



Search, unemployment, and the Beveridge curve: Experimental evidence[☆]

John Duffy^{a,b,1}, Brian C. Jenkins^{a,*,1}

^a Department of Economics, University of California, Irvine, United States of America

^b ISEER Osaka University, Japan

ARTICLE INFO

JEL classification:

J63
J64
C92

Keywords:

Beveridge curve
Job search
Matching
Vacancies
Unemployment
Experimental economics

ABSTRACT

We report on a laboratory experiment testing the predictions of the Diamond–Mortensen–Pissarides (DMP) search-and-matching model, which is a workhorse, decentralized model of unemployment and the labor market. We focus on the job vacancy posting problem that firms face in the DMP model. We explore the model's comparative statics predictions concerning variations in the separation rate, the vacancy posting cost, and the firm's surplus earned per employee. Across all treatments, we find strong evidence for an inverse relationship between vacancies and unemployment, consistent with the Beveridge curve. We also find that the results of our various comparative statics exercises are in-line with the predictions of the theory.

1. Introduction

Following the work of Diamond (1982), Pissarides (1985), and Mortensen and Pissarides (1994), the Diamond–Mortensen–Pissarides (DMP) search-and-matching model has become a workhorse framework for studying unemployment and labor market dynamics and policy interventions. The DMP approach takes search and matching frictions in labor markets seriously. Job creation and destruction are costly and time-consuming and worker–firm matches obtain in a decentralized and uncoordinated fashion. The core microeconomic foundation of the DMP model is a matching function specifying that the number of workers hired by firms over a given time period is increasing in both the number of unemployed workers looking for work and in the number of job vacancies that firms have posted and are seeking to fill.² A major aim of this literature is to provide theoretical foundations for the empirical *Beveridge curve*, the well-known inverse relationship between job vacancies and unemployment. Fig. 1 depicts the empirical Beveridge curve for the US from the start of the Great Depression.

The Beveridge curve relationship plays an important role in labor market analysis and policy (see, e.g., Elsby et al. (2015)) and is used to assess the health of and detect structural changes in the job market.

Movements along the Beveridge curve reflect the normal, negative relationship between unemployment and vacancies over the course of the business cycle. A downward shift of the Beveridge curve indicates improved job matching efficiency or increased labor market dynamism, while an outward shift of the Beveridge curve indicates the presence of mismatches in skills or other impediments to employment such as changes in worker search efforts. By assessing such shifts in the Beveridge curve, policymakers can act to address those impediments.

While there is not yet a consensus on what precise functional form the matching function in the DMP model ought to take, there is substantial empirical evidence supporting the general use of the matching function approach of the DMP model (Petrongolo and Pissarides, 2001). In this paper we provide further empirical support for the DMP model, but we follow a different path from that taken in the DMP literature. Rather than looking for empirical support for the matching function in the available economic data, we look for causal evidence as to whether the *decisions* of agents, interacting with the matching function of the theory and with complete information about other model parameters, produce aggregate outcomes consistent with the symmetric rational expectations (RE) equilibrium on the predicted Beveridge curve. Such

[☆] We thank the editor and two referees for helpful comments and suggestions on an earlier draft. We gratefully acknowledge funding for this project from the UC Irvine Academic Senate CORCL, United States of America. The human subjects data reported on in this paper was collected under a protocol approved by the UC Irvine Institutional Review Board (HS #2011-8378).

* Corresponding author.

E-mail addresses: duffy@uci.edu (J. Duffy), bcjenkin@uci.edu (B.C. Jenkins).

¹ Mailing address: 3151 Social Science Plaza, Irvine, CA 92697-5100.

² This matching function is analogous to a production function in the standard theory of the firm. Unemployed workers and vacant jobs are “inputs” and new hires are the “output”.




Fig. 1. Beveridge curve for the US. Monthly data from April 1929 through October 2023. See Petrosky-Nadeau and Zhang (2013) for a description of data sources and preparation.

causal evidence is not generally available using field data unless there has been some type of natural experiment.

In our experiment, we reduce the complexity of the model and focus on the behavior of firms, taking the behavior of workers as given (workers are robot players in our experiment). Thus, each subject in our experiment plays the role of a *firm* that must choose how many job vacancy notices to post in each period. In making these choices, subjects have full information about the aggregate level of unemployment, the vacancy posting cost, the expected surplus flow from a filled job and the job separation rate. The vacancy posting decisions, made simultaneously by the firms, determine the total number of vacancies, which together with the unemployment level determines the equilibrium vacancy filling rate.

An advantage of our approach is that the total number of vacancies is observable, along with the number of employed and unemployed persons. In the field, by contrast, unemployment and vacancy numbers are measured based on survey data and are only known with some lag. While there is some agreement about the measurement of unemployment status, there is less agreement about the measurement of vacancies in the economy. By contrast, in the laboratory, we have complete control over the measurement of such variables, and we provide subjects with a complete information setting in which to make decisions, thereby providing an ideal environment for evaluating the comparative static predictions of the theory.

Nevertheless, in the theoretical model, solving for the symmetric, rational expectations equilibrium choice of vacancies for a firm to post is not a trivial problem. The vacancy choice problem in our experimental setting is even more complicated than theory imagines because there is strategic uncertainty regarding the choices of other participants that is absent from the theoretical, representative agent framework. Subjects are likely heterogeneous in their ability to make the right vacancy posting choices and, to the extent that any subject deviates from the symmetric equilibrium, the optimal vacancy choice for the other subjects will be affected as well. Thus our study explores the extent to which human subjects, incentivized according to the theory and provided with all the information necessary to solve for the steady state, can learn to post vacancies in a way that, in the aggregate, is consistent with the symmetric information, rational expectations steady state of the DMP model.

The DMP theory predicts that changes in any one of several factors will move the location of the steady state on the Beveridge curve or will cause a shift of the Beveridge curve itself. We consider several

treatments in our experimental design that each entail a change in one key factor relative to a baseline parameterization. Specifically, we consider changes in: (1) the cost of posting a vacancy, (2) the surplus flow to firms from a filled job and (3) the rate at which workers separate from their jobs each period. Observations from the various treatments allow us to examine whether the locus of equilibrium points from each treatment condition corresponds to the predicted equilibrium along an appropriate Beveridge curve. If subjects behave as the theory predicts, then observations from our different treatments should reveal both *movements along* a Beveridge curve and, in the case of the treatment where the separation rate is changed, a *shift* to a new Beveridge curve.

We find that across all treatments, average outcomes for unemployment and vacancies tend to lie close to, if not precisely on the predicted Beveridge curve and are often close to the steady state equilibria predicted by theory. For the treatments that are predicted to produce movements along a Beveridge curve, we find that the data do in fact reveal movements in the correct direction. For the one treatment that should produce an equilibrium along a new Beveridge curve, we also find that the experiment predicts exactly that. Thus, our results provide strong evidence in support of the comparative statics predictions of the DMP model with respect to the behavior of firms.

What do we learn from such an experiment? First, we learn that subjects do a reasonably good job solving the non-trivial optimization problem required by the DMP framework, even in the face of strategic uncertainty about the decisions made by other participants. Such strategic uncertainty is an inherent property of real world vacancy posting decisions made by firms, but as noted earlier, it is generally missing from the theory. Second, and related to the first point, we consider whether human vacancy posting decisions are more deliberative than are the choices made by automated “noise traders” who make random vacancy choices (within certain limits). We find that human subjects clearly outperform these noise traders and so we learn that the matching function *by itself* does not explain why the experimental comparative statics results are consistent with theoretical predictions. That is, the findings we observe are not mechanically driven, but rather involve deliberation on the part of subjects in the role of firms. Finally, we learn that the comparative statics predictions of the DMP model are robust to the observed heterogeneity in subject behavior.

To our knowledge, the DMP model has not been studied in the controlled conditions of the laboratory and Beveridge curves have not been reproduced in any experimental setting. We view our research into this question as part of a larger agenda aimed at evaluating the assumptions

and predictions of macroeconomic models using laboratory methods; see Duffy (2016) for a survey of this literature.

There is an experimental literature evaluating the equilibrium predictions of labor search models, but these studies focus on the *worker* side of the market, e.g., whether workers will accept job offers conditional on their reservation wages. See, for example, Schotter and Braunstein (1981), Braunstein and Schotter (1982), Cox and Oaxaca (1989, 1992), Harrison and Morgan (1990) and Brown et al. (2011).

There is also some field experimental work assessing causality in labor market outcomes that is relevant to the literature on labor search. For instance, Kroft et al. (2013), Ghayad (2013) and Eriksson and Rooth (2014) explore the effect of unemployment duration on employer callback rates using factitious job applications that manipulate the duration of an applicant's unemployment. These studies find a negative relationship between unemployment duration and callback rates. Other researchers, e.g. Van Ours and Vodopivec (2006) have used natural experiments of reductions in the duration of unemployment benefits for some workers and not others to assess the causal impact on exit rates from unemployment; they find that a reduction in the duration of unemployment benefits raises the unemployment exit rate.

By contrast with all of this literature, our focus is on the *firm side* of the labor market. We study the firm's problem of posting costly job vacancy notices when each firm randomly loses some employees in every period and there is uncertainty about the vacancy posting choices made by rival firms operating in the same labor market, which affects the degree of labor market tightness and the likelihood that vacancy notices will be filled. We are not aware of any empirical or experimental research that directly addresses such firm behavior. There are some related experimental studies of firm investments in *market capacity*, but there the focus is on deterring entry by rival firms, see, e.g., Brandts et al. (2007) and Brandts and Giritligil (2008). These studies differ from the DMP framework we use, where firms do not choose production capacity or decide sequentially whether to enter an industry. Instead, firms in our setting simultaneously decide on how many costly vacancy notices to post, which has some (stochastic) relationship to the number of employees working at each firm.

The closest paper to this one is by Korenok and Munro (2021), who use the matching model of Diamond (1982) to explore how subjects bargain with firms over wages and how these bargaining outcomes are affected by changes in unemployment insurance and in the unemployment rate. Our study complements Korenok and Munro (2021) in that we also consider the impact of the unemployment rate on labor market outcomes, but instead we focus on firms' job vacancy posting decisions, with wages determined exogenously.

2. Theory

We adopt a simplified version of the DMP model described by Bhat-tacharya et al. (2018). In this version, the DMP model can be expressed without Bellman equations, making the problem of finding the optimal number of job vacancies to post more tractable for subjects in our experiment.

2.1. The matching function and the Beveridge curve

There is a fixed number of workers, L . In each period t , E_t workers are employed and the remaining $L - E_t = U_t$ workers are unemployed. There is also a fixed number of firms, N , each of whom posts vacancy notices. Let $V_t = \sum_{i=1}^N V_{it}$ denote the total number of vacancies posted in each period. Total new hires H_t are formed according to the matching function:

$$H_t = A\sqrt{U_t V_t}. \tag{1}$$

We chose this parameterization for the matching function to make the model easier to communicate to subjects, but we note that the elasticities of new hires with respect to unemployment and vacancies




Fig. 2. Beveridge curve. Steady state relationship between the vacancy rate v and the unemployment rate u implied by the model.

are consistent with many of the estimates reported in Petrongolo and Pissarides (2001).

The matching function is used to derive the probabilities that vacancies are filled. The ratio of new hires to vacancies is the *vacancy-filling rate* q_t :

$$q_t = \frac{H_t}{V_t} = A\sqrt{\frac{U_t}{V_t}} = A\frac{1}{\sqrt{\theta_t}}, \tag{2}$$

where $\theta_t \equiv V_t/U_t$ denotes *market tightness*. The law of motion for the number of unemployed persons is:

$$U_{t+1} = U_t - q_t V_t + \lambda E_t, \tag{3}$$

where λ is the exogenous separation rate or the probability that a given job is destroyed. Dividing both sides of (3) by L to express the law of motion in terms of the unemployment rate, $u_t = U_t/L$, and substituting for q_t using (2), we have:

$$u_{t+1} = (1 - A\sqrt{v_t/u_t} - \lambda)u_t + \lambda, \tag{4}$$

where $v_t = V_t/L$ is the vacancy rate. In the steady state, the model implies a Beveridge curve relationship:

$$u = \frac{\lambda}{\lambda + A\sqrt{v/u}}. \tag{5}$$

The Beveridge curve implied by Eq. (5) is depicted in Fig. 2.

2.2. Vacancy creation

Let $n = 1, \dots, N$ index individual firms. Firm n enters period t with E_{n-1} jobs filled and chooses how many vacancies V_{nt} to post in the current period. The cost of posting a vacancy notice is κ per period. A filled job provides the firm with a surplus flow of π per period. The surplus flow captures the difference between the productivity flow of the worker in the filled job and the wage paid to that worker. In most versions of the DMP model, productivity is exogenous while the wage is determined endogenously through some kind of bargaining process. We decided to fix the wage here in order to keep our experimental environment simple. With a fixed wage, there is no need to distinguish between the wage and productivity and so we just make the surplus exogenous.

We consider an indefinite horizon that proxies for the infinite horizon of the theory. The probability that the economy continues to operate from one period to the next is denoted by $\delta \in (0, 1)$ and so with probability $1 - \delta$ the current period is the final period for the economy.³

³ This random termination process, originating in the work of Roth and Murnighan (1978), is a standard mechanism for inducing both discounting and the stationarity associated with an infinite horizon.

Each period, firm n earns a profit flow equal to the difference between the surplus earned from current employees, and the cost of advertising new vacancies:

$$\text{Profit}_{nt} = \pi E_{nt} - \kappa V_{nt}. \tag{6}$$

The firm chooses vacancies V_{nt} to maximize the present discounted expected value of current and future profit flows. Therefore, in period t , firm n solves:

$$\max_{V_{nt}} E_t \sum_{s=0}^{\infty} \delta^s (\pi E_{nt+s} - \kappa V_{nt+s}), \tag{7}$$

subject to Eqs. (3) and (2). By substituting Eqs. (3) and (2) into the objective function and simplifying, the objective function can be written as a sum with two terms that depend on current vacancy posting and a term that is independent of V_{nt} :

$$\max_{V_{nt}} q(\theta_t) \underbrace{\left(\frac{\delta \pi}{1 - \delta(1 - \lambda)} \right)}_{\substack{\text{Expected PDV} \\ \text{of surplus flow} \\ \text{per employee}}} V_{nt} - \kappa V_{nt} + \xi_{nt}, \tag{8}$$

where ξ_{nt} includes all factors affecting current and future profit flows that do not depend on the vacancy posting choice V_{nt} and $q(\theta_t)$ depends on the choice of V_{nt} :

$$q(\theta_t) = A \sqrt{U_t / \sum_{j=1}^N V_{jt}}. \tag{9}$$

Note that aggregate unemployment U_t is predetermined via Eq. (3) and is known to firms when choosing V_{nt} and that the firm's current number of employees does not affect its vacancy posting decision.

In the symmetric rational expectations equilibrium, the choice of vacancies for firm n is given by:

$$V_{nt} = \frac{U_t}{N} \left[\left(\frac{A}{\kappa} \right) \left(\frac{2N-1}{2N} \right) \left(\frac{\delta \pi}{1 - \delta(1 - \lambda)} \right) \right]^2, \tag{10}$$

so total vacancies are proportional to unemployment:

$$V_t = U_t \left[\left(\frac{A}{\kappa} \right) \left(\frac{2N-1}{2N} \right) \left(\frac{\delta \pi}{1 - \delta(1 - \lambda)} \right) \right]^2. \tag{11}$$

Eqs. (5) and (11) jointly determine the unemployment rate and vacancies. We use the rational expectations equilibrium as the benchmark for comparison with the outcomes of the experimental economy.

3. Experimental design and implementation

The experiment was computerized and programmed using the oTree software (Chen et al., 2016). Each experimental session consisted of $N = 10$ subjects with no prior experience with the study; subjects were only allowed to participate in a single session/treatment. Our subjects were undergraduate students drawn from a broad set of majors at UC Irvine. They play the role of firms choosing how many vacancy notices to post per round.⁴

Each session consisted of a number of sequences or “supergames” (repeated dynamic games). Each sequence consisted of an indefinite number of rounds to capture the infinite horizon nature of the model environment. We set the probability δ that a sequence continues from one round to the next to be 0.8 and we pre-drew five sets of indefinite

⁴ While we use student subjects to play the role of firms, the model setting we study is sufficiently abstract that it would seem unlikely that “professional” subjects (if they could be identified) would perform differently. Indeed, in the few experimental studies that do compare the behavior of students versus “professionals” the overwhelming finding is that there is no difference in performance between these two different sets of subjects, see e.g., Fréchet (2015).

Table 1

Number of sequences, durations of each sequence and total number of rounds from the 5 sessions of each treatment.

Session	Sequences	Durations	Total rounds
1	9	1,10,9,1,6,3,5,7,6	48
2	9	7,6,1,4,2,5,1,6,13	45
3	9	4,2,5,7,5,1,11,8,2	45
4	10	2,7,4,6,5,10,4,2,1,6	47
5	8	7,2,1,7,8,10,4,5	44
All	45	–	229

Table 2

Parameterizations used in each treatment and rational expectations (RE) equilibrium vacancies, unemployment rate, vacancy rate, market tightness, and vacancy filling rate. Relative to the baseline treatment, the other three treatments had one different parameter value. The changed parameter value is boxed.

	Treatment			
	Baseline	Incr. κ	Incr. π	Decr. λ
Matching productivity (A)	0.5	0.5	0.5	0.5
Labor force (L)	120	120	120	120
Continuation probability (δ)	0.8	0.8	0.8	0.8
Vacancy posting cost (κ)	19	38	19	19
Surplus per employee (π)	24	24	48	24
Separation rate (λ)	0.5	0.5	0.5	0.25
RE vacancies (V)	42.67	13.71	118.15	50.82
RE unemployment rate (u)	0.56	0.71	0.38	0.29
RE vacancy rate (v)	0.36	0.11	0.98	0.42
RE market tightness (θ)	0.64	0.16	2.56	1.44
RE vacancy-filling rate (q)	0.62	1.00	0.31	0.42

sequence lengths using this continuation probability as reported in Table 1. For each of our four treatments, we conducted 5 sessions each with one of the five sets of sequence lengths and total number of rounds, also reported in Table 1.⁵ Thus, we have 5 sessions of 10 subjects for each of our four experimental treatments or results from $5 \times 4 \times 10 = 200$ experimental subjects.

Table 2 reveals how our four experimental treatments differ from one another in terms of the parameterization of the model. The bottom part of Table 2 shows, for each parameterization, the RE predictions for the four primary outcome variables of our analysis: total vacancies V , the unemployment rate, u , the vacancy rate v , and the vacancy filling rate q .

Note first that we hold several model parameters constant across all treatments. Specifically, the matching productivity parameter A is always set equal to 0.5; the size of labor force L is always set to 120 workers; and the continuation probability δ is always set to 0.8. The four treatments vary one of the other three parameters of the model relative to the parameterization given in what we label the “Baseline” treatment. In the “Increasing κ ” treatment, we double κ , the vacancy posting cost, from 19 per notice in the Baseline to 38. In the “Increasing π ” treatment, we double the surplus that firms earn per employee π from 24 in the Baseline to 48. Finally, in the “Decreasing λ ” treatment, we decrease by half the job separation rate λ , from 50% in the Baseline to 25%.

The treatment parameterizations were not chosen to be the most realistic. Instead they were chosen because they produce rational expectations equilibria that facilitate testing of the model in two important ways. First, a primary goal of the experiment is to test the comparative statics predictions of the theory and for that purpose it is necessary that changes in parameter values be large enough to be detectable in the (typically) noisy experimental data. Second, each parameterization was chosen so that the rational expectations equilibrium unemployment

⁵ This design allows for variations in the duration of sequences across sessions, but it further allows for control in the amount of noise in sequence durations across all four treatments.

rate would not be near the lower and upper bounds of 0 and 1 in order to reduce the chances of encountering those boundaries during experimental sessions.

At the beginning of an experimental session, subjects are given written instructions with a description of the environment. The instructions are reproduced in Appendix A. Since the fundamental aim of the experiment is to test whether subjects can comprehend the optimization problem they face and find the optimal number of vacancies to post all while interacting and coordinating with other subjects who face the same task and who jointly determine the value of q , we provided subjects with *complete information* about this task. Specifically, subjects are told:

1. All relevant parameter values for the environment. See Table 2 for the parameterizations of the various treatments.
2. The job separation and vacancy filling processes described in Section 2.1 including the unemployment law of motion in Eq. (3). We do not provide subjects with the explicit optimization problem in expression (8), nor do we provide them with the solution to the problem, because we are interested in whether and how they solve that problem.
3. The formula for the vacancy filling rate as a function of U_t and V_t in Eq. (2). We also provide subjects with a table of q values for different combinations of U_t and V_t and a contour plot to help them to visualize the relationship. These are reproduced in the Appendix — See Table A.1 and Figure A.1.
4. How points are accumulated and the point-to-dollar conversion rate for compensation.

The instructions are read aloud at the beginning of each experimental session and subjects have the opportunity to ask questions. All subjects then complete a comprehension quiz to make sure that they understand the environment. We verify that each subject with incorrect answers learns the correct answers before moving on to the decision-making part of the experiment.

Following the reading of the instructions, the experiment proceeds as follows. Subjects complete 8 to 10 independent sequences of the model environment, but the number of sequences is not known to them. The number of rounds in each sequence was chosen randomly based on the 80 percent continuation probability and so the length of each sequence is also not known in advance by the subjects. At the start of each sequence, all 10 subjects have the same initial number of employees which is set close to, but not exactly equal to, the symmetric rational expectations equilibrium value for the treatment. At the beginning of each round, subjects learn:

1. The total vacancies posted by all subjects in the previous round
2. How many of their own vacancies from the previous round were filled
3. How many of their own employees from the previous round separated from their firm
4. Their own payoff from the previous round which is equal to the total surplus from all filled jobs less the cost of advertising new vacancies
5. The new total level of unemployment for the current round

Points accumulated in a round are the difference between the total surplus earned from filled jobs and the total cost of advertising new vacancies:

$$\text{New points} = \pi E_{nt} - \kappa V_{nt} \tag{12}$$

Note that Eq. (12) is analogous to Eq. (6). Subjects are endowed with 100 points at the start of each new sequence; there is no carry over of points from one sequence to the next. Subjects were instructed that total

Table 3
Sessions, subject numbers, and average subject earnings across treatments.

Treatment	# Sessions	# Subjects	Average earnings
Baseline	5	50	\$24.72
Increasing κ	5	50	\$15.54
Increasing π	5	50	\$49.84
Decreasing λ	5	50	\$39.49

expenditures on vacancy postings were not to exceed their available point balance in a sequence.⁶

At the end of the experimental session, we randomly select two sequences from all sequences played. The subjects' final round point totals from those two chosen sequences are summed together and this sum is converted into dollars at the rate of 1 point = \$0.03.⁷ Subjects are guaranteed a minimum of \$7 for showing up and completing the experimental session.

4. Hypotheses

Our experiment was designed to test four main hypotheses which we evaluate using our experimental data.

Hypothesis H1. The relationship between the observed vacancy rate and the unemployment rate in the experiment is well approximated by the relevant Beveridge curve for each treatment.

Hypothesis H2. Relative to the Baseline treatment, an increase in the vacancy posting cost, κ , reduces the vacancy rate and increases the unemployment rate along the same Beveridge curve.

Hypothesis H3. Relative to the Baseline treatment, an increase in the surplus from a filled job, π , increases the vacancy rate and decreases the unemployment rate along the same Beveridge curve.

Hypothesis H4. Relative to the Baseline treatment, a decrease in the separation rate, λ , results in an increase in the vacancy rate and a decrease in the unemployment rate along a Beveridge curve that is shifted down relative to the Beveridge curve for the Baseline treatment.

5. Results

As noted we have data on the behavior of 200 subjects participating in 5 sessions of each of our 4 treatments (20 sessions total, 10 subjects per session). Table 3 reports on the number of sessions, the total number of subjects and their average earnings across the four treatments. Following standard practice in experimental economics, we held constant the conversion rate from points into money across treatments. Consequently, certain treatments, e.g., the Increasing π or Decreasing λ treatments yielded greater earnings opportunities than others, e.g., the Baseline parameterization. However, if we had changed the conversion rate of points into money across treatments, this would have presented a second confounding factor, in addition to the treatment change itself.

As one aggregate measure of behavior, we report treatment-level means (over all sequences of all sessions of each treatment) for five

⁶ Out of the 20 experimental sessions, there were three for which our program did not enforce this budget constraint. In these three sessions only, there were five subjects for which point balances fell below zero for just one round each: Two from a single Decreasing λ treatment session, two from a single Increasing π treatment session, and one from an Increasing κ treatment session. In the same Increasing κ treatment session, a sixth subject was allowed to consistently over-post vacancies and had large negative point balances throughout the session.

⁷ This payment procedure for indefinitely repeated games is advocated by Sherstyuk et al. (2013).




Fig. 3. Increasing vacancy posting cost κ and baseline treatments. Experimental data points are individual sequence averages for all groups of each treatment. Data are shown relative to the theoretical Beveridge curve (BC). Overall data averages and RE steady states (ss) are indicated for each treatment.

key variables: the number of vacancies posted, V , the unemployment rate, u , the vacancy rate, v , the measure of market tightness, θ and the vacancy filling rate, q . These means are reported in Table 4 along with standard deviations and predicted RE equilibrium values. Additionally, Table 4 reports the Root Mean Squared Error (RMSE) between the experimental data and the RE equilibrium prediction both in un-normalized form and in normalized form which involves dividing the un-normalized RMSE by the RE equilibrium prediction. The latter measure provides us with a unitless measure of the fit of the data to the RE equilibrium predictions enabling comparisons across treatments, with 0 indicating a perfect fit to the RE predictions.

As a second measure of the fit of our data to the RE predictions, we consider the mean value of the unemployment and vacancy rates over all rounds of each indefinite supergame or a “sequence” within a session. Recall from Table 1 that we have 45 such sequence-level observations per treatment and these are the data points shown in Figs. 3, 4, and 5. These figures show how these unemployment and vacancy rates align with the predicted Beveridge curve for each treatment. For comparison purposes, in each figure we always show the outcomes for the Baseline treatment (indicated by circles). For instance, Fig. 3 compares outcomes from the Baseline treatment with outcomes from the Increasing κ treatment, showing the average vacancy and unemployment rates for each of the 45 sequences across all sessions of these two treatments. The 45 sequence averages are represented by the small markers. The solid line is the theoretical Beveridge curve along which the equilibria for both treatments are predicted to lie. The large open markers are the rational expectations equilibria for the two treatment parameterizations. The large solid black markers are the average actual combinations (over all sequences and sessions) of the vacancy rate, v , and the unemployment rate, u , for each treatment.

We first address Hypothesis 1, using the aggregate results reported in Table 4 and the sequence average data shown in Figs. 3, 4, and 5. We start with the Baseline treatment. As Table 4 reveals, the RE equilibrium unemployment rate in the Baseline treatment is 56 percent. The actual

Table 4

Summary statistics for select group-level variables by treatment. Root-mean-square error (RMSE) of a quantity x is computed relative to the rational expectations (RE) equilibrium steady state: $\sqrt{N^{-1} \sum_{i=1}^N (x_i - x_{RE})^2}$, where x_i is the average value of the quantity for a sequence and x_{RE} is the RE steady state value. The final column reports RMSEs normalized by the RE steady state: $x_{RE}^{-1} \sqrt{N^{-1} \sum_{i=1}^N (x_i - x_{RE})^2}$.

		RE	Mean	Std. dev.	RMSE	
					Unnormalized	Normalized
Baseline	V	42.67	57.04	7.12	16.00	0.38
	u	0.56	0.51	0.02	0.05	0.10
	v	0.36	0.48	0.06	0.15	0.43
	θ	0.64	0.95	0.15	0.34	0.54
	q	0.62	0.52	0.04	0.11	0.18
Increasing κ	V	13.71	25.12	6.15	12.92	0.94
	u	0.71	0.66	0.03	0.06	0.09
	v	0.11	0.21	0.05	0.54	4.76
	θ	0.16	0.32	0.09	0.18	1.14
	q	1.25	0.91	0.11	0.36	0.28
Increasing π	V	118.15	105.63	24.06	26.89	0.23
	u	0.38	0.38	0.04	0.04	0.11
	v	0.98	0.88	0.20	0.60	0.61
	θ	2.56	2.37	0.66	0.67	0.26
	q	0.31	0.34	0.06	0.07	0.21
Decreasing λ	V	50.82	63.51	10.21	16.21	0.32
	u	0.29	0.26	0.02	0.04	0.14
	v	0.42	0.53	0.09	0.17	0.40
	θ	1.44	2.11	0.48	0.82	0.57
	q	0.42	0.36	0.04	0.07	0.17

mean unemployment rate across all sessions of the Baseline treatment is 51 percent. The RE equilibrium vacancy rate in the Baseline treatment is 36 percent and the mean vacancy rate across all Baseline treatment sessions is 48 percent. The normalized RMSEs are .10 and .43 for u and v respectively. As Fig. 3 reveals, the observations from the Baseline treatment sessions (circles) lie on the predicted Beveridge curve and




Fig. 4. Increasing surplus π and baseline treatments. Experimental data points are individual sequence averages for all groups of each treatment. Data are shown relative to the theoretical Beveridge curve (BC). Overall data averages and RE steady states (ss) are indicated for each treatment.




Fig. 5. Decreasing separation rate λ and baseline treatments. Experimental data points are individual sequence averages for all groups of each treatment. Data are shown relative to the two different theoretical Beveridge curves (BC). Overall data averages and RE steady states (ss) are indicated for each treatment.




Fig. 6. Experimental outcomes from all four treatments combined. Experimental data points are individual sequence averages for all groups of each treatment. Data are shown relative to the theoretical Beveridge curves (BC). Overall data averages and RE steady states (ss) are indicated for each treatment.

are close to the RE steady state prediction. However, the actual mean unemployment rate is lower and, correspondingly, the actual mean vacancy rate is higher relative to the RE steady state prediction, indicating that subjects in the baseline treatment had a tendency to over-post vacancies, on average.

Next, we compare the results from the Baseline treatment to the Increasing κ treatment. The RE equilibria for both of those treatments lie on the same Beveridge curve as shown in Fig. 3. As in the Baseline treatment, the observations from the Increasing κ treatment are also close to the rational expectations equilibrium steady state, though again the average unemployment rate is lower than predicted in the RE equilibrium while the vacancy rate is higher than predicted. From Table 4 we see that the mean unemployment rate in the Increasing κ treatment is 66 percent compared with the rational expectations equilibrium value of 77 percent. Further, the mean vacancy rate is 21 percent compared to the RE value of 11 percent. Using the normalized RSME, the vacancy rate, v , is the furthest from the RE prediction in the Increasing κ treatment.

Fig. 4 shows the average outcomes for the Increasing π treatment relative to the Baseline treatment. The outcomes for the Increasing π treatment are predicted to lie on the same Beveridge curve as in the Baseline treatment and that indeed appears to be approximately the case. Table 4 reveals that the average actual unemployment rate across all sessions of the Increasing π treatment is exactly equal to the RE prediction of 38 percent. The RE equilibrium vacancy rate is 98 percent while the actual mean vacancy rate across all sessions of this treatment is lower, at 88 percent. As Fig. 4 makes clear, the standard deviation of the vacancy rate, v , as well as the number of vacancy postings, V , is higher in the Increasing π treatment than in the other three treatments. However, the normalized RMSE measures for these and the other outcome variables for the Increasing π treatment are not greatly different from the Baseline treatment values.

Finally, Fig. 5 compares outcomes from the Decreasing λ treatment with the Baseline treatment. In this case, the reduction in the separation

rate, λ , by one half is predicted to *shift* the Beveridge curve. The dashed line in Fig. 5 is the Beveridge curve associated with the lowered value for λ . Again, qualitatively, outcomes for the Decreasing λ treatment match the predictions: they lie on the correct, shifted Beveridge curve and clearly show a higher vacancy rate and lower unemployment. Quantitatively, the average outcomes involve the combination of a lower unemployment rate and a higher vacancy rate relative to the RE predictions. Table 4 reveals that the mean unemployment rate is 26 percent which is only slightly lower than the RE prediction of 29 percent. At 53 percent, the mean vacancy rate is above the RE prediction of 42 percent, indicating that subjects in the Decreasing λ treatments were apparently also prone to over-posting vacancies. However, the normalized RMSEs for the Decreasing λ treatment are very similar to those for the Baseline treatment.

Fig. 6 plots the mean vacancy and unemployment rates from each experimental sequence for all four treatments (20 sessions total) in a single graph. Three clear insights emerge from examining this figure. First, the results from each experimental sequence are clustered around the theoretical Beveridge curve corresponding to the treatment-specific parameterization of the model. That is, the data generated by the experiment do in fact appear to lie along the correct Beveridge curves. Second, the mean values generated by the sequences in each treatment (large open markers) are generally close to the rational expectations equilibrium predictions (large solid markers). Finally, because the average values generated by the sequences in the treatment are close to the respective rational expectations equilibria, the experimental data produces comparative statics results that are consistent with the theory.

Thus, with respect to Hypothesis 1, the Beveridge curve does appear to be a good approximation of the average outcomes of the experimental sessions. Further, in line with Hypotheses 2, 3, and 4, changes in model parameters relative to the baseline treatment follow the comparative static predictions of theory. On the other hand, as Fig. 6 makes clear, there is excessive vacancy posting in the baseline, increasing κ and decreasing λ treatments, and there is under-posting of

vacancies and a higher variance in vacancy postings and vacancy rates in the increasing π treatment.

5.1. Evidence of deliberative activity

Before providing a more detailed analysis of the experimental results, we briefly consider what the outcome for the experiment would have been if subjects were not thinking too hard about how many vacancies to choose and were instead randomly choosing the number of vacancies they posted up to certain feasible limits. Here, we are interested in establishing that subjects in our experiment were not random in their behavior, but were instead *deliberately engaged* in the task of determining how many vacancies to post given the benefits and costs of those decisions.

For this purpose, we simulate the DMP model for each of our four treatment parameterizations under the assumption that at the beginning of each sequence, each of the 10 firms randomly and independently chooses a number of vacancy postings and posts that number for each remaining round of the sequence. We report results from simulations under the assumption that the number of vacancies posted by each firm is randomly drawn from a uniform distribution with a lower limit of 0 and an upper limit of 75% of the maximum vacancies allowed by firms' beginning of sequence point balance.⁸ We use the same sequence lengths that we used in the actual experiment. Fig. 7 shows the results of this random vacancy choice simulation. The top panel compares the outcomes for the randomized Increasing κ treatment with the randomized Baseline treatment, the middle panel compares the randomized Increasing π treatment with the randomized Baseline treatment, and the bottom panel compares the randomized Decreasing λ treatment with the randomized Baseline treatment. Relative to the experimental data plotted in Fig. 6, the averages for the vacancy rate and, especially, the unemployment rate are considerably more volatile in the data with randomly chosen vacancies.

In some cases, like the Decreasing λ treatment, the outcomes from the randomization exercise appear centered on the RE steady state, but that is in fact an artifact of the distributional assumption for the randomized vacancy drawings. Indeed, by adjusting the bounds of the distribution from which vacancies are drawn, one can bring the central tendency of the random data in line with that of the experimental data, but the volatility remains higher than in the experimental data. The simulations with randomized vacancy postings reveal that the model does not mechanically produce outcomes that are near the RE steady state. It also does not guarantee that outcomes will be tightly clustered around a given Beveridge curve. We interpret the findings from these randomization exercises as providing strong evidence that subjects in our experiment were *not* randomly choosing the number of vacancy postings, but are instead proceeding in a more deliberative manner in line with the rational actor assumption.

5.2. Aggregate vacancy postings

In this section, we examine more precisely how the vacancy posting behavior of subjects in our experiment conforms to the rational expectations equilibrium of the DMP model. From Eq. (10), the log rational

⁸ There is no obvious best choice for what the bounds on randomized vacancy postings ought to be for this simulation exercise. In our view, the bounds used to produce the results in the main body of our paper reasonably balance our desire to create variability in vacancy postings while also excluding outlandishly large vacancy postings that would deplete all or most of firms' point balances. Figures B.1 through B.4 in Appendix B depict simulations of random vacancy posting outcomes for each of the four treatments under alternative assumptions about the upper bounds on randomly drawn vacancies.

Table 5

Regression analysis: Log market tightness regressed on treatment indicators together with RE predictions for comparison purposes. Standard errors are clustered at the session level.

	Estimated	Predicted
Intercept (β_0)	-0.0864 (0.0655)	-0.4463
Indicator: Surplus treatment (β_π)	0.9231*** (0.1515)	1.3863
Indicator: Vacancy cost treatment (β_κ)	-1.1220*** (0.1054)	-1.3863
Indicator: Separation rate treatment (β_λ)	0.8116*** (0.0995)	0.8109
R-squared	0.8751	
R-squared Adj.	0.8751	
N	9160	

expectations equilibrium vacancy postings by firms implies that market tightness is a function of N , A , δ , π , κ , and λ :

$$\log \theta = 2 \log \left[\frac{A\delta\pi}{\kappa(1-\delta(1-\lambda))} \left(\frac{2N-1}{2N} \right) \right]. \quad (13)$$

Thus, we estimate the following linear model:

$$\log \theta_{it} = \beta_0 + \beta_\pi D_{it}^\pi + \beta_\kappa D_{it}^\kappa + \beta_\lambda D_{it}^\lambda + \epsilon_{it} \quad (14)$$

where i denotes the group and t denotes the round within the session and where the treatment indicator variables are defined as follows:

$$D_{it}^\pi = \begin{cases} 0, & \pi = 24 \\ 1, & \pi = 48 \end{cases} \quad (15)$$

$$D_{it}^\kappa = \begin{cases} 0, & \kappa = 19 \\ 1, & \kappa = 38 \end{cases} \quad (16)$$

$$D_{it}^\lambda = \begin{cases} 0, & \lambda = 0.5 \\ 1, & \lambda = 0.25 \end{cases} \quad (17)$$

Note that the predicted coefficients are:

$$\beta_0 = 2 \log \left[\frac{A\pi\delta}{\kappa(1-\delta(1-\lambda))} \left(\frac{2N-1}{2N} \right) \right] = -0.4463 \quad (18)$$

$$\beta_\pi = 2 \log(\pi'/\pi) = 2 \log 2 = 1.3863 \quad (19)$$

$$\beta_\kappa = -2 \log(\kappa'/\kappa) = -2 \log 2 = -1.3863 \quad (20)$$

$$\beta_\lambda = -2 \log \left(\frac{1-\delta(1-\lambda)}{1-\delta(1-\lambda')} \right) = -2 \log(2/3) = 0.8109 \quad (21)$$

We estimate the model using OLS and present the results in Table 5. The constant of the regression is estimated to be a small negative number that is slightly larger than the RE equilibrium prediction. The estimate is consistent with our observation of excessive vacancy postings in the Baseline treatment relative to the RE prediction.

The coefficients on the treatment indicator variables all have the signs that are predicted by theory and the coefficients on the indicator variables for the Increasing κ and the Decreasing λ treatments are indeed remarkably close to what is predicted by theory. Even though we observe excessive vacancy posting in these treatments relative to the RE predictions, the *differences* in vacancy postings between the Increasing κ and Baseline treatments and the Decreasing λ and Baseline treatments are in-line with theoretical predictions.

The coefficient on the indicator variable for the π treatment has the proper sign, but it is only about 67 percent of the value predicted by theory. The *difference* in vacancy postings between the Increasing π treatment and the Baseline treatment is lower than the theoretical prediction. Firms in the Increasing π treatment under-responded to the higher surplus conditional on over-responding to the other parameters that make-up β_0 .

5.3. Firm-level vacancy postings

In this section we consider *firm* (or subject)-level vacancy posting behavior. We allow for the possibility that firms use information about

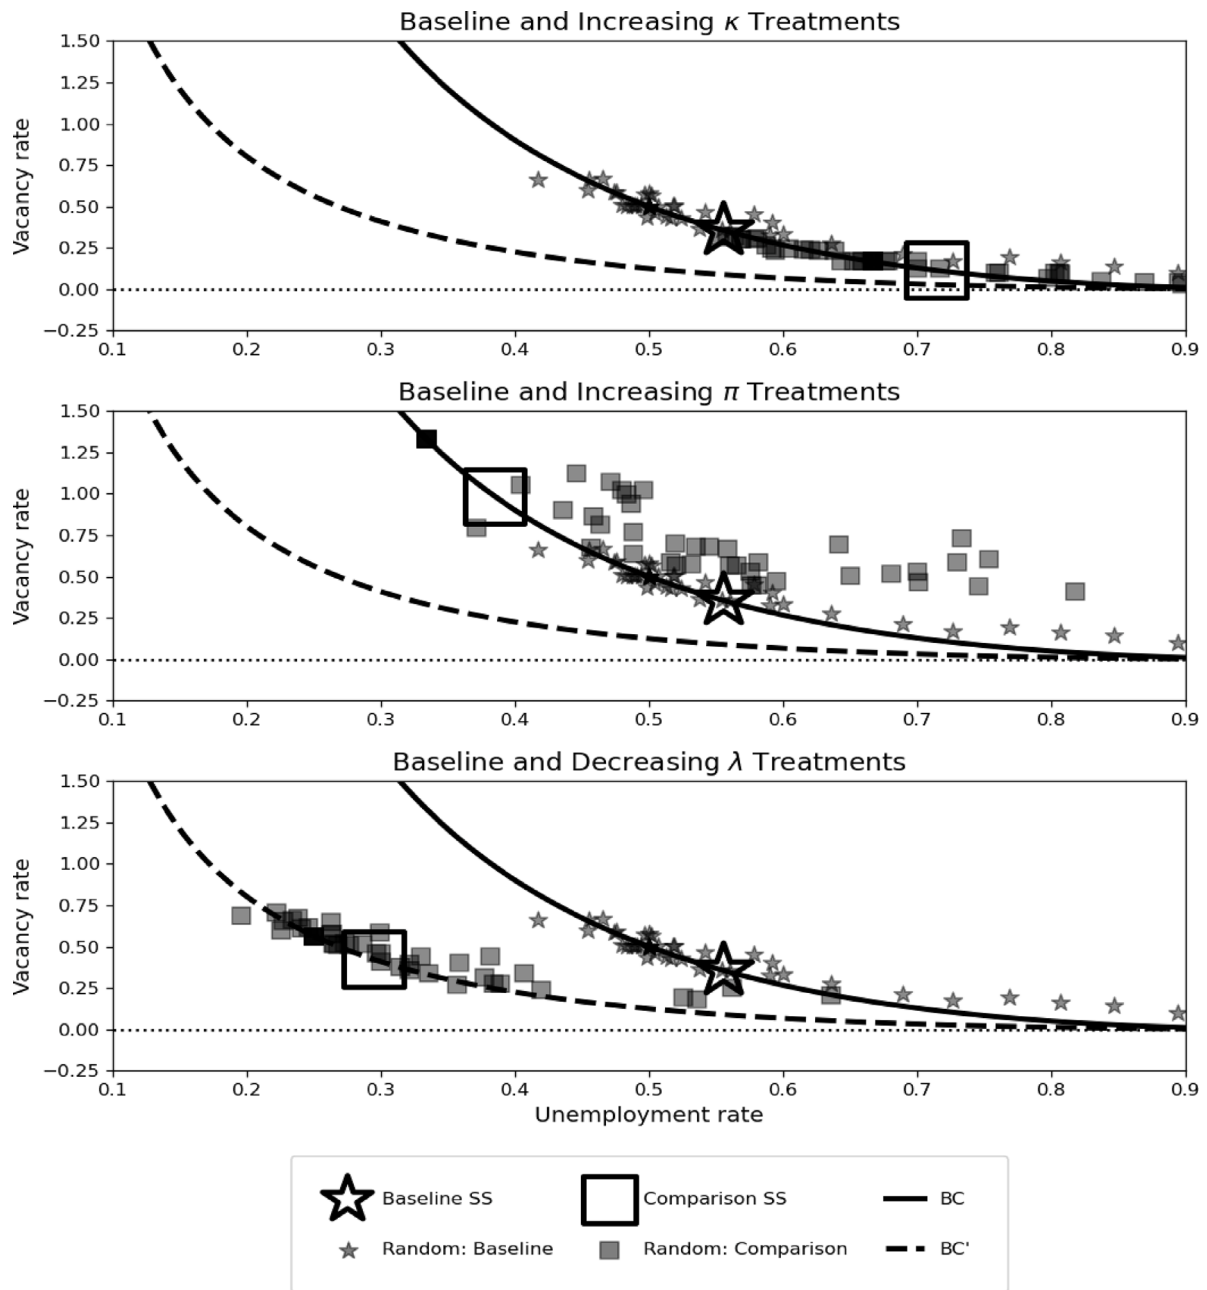


Fig. 7. Outcomes of a randomized vacancy posting exercise comparing the Baseline treatment (Random: Baseline) with the Increasing κ , Increasing π , and Decreasing λ treatments (Random: Comparison). The random data are constructed assuming that firms randomly draw vacancy postings from a discrete uniform distribution with a lower limit of 0 and an upper limit equal to 75% of the maximum vacancies allowed by subjects' point endowments in the first round of each sequence. The simulation assumes each firms post that same number of randomly drawn vacancies in every round of a given sequence. The relevant RE steady state equilibria and Beveridge curves are also shown.

their own past vacancy postings and employment changes in addition to responding to the different treatment conditions of the experiment. While theory predicts that vacancy posting decisions are not affected by either past vacancy postings or employment changes, we include these variables in the analysis to see if reaction to these quantities accounts for some of the difference between the experimental outcomes and the predictions from theory. Specifically, we estimate the following equation using data from all treatments and sessions:

$$V_{it} = \beta_0 + \beta_V V_{it-1} + \beta_E \Delta E_{it-1} + \beta_\pi D_{it}^\pi + \beta_\kappa D_{it}^\kappa + \beta_\lambda D_{it}^\lambda + \epsilon_{it} \quad (22)$$

Here V_{it} denotes the vacancy postings of firm i in period t and ΔE_{it-1} denotes the change in the number firm i 's employees (hires minus separations) in the previous period. The other regressors are dummy

indicator variables for treatment conditions. Table 6 reports our estimation of Eq. (22). We find a strong degree of autocorrelation in vacancy postings (around 0.9). In the symmetric rational expectations equilibrium, we would expect to find an autocorrelation coefficient of 1; that is not because it would be optimal for firms to respond to past vacancies, but rather because the optimal vacancy postings would change minimally between rounds and would be identical across firms. We find no evidence that firms are using the net change in employees in the previous round to determine their vacancy postings.

Next, we look at the distribution of vacancy postings across treatments. Fig. 8 plots three distributions for each treatment: (1) the RE predicted cumulative density functions (CDFs) for firm-level vacancy postings (2) the simulated RE firm-level prediction for this same density taking into account actual sequence lengths and the same number of

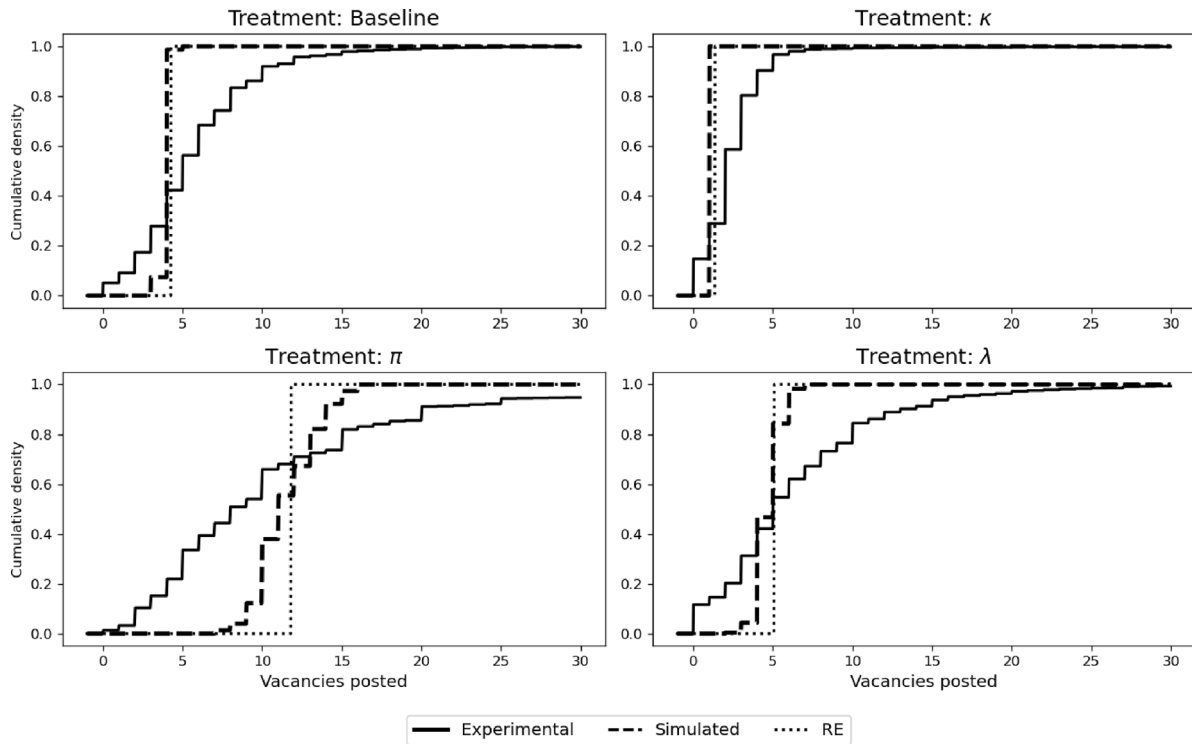


Fig. 8. Experimental and simulated cumulative distributions of vacancies posted by firms as compared with RE predictions.

Table 6

Regression analysis: Firm-level vacancy postings regressed on lagged firm vacancy postings, lagged net new hires and treatment indicators. Standard errors are clustered at the session level.

	V_{it}	V_{it}	V_{it}
Indicator: Vacancy cost treatment (β_κ)	-3.2984*** (0.3801)	-0.5634** (0.2360)	-0.5850** (0.2425)
Indicator: Surplus treatment (β_π)	6.0179*** (1.0656)	1.4898*** (0.3447)	1.6108*** (0.4036)
Indicator: Separation rate treatment (β_λ)	1.0310** (0.4889)	0.4125*** (0.1367)	0.5181** (0.1840)
Lagged vacancies (β_ν)		0.8436*** (0.0674)	0.9178*** (0.0860)
Lagged net new hires (β_μ)			0.1083 (0.0962)
Intercept (β_0)	5.6609*** (0.3267)	0.9016** (0.3873)	0.7781* (0.3867)
R-squared	0.1839	0.6592	0.6596
R-squared Adj.	0.1836	0.6591	0.6594
N	7360	7360	7360

players as in the experiment and finally (3) the actual “empirical” CDFs of vacancy postings for each treatment. Note that the firm level predictions are equal to the aggregate vacancy postings divided by 10 (the number of firms). In particular, for the Baseline treatment, the RE number of vacancy postings per firm is 4.267, for the Increasing κ treatment it is 1.371, for the Increasing π treatment it is 11.815, and for Decreasing λ treatment it is 5.082. Fig. 9 shows the actual CDFs of vacancy postings from all four treatments combined together in a single graph.

As Fig. 8 reveals, the cumulative distributions of firm level vacancy postings are more dispersed than theory or RE simulations would predict and are not degenerate at the RE predictions. Note that the simulated RE predictions are also not always degenerate at those point predictions either. More importantly, however, these distributions do shift with changes in the treatment conditions as is clearly revealed by Fig. 9. Specifically, the distribution of firm-level vacancy choices for the Increasing κ treatment lies to the left of the Baseline treatment

Table 7

D statistics and p-values from Kolmogorov–Smirnov tests on the equality of the distributions for individual vacancy postings for all treatment pairs. The alternative hypothesis is that the treatment on the left side of the “Treatment pairing” had lower vacancy posting values than the treatment on the right side of the pair.

Treatment pairing	D statistic	p-value
κ , Baseline	-0.5249	0.000
Baseline, π	0.3253	0.000
Baseline, λ	0.1022	0.000
κ , π	0.6843	0.000
κ , λ	0.4913	0.000
λ , π	0.2284	0.000

while the distribution of firm-level vacancy choices for the Increasing π and Decreasing λ treatments lie to the right of the Baseline treatment. Pairwise Kolmogorov–Smirnov tests of the null hypothesis of no difference between pairs of distributions of vacancy postings per firm are reported in Table 7 for all six pairwise comparisons. This table confirms the impression given by Fig. 9 that the empirical distributions are all significantly different from one another.

On the other hand, as noted, the distributions are not degenerate at the steady state point predictions, nor are they so close to the simulated distributions as is revealed in further Kolmogorov–Smirnov tests reported on in Table 8. This table reveals that the vacancy postings are in fact significantly different from both the RE steady state and the Simulated RE distributions.

5.4. Convergence in vacancy postings

In this section we study whether the vacancies posted by subjects (firms) in each treatment are tending toward some long-run value and whether that long-run value is statistically close to the rational expectations equilibrium. Duffy (2016) suggests estimating the following model:

$$v_{j,t} = (1 - \rho_j)\mu_j + \rho_j v_{j,t-1} + \epsilon_{j,t}, \quad (23)$$




Fig. 9. Experimental cumulative frequency distributions of vacancy postings by subject.

Table 8

D statistics and p-values for two-sided Kolmogorov–Smirnov tests on the equality of the distributions for individual vacancy postings from each treatment pairs against the respective RE equilibrium distribution and the simulated data.

Treatment	Empirical vs.	D statistic	p-value
Baseline	RE	0.5769	0.000
	Simulated	0.5638	0.000
κ	RE	0.7109	0.000
	Simulated	0.7109	0.000
π	RE	0.6803	0.000
	Simulated	0.4694	0.000
λ	RE	0.5467	0.000
	Simulated	0.3624	0.000

where $v_{j,t}$ is the average vacancies posted by firms in treatment j in round t . If $|\rho_j| < 1$, then average vacancy postings for the treatment is a stationary process that converges toward the average value μ_j . Further, one can evaluate this long-run average value relative to the RE steady state prediction.

We estimate Eq. (23) using maximum likelihood and report the results in Table 9. For all treatments, the point estimates of ρ_j are less than 1 in absolute value suggesting evidence for convergence. However, in all cases, the upper-bounds of the 95% confidence intervals exceed 1. The 95% confidence intervals for the process mean μ_j contain the RE equilibrium vacancy posting predictions for the Increasing π and Decreasing λ treatments. However, for the Baseline and Increasing κ treatments, the RE equilibrium vacancy posting predictions lie below and outside the 95% confidence interval, reflecting the over-posting of vacancies in these treatments.

In our experiment, each round is initialized with an unemployment rate that is close to, but not equal to, the RE steady state value. Therefore, the analysis above provides some perspective on the transitional dynamics of the DMP model. However, to study the transitional dynamics from one RE steady state to another following a parameter

change, our experimental implementation of the DMP would need to be modified to a “within subjects” design with the theory augmented to allow for how the potential for changing parameters would affect firms’ expectations when setting vacancies.

5.5. Behavioral explanations for results

We have demonstrated that for each treatment, the average vacancy rate and unemployment rates for each sequence lie near the predicted Beveridge curves and close the rational expectations point predictions. However, as noted there are deviations in the experimental results from the RE predictions that warrant further investigation. For the Baseline, Increasing κ , and Decreasing λ treatments, the average vacancy rate is sufficiently above the rational expectations equilibria to suggest a pattern of vacancy over-posting. Additionally, for all treatments, the average vacancy rate per sequence varies substantially with the greatest variation in the Decreasing λ and Increasing π treatments. We note that in all experimental sessions there is considerable variability in the number of vacancies posted by subjects each round. We suspect that subject heterogeneity may be an important factor explaining our results. In this section we explore the possibility that there might be subjects of different “types” operating according to different decision processes. We address this possibility by estimating a finite mixture model (FMM) with several latent classes (see, e.g., Moffatt (2016)).

Let v_i denote the average vacancies posted by subject i . With probability p_j the pdf of v_i is $f_j(v_i|\mu_j, \sigma_j^2)$, where f_j is the density function for a normal distribution truncated below at zero and with $\sum_j p_j = 1$. We estimate the FMM model assuming three latent classes.⁹

⁹ We estimate the FMM for 1 to 5 latent classes for each treatment. The AIC suggests three latent classes for the Decreasing λ and Increasing π treatments, two classes for the Baseline treatment, and one class for the increasing κ treatment. Since three latent classes is the modal number minimizing the AIC, we report results for the three latent class model for all treatments so that the results are comparable.

Table 9

Estimates of autocorrelation and mean for average vacancies posted per round by treatment. The final column is rational expectations equilibrium vacancies posted per firm.

	ρ_j			μ_j			
	Estimate	S.E.	95% C.I.	Estimate	S.E.	95% C.I.	RE
Baseline	0.70	0.45	[-0.18,1.58]	6.00	0.45	[5.11,6.89]	4.27
Increasing κ	-0.28	1.04	[-2.31,1.76]	2.44	0.09	[2.26,2.62]	1.37
Increasing π	0.87	0.30	[0.29,1.45]	12.75	2.41	[8.04,17.47]	11.82
Decreasing λ	0.61	0.56	[-0.48,1.70]	6.20	1.80	[2.68,9.72]	5.08

Table 10

Finite mixture model estimates by class and treatment. Three latent classes. The table reports estimated mean vacancy postings and variances for each class, denoted by μ and σ^2 , and proportions, p , of subjects in each class. Standard errors are in parentheses.

Treatment	RE	Class 1			Class 2			Class 3		
		μ	σ^2	p	μ	σ^2	p	μ	σ^2	p
Baseline	4.27	2.60 (0.19)	0.08 (0.06)	0.09 (0.06)	5.44 (0.35)	3.02 (0.92)	0.81 (0.12)	10.84 (4.99)	13.20 (16.51)	0.09 (0.11)
κ	1.37	0.80 (0.04)	0.01 (0.00)	0.09 (0.04)	2.33 (0.11)	0.42 (0.12)	0.84 (0.06)	4.02 (0.09)	0.02 (0.02)	0.07 (0.04)
π	11.82	5.96 (0.76)	4.90 (2.46)	0.47 (0.16)	13.60 (2.06)	23.59 (10.89)	0.49 (0.16)	38.07 (8.28)	54.46 (94.58)	0.04 (0.03)
λ	5.08	2.20 (0.06)	0.01 (0.01)	0.10 (0.05)	4.86 (0.22)	0.35 (0.25)	0.27 (0.09)	7.87 (0.95)	23.43 (6.24)	0.62 (0.10)

Table 10 reports the estimated parameters of the finite mixture model for the three classes in each treatment. In general we find evidence of a “low mean” group of vacancy posters in each treatment, a group of vacancy posters reasonably close to the RE equilibrium value, and a group of “high mean” posters. The share of subjects estimated to be from each class varies by treatment.

For the Baseline, Increasing κ , and Decreasing λ treatments, we find that about 10 percent of subjects belong to the “low mean” group, while nearly 50 percent of the subjects in the Increasing π treatment are from the “low mean” group. In the Baseline and Increasing κ treatments, slightly more than 80 percent of subjects are estimated to belong to the “near RE” group, while the estimates for Increasing π and Decreasing λ treatments are about 50 percent and 27 percent. For the Baseline, Increasing κ , and Increasing π treatments, the share of subjects in the “high mean” group ranges from 4 to 9 percent and, by contrast, about 62 percent of the subjects in the Decreasing λ treatment are estimated to be “high mean” subjects.

The FMM estimates suggest that each treatment contained subjects with different propensities to post vacancies. We regard the “high mean” subjects as those who possibly perceived there to be greater uncertainty about the number of workers they would be able to employ and so they posted more than the theoretically optimal number of vacancies on average. The excessive vacancy posting of the “high mean” subjects affects the other subjects because the vacancy filling rate is driven down for all by their behavior. In our view, members of the “low mean” group are subjects who respond strategically to the excessive posting of the “high mean” group by posting less than the RE equilibrium number of vacancies. Alternatively, the causality could run in the other direction.

6. Conclusion

We have designed and reported on a laboratory experiment in which subjects, in the role of firms, choose how many job vacancies to post in a decentralized matching model of unemployment and the labor market with real search and matching frictions. The total number of job vacancy postings affects the vacancy filling rate and therefore the economy-wide vacancy and unemployment rates.

To our knowledge, this is the first experimental test of the workhorse Diamond–Mortensen–Pissarides model. In essence we are testing the empirical relevance of that model for generating a Beveridge

curve, which is the inverse relationship between unemployment and the vacancy rate, and for producing comparative statics predictions in response to changes in labor market conditions. We also consider the role played by strategic uncertainty in firms’ vacancy posting decisions by relaxing the representative firm assumption. The advantage of the experimental method is that we induce a matching function and we know the vacancy posting costs, the benefit value of employees to firms, the employee separation rate and other model parameters so that we can assess the empirical relevance of the theory at a more detailed, micro-level than would be possible using field data alone. In other words, we have good *internal validity*. A disadvantage, of course, is that one has to exercise caution in making claims about the external validity or generalizability of experimental findings outside of the laboratory. However, that same criticism also applies to the theory itself.

Overall, we find strong evidence in our experimental data for a downward-sloping Beveridge curve across all treatments. We further find that the comparative statics resulting from changes in model parameters are consistent with RE theoretical predictions despite heterogeneity in subject behavior. Our data suggest that subjects playing the role of firms vary their vacancy postings numbers in response to the different treatments that we explore in a manner that is consistent with the theory but that there is often an upward-bias in vacancy postings which we explain as being primarily the result of a small number of subjects posting far more than the optimal number of vacancies.

The experimental exercise reveals that, as a whole, subjects do a good job of approximately solving the complicated optimization problem that the DMP model requires. The experiment further reveals that the support for the comparative statics predictions of the theory is the result of deliberate behavior on the part of subjects and cannot be explained by random vacancy postings in combination with the matching function. Indeed, the matching function does not mechanically generate movements of equilibria along a Beveridge curve. Finally, the experiment demonstrates that the predictions of the representative agent DMP model are robust to heterogeneous subject behavior.

In future research, it would be of interest to explore whether and how the same group of subjects react to changes in the model’s parameters as in a “within subjects” design. Such a design would enable study of the model’s predicted transitional dynamics between steady states. One could also have subjects play the role of both firms and workers, which would increase the complexity of the environment. Further, it would be of interest to use our framework to study the Shimer

puzzle (Shimer, 2005): that vacancy fluctuations due to productivity shocks are too large to be explained by the DMP model. While we did not design our experiment to test the Shimer puzzle, we think that our experimental framework could provide a fruitful avenue for understanding that puzzle. We leave these and other extensions to future research.

CRedit authorship contribution statement

John Duffy: Conceptualization, Methodology, Writing – original draft, Writing – review & editing. **Brian C. Jenkins:** Data curation, Formal analysis, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.labeco.2024.102518>.

References

- Bhattacharya, Arghya, Jackson, Paul, Jenkins, Brian C., 2018. Revisiting unemployment in intermediate macroeconomics: A new approach for teaching Diamond-Mortensen-Pissarides. *J. Econ. Educ.* 49 (1), 22–37.
- Brandts, Jordi, Giritligil, Ayça Ebru, 2008. Entry and market selection of firms: A laboratory study. *J. Econ. Behav. Organ.* 68 (3), 593–612.
- Brandts, Jordi, Giritligil, Ayça Ebru, Cabrales, Antonio, Charness, Gary, 2007. Forward induction and entry deterrence: an experiment. *Econom. Theory* 33 (1), 183–209.
- Braunstein, Yale M., Schotter, Andrew, 1982. Labor market search: An experimental study. *Econ. Inq.* 20 (1), 133–144.
- Brown, Meta, Flinn, Christopher J., Schotter, Andrew, 2011. Real-time search in the laboratory and the market. *Amer. Econ. Rev.* 101 (2), 948–974.
- Chen, Daniel L., Schonger, Martin, Wickens, Chris, 2016. oTree—An open-source platform for laboratory, online, and field experiments. *J. Behav. Exp. Finance* 9, 88–97.
- Cox, James C., Oaxaca, Ronald L., 1989. Laboratory experiments with a finite-horizon job-search model. *J. Risk Uncertain.* 2 (3), 301–329.
- Cox, James C., Oaxaca, Ronald L., 1992. Direct tests of the reservation wage property. *Econ. J.* 102 (415), 1423–1432.
- Diamond, Peter A., 1982. Wage determination and efficiency in search equilibrium. *Rev. Econom. Stud.* 49 (2), 217–227.
- Duffy, John, 2016. Macroeconomics: A survey of laboratory research. In: Kagel, John H., Roth, Alvin E. (Eds.), *The Handbook of Experimental Economics, Volume 2*. Princeton University Press, pp. 1–90.
- Elsby, Michael W.L., Michaels, Ryan, Ratner, David, 2015. The beveridge curve: A survey. *J. Econ. Lit.* 53 (3), 571–630.
- Eriksson, Stefan, Rooth, Dan-Olof, 2014. Do employers use unemployment as a sorting criterion when hiring? Evidence from a field experiment. *Am. Econ. Rev.* 104 (3), 1014–1039.
- Fréchette, Guillaume R., 2015. Laboratory experiments: Professionals versus students. In: Fréchette, Guillaume R., Schotter, Andrew (Eds.), *Handbook of Experimental Economic Methodology*. Oxford University Press, pp. 360–390.
- Ghayad, Rand, 2013. The Jobless Trap. Working paper, Northeastern University.
- Harrison, Glenn W., Morgan, Peter, 1990. Search intensity in experiments. *Econ. J.* 100 (401), 478–486.
- Korenok, Oleg, Munro, David, 2021. Wage bargaining in a matching market: Experimental evidence. *Labour Econ.* 73, 1–13.
- Kroft, Kory, Lange, Fabian, Notowidigdo, Matthew J., 2013. Duration dependence and labor market conditions: Evidence from a field experiment. *Q. J. Econ.* 128 (3), 1123–1167.
- Moffatt, Peter G., 2016. *Experimentics: Econometrics for Experimental Economics*. Red Globe Press.
- Mortensen, Dale T., Pissarides, Christopher A., 1994. Job creation and job destruction in the theory of unemployment. *Rev. Econom. Stud.* 120 (5), 397–415.
- Petrongolo, Barbara, Pissarides, Christopher A., 2001. Looking into the black box: A survey of the matching function. *J. Econ. Lit.* 39 (2), 390–431.
- Petrosky-Nadeau, Nicolas, Zhang, Lu, 2013. Empirical Matching Functions: Estimation and Interpretation Using Disaggregate Data. Fisher College of Business Working Paper No. 2014-03-11.
- Pissarides, Christopher A., 1985. Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *Am. Econ. Rev.* 75 (4), 676–690.
- Roth, Alvin E., Murnighan, J. Keith, 1978. Equilibrium behavior and repeated play of the prisoner's dilemma. *J. Math. Psych.* 17 (2), 189–198.
- Schotter, Andrew, Braunstein, Yale M., 1981. Economic search: an experimental study. *Econ. Inq.* 19 (1), 1–25.
- Sherstyuk, Katerina, Tarui, Nori, Saijo, Tatsuyoshi, 2013. Payment schemes in infinite-horizon experimental games. *Exp. Econ.* 16 (1), 125–153.
- Shimer, Robert, 2005. The cyclical behavior of equilibrium unemployment and vacancies. *Amer. Econ. Rev.* 95 (1), 25–49.
- Van Ours, Jan C., Vodopivec, Milan, 2006. How shortening the potential duration of unemployment benefits affects the duration of unemployment: Evidence from a natural experiment. *J. Labor Econ.* 24 (2), 351–378.