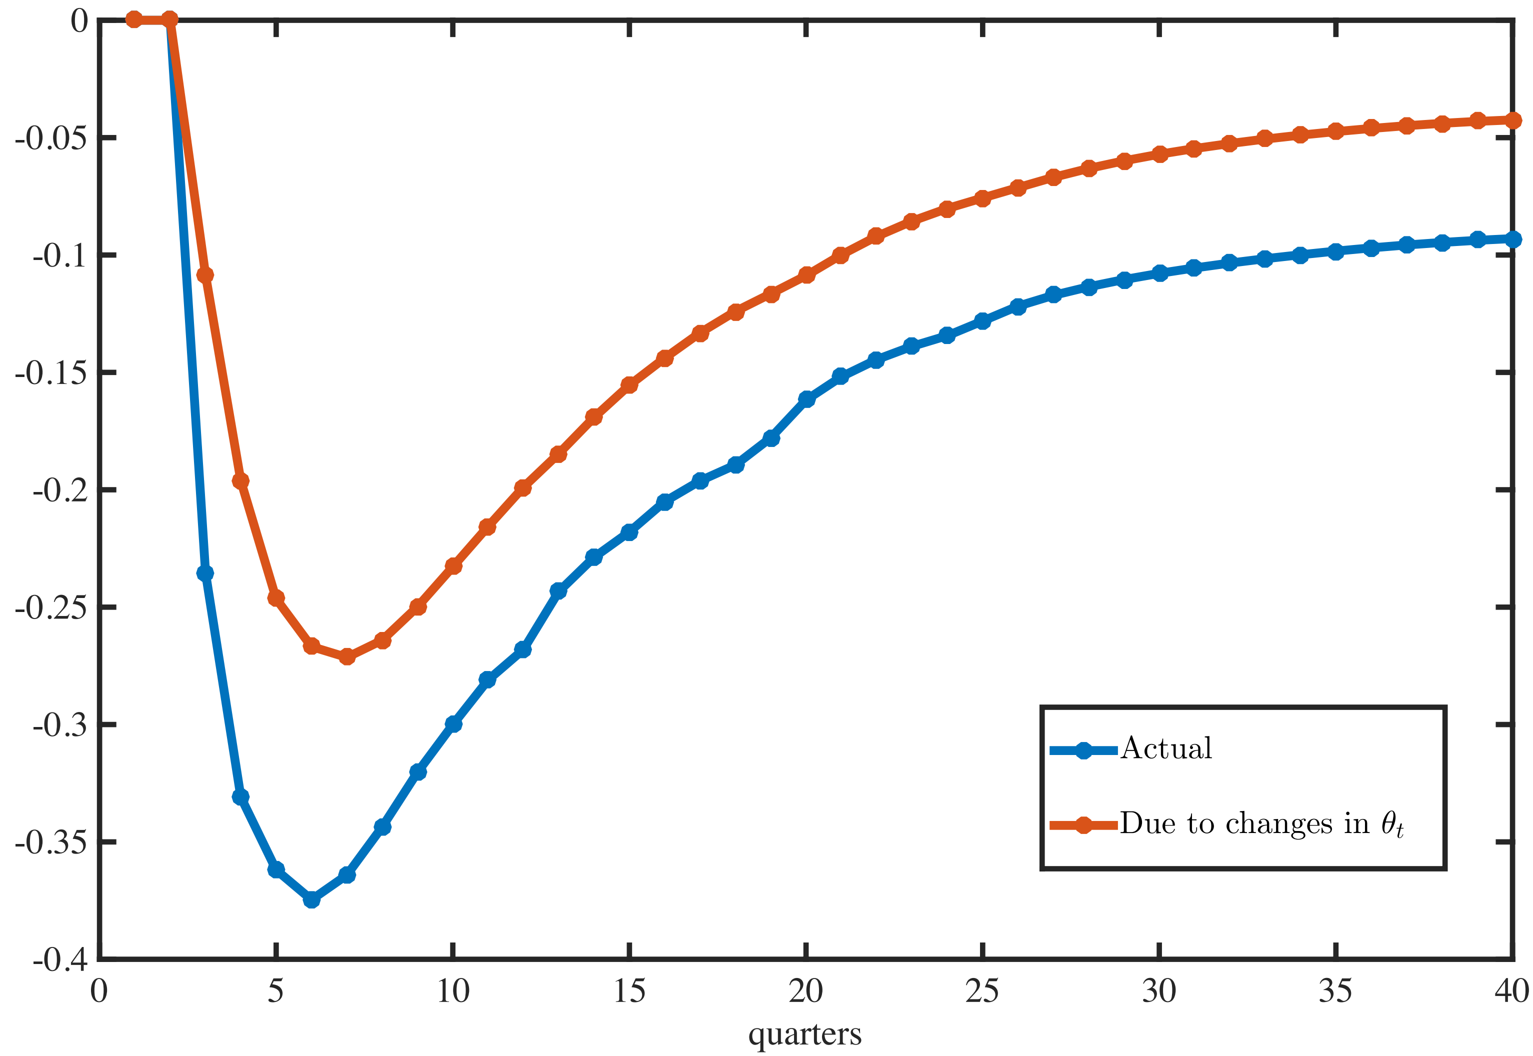

Figure 3: Employment Responses as a function of the EIS

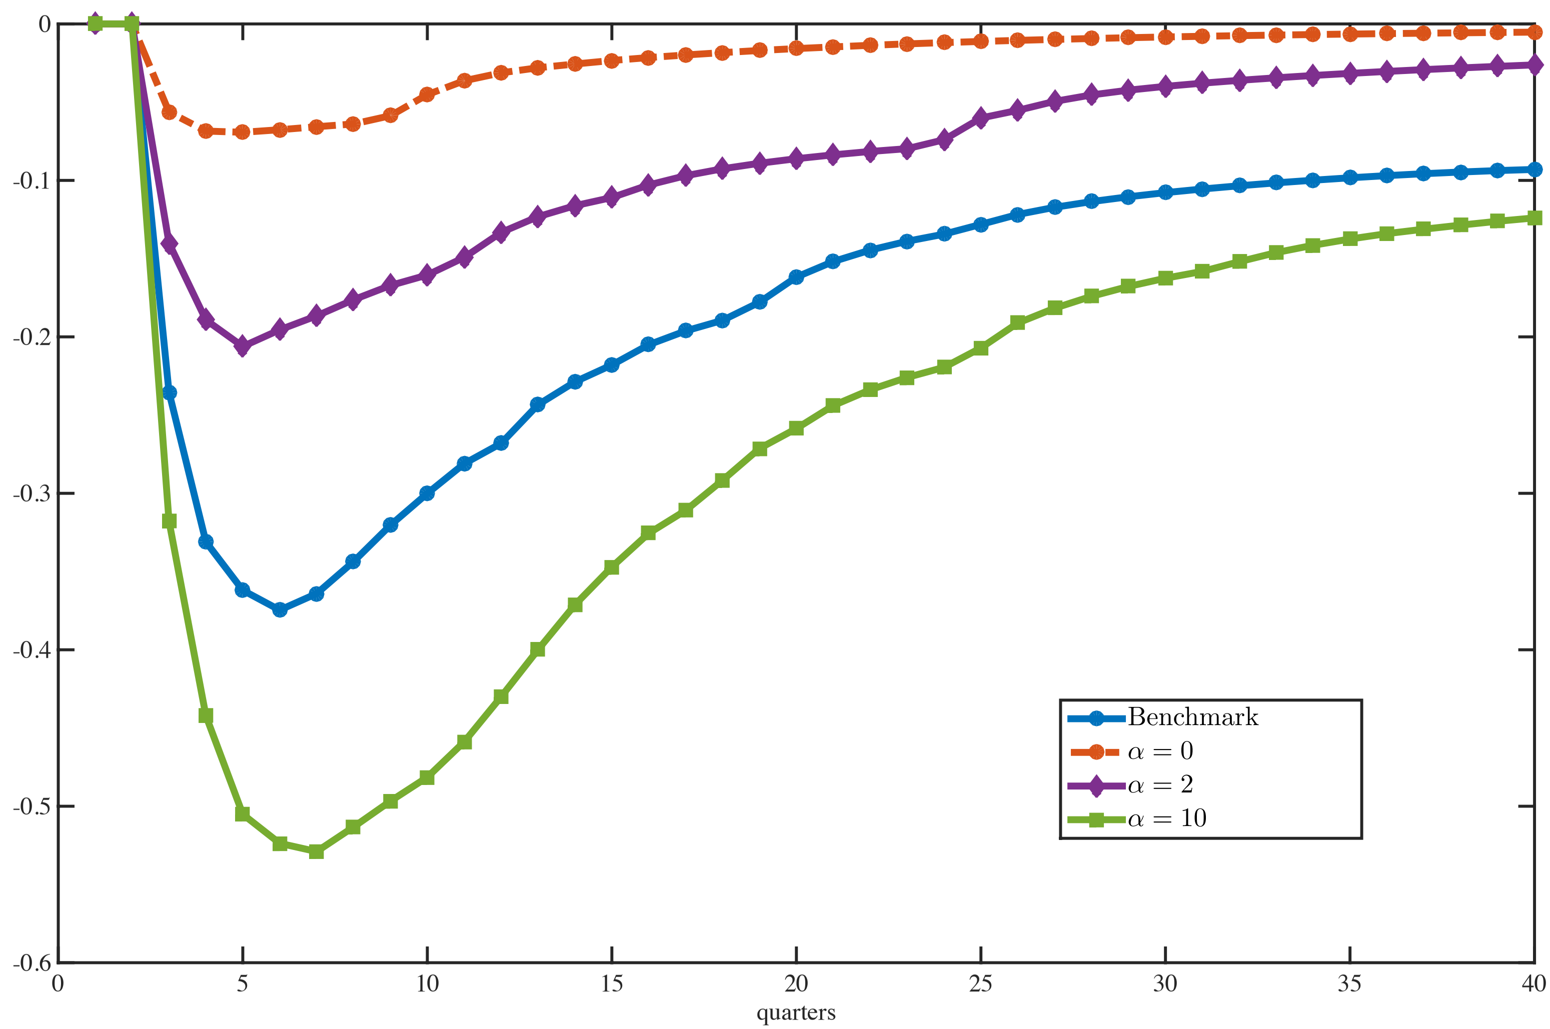

Figure 4: Employment Responses Absent Returns to Work

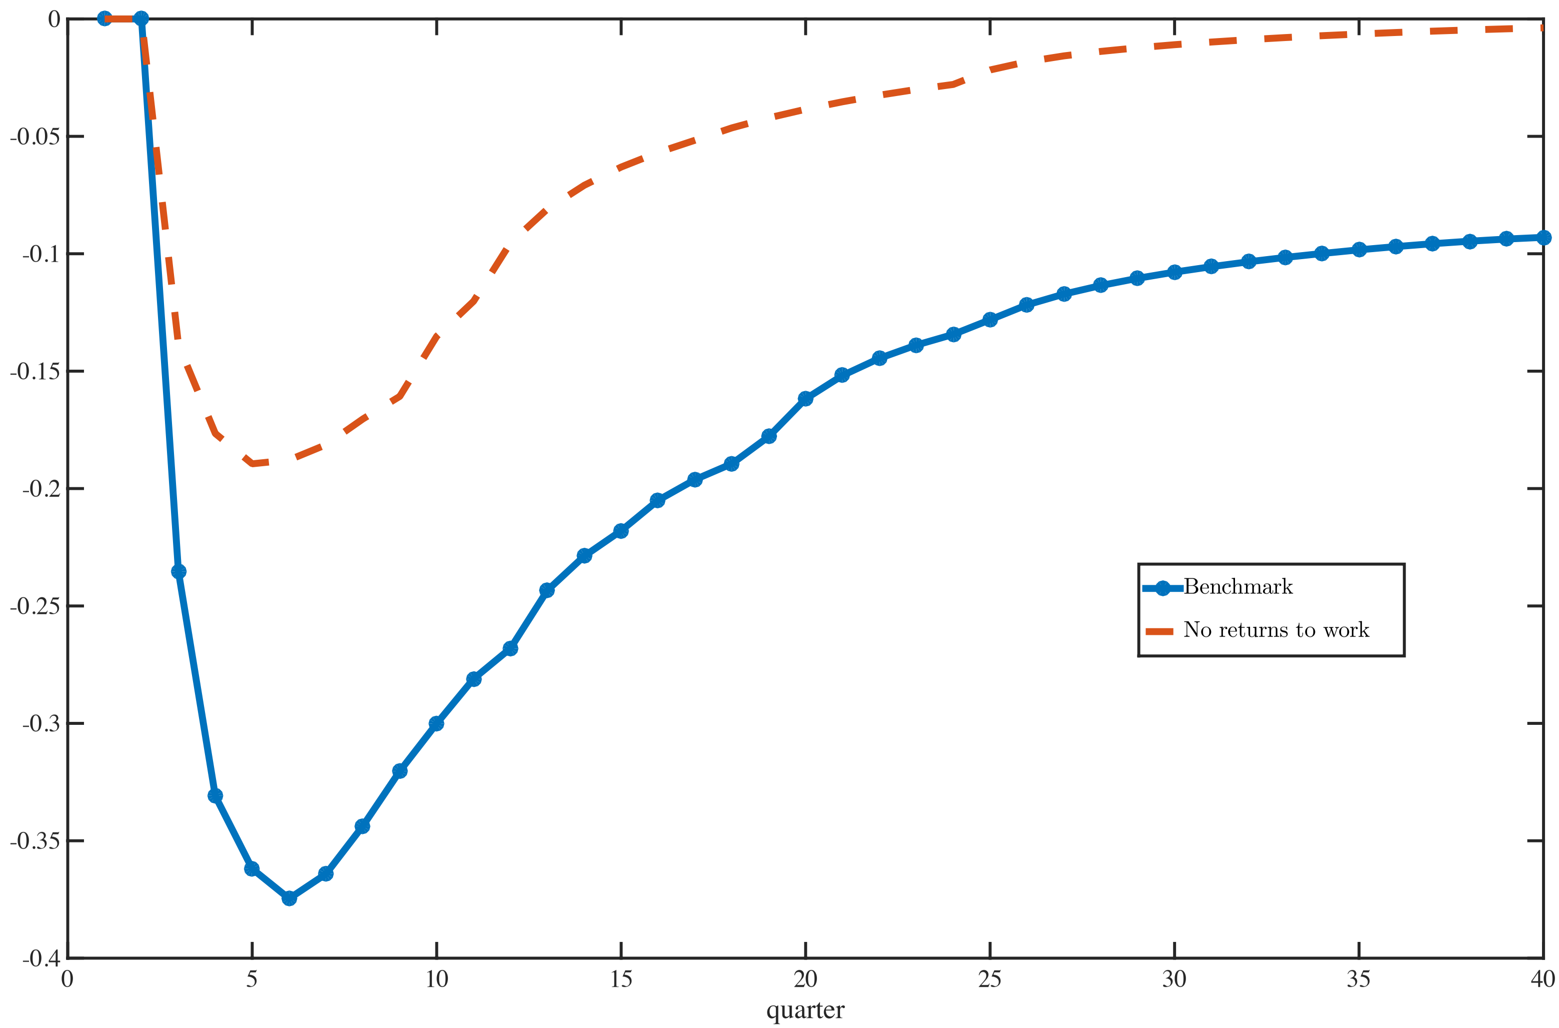

Figure 5: Employment Responses in Economy with Proportional Benefits