Unemployment and Labor Productivity Co-movement: the Role of Firm Exit

Miroslav Gabrovski[b], Mario Rafael Silva[a]

[a] *Hong Kong Baptist University*
[b] *University of Hawai'i Manoa*

Abstract

The Diamond-Mortensen-Pissarides model implies a nearly perfect correlation between labor productivity and unemployment/vacancies, yet the relationship in the data is mild. We show that incorporating sunk entry costs and vacancy creation in an otherwise standard setup can reconcile the discrepancy. Sunk costs cause vacancies to be a positively valued, predetermined variable. If the destruction shock is infrequent, then most vacancies were created in the past, and hence the number of vacancies in the market correlate more closely with past than current labor productivity. Provided the destruction shock is calibrated to match either micro-level evidence on product destruction and firm exit rates or commonly used values in the growth literature, the model reproduces the empirically observed mild correlation between productivity and unemployment without breaking the strong negative co-movement between unemployment and vacancies.

JEL Classification: E24, E32, J63, J64

Key words: job destruction; entry costs; unemployment; aggregate fluctuations

[1]We would like to thank Michael Choi, Ioannis Kospentaris, Victor Ortego-Marti, Guillaume Rocheteau, and other participants of the UC Irvine Macroeconomics Workshop.

1. Introduction

The Diamond-Mortensen-Pissarides (DMP henceforth) has been the primary workhorse for studying the business cycle properties of unemployment, labor market tightness, and vacancies. Following Shimer (2005), much of the literature has focused on reproducing the empirically observed volatilities of these three labor market variables.[2] The majority of these past studies follow the tradition of the business cycle literature and utilize technology shocks as the fundamental driving force behind business cycles.[3] As a consequence, the model predicts near-perfect cross-correlation between productivity and the labor market variables. This is in sharp contrast with the data which features a mild correlation between productivity and the labor market variables.[4] Yet the reasons behind this discrepancy between the model and the data have been under-explored in the literature.[5]

Explaining the co-movement between productivity and other real macroeconomic aggregates has been a major objective of real business cycle theory. One major drawback in the workhorse real business cycle model is that all unemployment is voluntary and there is no room for the co-existence of unemployed workers and vacant jobs at the same time. To address this shortcoming, search-and-matching models were built to better capture the dynamics of the labor market. Whereas much effort has been devoted to addressing inconsistencies between the model predicted volatility of labor market variables and their empirical counterparts, the discrepancy in the theoretical and empirical cross-correlations has been relatively understudied. Given the original motivation behind studying real business cycle theory, and the focus of search-and-matching models on the dynamics of the labor market, we see the discrepancy between empirical and theoretical cross-correlations of labor market

[2]Some notable examples include Hagedorn and Manovskii (2008), Hall and Milgrom (2008), Pissarides (2009), Ljungqvist and Sargent (2017).

[3]There are some notable exceptions. For example, Barnichon (2010) features a demand shock as well and both Shimer (2005) and Coles and Moghaddasi Kelishomi (2018) also feature a separation shock.

[4]In fact, Barnichon (2010) documents that the unemployment-productivity relationship shifts from moderately negative between 1948-1984 to moderately positive thereafter and explains the shift in terms of a change in the composition of technology and non-technology shocks as well as nominal rigidity.

[5]Notable exceptions include Barnichon (2010) and the papers cited in the “Related Literature” section of this paper.

variables as a natural research question that should be addressed.

There are two major possibilities as to why the baseline DMP model falls short when it comes to reproducing the cross correlation between productivity and the labor market variables. First, the dynamics of the labor market variables may be more closely connected to other types of shocks, such as demand shocks. If this is the case, then it may be more prudent to analyze the DMP model incorporating demand shocks as the fundamental source of uncertainty in the model. Second, it may be the case that unemployment/vacancy/market tightness dynamics are mainly driven by productivity shocks, but there exist a mechanism in the economy which endogenously decreases the correlation of these variables with productivity. If this mechanism is missing from the model, the predictions from the theory would not be able to match the empirical cross correlations. More importantly, the model would miss key labor market dynamics and potentially lead to false policy recommendations.

In this paper, we explore the latter possibility and show that the discrepancy can be reconciled in an extension of the DMP model with sunk vacancy creation costs. The sunk costs cause vacancies to be a positively valued asset in equilibrium. Moreover, the vacancy stock is predetermined and depends on the likelihood of a destruction/obsolescence shock as well as the expectations of future profits. Under a high destruction rate, vacancies do not survive as long, so most of the vacancy pool is composed of newly created job opportunities. As a result, vacancies — and by extension the market tightness and unemployment – are highly correlated with current productivity, which governs the expectations for future profits. Alternatively, if the shock is less frequent, then most of the current vacancies were created in the past. In that case, labor market variables co-move relatively more with past productivity and relatively less with the current state of technology. Our numerical exercises show that if the destruction shock is calibrated to be consistent with (i) micro-level evidence on product destruction and firm exits or with (ii) values usually used in the growth literature the model can reproduce the empirically observed mild correlation between productivity and the labor market variables (-0.41 in the data versus -0.33 in the benchmark calibration), while still maintaining the high cross-correlation between labor market variables.

The idea of introducing sunk vacancy creation costs into the DMP model is not novel to our paper. For example Diamond (1982), Fujita and Ramey (2007), Shao and Silos (2013), Coles and Moghaddasi Kelishomi (2018), and Mercan and Schoefer (2020) have all studied a search model à la DMP with sunk costs of vacancy creation. We contribute to the literature via a further exploration of the numerical properties of this existing class of models. Given this research question, we keep our theoretical environment close to the existing literature to facilitate comparison. In particular, our environment follows closely Coles and Moghaddasi Kelishomi (2018). Creating a new vacancy requires an up-front investment in a new technology, with the cost of this sunk investment being a random draw from a known exogenous distribution.[6] Consequently, each vacancy carries a positive asset value. As a result, following a separation with a worker, the firm strictly prefers to keep its vacancy open. We distinguish between this separation shock and a destruction shock. Following the latter, the vacancy is lost alongside the job. We interpret this as a shock that makes the firm's product obsolete or as a shock that destroys the firm's business opportunity. That is, if the destruction shock hits, the firm exits the labor market altogether. The only difference between our model and Coles and Moghaddasi Kelishomi (2018) is that in theirs all job loss is due to destruction. Thus, our theoretical environment nests theirs if we set the separation shock to zero. As compared to the model in our paper, the one in Shao and Silos (2013) differs mostly due to the existence of capital in their economy as the driving source of congestion in vacancy creation — the more firms enter, the higher the demand for capital which increases its rental price and, as a consequence, the equilibrium cost of vacancy creation. Similarly, the model in Mercan and Schoefer (2020) incorporates job-to-job flows, which we abstract from.

The only source of uncertainty in the economy is a productivity shock. A negative shock lowers future expected profits, which dampens vacancy creation. As a result, unemployment rises and market tightness falls. In the standard DMP setup this is the only effect of the

[6] Thus specified entry in the model behaves similarly to an environment where there is entry in congestion. That is, the entry cost each firm has to pay in order to open a vacancy is increasing in the number of entrants. As Fujita and Ramey (2007) show congestion in entry leads firms to spread out their response to productivity shocks over several periods. This gradual propagation of shocks also serves to reduce the contemporaneous correlation between productivity and the labor market variables. Although, on its own, it is not enough to reproduce the mild correlations observed in the data.

shock on labor market variables, thus they are near-perfectly correlated with productivity. In the current setting, however, vacancies are long-lived assets with a positive value, i.e they are a stock variable. Thus, the number of vacancies in the market is correlated not only with the current technology shock, but also with past ones that may have (dis)incentivized entry in previous periods. The longer the expected life of a vacancy, the longer the history of past shocks that affects the pool of vacancies today, the larger the fraction of vacancies in the pool that were created in previous periods, and the lower the correlation of vacancies with current productivity. Thus, the magnitude of the cross-correlations between labor market variables and productivity in the model depends on how long-lived vacancies are, i.e. the size of the destruction shock.

Consequently, disciplining the size of the destruction shock plays a central role. In our model, initial vacancy creation requires an up-front sunk investment on the part of the firm in order to capture a business opportunity. The destruction shock is therefore more closely linked to the loss of the business opportunity rather than alternative reasons for which firm and worker may separate. It is thereby appropriate to calibrate the destruction shock with data on product obsolescence or firm exit and the separation shock with data on job loss and separations. Furthermore, there is empirical evidence that product turnover and firm exit are tightly linked to employment turnover, e.g. Bernard et al. (2010), Lee and Mukoyama (2015). Indeed, this is quite intuitive: when firms retire production lines and products this causes an organizational change, which induces a shift in the labor needs of the company. This change would likely both reshuffle workers to different roles within the company and cause job and vacancy destruction. Using micro-level evidence Broda and Weinstein (2010) find an annual destruction rate of 3%. This value is consistent with the obsolescence rate of 3% that Comin and Gertler (2006) calibrate using balanced growth path restrictions from the U.S. data. This value is also within the ball park of the $5\% - 6\%$ annual destruction rate implied by the estimates in Bernard et al. (2010).[7] Alternatively one may turn to the data on firm exit, which the growth literature has commonly used.[8] In particular, Broda and Weinstein (2010)

[7] In particular, the authors estimate that between 26% and 29% of firm output is accounted for by products about to be dropped in the next five years.

[8] See, for example Bilbiie et al. (2012) and Gabrovski (2019).

find firm exit to be about 10% annually. This number is also consistent with the estimates of job destruction in Lee and Mukoyama (2015), which the authors obtain from analyzing data on plant-level entry and exit over the business cycle. Since destruction is likely to be caused by both firm exit and product turnover we take the mid point of the estimates and calibrate our benchmark to be consistent with the 6% annual destruction that Bernard et al. (2010) find in the data. We also perform a robustness check utilizing the 10% rate implied by the evidence on firm exit alone. In both cases the model reproduces the mild correlation between productivity and all three of the labor market variables of interest (vacancies, the market tightness, and unemployment) reasonably well. Moreover, the cross-correlations between the labor market variables themselves remain strong.

Related Literature. Existing studies within the labor search literature have targeted job destruction rates in alternative ways. For example, Coles and Moghaddasi Kelishomi (2018) set the separation rate to zero and attribute all job separations in the data to destruction. Fujita and Ramey (2007), on the other hand, identify the destruction and separation rates using evidence on total job losses from the Business Employment Dynamics (BED) program, coupled with the usual moments of the job-finding rate and steady state unemployment. A similar strategy is used by Shao and Silos (2013) who pin down the destruction rate using evidence on total job losses from Shimer (2005) and a moment restriction for the steady state level of unemployment. Mercan and Schoefer (2020) use survey evidence from German employers to distinguish between hiring aimed at replacing workers who have quit (replacement hiring) and new job creation to separately calibrate their destruction and separation rates.

We regard our calibration as a viable alternative to address several shortcomings in the previous literature. In particular, we follow Coles and Moghaddasi Kelishomi (2018) in our theoretical framework and allow the congestion of vacancy creation in the model to follow a flexible functional form that we calibrate using empirical evidence. The model, with this additional degree of freedom, can match the aggregate job loss rate, unemployment, and job finding rate in the steady state. There is no need to restrict the destruction rate unlike Shao and Silos (2013). Compared to Coles and Moghaddasi Kelishomi (2018) our calibration strategy utilizes evidence on product/firm destruction; we do not have to attribute all job separations due to destruction, as they do.[9] Indeed, as we show in this paper, the share of job loss attributed to the destruction shock has a first-order effect on the numerical properties of the model. Relative to Mercan and Schoefer (2020) our approach uses only data from the U.S. Although they can very accurately back out the destruction rate implied by their model using data on new jobs vs re-hires, their data is from Germany. This is a drawback for the purposes of this paper because it is likely the German and U.S. labor markets differ in important aspects that affect destruction and separation rates as well as their aggregate business cycle properties. Moreover, the calibration strategy in Mercan and Schoefer (2020) requires the use of a lot of data that is difficult to obtain. In contrast, ours uses only estimates and data moments that are easily obtainable. Fujita and Ramey (2007) use the evidence on job destruction from the BED program reported in Faberman et al. (2004) for their calibration. Because job destruction is defined as "the gross number of jobs lost at establishments either closing down or contracting their workforce" (Faberman et al., 2004, p.1), identification of the destruction rate in their model requires one to back it out from the equilibrium in the model using information on job separation and job finding rates. In contrast, our calibration strategy is able to identify the destruction rate using evidence on job destruction alone.

Although the literature has mainly focused on exploring the DMP model's ability to reproduce the empirically observed relative volatility of unemployment, several studies have stressed the correlation between productivity and unemployment. Notably, Barnichon (2010) uncovers the stylized fact that these two series are only mildly correlated, using a variety of productivity measures. Furthermore, the author finds the cross-correlation to be negative pre-1984 and positive post-1984 (although in both sample periods its magnitude is mild). Barnichon (2010) focuses on explaining the empirically observed sign change of the correlation and does not focus on its magnitude. Moreover, the model introduces nominal rigidity into the DMP framework, following the New Keynesian literature. Nevertheless, it is able to reproduce the sign and magnitude of the empirically observed correlation between unemploy-

[9]To be precise, Coles and Moghaddasi Kelishomi (2018) specify their model so that all separations lead to the loss of a vacancy. Thus, in contrast to their framework we make the distinction between separations and destruction.

ment and productivity even though it requires nominal rigidity and two sources of business cycle fluctuations: a demand shock as well as a technology shock.

The key way our theoretical model departs from the baseline DMP is by featuring finitely elastic vacancy supply/congestion. In that respect our modeling environment follows Coles and Moghaddasi Kelishomi (2018) and we assume a flexible congestion function. In a recent paper Potter (2022) uses the same technique to study labor market dynamics when vacancy supply is finitely elastic and modern search technologies allow workers to monitor and quickly apply to newly posted jobs. Beaudry et al. (2018) use a similar specification of vacancy creation congestion to provide an estimate of the labor demand curve. Fujita and Ramey (2007) study the dynamics of vacancy creation using a functional form that is nested in our specification.[10] Shao and Silos (2013) also study a model with congestion in vacancy creation. However, in their setting the congestion works through the rental rate on capital: newly created vacancies require capital, so the more vacancies are created the higher the demand for capital, which raises its rental rate. Mercan and Schoefer (2020) employ a similar, but more general, vacancy congestion specification to Fujita and Ramey (2007) in order to study the implications of replacement hiring in the DMP model. In a model with long-lived vacancies, Haefke and Reiter (2020) analyze a vacancy depletion channel and the implications of vacancy longevity for the dynamics of the labor market and its efficiency properties. The vacancy creation mechanism is analogous to ours and the one in Coles and Moghaddasi Kelishomi (2018). However, they endogenize separations, which we do not. More importantly, although their calibration is able to reproduce the empirical properties of the data in several important dimensions, the correlation between unemployment and productivity is still reported to be almost perfect.

2. Model

Environment. Our environment closely follows a conventional equilibrium unemployment model in discrete time, e.g. Pissarides (2000). The only point of departure is that we follow Coles and Moghaddasi Kelishomi (2018) and assume vacancy creation is finitely

[10] See Section 2 for details.

elastic, i.e. there is congestion in vacancy creation. Unemployed workers search for jobs, firms search for workers to fill their vacancies, and matches are formed according to a matching function. Once matched, workers and firms decide on wages using Nash Bargaining and the match persists until the pair exogenously separates. There is a fixed measure $F > 0$ of firms that can create vacancies. In each period firms receive access to a new independent business opportunity, which can be undertaken by paying an investment cost x. This cost, for example, can reflect R&D and taking a new product to the production phase. Let Q_t denote the value of posting a vacancy at time t and suppose each firm draws an investment cost from a known distribution H. Given that firms undertake their business opportunity if and only if $x \leq Q_t$, the aggregate amount of new vacancy creation is $E_t = FH(Q_t)$.[11]

Figure 1 below summarizes the timing in our model. At the beginning of each period the aggregate productivity shock is realized and agents observe the current productivity level p_t. Next, firms receive their investment opportunities and make entry decisions. Third, production takes place: firm-worker pairs that are matched produce p_t, workers are paid a wage w_t, and unemployed workers receive benefits b_t. Furthermore, at this stage firms which have an unfilled vacancy pay the vacancy posting cost γ. Fourth, matching take place. Specifically, each period a number $M(u_t, v_t)$ of firm-worker pairs are formed, where u_t denotes unemployment and v_t the number of vacancies on the market. We make the standard assumptions on the matching function: it is concave in each of its arguments, constant returns to scale, and increasing in both its arguments. Thus, the job-filling rate for firms is $q(\theta_t) \equiv M(u_t, v_t)/v_t = M(\theta_t^{-1}, 1)$ and the job-finding rate is $f(\theta_t) \equiv M(u_t, v_t)/u_t = M(1, \theta_t)$. Fifth, separations take place: with some probability s the firm-worker pair separates. In that event, the firm keeps its business opportunity, but must hire a new worker for the job. Lastly, firms and matches are subject to a destruction shock — with probability δ the firm's business opportunity becomes obsolete. In that event the match is dissolved and the existing firm is destroyed. If a firm has an unfilled vacancy that suffers a destruction shock, it is destroyed

[11] Note that our entry condition implies there is congestion in vacancy creation — the more firms post vacancies each period, the higher the average investment firms must make. We can arrive at an analogous congestion mechanism if we instead assume that firms compete for new business opportunities that come in the form of innovations. See, for example, Gabrovski (2019) and Gabrovski (2022).

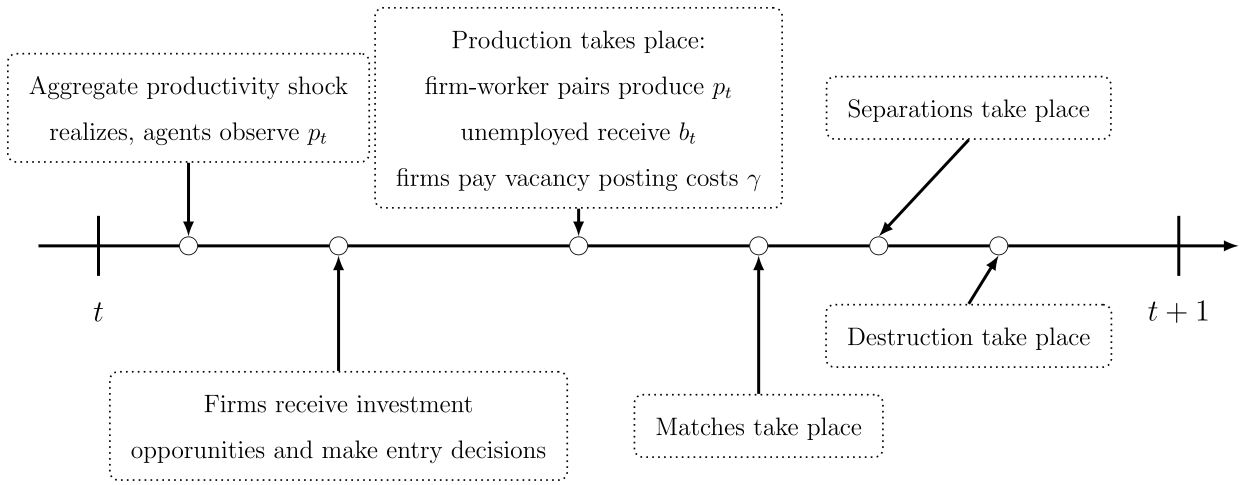

Figure 1: Labor Market Timing

as well. We assume that a match formed in period t is not subject to a separation shock at time t, but may be hit with a destruction shock. Moreover, newly created vacancies are also at risk of destruction. The distinction between a separation shock, s, and a destruction shock, δ, is our point of departure from Coles and Moghaddasi Kelishomi (2018) (where all job loss be due to destruction). In the event of a separation shock, the firm keeps its vacancy and begins searching for a worker right away. In the event of a destruction shock, however, the firm has to make an entry decision. As we show in the next section, this distinction has a major effect on the quantitative predictions of the model. If the destruction shock dominates, then dynamics are close to that of the baseline Pissarides (2000) model and vacancy creation is mainly determined by the current labor market conditions; otherwise, separation shock dominates and the dynamics are very different because the mass of vacancies is mainly determined by *past* labor market conditions.

Bellman Equations. The value of a vacancy, Q_t, comprises several terms. First, firms must pay a vacancy posting cost γ in order to search for workers in the labor market. If they are matched with a worker, which occurs with probability $q(\theta_t)$, their vacancy transitions to a filled job with value J_t. If not, they keep the opportunity to search for a worker next period. Finally, firms discount the future with a factor β and expect their business opportunity will not become obsolete with probability $(1-\delta)$:

$$Q_t = -\gamma + \beta(1-\delta)\mathbb{E}_t \left[q(\theta_t)J_{t+1} + (1-q(\theta_t))Q_{t+1} \right]. \tag{1}$$

A filled job has productivity p_t. Firms pay workers a wage w_t, so the per-period profits are $p_t - w_t$. With probability $(1-s)(1-\delta)$ the firm and the worker do not separate and the business opportunity does not become obsolete, so the firm keeps its filled job next period. There is a chance $s(1-\delta)$ that the firm-worker pair dissolves due to separation (rather than destruction). In that event, the firm keeps a vacancy with expected value $\mathbb{E}_t Q_{t+1}$ and can search for a new worker next period. Lastly, in the event that the business opportunity is destroyed, the firm-worker pair dissolves and the firm exits the market. Thus,

$$J_t = p_t - w_t + \beta \left[(1-s)(1-\delta)\mathbb{E}_t J_{t+1} + s(1-\delta)\mathbb{E}_t Q_{t+1} \right]. \tag{2}$$

An unemployed worker receives benefits b. With probability $f(\theta_t)$ she is matched with a firm and, conditional on that, there is a $(1-\delta)$ chance the job survives until the beginning of next period. In that event the worker becomes employed. Otherwise, the worker remains unemployed. Thus, the value of unemployment, U_t, satisfies

$$U_t = b + \beta \left[(1-\delta) f(\theta_t)\mathbb{E}_t W_{t+1} + [1-(1-\delta)f(\theta_t)]\mathbb{E}_t U_{t+1} \right]. \tag{3}$$

An employed worker receives wages w_t. She keeps the job whenever the firm-worker pair do not separate and the firm survives the destruction shock, i.e. with probability $(1-\delta)(1-s)$. Otherwise, the worker loses the job and transitions to unemployment next period. Hence, the value of the job for the worker, W_t, satisfies

$$W_t = w_t + \beta \left[(1-\delta)(1-s)\mathbb{E}_t W_{t+1} + [1-(1-\delta)(1-s)]\mathbb{E}_t U_{t+1} \right]. \tag{4}$$

Wages, laws of motion, and entry. As is standard in the literature, wages are determined according to Nash bargaining:

$$w_t = \arg\max_{w_t} [J_t - Q_t]^{1-\alpha}[W_t - U_t]^{\alpha},$$

where α is the worker's bargaining power. Thus, the solution for the wage satisfies

$$\alpha(J_t - V_t) = (1-\alpha)(W_t - U_t). \tag{5}$$

Lastly, we close the model by specifying the laws of motion and the entry decision of firms. Now that entry into the labor market requires firms to make a sunk investment cost, vacancies are a state variable. In particular, the number of vacancies today v_t is the sum of three terms. First, all vacancies last period that were not matched with a worker and survived the destruction shock, $(1-\delta)[(1-q(\theta_{t-1}))v_{t-1}$, remain in the pool of vacancies today. Second, all filled jobs that experienced a separation shock, but survived the destruction shock last period, $(1-\delta)s(1-u_{t-1})$, transform into vacancies today. Lastly, all new entrants, e_t, have a vacancy. Hence,

$$v_t = (1-\delta)[(1-q(\theta_{t-1}))v_{t-1} + s(1-u_{t-1})] + e_t. \tag{6}$$

In turn the amount of new entrants is determined by the free entry condition $e_t = FH(Q_t)$. We follow Coles and Moghaddasi Kelishomi (2018) and postulate a parsimonious power law distribution: $H(Q_t) = Q_t^{\xi}$.[12] Note that under this specification the value of the vacancy in equilibrium is $Q_t = (e_t/F)^{1/\xi}$. Thus, the congestion elasticity is $1/\xi$ and is both positive and finite. In contrast, under the baseline DMP model with no congestion this elasticity is infinite.

Due to the destruction shock, the law of motion for unemployment differs from the one in the Pissarides model as well. In particular, worker-firm pairs are subject to a destruction shock immediately after the match, before production takes place. Thus, only a fraction $(1-\delta)f(\theta_{t-1})$ of unemployed workers the previous period transition to employment this period. The rest remain unemployed. Second, the effective separation rate for workers is $\tau \equiv 1-(1-\delta)(1-s)$, so a number of $\tau(1-u_{t-1})$ previously employed workers join the unemployment pool today. Therefore,

$$u_t = [1-(1-\delta)f(\theta_{t-1})]u_{t-1} + \tau(1-u_{t-1}). \tag{7}$$

Because our focus is on business cycle dynamics, we follow the literature and assume the natural logarithm of productivity follows an AR(1) process

$$\log(p_t) = \rho \log(p_{t-1}) + \epsilon_t, \tag{8}$$

where ρ is a persistence parameter, and $\epsilon_t \sim N(0,\sigma)$ is a white noise process.

[12] Beaudry et al. (2018) and Potter (2022) use similar functional forms.

Equilibrium. Combining the free entry condition, $Q_t = (e_t/F)^{1/\xi}$, and the Bellman equations for vacancies (1) and for filled jobs (2) yields the job creation condition in our economy:

$$\frac{\gamma + K_t}{q(\theta_t)} = \beta(1-\delta)\mathbb{E}_t\left[p_{t+1} - w_{t+1} - K_{t+1} + (1-s)\frac{\gamma + K_{t+1}}{q(\theta_{t+1})}\right], \tag{9}$$

where $K_t \equiv \mathbb{E}_t\left[e_t^{1/\xi} - \beta(1-\delta)e_{t+1}^{1/\xi}\right]/F^{\xi}$ is the expected flow entry cost: the difference between the entry cost the firms face today and the discounted expected entry cost tomorrow. The interpretation of (9) is analogous to that in the baseline DMP model: the left-hand side of the equation represents the expected costs of posting and maintaining a vacancy, whereas the right-hand side is the expected profit. The costs in our environment take into account the congestion that takes place in the economy. For example, if many firms are entering the market today the market is congested and an entrant might choose to delay her entry decision until tomorrow when she will face a lower entry cost. This force is captured by a higher K_t in the job creation condition above. When entry today is high firms have an incentive to wait, whereas when entry today is low firms have an incentive to expedite their entry. This smoothing mechanism yields a hump-shaped response in vacancies, a feature consistent with the empirical evidence provided by Fujita and Ramey (2007). In contrast, vacancies peak on impact in the standard DMP model.[13] The expected benefit from posting a vacancy is the expected discounted profits next period, $p_{t+1} - w_{t+1}$, plus the continuation value of the vacancy in the event the firm and worker separate, $(1-s)[\gamma + K_{t+1}]/q(\theta_{t+1})$, net of the expected flow costs, K_{t+1}.

Next, we turn our attention to the wage equation. Using the surplus sharing rule (5) together with the four Bellman equations (1), (2), (3), (4) one can obtain an expression of the wage that is similar to the one in the baseline DMP model:

$$w_t = \alpha\left[p_t - K_t + \frac{\theta_t}{1-\delta}(\gamma + K_t)\right] + (1-\alpha)b. \tag{10}$$

The wage is a weighted average between the flow payoff from the match and the worker's outside option of receiving unemployment benefits. The benefit of the match constitutes

[13] Fujita and Ramey (2007) study this mechanism in detail and show that it helps the model match key properties of vacancies in the data. Unlike us, however, Fujita and Ramey (2007) do not focus on the mild procyclicality of vacancies. Indeed in their model vacancies and productivity are highly correlated.

the match output together with the search costs the firm saves from not having a vacancy, $p_t + \frac{\theta_t}{1-\delta}(\gamma + K_t)$, just as in the baseline DMP model with the only difference that now the job opportunity has a finite expected life of only $1/(1-\delta)$ and that the vacancy costs include both γ and the flow entry cost K_t. Additionally, vacancies are an asset with a positive value and are subject to congestion in our economy. Consequently, the wage includes an additional term: $-K_t$. This term reflects the fact that the outside option of firms in the bargaining game has a flow value K_t. If vacancies are expected to be more valuable tomorrow, then the firm will have an easier time recruiting a worker tomorrow, which raises the attractiveness of her outside option and subsequently reduces the wage. Alternatively, if the expected value of a vacancy tomorrow is lower the firm is more eager to match with a worker today and is thus willing to offer a higher wage.

We are ready to define equilibrium. Formally,

Definition 1. *An equilibrium is an infinite, bounded sequence of productivity, wages, market tightness, entrants, vacancies, and unemployment* $\{p_t, w_t, \theta_t, e_t, v_t, u_t\}_0^\infty$ *such that (i) firms set entry optimally according to* (9)*; (ii) the wage solves the Nash Bargaining problem between the firm and the worker as in* (10)*; (iii) vacancies follow the low of motion* (6)*; (iv) unemployment follows the law of motion* (7)*; (v) productivity follows the AR(1) process defined in* (8)*.*

3. Numerical Results

The appendix details the numerical procedure we use to simulate the model and to derive the impulse response functions. In this section, we outline the calibration and focus on our results, and the mechanism behind them. Namely, we show that the model-implied cross-correlation between productivity on the one hand and vacancies, unemployment, and the market tightness on the other depends crucially on the calibrated value of the destruction shock. This is true even though the labor market variables themselves remain nearly perfectly correlated with each other for any value of δ. Intuitively, the model can reproduce the mild cross-correlation for two reasons. First, vacancies are a predetermined variable, which evolve according to the law of motion (6). As is evident in the equation, the pool of vacancies comprises new entrants and surviving vacancies created in previous periods. This feature is in contrast to the standard DMP model where new entrants constitute all vacancies. As a result, in the baseline DMP model vacancies are a choice variable and adjust instantaneously in response to changes in labor productivity. In contrast, when there is a fixed cost of entry, only a fraction of the vacancy pool, e_t, is determined by current labor market conditions whereas the remaining fraction of vacancies, $(1-\delta)[(1-q(\theta_{t-1}))v_{t-1}+s(1-u_{t-1})]$, is correlated with past labor productivities. When the destruction shock δ is relatively high, the contribution of entrants to the mass of vacancies is relatively high, so the correlation between vacancies and labor productivity is relatively high as well. On the other hand, if the shock is less frequent, most vacancies from the previous period survive which implies only a small part of current vacancies are new entrants. Second, the equilibrium of our economy features congestion in entry: the more entrants there are the higher the average costs of entry. This generates a smoothing mechanism similar to the one in Fujita and Ramey (2007). As the authors show, this mechanism incentivizes firms to smooth out their response to productivity shocks over several periods so as to avoid excessively high entry costs. Fujita and Ramey (2007) highlight this mechanism and show it reproduces the hump-shaped response of vacancies to productivity shocks that we observe in the data.

3.1. Calibration

Our paper highlights the implications of sunk investment in vacancy creation for the DMP model's ability to reproduce the mild cross-correlation between productivity and the labor market variables. In this theoretical context creating a vacancy is closely tied to a business opportunity. Thus, job loss in the economy is the result of one of two distinct processes: (i) a separation between the firm and the worker which dissolves the match but leaves the firm's business opportunity intact and (ii) a destruction of the business opportunity which makes the firm's product no longer viable, i.e. the firm exits the labor market altogether. The first source of job loss is closely tied to the view of separations in the standard DMP model: the match might dissolve because the worker moves to a new city, decides to change jobs because of conflict with management, is fired for cause, etc. The second source of job loss is more closely related to the business environment the firm operates in: the match might dissolve because a competitor has stolen the firm's market share, because the product the firm makes is now obsolete, because the firm went out of business, etc. Hence, we pay special attention to calibrating the destruction shock in a way that captures the firm's lost business opportunity. To facilitate comparison with the existing literature we keep our calibration as close as possible to Coles and Moghaddasi Kelishomi (2018). The only departure is our strategy for calibrating the destruction and separation shocks. In the benchmark we use data on product destruction and firm exit to calibrate δ. We then present the model's implied moments for different values of δ used in the literature so as to highlight the dependence of the numerical results on the calibrated value of the destruction shock.

The model is calibrated at monthly frequency. The discount factor $\beta = 0.9967$ is chosen so that it matches an annual discount rate of 4%. Further, we set the vacancy posting costs γ to zero, $b = 0.7$, $\alpha = 0.6$, and $\xi = 1$, so that the distribution $H(Q_t)$ is uniform, following Coles and Moghaddasi Kelishomi (2018). Next, we set $\delta = 0.0051$ to match a 6% annual destruction rate following the evidence in Bernard et al. (2010). To match the empirical job-filling and job-finding rates we set $1/[(1-\delta)q(\theta)] = 0.75$ months and $1/[(1-\delta)f(\theta)] = 2.2$ months. Further, the mean job separation probability is set to $\tau = 3.4$ percent per month, thus steady state unemployment is $u = 0.07$. Given δ and ξ these three moments jointly yield $\nu = 1.47$, $s = 0.0258$, and $F = 0.000235$. The only source of uncertainty in our model is the productivity shock. We postulate that productivity follows an AR(1) process with autoregression coefficient $\rho = 0.979$ and variance $\sigma = 0.007$, similar to the values used by Coles and Moghaddasi Kelishomi (2018). Table 1 below summarizes the calibration.

Table 1: Calibration

Preferences/Technology	Parameter	Value	Calibration Strategy
Vacancy posting cost	γ	0	Coles and Moghaddasi Kelishomi (2018)
Bargaining power	α	0.6	Coles and Moghaddasi Kelishomi (2018)
Unemployment benefits	b	0.7	Coles and Moghaddasi Kelishomi (2018)
Matching function elasticity	ν	1.575	Job-finding rate
Discount factor	β	0.997	4% annual discount rate
Separation rate	s	0.0258	3.4% monthly match dissolution probability
Destruction rate	δ	0.0051	6% annual destruction rate
Population of firms	F	0.000235	Job-filling rate
Cost distribution parameter	ξ	1	Coles and Moghaddasi Kelishomi (2018)

3.2. *Impulse Response Functions and Mechanism*

To highlight the intuition behind our numerical results we first turn to the impulse response functions generated by the model. To this end we graph the economy's response to a one standard deviation negative technology shock in Figure 2. On impact all vacancies, entry, and the market tightness all decrease. The response of entry is largest on impact and begins to recover in subsequent periods, but slowly. This sluggish response is due to the congestion in the model: increasing entry faster leads to larger vacancy creation costs, which incentivizes firms to delay entry to future periods. Since vacancies are an asset with a positive value firms do not voluntarily exit the labor market. Instead, they keep their surviving vacancies and the only change in the pool comes from the reduced entry. Thus, vacancies respond very little on impact and achieve their maximum response only around 5 years after the initial shock. The market tightness and unemployment follow a similar pattern because their behavior is a direct consequence of firm's entry decisions. As highlighted by Fujita and Ramey (2007) this sluggishness (i) is absent from the baseline DMP model without congestion in entry and (ii) is well-supported empirically. In what follows we highlight the importance of this sluggishness for the model's ability to match the cross-correlation between productivity and the labor market variables observed in the data.

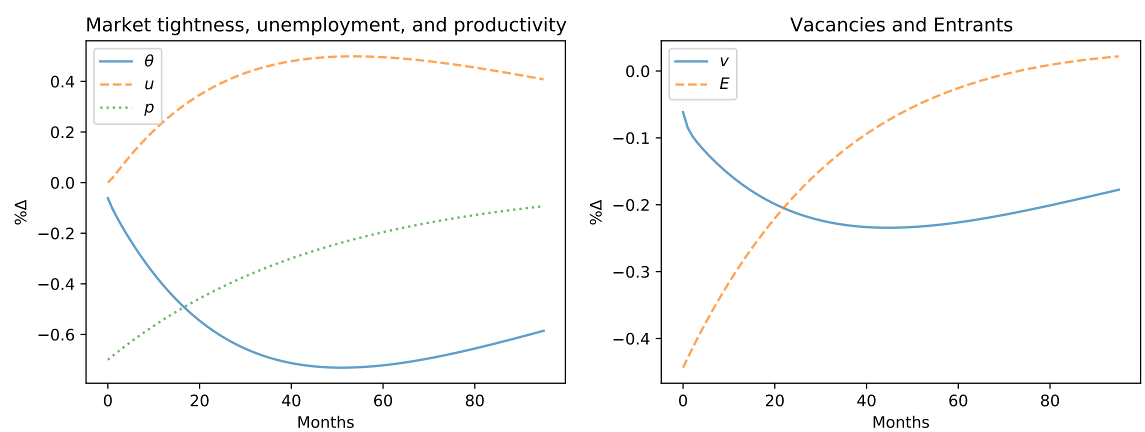

Figure 2: Impulse response functions in the benchmark calibration with $\delta = 0.0051$.

Next, we focus on the impulse response functions generated by the model when it is calibrated to match the data on firm destruction. The destruction shock is set to $\delta = 0.00874$

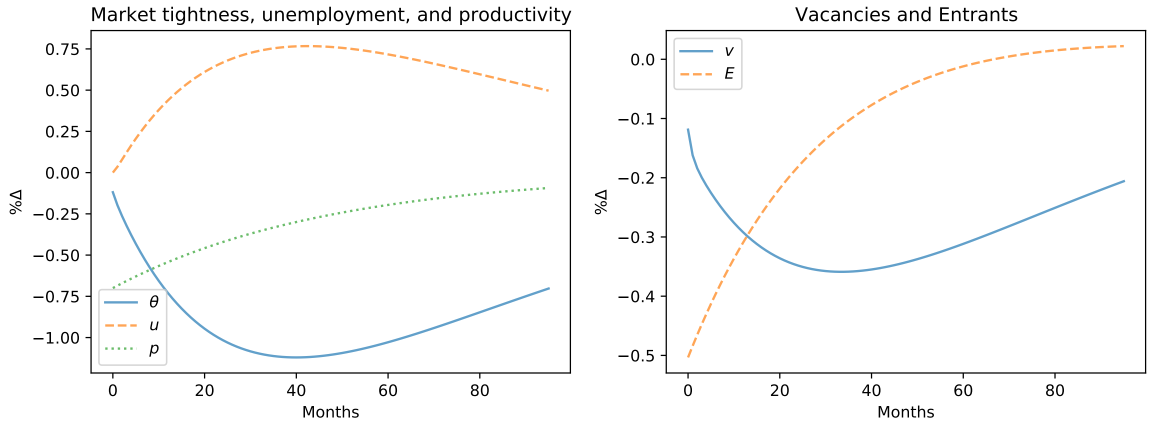

Figure 3: Impulse response functions for model with $\delta = 0.0087$.

so that the annual rate of firm exit is 10% following the evidence in Broda and Weinstein (2010). All other moments are kept the same as in the benchmark calibration so as to isolate the model's response to a more frequent destruction shock. Figure 3 presents the impulse response functions. The qualitative behavior of entry, vacancies, market tightness, and unemployment is the same as in the benchmark model: entry responds most on impact and then slowly recovers, whereas vacancies, tightness, and unemployment respond sluggishly to the shock reaching their peak response many periods after the shock. The difference between this calibration and the benchmark is in the magnitude of the responses and the speed of the response: the peak response in vacancies, tightness, and unemployment is more than twice that in the benchmark calibration and their response is not as sluggish, i.e. the peak response occurs only after about 3 years.

Figure 4 graphs the impulse response functions generated by the model when all job losses are calibrated to be due to destruction, i.e. $s = 0$ and $\delta = 0.0342$. This calibration corresponds to the one in Coles and Moghaddasi Kelishomi (2018). The pattern we have seen emerge from the previous two figures is still present: following a productivity shock, the model predicts a sluggish response of vacancies, the market tightness, and entry. The extent of this sluggishness and the peak response in the vacancies and market tightness, however, are tightly linked to the calibrated value of the destruction shock: the larger δ is, the bigger the peak response and the faster the converges of the vacancies and the market tightness to

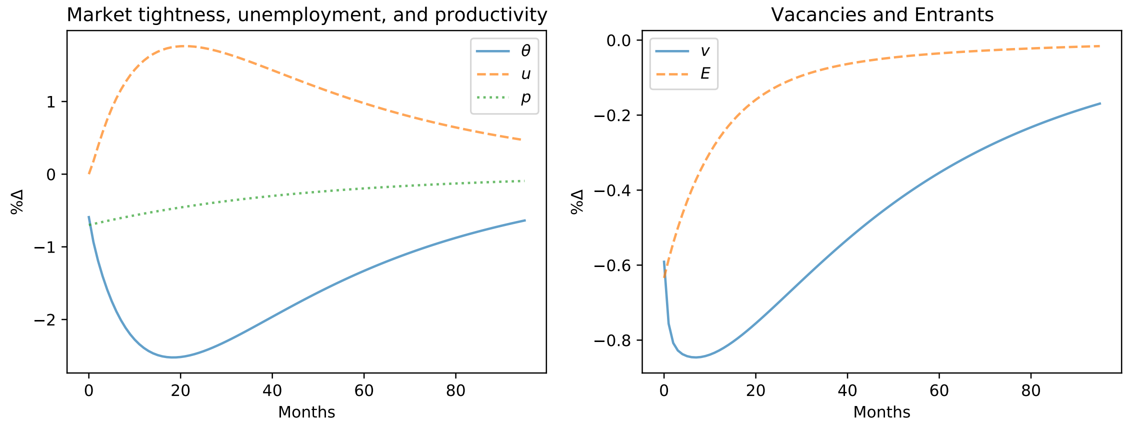

Figure 4: Impulse response functions for model with $\delta = 0.0342$.

the steady state. It is thus evident that the calibrated value of the destruction shock is a key determinant of the model's numerical predictions. A lower destruction rate implies that most vacancies in the pool are surviving business opportunities from previous periods. As a result, new entrants comprise a small fraction of the pool. Since this is the portion of vacancies that responds to aggregate conditions the reduction of vacancies is small on impact. At the same time both the flows in (entry and separations) and out of vacancies (matching and destruction) are relatively small as compared to the entire pool which contributes to a sluggish response of vacancies. Because the behavior of the market tightness and unemployment are directly determined by the behavior of vacancies their response is qualitatively the same. At the same time Figure 4 highlights the second reason behind sluggishness in our model: congestion in entry induces firms to smooth out their response to productivity shocks over several periods.

When vacancies respond sluggishly to shocks, the labor-market variables comove more strongly with past than with current levels of productivity. As a result, the absolute value of the cross-correlation between productivity and the labor market variables is lower. This happens for two reasons. First, sluggish response in vacancies implies lower entry, so a higher fraction of the vacancy pool are vacancies that were created in previous periods. In those previous periods firms were looking at past levels of productivity to make their entry decisions, so those vacancies are correlated to these past productivity values. Second, a more

Table 2: Moments

Variable X	Data	Benchmark ($\delta = 0.0051$)	10% Destruction Rate ($\delta = 0.00874$)	CM ($\delta = 0.0342$)	FR/MS ($\delta = 0.0222$)	SS ($\delta = 0.018$)
		$Corr(X, p)$				
u	-0.408	-0.329	-0.462	-0.77	-0.686	-0.648
θ	0.396	0.419	0.554	0.861	0.78	0.736
v	0.364	0.593	0.721	0.975	0.924	0.887

sluggish response in vacancies is tied to lower destruction rates. Thus, more vacancies in the pool were created in previous periods, but also the average age of the vacancy in the pool is older. Since productivity are less correlated over a longer horizon, this also helps lower the cross-correlation between current productivity and the labor market variables as compared to the correlation between past productivity and the labor market variables.

Table 2 highlights this point by presenting the cross-correlation of productivity with vacancies, unemployment, and the market tightness. The second column shows the model predicted moments in the benchmark calibration and the third column reports the moments when the destruction shock is calibrated to match the empirically observed exit rate for firms. In Column 4, referred to as CM, we calibrate the model to match the destruction rate used in Coles and Moghaddasi Kelishomi (2018). Column 5, referred to as FR/MS, matches the destruction rates used in Mercan and Schoefer (2020) and Fujita and Ramey (2007).[14] The last column, referred to as SS, sets $\delta = 0.018$, which matches the calibration in Shao and Silos (2013).

The benchmark calibration in Column 2 is able to reproduce the mild cross-correlation of the labor market variables with productivity. For unemployment and the market tightness the correlation is slightly lower than that in the data, whereas for vacancies it is slightly higher. Turning to Column 3 and a destruction shock calibrated to match the data on firm

[14] To be precise, Fujita and Ramey (2007) calibrate $\delta = 0.021$ but this small difference does not change the simulated moments in a meaningful way, so we group the two calibrations togeter.

exit, the model is able to replicate a mild cross-correlation for unemployment and the market tightness as well. The correlation between vacancies and productivity is somewhat high at 0.72, yet still significantly lower than what the baseline DMP model predicts and the values implied by models in the existing literature. In particular, for the last three columns vacancies are strongly pro-cyclical. A similar pattern emerges for unemployment and market tightness: when we calibrate the destruction shock to values from the existing literature, the cross-correlation is much higher than that in the data. Of course, it is not perfect because even in the case when all separations are due to job destruction (Column CM) the model still features the vacancy smoothing mechanism from Fujita and Ramey (2007). It is worth noting that, for low δ, the model is able to reproduce the mild cross-correlation of the labor variables with productivity while maintaining a strong correlation between the labor variables themselves. Specifically, in the benchmark calibration $\text{Corr}(u, v) = -0.96$ and $\text{Corr}(u, \theta) = -0.99$.

Our benchmark calibration uses the value of $\xi = 1$. This value is also used in Fujita and Ramey (2007). Coles and Moghaddasi Kelishomi (2018) also consider a value of $\xi = 0.265$, whereas Haefke and Reiter (2020) use 15.878 for their benchmark. Although values of ξ are not directly comparable across models, we report our main results for different calibrated ξ in Table 3 below. As expected, the lower ξ is the more correlated vacancies are with productivity. However, even for values as low as $\xi = 0.25$ the correlation remains mild.

Table 3: Alternative values of ξ

Value of ξ	0.25	0.5	1	2	4	8
$\text{Corr}(u_t, p_t)$	-0.515	-0.411	-0.329	-0.255	-0.196	-0.151

3.3. Dynamic Correlations

When the destruction rate is low, the stock of vacancies is relatively more skewed towards vacancies created in past periods, which correlates more closely with past values of productivity. We observe this pattern for both the benchmark calibration and the calibration with 10% annual destruction rate. The same is not true for the values of δ taken from the previous literature where the correlation between vacancies and productivity peters out with increasing lags. Of course, even for low values of the destruction shock, eventually vacancies are less strongly correlated with past productivity than current productivity, as we consider productivity values further in the past. Specifically, the correlation seems to be strongest around 2 quarters after the shock. Table 4 shows the results for the different calibrated values of δ.[15]

Table 4: Moments

$Corr(v_t, p_{t-i})$					
Lagged Productivity	Benchmark ($\delta = 0.0051$)	10% Destruction Rate ($\delta = 0.00874$)	CM ($\delta = 0.0342$)	FR/MS ($\delta = 0.0222$)	SS ($\delta = 0.018$)
p_{t-1}	0.667	0.784	0.95	0.933	0.917
p_{t-2}	0.709	0.809	0.87	0.902	0.898
p_{t-3}	0.735	0.816	0.78	0.849	0.861
p_{t-4}	0.749	0.81	0.69	0.785	0.812

Figure 5 depicts the dynamic correlations in greater detail. For conciseness we focus on the benchmark calibration and the one in which the destruction rate is calibrated as by Coles and Moghaddasi Kelishomi (2018). The first panel of the figure shows the dynamic correlations between unemployment and productivity. Here the benchmark calibration does better again in terms of being able to match both the correlations and the peak correlation lag more closely. The second panel shows the dynamic correlations between vacancies and productivity. The curves corresponding to the two calibrations confirm our results from Table 4: the correlation peaks in the current period and peters off monotonically as lags increase for the calibration with $\delta = 0.0342$, whereas the peak correlation under the benchmark calibration is when productivity is lagged several months. Importantly, this is the behavior of the data as well: vacancies are more strongly correlated with past levels of productivity than current ones. Lastly, the third panel shows that the two calibrations produce comparable correlations between unemployment and vacancies.

[15] We only show the correlations for vacancies, but those for unemployment and the market tightness follow a similar pattern.

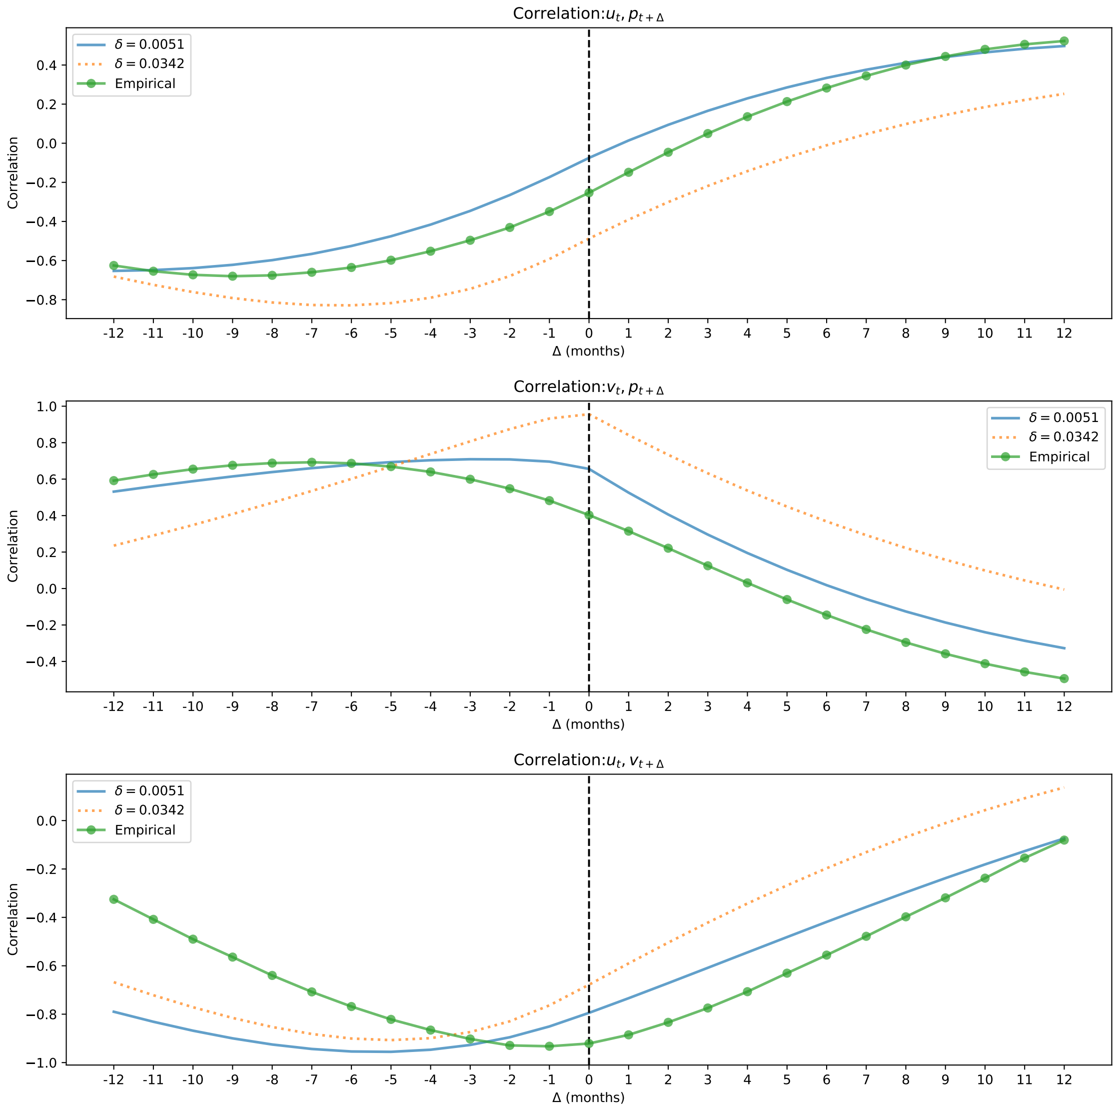

Figure 5: Dynamic correlations. The horizontal axis in each period depicts the time-shift Δ, measured in months, the vertical axis the correlation coefficient. For each pairwise combination of variables, we consider the data alongside the model calibrated under $\delta = 0.0051$ and $\delta = 0.0342$.

4. Conclusion

The DMP model has proved to be a very useful theoretical framework to analyze labor market dynamics, with several extensions of the baseline model being able to reproduce the relative volatility of unemployment, vacancies, and market tightness. Yet the ability of the model to reproduce the empirical mild correlation between productivity and labor market variables remains under-explored. We show that an extension of the baseline model with sunk vacancy creation costs and congestion in entry can help the model reproduce these empirical moments. Under sunk creation costs, vacancies become a positively valued state variable. Consequently, the model distinguishes between a separation shock — the firm-worker match dissolves, but the firm keeps its business opportunity — and a destruction shock, in which the firm-worker match dissolves due to a lost business opportunity. In the former case, the firm can look for a replacement worker without incurring additional vacancy creation costs, whereas in the latter the firm must seek a new business opportunity to re-enter the labor market.

The model eliminates the near-perfect correlation between the labor market variables and productivity for two reasons. First, some vacancies are created by new entrants whereas others originate from previous periods. Because firms have no incentive to voluntarily destroy their vacancies, only the newly form vacancies respond to a technology shock. That is, only a fraction of vacancies are determined by the current macroeconomic conditions, whereas the remaining fraction reflects past levels of productivity. The smaller the destruction shock, the smaller the pool of new entrants, and ultimately the less correlated vacancies are with current productivity. Moreover, congestion in entry induces firms to smooth out entry in response to a shock, which serves to further reduce the correlation between vacancies and productivity. Our analysis shows that calibrating the economy with a destruction rate that matches (i) the micro-level evidence on product destruction and firm exits or (ii) the commonly used values in the growth literature allows the model to reproduce the mild cross-correlation between vacancies, unemployment, an the market tightness on the one hand and productivity on the other. The model achieves this feature while retaining the strong correlation between the labor market variables themselves.

Appendix A. Numerical Algorithm

The algorithm for computing the dynamic stochastic equilibrium is an Euler-equation based method described in detail in Coleman et al. (2021). The unknown policy functions are approximated using complete quadratic monomials in terms of the states with coefficients Θ. There is one exogenous state variable p_t and the two endogenous state variables: unemployment u_t and predetermined vacancies $v_{pret,t} = (1-\delta)[(1-q(\theta_{t-1}))v_{t-1} + s(1-u_{t-1})]$. The aggregate state space is thus $\mathcal{S}_t = (u_t, v_{pret,t}, p_t)$. We use a quasi-random grid (Sobol) on a fixed hypercube to discretize the state space. We approximate the flow entry cost K_t and entrants e_t.

The steady state is useful as a precursor to computing the stochastic equilibrium for at least three reason: (1) it can be used to initialize the unknown functions at the steady state, (2) it is essential to express impulse responses in percentage deviations from steady state, and (3) it can be used to appropriately set bounds of the endogenous state variables.

1. Step 1: Initialization
   (a) Choose (u_t, F_t, p_t) and T
   (b) Choose approximating functions $H \approx \hat{H}(;\Theta)$
   (c) Make initial guesses on Θ: set equal to steady-state values
   (d) Choose integration nodes $\{\varepsilon_{x,j}\}_{j=1}^{J}$ and weights $\{\omega_j\}_{j=1}^{J}$
   (e) Construct a grid $\Gamma = \{u_m, v_{pret,m}, p_m\}_{m=1}^{M} \equiv \{X_m\}_{m=1}^{M}$
   (f) Choose termination criterion $crit = 1e-6$
2. Step 2: Computation of a solution for H
   (a) At iteration i, for $m = 1, \ldots, M$, compute
      - $K_m = \hat{H}_1(X_m), e_m = \hat{H}_2(X_m)$
      - $v_m = v_{pret,m} + e_m$
      - $\theta_m = v_m/u_m$
      - $q_m = q(\theta_m)$
      - $f_m = \theta_m q_m$

- Update stocks

$$u'_m = (1-(1-\delta)f_m)u_m + \tau(1-u_m)$$
$$v'_{pret,m} = (1-\delta)((1-q_m)v_m + s(1-u_m))$$

- If necessary, constrain u'_m and $v_{pret,m}{}'$ to lie within admissible bounds
- Update state vector with future productivity nodes:

$$X_{mj} = (u'_m, v'_{pret,m}, p'_{mj}) \quad \forall j$$

- Interpolation of function

$$H'_{mj} = \hat{H}(X'_{mj};\Theta)$$

- Repeat steps above to get x'_{mj}, θ'_{mj}
- Numerical integration

$$EJ_m = \beta \sum_{j=1}^{J} \omega_j \left[(1-\alpha_L)(x'_{mj} - b - K'_{mj}) - \alpha_L(\theta'_{mj}\gamma + K'_{mj})/(1-\delta) + (1-s)\frac{\gamma + K'_{mj}}{q'_{mj}} \right]$$

$$\hat{H}_{1,m} = q_m EJ_m - \gamma$$
$$\hat{H}_{2,m} = \beta \sum_{j=1}^{J} \omega_j \left[(K_m F^{\xi} + \beta(1-\delta)e'^{1/\xi}_{mj})^{\xi} \right]$$

(b) Find b that solves the system in (2a)

- Use ordinary least squares

$$\hat{\Theta}^g \equiv \arg\min \sum_{m=1}^{M} ||\hat{H}_m - \hat{H}(X_m;\Theta)||$$

- Dampening: weight η on new coefficients

$$\Theta^{(i+1)} = (1-\eta)\Theta^{(i)} + \eta\hat{\Theta}^g$$

- Check for convergence: end Step 2 if

$$\frac{1}{M}\left\{ \sum_{m=1}^{M} |\frac{(H_m)^{i+1} - (H_m)^i}{(H_m)^i}|, \right\} < crit$$

(c) Iterate on Step 2 until convergence

The coefficients Θ give us approximate solutions to the policy functions. it is straightforward to simulate data, construct moments, and generate impulse responses.

References

Barnichon, R., 2010. Productivity and unemployment over the business cycle. Journal of Monetary Economics 57, 1013–1025.

Beaudry, P., Green, D.A., Sand, B.M., 2018. In search of labor demand. American Economic Review 108, 2714–57.

Bernard, A.B., Redding, S.J., Schott, P.K., 2010. Multiple-product firms and product switching. American Economic Review 100, 70–97.

Bilbiie, F.O., Ghironi, F., Melitz, M.J., 2012. Endogenous entry, product variety, and business cycles. Journal of Political Economy 120.

Broda, C., Weinstein, D.E., 2010. Product creation and destruction: Evidence and price implications. American Economic Review 100, 691–723.

Coleman, C., Lyon, S., Maliar, L., Maliar, S., 2021. Matlab, python, julia: What to choose in economics? Computational Economics 58, 1263–1288.

Coles, M.G., Moghaddasi Kelishomi, A., 2018. Do job destruction shocks matter in the theory of unemployment? American Economic Journal: Macroeconomics 10, 118–36.

Comin, D., Gertler, M., 2006. Medium-term business cycles. The American Economic Review 96, 523.

Diamond, P.A., 1982. Aggregate demand management in search equilibrium. Journal of political Economy 90, 881–894.

Faberman, R.J., et al., 2004. Gross Job Flows over the Past Two Business Cycles: Not all'Recoveries' are Created Equal. Technical Report.

Fujita, S., Ramey, G., 2007. Job matching and propagation. Journal of Economic dynamics and control 31, 3671–3698.

Gabrovski, M., 2019. Simultaneous innovation and the cyclicality of r&d. Review of Economic Dynamics, forthcoming .

Gabrovski, M., 2022. Coordination frictions and economic growth. Macroeconomic Dynamics forthcoming.

Haefke, C., Reiter, M., 2020. Long live the vacancy .

Hagedorn, M., Manovskii, I., 2008. The cyclical behavior of equilibrium unemployment and vacancies revisited. American Economic Review 98, 1692–1706.

Hall, R.E., Milgrom, P.R., 2008. The limited influence of unemployment on the wage bargain. American economic review 98, 1653–74.

Lee, Y., Mukoyama, T., 2015. Entry and exit of manufacturing plants over the business cycle. European Economic Review 77, 20–27.

Ljungqvist, L., Sargent, T.J., 2017. The fundamental surplus. American Economic Review 107, 2630–65.

Mercan, Y., Schoefer, B., 2020. Jobs and matches: Quits, replacement hiring, and vacancy chains. American Economic Review: Insights 2, 101–24.

Pissarides, C.A., 2000. Equilibrium unemployment theory. MIT press.

Pissarides, C.A., 2009. The unemployment volatility puzzle: Is wage stickiness the answer? Econometrica 77, 1339–1369.

Potter, T., 2022. Destabilizing search technology .

Shao, E., Silos, P., 2013. Entry costs and labor market dynamics. European Economic Review 63, 243–255.

Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. American economic review 95, 25–49.