Journal Pre-proof

Nonlinear unemployment effects of the inflation tax

Mohammed Ait Lahcen, Garth Baughman, Stanislav Rabinovich,
Hugo van Buggenum

PII: S0014-2921(22)00146-5
DOI: https://doi.org/10.1016/j.euroecorev.2022.104247
Reference: EER 104247

To appear in: *European Economic Review*

Received date : 19 October 2021
Revised date : 1 July 2022
Accepted date : 9 July 2022

Please cite this article as: M.A. Lahcen, G. Baughman, S. Rabinovich et al., Nonlinear unemployment effects of the inflation tax. *European Economic Review* (2022), doi: https://doi.org/10.1016/j.euroecorev.2022.104247.

Nonlinear Unemployment Effects of the Inflation Tax*

Mohammed Ait Lahcen†
Qatar University

Garth Baughman
Federal Reserve Board of Governors

Stanislav Rabinovich‡
University of North Carolina at Chapel Hill

Hugo van Buggenum§
ETH Zürich

July 1, 2022

Abstract

Long-run inflation has nonlinear and state-dependent effects on unemployment, output, and welfare. We show this using a standard monetary search model with two shocks – productivity and monetary – and frictions in both labor and goods markets. Inflation lowers the surplus from a worker-firm match, in turn making it more sensitive to both productivity shocks and further increases in inflation. We calibrate the model to match key aspects of the US labor market and monetary data. The calibrated model is consistent with a number of empirical correlations, which we document using panel data from the OECD: (1) there is a positive long-run relationship between anticipated inflation and unemployment; (2) there is also a positive correlation between anticipated inflation and unemployment volatility; (3) the long-run inflation-unemployment relationship is stronger when unemployment is higher. The key mechanism through which the model generates these results is the negative effect of inflation on measured output per worker, which is likewise consistent with cross-country data. Finally, we show that the welfare cost of inflation is nonlinear in the level of inflation and is amplified by the presence of aggregate uncertainty.

*The views expressed are those of the authors and do not reflect the official position of the Federal Reserve System or the Board of Governors. We are very grateful to the editor, Peter Rupert, and an anonymous referee for numerous constructive comments that significantly improved the paper. We would also like to thank Lukas Altermatt, Aleks Berentsen, Zach Bethune, Paul Jackson, Kuk Mo Jung, Athanasios Geromichalos, Pedro Gomis-Porqueras, Serene Tan and seminar participants at the University of Basel, the National University of Singapore, Qatar University, Doha Institute for Graduate Studies, the Federal Reserve Board and the Sogang Search and Matching Workshop for valuable comments and suggestions.

†Department of Finance and Economics, Qatar University, P. O. Box: 2713, Doha, Qatar. e-mail: m.aitlahcen@gmail.com

‡Department Economics, University of North Carolina at Chapel Hill, 107 Gardner Hall, CB 3305. Chapel Hill, NC 27599-3305 Email: srabinov@email.unc.edu

§ETH Zürich, Zürichbergstrasse 18, 8092 Zürich, Switzerland. e-mail: hugo.buggenum@gmail.com

1 Introduction

By acting as a tax on cash-intensive activity, long-run inflation leads to reduced output per worker, and, in the presence of labor market frictions, increased unemployment. This result is a robust prediction of the monetary search literature that explicitly models both goods and labor market frictions, such as Berentsen et al. (2011), as well as much of the literature preceding it that models money in reduced form, e.g. Cooley and Hansen (1989) and Lucas (2000). However, research on the effects of anticipated inflation has focused predominantly on its effect on average outcomes, largely abstracting from whether or how it affects the business cycle. In this paper, we argue that long-run inflation has significant effects on short-run unemployment volatility; more generally, the nonlinear and state-dependent effects of inflation should be taken into account when evaluating its effects on employment, output and welfare. We motivate this empirically, using cross-country data. We then show the importance and sources of nonlinearities quantitatively, using a standard monetary search framework with goods and labor market frictions.

We first use cross-country panel data from the Organization for Economic Cooperation and Development (OECD) to document three empirical correlations. First, as already expected from the existing literature and confirmed e.g. by Berentsen et al. (2011) in US data, the long-run correlation between anticipated inflation and unemployment is positive.[1] Second, there is also a positive correlation between anticipated inflation and unemployment volatility. Third, the relationship between anticipated inflation and the level of unemployment is also state-dependent and nonlinear: the positive correlation is stronger when unemployment is higher.

We then explore whether, and through what channels, a standard monetary search model can rationalize these correlations. Our model environment closely follows the model of Berentsen et al. (2011), which combines labor and goods market frictions. Firms hire workers in a frictional labor market, and then sell some of their production to households in a frictional goods market where money is essential. We extend this framework by introducing both monetary and productivity shocks, so that our model allows for the analysis of unemployment volatility in response to these shocks, as well as the interaction between them.

Theoretically, several forces might amplify shocks and drive nonlinear behavior in this framework. To start with, insights from the analysis of unemployment volatility in the Diamond-Mortensen-Pissarides (DMP) framework (Shimer (2005), Hagedorn and Manovskii (2008), Ljungqvist and Sargent (2017)) carry over naturally to the Berentsen et al. (2011) environment. A small surplus from an employment match amplifies the responsiveness of vacancy creation, and hence unemployment, to shocks. This well-known result turns out to have a number of novel implications once both productivity shocks and monetary shocks are present and interact. By reducing

[1] Throughout, we use long-term nominal interest rates as a proxy for anticipated inflation; see the discussion in section 3.

monetary trade in the goods market, inflation lowers the surplus from a job match, thus making unemployment more responsive to productivity shocks and amplifying unemployment volatility. Through the same channel, a higher level of inflation also makes unemployment more sensitive to further increases in inflation, implying a nonlinear effect of inflation on unemployment levels. Conversely, a decrease in productivity likewise makes unemployment more sensitive to changes in inflation. This two-way interaction between real and monetary shocks has the potential to rationalize the empirical patterns described above and generate significant nonlinearities. Moreover, the nonlinearity of the matching technology itself may drive some of the model's dynamics, as emphasized in Bernstein et al. (2021).

Because the framework upon which we build, Berentsen et al. (2011), features both monetary trade and goods market frictions, it therefore contains mechanisms absent from the benchmark DMP model. Endogenous money demand implies that the relative size of the monetary sector, an important factor for the aforementioned amplification effects, depends itself on inflation. Moreover, goods market frictions generate a feedback effect from employment to match output. An increase in employment lowers a firm's probability of finding a buyer in the goods market, thus lowering its profits from a match, all else equal; on the other hand, it also raises a buyer's probability of finding a seller, thereby raising a buyer's incentive to carry real balances. Goods market frictions may thus act as either an amplifying or a dampening mechanism, depending on parameters. As a result, there are multiple mechanisms potentially contributing to the model's dynamics; both the signs of these channels and their relative importance are a priori unclear, motivating our quantitative analysis.

To assess the model's performance as well as its implications for volatility and welfare, we simulate the model numerically. We discipline the model by calibrating parameters to match salient features of both money demand data, such as velocity, and labor market data, such as labor market flows and unemployment volatility. We validate the model based on its ability to replicate the correlation between inflation and unemployment volatility, as well as the nonlinear correlation between inflation and unemployment levels that we find in the data. We find that the model matches these regularities well. In line with the aforementioned match surplus channel, it also predicts a negative correlation between anticipated inflation and measured labor productivity, which we likewise confirm in the cross-country data. We illustrate the nonlinear and state-dependent dynamics of the model by computing generalized impulse response functions following a negative productivity shock. The response of employment, output, and monetary trade to shocks is stronger when trend inflation is high. For example, the average increase in unemployment on impact, in response to a one standard-deviation productivity shock, is on average 1.5 times stronger when trend inflation is 8% than when it is 3%.

We also examine in more detail the sources of nonlinear behavior. As mentioned above, the obvious candidate is the conventional small-surplus channel: inflation lowers the match surplus, thereby amplifying its response to both productivity shocks and further changes in inflation. As pointed out by Bernstein et al. (2021), the shape of the matching function itself implies that shocks have a larger effect on unemployment when unemployment is already high, further contributing to nonlinear dynamics. However, goods market frictions specific to the Berentsen et al. (2011) introduce additional feedback effects from employment to match surplus, due to either buyers' money demand responses or congestion effects in the goods market. In Section 7 we argue, via a steady-state decomposition exercise, that nonlinear effects in the calibrated model are driven primarily by the small-surplus channel, though the nonlinearity of the matching function also plays an important role. On the other hand, feedback effects operating through goods market frictions are less significant, both because these feedback effects consist of opposing channels that offset each other, and because the fraction of total trade taking place in the frictional goods market is small in the calibrated model. These claims are also confirmed by robustness checks in Section 8.4, which show that it is the labor market parameters (notably the value of unemployment consumption and the elasticity of the matching function) that are of primary importance for the model's amplification mechanism.

Finally, we use the model to evaluate the welfare cost of inflation and the extent to which its interaction with volatility matters. We find that increasing the trend inflation rate from -2.75% (which corresponds to a 0% nominal interest rate) to 10% leads to a welfare loss of 5.28% in consumption-equivalent terms. Decomposing this number reveals that the welfare cost is nonlinear in the level of inflation. For example, increasing inflation from -2.75% to 5% reduces welfare by 2.54 percentage points, while increasing it from 5% to 10% leads to an additional welfare loss of 2.75 percentage points. We then analyse the contribution of aggregate uncertainty to the cost of inflation. An identical economy without aggregate shocks implies a welfare loss of only 4.88%, compared to 5.28% in the baseline economy. This amplification effect matters only at high levels of inflation. Our results thus suggest that the interaction of high inflation and aggregate uncertainty is important for assessing welfare effects – both because inflation amplifies the responsiveness of unemployment to further increases in inflation, and because inflation amplifies unemployment volatility in response to productivity shocks.

Section 2 positions our paper in the context of the existing literature on monetary search models, the effects of inflation, and unemployment dynamics. Section 3 describes our empirical analysis and findings using OECD data. Section 4 lays out the model environment, and Section 5 characterizes its equilibrium conditions. Section 6 describes our calibration strategy. In Section 7, we inspect the mechanisms driving the model's nonlinearities. We do so by analytically characterizing the model's comparative statics properties in steady state, and then provide a numerical decomposition of these comparative statics in the calibrated model. Section 8 reports the main quantitative results, verifying that the model is consistent with the documented empirical correlations, and providing additional robustness checks. Section 9 draws implications for the welfare cost of inflation. Section 10 concludes.

2 Relationship to literature

Our paper contributes to the research on the welfare cost of inflation in micro-founded models of money demand, starting with the work by Lagos and Wright (2005). Our work complements a growing literature combining goods and labor market frictions to study the effects of monetary policy on unemployment, starting with Berentsen et al. (2011) and developed in work such as Gomis-Porqueras et al. (2013), Rocheteau and Rodriguez-Lopez (2014), Bethune et al. (2015), Gu et al. (2021), Bethune and Rocheteau (2017), Dong and Xiao (2019), Ait Lahcen (2020), Gomis-Porqueras et al. (2020), He and Zhang (2020), and Jung and Pyun (2020), among others. As mentioned above, this literature has largely viewed short-run volatility as orthogonal to the long-run welfare effect of anticipated inflation. Perhaps this view stems from the conclusion of Cooley and Hansen (1989) that, in a *frictionless* environment, inflation has no first-order effect on the business cycle. Our findings imply that this view is not innocuous and that this conclusion is overturned in a frictional economy. To our knowledge, this is the first paper to analyze the effects of inflation in a fully stochastic model with both real and monetary shocks, and both goods and labor market frictions. There is, of course, a rich literature studying the short-run effects of monetary policy in the presence of nominal rigidities. Moreover, an important recent line of research, exemplified by the recent findings of Ascari et al. (2022) and surveyed in Ascari and Sbordone (2014), shows that trend inflation matters in New Keynesian models as well. We see our work as complementary: we deliberately abstract from any consideration of nominal rigidities emphasized in that literature, since we want to demonstrate that inflation matters for volatility, and its effect on volatility matters for welfare, in a fully flexible-price economy where money serves a medium-of-exchange role.

Methodologically, we see our results as providing a quantitative validation of the Berentsen et al. (2011) framework, which, as demonstrated by the above literature, has become a workhorse model for thinking about monetary-labor market interactions. Moreover, relative to the previous literature, we provide a more detailed inspection of which mechanisms are quantitatively dominant in this framework. By explicitly modeling both goods and labor market frictions, the Berentsen et al. (2011) model generates feedback mechanisms above and beyond the well-known amplification channels in the standard DMP framework: e.g., as mentioned above, households' money demand depends on the employment level through the probability of meeting a seller and, conversely, a firm's profits from a match are endogenous and depend on the employment level as well. We nevertheless show that the conventional small-surplus channel, rather than these goods-labor market interactions, is the dominant force quantitatively in driving the model's nonlinearities.

Our paper also contributes to the research on unemployment volatility, by examining how goods market frictions and monetary trade impact unemployment fluctuations. In fact, the analysis in this paper directly builds on the well-known insight from that literature that a small job match surplus leads to high unemployment volatility (Shimer (2005), Hagedorn and Manovskii (2008), Ljungqvist and Sargent (2017)). To our knowledge, we are the first to draw implications for the effect of the inflation tax. The highly nonlinear behavior of our stochastic model is also directly related to similar findings in Petrosky-Nadeau and Zhang (2017, 2020), and Petrosky-Nadeau et al. (2018). Like them, we employ a global solution method in order to accurately characterize the model-implied dynamics. Our finding that the shape of the matching technology contributes substantially to nonlinear behavior is also in line with the recent results of Bernstein et al. (2021). More broadly, the interaction of goods and labor market frictions also connects our work to the research by Petrosky-Nadeau and Wasmer (2015), who likewise find that such interaction has important effects for the resulting dynamics, though trade is non-monetary in these frameworks and hence they do not speak to the effects of the inflation tax.

3 Empirical evidence

In this section we present some suggestive cross-country evidence concerning the long-run relationship between anticipated inflation, unemployment, and unemployment volatility. We make no claims of causality here: the correlations we display are meant to be illustrative and serve as motivating evidence for the mechanism we highlight. We will also verify, in our numerical analysis, that our calibrated model can reproduce the correlations found here.

We use quarterly data from the Main Economic Indicators database published by the OECD, covering 35 developed countries. The start date of the period covered varies between countries and ends in 2019. On the labor market side, we use the harmonised unemployment rate series in order to ensure that the data are consistent across countries. We use the long-term nominal interest rate as our proxy for the opportunity cost of holding nominal balances; this series consists mostly of yields on government bonds with a 10-year maturity. Theoretically the measure relevant for our mechanism is anticipated inflation, rather than realized inflation; hence our focus on the long-term nominal interest rate. We extract the trend component of each data series using the HP filter (Hodrick and Prescott, 1997) with a smoothing parameter value of 1,600. Appendix A shows that using the 5-year moving average, instead of the HP filter, yields very similar results.

3.1 A positive relationship between $\bar{u}$ and $\bar{\iota}$

We first regress the trend component of unemployment on the trend component of the long-term nominal interest rate. Table 1 shows a positive and significant relationship. In particular, Column

(1) presents the results of the pooled OLS regression

$$\bar{u}_{jt} = \alpha + \beta\bar{\iota}_{jt} + \varepsilon_{jt}, \tag{1}$$

where $\bar{u}_{jt}$ and $\bar{\iota}_{jt}$ represent the trend components of unemployment and the long-term nominal interest rate in country j at quarter t. It indicates that a 1 percentage point (pp) increase in $\bar{\iota}$ is associated with a 0.35pp increase in $\bar{u}$. Columns (2), (3) and (4) present results from fixed effects panel regressions of the type

$$\bar{u}_{jt} = \alpha + \beta\bar{\iota}_{jt} + \gamma_j + \delta_t + \varepsilon_{jt}, \tag{2}$$

where γ_j and δ_t represent, respectively, country and time fixed effects. The results of all three panel regressions confirm the positive relationship between $\bar{u}$ and $\bar{\iota}$. In particular, the estimates presented in Column (4), which include both country and time fixed effects, indicate that a 1pp increase in $\bar{\iota}$ is associated with a 0.77pp increase in $\bar{u}$, an economically and statistically significant relationship.

To verify the robustness of our finding, Table 14 in Appendix A reports the results of running the same regression specifications on trend components of unemployment and the long-term nominal interest rate extracted using a 5-year moving average instead of the HP filter. The results are similarly significant and strong. In particular, the specification in Column (4), which includes both country and time fixed effects, indicates that a 1pp increase in $\bar{\iota}$ is associated with a 0.915pp increase in $\bar{u}$. Finally, Tables 15 and 16 in Appendix A present the regression results using the logarithm of unemployment instead of its level with HP-filtered and 5-year moving average trends; the results are again very similar.[2] These cross-country regression results are thus consistent with the positive long-run relationship between unemployment and inflation in the US, documented by Beyer and Farmer (2007), Berentsen et al. (2011) and Haug and King (2014).[3]

3.2 A positive relationship between unemployment volatility and $\bar{\iota}$

Next, we seek to provide some evidence on the relationship between anticipated inflation and unemployment volatility. We use the HP-filtered cyclical component of the logarithm of unemployment. Our measure of volatility $\sigma_{u_{jt}}$ is the 5-year rolling window standard deviation of this cyclical component. Column (1) of Table 2 presents the results of the pooled OLS regression

$$\sigma_{u_{jt}} = \alpha + \beta\bar{\iota}_{jt} + \varepsilon_{jt}, \tag{3}$$

[2] In addition, we ran similar panel regressions of either the unfiltered or the cyclical components of unemployment on the short-term nominal interest rate. The resulting coefficients are either statistically not significant (for the unfiltered data) or significant and negative (for the HP-detrended data). These results are available on request.

[3] In additional robustness checks, available on request, we have also verified that the empirical correlations documented here are not driven by particular outlier countries.

Table 1: Regression of $\bar{u}$ on $\bar{\imath}$

	Trend unemployment			
	(1)	(2)	(3)	(4)
Constant	5.773***	6.035***	3.750***	3.507***
	(0.505)	(0.362)	(1.366)	(1.205)
Trend long-term rate	0.350***	0.302***	0.725**	0.770***
	(0.091)	(0.067)	(0.288)	(0.223)
Observations	4,007	4,007	4,007	4,007
R^2	0.086	0.140	0.121	0.135
F-Statistic	376.96***	646.04***	513.26***	581.27***
Country fixed effects	No	Yes	No	Yes
Time fixed effects	No	No	Yes	Yes
Clustered errors (country level)	Yes	Yes	Yes	Yes

*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parentheses. Data are from the OECD. Both unemployment and long-term nominal interest rate series for each country are filtered out of high frequency variations using the HP filter with $\lambda = 1600$.

while columns (2), (3) and (4) present the results of the fixed effects panel regression

$$\sigma_{u_{jt}} = \alpha + \beta\bar{\imath}_{jt} + \gamma_j + \delta_t + \varepsilon_{jt}, \tag{4}$$

with country and/or time fixed effects. Our preferred specification in Column (4) shows that an increase of 1pp in $\bar{\imath}$ is associated with a 0.011 increase in the volatility of unemployment, which corresponds to about 12% of the average volatility. Table 17 in Appendix A presents very similar results obtained by running the same panel regressions using the level of unemployment instead of the logarithm of unemployment.

In Appendix A, we also compute unemployment volatility based on the cyclical component of the logarithm of unemployment extracted assuming a 5-year moving average as the trend, and then regress it on the 5-year moving average of the long-run interest rate. Table 18 presents the results. Column (4) with country and time fixed effects shows that an increase of 1pp in $\bar{\imath}$ is associated with a 0.026pp increase in unemployment volatility. This increase corresponds to about 18% of average unemployment volatility. We also run the same panel regressions using the level of unemployment and the results, presented in Table 19 in Appendix A, are very similar.

Table 2: Regression of unemployment volatility on $\bar{\iota}$

	log unemployment volatility			
	(1)	(2)	(3)	(4)
Constant	0.058***	0.052***	0.060***	0.031
	(0.005)	(0.007)	(0.011)	(0.023)
Trend long-term rate	0.005***	0.006***	0.005**	0.011**
	(0.001)	(0.001)	(0.002)	(0.005)
Observations	3,616	3,616	3,616	3,616
R^2	0.079	0.114	0.031	0.062
F-Statistic	309.20***	462.40***	109.11***	221.28***
Country fixed effects	No	Yes	No	Yes
Time fixed effects	No	No	Yes	Yes
Clustered errors (country level)	Yes	Yes	Yes	Yes

*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parenthesis. Data are from the OECD. Unemployment volatility is measured as the 5-year rolling window standard deviation of HP-detrended log unemployment. Long-term nominal interest rate series for each country is filtered out of high frequency variations using the HP filter with $\lambda = 1600$.

3.3 A state-dependent relationship between $\bar{u}$ and $\bar{\iota}$

Finally, we want to examine whether the positive long-run relationship observed between unemployment and anticipated inflation varies with the level of unemployment. To do so, we use a quantile regression specification. As opposed to the linear regression, which approximates the conditional expectation function, the quantile regression approximates the conditional quantile function at quantile q by the linear relationship

$$\mathcal{Q}_q(\bar{u}_{jt}|\bar{\iota}_{jt}) = \alpha_q + \beta_q \bar{\iota}_{jt} + \varepsilon_{qjt}. \tag{5}$$

By estimating the above regression for different quantiles of the distribution of $\bar{u}$, we can check whether the relationship between the low frequency components of unemployment and nominal interest rates varies with the level of unemployment. Figure 1 plots, in solid blue, the values of β_q for various values of q estimated by pooling together all observations in our sample, with the dashed blue lines depicting the 95% confidence interval. The coefficient of the pooled OLS regression discussed above is depicted as the flat red line for comparison. The relationship is clearly nonlinear. At the 5th percentile of the distribution of trend unemployment, i.e. $q = 5\%$, a

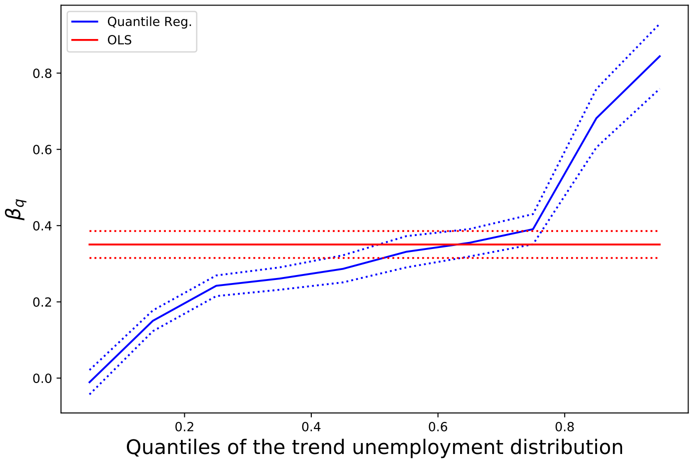

Figure 1: Quantile regression coefficients of $\bar{u}$ on $\bar{\iota}$ for various quantiles of $\bar{u}$.
Notes: The dotted lines represent the 95% confidence intervals. Data are from the OECD. Both unemployment and long-term nominal interest rate series for each country are filtered out of high frequency variations using the HP filter with $\lambda = 1600$.

1pp increase in the trend long-term nominal interest rate is associated with almost no change in trend unemployment, whereas at the 95th percentile the coefficient reflects an increase of 0.84pp in trend unemployment. The quantile regression depicted in Figure 8 in Appendix A uses 5-year moving averages and the ones depicted in Figures 9 and 10 use the logarithm of unemployment with 5-year moving averages and HP-filtered trends, respectively. The results are very similar.

3.4 Summary of the empirical results

We have provided illustrative cross-country evidence that there is a positive long-run correlation between anticipated inflation, as proxied by nominal interest rates, and unemployment; that there is also a positive correlation between anticipated inflation and unemployment volatility; and that the inflation-unemployment correlation is stronger at high unemployment rates. The rest of the paper explores the mechanisms through which a monetary search model can generate these patterns and the implications of these mechanisms.

4 Model

The model environment closely follows Berentsen et al. (2011), augmented with aggregate productivity and monetary shocks. Time is discrete, and the time horizon is infinite. The economy is populated by a measure 1 of infinitely-lived households and a large measure of infinitely-lived firms. Households have preferences

$$\mathcal{U} = \sum_{t=0}^{\infty} \beta^t \left[h_t + \mathrm{u}\left(x_t\right)\right], \tag{6}$$

where h_t is consumption of a general good (taken as the numeraire), and x_t is consumption of a specialized good. Firms consume the general good and have utility

$$\mathcal{U} = \sum_{t=0}^{\infty} \beta^t h_t. \tag{7}$$

Each period is divided into three sub-periods. In the first sub-period (LM), there is a frictional labor market, in which households and firms match pairwise to produce the specialized good. In the second sub-period (DM), there is a frictional goods market, in which productive firms sell the specialized good to other households. In the third sub-period (CM), a frictionless centralized market convenes, in which firms pay wages and liquidate unsold inventories, and all agents rebalance their portfolios.

Matching and production (LM). At each point in time, each household is either employed or unemployed. Matching of unemployed households (workers) and firms (employers) in the LM is random. Firms create vacancies at cost κ. If the measure of vacancies is v_t and the measure of unemployed households at the beginning of period t is $1 - n_{t-1}$, the measure of new matches created in period t is $\mathcal{M}\left(v_t, 1 - n_{t-1}\right)$, where $\mathcal{M}$ is a constant returns to scale function satisfying the standard assumptions. The job-finding probability for an unemployed household is then

$$\frac{\mathcal{M}\left(v_t, 1 - n_{t-1}\right)}{1 - n_{t-1}} = \mathcal{M}\left(\theta_t, 1\right) \equiv f\left(\theta_t\right), \tag{8}$$

where $\theta_t = v_t / \left(1 - n_{t-1}\right)$ is the market tightness. Similarly, the probability for a firm of filling its vacancy is

$$\frac{\mathcal{M}\left(v_t, 1 - n_{t-1}\right)}{v_t} = \mathcal{M}\left(1, \frac{1}{\theta_t}\right) \equiv q\left(\theta_t\right). \tag{9}$$

Existing matches are destroyed every period with an exogenous probability δ. The above implies that the employment level upon exiting the period-t LM is

$$n_t = \left(1 - \delta\right) n_{t-1} + f\left(\theta_t\right)\left(1 - n_{t-1}\right). \tag{10}$$

A job match between a household and a firm produces y_t units of the general good in the LM, to be sold in the subsequent DM (see below). The productivity of a match, y_t, is subject to aggregate shocks and follows an AR(1) process

$$\log y_t = (1-\rho_y)\log \bar{y} + \rho_y \log y_{t-1} + \sigma_y \varepsilon_{y,t}, \quad \varepsilon_{y,t} \sim \mathcal{N}(0,1). \tag{11}$$

Trade in the goods market (DM). The DM is a decentralized goods market, in which productive firms (i.e. firms who have a worker in the LM) match pairwise with households in order to convert their general good into specialized goods and sell them. A firm with a worker in the LM has a technology for converting general goods into specialized goods: producing x_t units of the specialized good costs $\mathrm{c}(x_t)$ units of the general good. Matching is random and governed by a constant returns to scale matching function. Specifically, since there are n_t productive firms seeking to sell and a measure 1 of households seeking to buy, each household meets a firm with probability $\alpha(n_t)$, and a firm meets a household with probability $\alpha(n_t)/n_t$, where $\alpha' > 0$, $\alpha'' < 0$, and $\alpha(n) \leq n$. Terms of trade between a firm and a household in the DM are determined by proportional (Kalai) bargaining and are described in detail below.

Meetings in the DM are anonymous, there is no record-keeping of past transactions, and agents cannot commit to repay loans. This rules out credit in the DM, making a medium of exchange necessary for trade. This role is served by fiat money. The supply of fiat money is augmented stochastically via lump-sum transfers T_t to households in the CM. We find it convenient to think of monetary policy in terms of the nominal interest rate $\iota_t = (1+\pi_t)/\beta - 1$, where π_t is the rate of inflation. The nominal interest rate ι_t follows a Markov process, which will be described in detail below in Section 6.

Settlement and rebalancing (CM). The CM is a frictionless market, in which wages are paid out to employed households, and agents decide on money holdings for the following period. Wages w_t to be paid out in the CM are determined in the previous LM by Nash bargaining between the household and the firm (see below). Unemployed households receive an exogenous amount b of the general good, which can be interpreted as unemployment benefits, a value of leisure, or home production.

5 Equilibrium

We define equilibrium recursively. The aggregate state of the economy is $\Omega = (n^-, \iota, y)$, where n^- is the start-of-period employment level, ι is the nominal interest rate, and y is productivity. Throughout, "$^-$" superscripts are used to denote the previous period's variables, and "$^+$" superscripts are used to denote the next period's variables.

The timing is as follows. The ι and y shocks realize at the beginning of the CM. The realized ι shock determines the amount of lump-sum transfers $T(\iota)$ that will be paid in the following CM. Agents then make portfolio choices, and at the same time firms post vacancies v at cost κ. In the subsequent LM, matching and separations take place. The market tightness is

$$\theta = \frac{v}{1-n^-} \tag{12}$$

and the law of motion for employment is

$$n = (1-\delta)\, n^- + f(\theta)\left(1-n^-\right). \tag{13}$$

Wages are negotiated at this stage. After that, worker-firm matches produce and go on to the DM, where they sell their output in exchange for money. The buyer's matching probability in the DM is then $\alpha(n)$. A firm with productivity y selling x units of output incurs a cost of $\mathrm{c}(x)$. A firm who did not trade receives y units of output. Finally, in the CM either wages (w) or unemployment benefits (b) are paid, along with transfers $T(\iota)$.

5.1 Workers

We start by writing the CM value function for a worker of each employment status, denoted by $j \in \{e, u\}$. If a worker enters the CM employed ($j = e$), with promised wage w, and with real balances z, their value function is

$$V^e_{CM}(z, w, \Omega) = \max_{c, z^+} c + \beta V^e_{LM}\left(z^+, \Omega\right), \tag{14}$$

subject to the budget constraint

$$c + (1+\pi) z^+ = w + T(\iota). \tag{15}$$

In the CM, the worker receives their promised wage, as well as state-dependent transfers. They then choose how much to allocate to consumption, c, and to real balances for the following period, z^+. Similarly, if a worker enters the CM unemployed ($j = u$ with real balances z, their value function is

$$V^u_{CM}(z, \Omega) = \max_{c, z^+} c + \beta V^u_{LM}\left(z^+, \Omega\right), \tag{16}$$

subject to the budget constraint

$$c + (1+\pi) z^+ = b + T(\iota). \tag{17}$$

The unemployed worker's budget constraint differs from the employed in that, rather than receiving a wage, they receive fixed unemployment benefits, b. The outcome of the worker's maximization problem gives us the real balances policy function $z^+(\Omega)$, which, by quasi-linearity, is independent of z as well as employment status. For the same reason, it follows that the worker's CM value function is linear in real balances, with

$$\frac{\partial}{\partial z}V_{CM}^{e}\left(z,w,\Omega\right)=\frac{\partial}{\partial z}V_{CM}^{u}\left(z,\Omega\right)=1.$$

Now, consider the LM. An employed worker ($j=e$) stays employed with probability $1-\delta$. An unemployed worker ($j=u$) finds a job with probability $f\left(\theta(\Omega)\right)$. Therefore,

$$V_{LM}^{e}\left(z,\Omega\right) = \left(1-\delta\right)V_{DM}^{e}\left(z,w\left(\Omega\right),\Omega\right)+\delta V_{DM}^{u}\left(z,\Omega\right), \tag{18}$$

$$V_{LM}^{u}\left(z,\Omega\right) = f\left(\theta(\Omega)\right)V_{DM}^{e}\left(z,w\left(\Omega\right),\Omega\right)+\left(1-f\left(\theta(\Omega)\right)\right)V_{DM}^{u}\left(z,\Omega\right). \tag{19}$$

Next, consider the DM. Each worker meets a firm in the product market with probability $\alpha\left(n\right)$, where $n=n\left(\Omega\right)$ is the end-of-period employment level as given by (13). If meeting a firm, he buys some negotiated amount x of the special good in exchange for some negotiated amount $d\leq z$ of real balances. As we will verify below, x,d (determined by bargaining) depend on z but not on the worker's employment status, and will therefore be written as $\tilde{x}\left(z,\Omega\right)$, $\tilde{d}\left(z,\Omega\right)$. Thus:

$$\begin{aligned}V_{DM}^{e}\left(z,w,\Omega\right)=&\alpha\left(n\right)\left[\mathrm{u}\left(\tilde{x}\left(z,\Omega\right)\right)+\mathbb{E}V_{CM}^{e}\left(z-\tilde{d}\left(z,\Omega\right),w,\Omega^{+}\right)\right]\\&+\left(1-\alpha\left(n\right)\right)\mathbb{E}V_{CM}^{e}\left(z,w,\Omega^{+}\right)\end{aligned}, \tag{20}$$

$$\begin{aligned}V_{DM}^{u}\left(z,\Omega\right)=&\alpha\left(n\right)\left[\mathrm{u}\left(\tilde{x}\left(z,\Omega\right)\right)+\mathbb{E}V_{CM}^{u}\left(z-\tilde{d}\left(z,\Omega\right),\Omega^{+}\right)\right]\\&+\left(1-\alpha\left(n\right)\right)\mathbb{E}V_{CM}^{u}\left(z,\Omega^{+}\right)\end{aligned}, \tag{21}$$

where Ω^{+} denotes the following period's aggregate state.

5.2 Firms

Consider first the CM. A firm with a worker, who has unsold output o, real balances z and a wage commitment w, has the value

$$J_{CM}^{e}\left(o,z,w,\Omega\right) = o+z-w+\beta J_{LM}^{e}\left(\Omega\right). \tag{22}$$

A firm without a worker decides whether to post a vacancy, at cost κ:

$$J_{CM}^{v}\left(z,\Omega\right) = z+\max\left\{0,-\kappa+\beta J_{LM}^{v}\left(\Omega\right)\right\}. \tag{23}$$

In the LM, a firm with a worker loses that worker with probability δ, and a firm with a vacancy fills it with probability $q(\theta(\Omega))$:

$$\begin{aligned} J^e_{LM}(\Omega) &= (1-\delta)\, J^e_{DM}(w(\Omega),\Omega) + \delta J^v_{DM}(\Omega), \qquad (24)\\ J^v_{LM}(\Omega) &= q(\theta(\Omega))\, J^e_{DM}(w(\Omega),\Omega) + (1-q(\theta(\Omega)))\, J^v_{DM}(\Omega). \qquad (25) \end{aligned}$$

In the DM, a firm with a worker produces and sells its output, getting

$$\begin{aligned} J^e_{DM}(w,\Omega) =& \frac{\alpha(n)}{n}\mathbb{E}J^e_{CM}\left(y - \mathrm{c}(x(\Omega)), d(\Omega), w, \Omega^+\right) \\ &+ \left(1 - \frac{\alpha(n)}{n}\right)\mathbb{E}J^e_{CM}\left(y, 0, w, \Omega^+\right) \end{aligned}, \qquad (26)$$

where the traded quantities $x(\Omega)$, $d(\Omega)$ are determined by bargaining (as described below) through

$$x(\Omega) = \tilde{x}(z(\Omega),\Omega),$$

$$d(\Omega) = \tilde{d}(z(\Omega),\Omega).$$

In other words, the firm takes as given the z of the worker it will meet when forecasting the traded quantities. A firm without a worker does not trade in the DM, so

$$J^v_{DM}(\Omega) = \mathbb{E}J^v_{CM}\left(\Omega^+\right). \qquad (27)$$

5.3 Goods market bargaining

We now turn to the determination of x, d. We assume that x and d are determined by Kalai bargaining between the worker and the firm, with $\varphi \in [0,1]$ denoting the workers's bargaining weight. Consider a DM meeting in which the worker is employed with promised wage w and has real balances z, and the firm has a wage commitment w'. Then the traded quantity and price x, d solve

$$\max_{x,d} \mathrm{u}(x) + \mathbb{E}\left[V^e_{CM}\left(z-d, w, \Omega^+\right) - V^e_{CM}\left(z, w, \Omega^+\right)\right], \qquad (28)$$

subject to $d \le z$ and

$$\begin{aligned} \mathrm{u}(x) +&\mathbb{E}\left[V^e_{CM}\left(z-d, w, \Omega^+\right) - V^e_{CM}\left(z, w, \Omega^+\right)\right] \\ &= \frac{\varphi}{1-\varphi}\mathbb{E}\left[J^e_{CM}\left(y - \mathrm{c}(x), d, w', \Omega^+\right) - J^e_{CM}\left(y, 0, w', \Omega^+\right)\right]. \end{aligned} \qquad (29)$$

Similarly, if an unemployed worker with real balances z meets a firm with wage commitment w', the traded amounts x, d solve

$$\max_{x,d} \mathrm{u}(x) + \mathbb{E}\left[V^u_{CM}\left(z-d,\Omega^+\right) - V^u_{CM}\left(z,\Omega^+\right)\right], \tag{30}$$

subject to $d \leq z$ and

$$\begin{aligned}\mathrm{u}(x) +& \mathbb{E}\left[V^u_{CM}\left(z-d,\Omega^+\right) - V^u_{CM}\left(z,\Omega^+\right)\right] \\ &= \frac{\varphi}{1-\varphi}\mathbb{E}\left[J^e_{CM}\left(y-\mathrm{c}(x), d, w', \Omega^+\right) - J^e_{CM}\left(y,0,w',\Omega^+\right)\right].\end{aligned} \tag{31}$$

By linearity of the CM value functions, it is straightforward that a worker's surplus from trading x, d is

$$\begin{aligned}\mathrm{u}(x) + \mathbb{E}\left[V^e_{CM}\left(z-d,w,\Omega^+\right) - V^e_{CM}\left(z,w,\Omega^+\right)\right] &= \mathrm{u}(x) + \mathbb{E}\left[V^u_{CM}\left(z-d,\Omega^+\right) - V^u_{CM}\left(z,\Omega^+\right)\right] \\ &= \mathrm{u}(x) - d,\end{aligned} \tag{32}$$

regardless of employment status and current real balances. Similarly, a firm's surplus from trading x, d is

$$\mathbb{E}\left[J^e_{CM}\left(y-\mathrm{c}(x), d, w', \Omega^+\right) - J^e_{CM}\left(y,0,w',\Omega^+\right)\right] = d - \mathrm{c}(x). \tag{33}$$

This means that we can write the bargaining problem as

$$\max_{x,d} \mathrm{u}(x) - d \quad s.t. \quad \mathrm{u}(x) - d = \frac{\varphi}{1-\varphi}\left[d - \mathrm{c}(x)\right] \quad \text{and} \quad d \leq z. \tag{34}$$

The solution is well known. Define $\mathrm{g}(x) = (1-\varphi)\,\mathrm{u}(x) + \varphi\mathrm{c}(x)$, and define x^* as the solution to $\mathrm{u}'(x) = \mathrm{c}'(x)$. The solution to the bargaining problem is

$$\begin{aligned}x &= \min\left\{x^*, \mathrm{g}^{-1}(z)\right\}, \\ d &= \min\left\{\mathrm{g}(x^*), z\right\}.\end{aligned} \tag{35}$$

5.4 Optimal choice of real balances

The DM bargaining solution described above gives us

$$\frac{\partial V^j_{LM}}{\partial z} = 1 + \alpha(n)\max\left\{0, \frac{\mathrm{u}'(x)}{\mathrm{g}'(x)} - 1\right\}. \tag{36}$$

In the CM, the first-order condition for z is

$$1 + \iota = \frac{\partial V^j_{LM}}{\partial z}. \tag{37}$$

Combining these, we get

$$\mathrm{u}'(x) = \left(1 + \frac{\iota}{\alpha(n)}\right)\mathrm{g}'(x) \tag{38}$$

and $z = \mathrm{g}(x)$.

5.5 Labor market bargaining

We next consider the determination of the wage. The worker's surplus from being employed at wage w is

$$V^e_{DM}(z, w, \Omega) - V^u_{DM}(z, \Omega) = V^e_{DM}(0, w, \Omega) - V^u_{DM}(0, \Omega) \equiv S^e_{DM}(w, \Omega).$$

This can be written recursively as

$$S^e_{DM}(w, \Omega) = w - b + \beta\mathbb{E}\left(1 - \delta - f\left(\theta\left(\Omega^+\right)\right)\right) S^e_{DM}\left(w\left(\Omega^+\right), \Omega^+\right). \tag{39}$$

The firm's surplus from having a worker at wage w is $J^e_{DM}(w, \Omega)$. Defining the firm's total output as

$$\begin{aligned}\mathcal{O}(\Omega) =& y + \frac{\alpha(n)}{n}(y + d(\Omega) - \mathrm{c}(x(\Omega)) - y) \\ =& y + \frac{\alpha(n)}{n}(1 - \varphi)(\mathrm{u}(x(\Omega)) - \mathrm{c}(x(\Omega))),\end{aligned} \tag{40}$$

we can write

$$J^e_{DM}(w, \Omega) = \mathcal{O}(\Omega) - w + \beta(1 - \delta)\mathbb{E}J^e_{DM}\left(w\left(\Omega^+\right), \Omega^+\right). \tag{41}$$

The surplus from an employment match is $\mathcal{S}(\Omega) = S^e_{DM}(w, \Omega) + J^e_{DM}(w, \Omega)$. We assume that the wage is determined by Nash bargaining, with worker bargaining weight equal to ξ, so that $w = w(\Omega)$ solves

$$S^e_{DM}(w(\Omega), \Omega) = \xi\mathcal{S}(\Omega). \tag{42}$$

Adding (39) and (41) and using (42), we obtain

$$\mathcal{S}(\Omega) = \mathcal{O}(\Omega) - b + \beta\mathbb{E}\left(1 - \delta - \xi f\left(\theta\left(\Omega^+\right)\right)\right)\mathcal{S}\left(\Omega^+\right). \tag{43}$$

5.6 Free entry and wages

Market tightness $\theta = \theta(\Omega)$ is determined by the free entry condition $\kappa = \beta J^v_{LM}(\Omega)$, which implies

$$\kappa = \beta q(\theta)(1 - \xi)\mathcal{S}(\Omega). \tag{44}$$

Combining (44) with the expression for the match surplus in (43), we also obtain the expression for the wage,

$$w(\Omega) = \xi\mathcal{O}(\Omega) + (1-\xi)\, b + \mathbb{E}\xi\kappa\theta\left(\Omega^{+}\right). \tag{45}$$

5.7 Equilibrium conditions

The equilibrium consists of functions $x(\Omega)$, $\mathcal{O}(\Omega)$, $\mathcal{S}(\Omega)$, $\theta(\Omega)$, $w(\Omega)$, and $n(\Omega)$ satisfying the law of motion (13), optimal choice of real balances (38), the output equation (40), the Bellman equation for the match surplus (43), the free entry condition (44), and the wage equation (45).

6 Calibration

Our calibration strategy is to match features of both labor market dynamics and monetary data in the US. As pointed out earlier, two features distinguish our environment from the standard labor search model. First, our model economy is subject to two shocks: productivity and nominal interest rates. Second, measured output per worker includes DM trade and is therefore endogenous. As a consequence, the productivity process cannot be backed out directly from output per worker as, e.g., in Shimer (2005) or Hagedorn and Manovskii (2008), and will instead be calibrated jointly with other model parameters.

The model is calibrated on a monthly basis using a mix of monthly and quarterly data. We use, whenever possible, data series covering the period from January 1948 to December 2019.[4] Except for the discount factor β, the job separation rate δ and the exogenous process of interest rate shocks, the model's remaining parameters are calibrated jointly such that a selected set of model-based simulated moments match their empirical counterparts. We set the monthly discount factor β externally to 0.998 to be consistent with an average monthly real interest rate of 0.23%.

Labor market. The exogenous process for y_t in the model is assumed to follow the AR(1) process

$$\log y_t = (1-\rho_y)\log\bar{y} + \rho_y \log y_{t-1} + \sigma_y \varepsilon_{y,t} \tag{46}$$

where $\varepsilon_y \sim \mathcal{N}(0,1)$ and $\bar{y}$ is normalized to 1. We calibrate the process parameters ρ_y and σ_y internally such that the cyclical component of the logarithm of total output per worker in the model, $\mathcal{O}_t$, matches the volatility and autocorrelation of output per worker in the data. We follow Shimer (2005) in using the Bureau of Labor Statistics' (BLS) monthly data series measuring real output per person in the non-farm business sector. We extract the cyclical component of the

[4] Our sample stops short of the onset of the COVID-19 pandemic and hence does not capture the wild movements in macroeconomic variables that occurred afterward.

logarithm of the quarterly observations using the HP filter with a smoothing parameter value of 1,600.

Similarly to Hagedorn and Manovskii (2008), we use the elasticity of wages with respect to labor productivity as a calibration target to identify the worker bargaining weight. We use the BLS series for labor productivity and labor income share to compute a series for the real wage as the product of the two. We then extract the HP-filtered cyclical component of the logarithm of the computed wage series and use it to estimate the elasticity of the wage to labor productivity. We include the wage elasticity in our list of targeted moments and add the bargaining power of workers ξ to the internal calibration to match it.

We calibrate the cost of posting vacancies κ to match the average labor market tightness θ, computed as the ratio of the vacancy and unemployment rates. The vacancy rate series is constructed as in Petrosky-Nadeau and Zhang (2020). From December 2000 to December 2019, the number of job openings is obtained from the Job Openings and Labor Turnover Survey (JOLTS) of the BLS which we divide by the civilian labor force to obtain the vacancy rate. From January 1951 to November 2000 we use the vacancy rate series from Barnichon (2010), which is based on a composite print and online help wanted index. We then splice Barnichon's series to the JOLTS series in December 2000 to obtain one series stretching from January 1951 to December 2019.

We assume the matching function

$$\mathcal{M}\left(v, 1-n\right) = \frac{v\left(1-n\right)}{\left(v^{\chi} + (1-n)^{\chi}\right)^{1/\chi}}, \tag{47}$$

similarly to Den Haan et al. (2000), as it ensures job finding and vacancy filling probabilities lie in the interval $[0, 1]$. The corresponding job finding probability as a function of θ is

$$f(\theta) = \frac{\theta}{(1+\theta^{\chi})^{1/\chi}}. \tag{48}$$

We add χ to our internal calibration to match the average job finding probability, which we compute using data on short-term unemployment as in Shimer (2012).[5] We calculate the job separation rate in the standard way by dividing the number of short-term unemployed workers by last period's employed population. We obtain an average monthly separation rate of 2.51%.[6] We set δ directly to that value.

The value of non-market activity b is calibrated to match the standard deviation of unemployment, which stands at 0.138. This statistic is measured in logs as the quarterly deviation from an HP-filtered trend with a smoothing parameter of 1,600.

[5]We correct for the 1994 CPS redesign following Shimer (2012).

[6]Correcting for time aggregation following Shimer (2005) yields a very similar value.

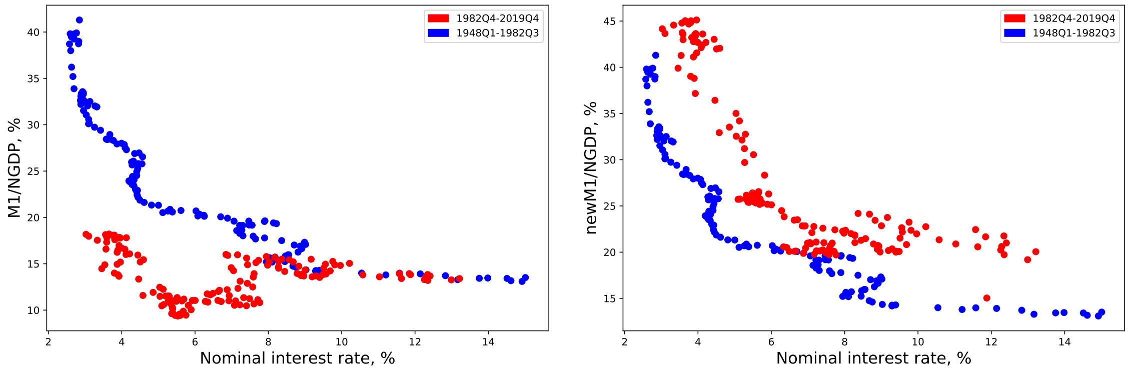

Figure 2: Measuring money demand: M1 v. NewM1.

Decentralized goods market. We assume the utility of DM good consumption takes the form

$$\mathrm{u}(x) = A\frac{x^{1-\gamma}}{1-\gamma}, \tag{49}$$

where $\gamma \in (0,1)$ and $A > 0$. For the DM cost function, we set

$$\mathrm{c}(x) = x. \tag{50}$$

For the DM matching function, we assume that the buyer's probability of finding a seller is

$$\alpha(n) = \zeta\frac{n}{1+n}, \tag{51}$$

where n is the measure of active sellers (i.e. firms with workers), 1 is the measure of active buyers (households), and ζ is a matching efficiency parameter.

Most of the parameters related to the DM are calibrated following Lagos and Wright (2005) and Berentsen et al. (2011). In particular, A and γ are calibrated to match both the average ratio of aggregate money supply to nominal GDP (i.e. the inverse of the velocity of money) and its elasticity with respect to the nominal interest rate.[7] To compute these moments, we use the NewM1 monetary aggregate following Lucas and Nicolini (2015). This measure adds to the standard M1 aggregate, published by the Federal Reserve Board, the total amount of Money Market Deposit Accounts held at commercial banks in the US starting from 1982. As discussed by Lucas and Nicolini (2015), money demand measures based on this aggregate perform much better than the ones based on the conventional M1 aggregate. This can clearly be seen in Figure

[7]Interest rate elasticity in the data is computed using raw quarterly observations. We run a log-log regression of money demand on interest rates and the resulting slope coefficient is used as a point estimate for the interest rate elasticity.

2, where the relationship between money demand and nominal interest rates is much more stable when using the NewM1 aggregate (right panel). In addition, we extend our NewM1 series back in time to January 1948 with the pre-1959 M1 series produced by Rasche (1987). The latter is consistent with the Board's post-1959 methodology.[8]

Regarding the opportunity cost of holding liquidity, we follow Lagos and Wright (2005) in using the monthly Moody's composite yield on Aaa-rated long-term US corporate bonds. Because this measure is non-stationary, we use the HP filter to decompose it into a trend component $\bar{\iota}_t$ and cycle component $\hat{\iota}_t$, i.e.

$$\iota_t = \bar{\iota}_t + \hat{\iota}_t. \tag{52}$$

The cyclical component is modeled as a stationary AR(1) process

$$\hat{\iota}_t = \rho_{\hat{\iota}} \hat{\iota}_{t-1} + \sigma_{\hat{\iota}} \varepsilon_{\hat{\iota},t}, \tag{53}$$

where $\varepsilon_{\hat{\iota}} \sim \mathcal{N}(0,1)$. Its parameters are estimated, at a monthly frequency, to be $\rho_{\hat{\iota}} = 0.939$ and $\sigma_{\hat{\iota}} = 0.0001$. The non-stationary trend component is modeled as a very persistent Markov chain with 5 states. The state values and the estimated transition probabilities are presented in Appendix B.[9]

Following Aruoba et al. (2011), the bargaining power of buyers φ is calibrated internally such that the average markup (i.e. price-to-marginal cost ratio) in DM transactions matches the average markup in the data. De Loecker et al. (2020) use financial statements of all publicly traded firms covering all sectors of the US economy over the period 1955-2016 to estimate an average net markup of 36%. We add their measure to the list of targeted moments. The DM matching efficiency parameter ζ is added to the internal calibration to match a monthly interest rate elasticity of unemployment of 0.297 in the data. Intuitively, this parameter is important for the slope of the inflation-unemployment relationship because of feedback effects between goods and labor markets. For example, an increase in inflation lowers money demand, which, in equilibrium, lowers firm profits and raises unemployment; however, this rise in unemployment lowers the probability of finding a seller in the DM and thus further lowers money demand, and the extent to which it does so depends on goods market frictions.

SMM calibration procedure. The above discussion leaves us with the set of 10 parameters $\{\kappa, b, \chi, \xi, \rho_y, \sigma_y, A, \gamma, \zeta, \varphi\}$ that we calibrate internally following a Simulated Method of Moments

[8] The pre-1959 M1 data series is available from the website of the Federal Reserve Bank of St. Louis (https://research.stlouisfed.org/aggreg/).

[9] Most of the results in our quantitative analysis below have to do with the effect of trend inflation: for example, comparing unemployment volatility in different $\bar{\iota}_t$ regimes. It is nonetheless important to include the cyclical component to generate sufficient variation in nominal interest rates in model simulations; this is essential both when calibrating the model and when comparing its predictions to cross-country regressions.

Table 3: SMM calibrated parameters

Parameter	Description	Value	Moment	Frequency	Data	Model
κ	Vacancy cost	1.471	Average θ	Monthly	0.634	0.634
b	Flow value of unemployment	0.990	Unemployment volatility	Quarterly	0.138	0.138
χ	Parameter of the LM matching fun.	1.269	Average JFP	Monthly	0.430	0.430
ξ	Worker bargaining weight	0.035	Elast. of wage to labor prod.	Quarterly	0.470	0.470
ρ_y	Persistence parameter of y_t process	0.967	Autocorr. of labor productivity	Quarterly	0.758	0.761
σ_y	Volatility parameter of y_t process	0.007	SD of labor productivity	Quarterly	0.013	0.013
A	Level parameter of DM utility	1.421	Average money demand	Quarterly	25.73%	25.72%
γ	Curvature parameter of DM utility	0.217	Elast. of money demand to ι	Quarterly	-0.594	-0.594
ζ	Parameter of the DM matching fun.	0.204	Elast. of u to ι	Monthly	0.297	0.297
φ	Buyer bargaining weight	0.320	Average price markup	Monthly	36.00%	36.00%

(SMM) procedure. The model is solved using a global solution method that preserves its strong nonlinearities (Petrosky-Nadeau and Zhang, 2017). Computational details are available in Appendix B.[10] Table 3 summarizes the results of the SMM calibration procedure. In particular, we match all the targeted moments.[11] Table 4 compares some US labor market statistics to those from the calibrated model. The model does a good job in capturing some business cycle properties of the US labor market.

Discussion. The primary aim of the quantitative analysis below is to examine the state-dependent and nonlinear effects of trend inflation, rather than its average long-run effect on unemployment, which has been discussed at length by e.g. Berentsen et al. (2011). As such, our calibration targets the slope of the long-run relationship between nominal interest rates and unemployment. The model will instead be validated based on its ability to match the state-dependence of this relationship, as evidenced in our empirical analysis.[12]

7 Understanding the mechanisms

As in any environment building on the DMP framework, the volatility of market tightness and therefore employment depends on the size of the match surplus. This result carries over to our setting with two shocks, and furthermore implies that the two shocks interact: inflation affects the size of the employment response to productivity, and vice versa. However, the amplification

[10]As in Berentsen et al. (2011), the model can exhibit multiple steady state equilibria. We focus our analysis on the high employment equilibrium. We show in Appendix C.1 that it is locally stable. We also plot the ergodic distributions of the endogenous variables to verify that the calibrated economy stays around the high employment equilibrium throughout the simulations.

[11]Our calibrated value of b is is 0.990, which corresponds to a ratio of the flow value of unemployment to average labor productivity of 0.913 – slightly below, e.g. Hagedorn and Manovskii (2008). Note that the value normalized to one is $\overline{y}$, which is less than average output per worker in the model, since the latter includes DM trade.

[12]Similarly, our calibration targets the level of unemployment volatility, and thus we do not view our analysis here as offering a novel explanation of the unemployment volatility "puzzle." How this volatility varies with the level of inflation, however, is a non-targeted prediction.

Table 4: Labor market statistics

		u	v	θ	$\mathcal{O}$
Quarterly US data, 1948-2019					
Standard deviation		**0.138**	0.138	0.257	**0.013**
Autocorrelation		0.895	0.902	0.903	**0.758**
Correlation matrix	u	1	-0.900	-0.950	-0.231
	v	-	1	0.982	0.363
	θ	-	-	1	0.296
	$\mathcal{O}$	-	-	-	1
Model simulations					
Standard deviation		**0.138**	0.633	0.745	**0.013**
Autocorrelation		0.843	0.431	0.635	**0.761**
Correlation matrix	u	1	-0.554	-0.787	-0.851
	v	-	1	0.904	0.639
	θ	-	-	1	0.756
	$\mathcal{O}$	-	-	-	1

Notes: All variables are reported in logs as deviations from an HP trend with $\lambda = 1600$. Model-based statistics are computed using 1,000 simulations of 1,000 months each. For each simulation, we burn the first 136 periods to match the length of the data series. The simulated series are then averaged quarterly. The reported statistics are averages over all simulations. The bold statistics are targeted in the calibration.

mechanism in our framework, following Berentsen et al. (2011), also contains two features absent from the benchmark DMP model. First, the endogenous choice of real balances implies that measured match output per worker is endogenous and, in particular, responds directly to the cost of liquidity. Second, goods market frictions generate an additional feedback effect from employment to match output. An increase in employment lowers a firm's probability of finding a buyer in the goods market, thus lowering its profits from a match, all else equal; on the other hand, it also raises a buyer's probability of finding a seller, thereby raising a buyer's incentive to carry real balances. Goods market frictions may thus act as either an amplifying or a dampening mechanism, depending on parameters.

To gauge the relative magnitude of the above mechanisms in our calibrated model, we consider the comparative statics of the model in steady state.[13] The free entry condition (44) in steady state becomes

$$\kappa = \underbrace{\frac{\beta q(\theta)(1-\xi)}{1-\beta(1-\delta-\xi f(\theta))}}_{\Upsilon(\theta)} (\mathcal{O} - b). \tag{54}$$

[13]This is similar to the analysis of Hagedorn and Manovskii (2008) and Ljungqvist and Sargent (2017) for the benchmark DMP model. Steady-state comparative statics are a good guide to the mechanisms driving the transitional dynamics of the model, since convergence to steady state is fast.

The steady-state employment level is then given by

$$n = \frac{f(\theta)}{\delta + f(\theta)}. \tag{55}$$

Equation (54) implicitly defines steady-state θ, and therefore n, as a function of match output $\mathcal{O}$. In turn, from (38) and (40), $\mathcal{O}$ is a function of n, ι and y. We can then write the elasticity of LM tightness with respect to y as

$$\varepsilon_{\theta,y} \equiv \frac{y}{\theta}\frac{\partial\theta}{\partial y} = \frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{O}}{\mathcal{O}-b}\varepsilon_{\mathcal{O},y}, \qquad \varepsilon_{\mathcal{O},y} \equiv \frac{y}{\mathcal{O}}\frac{\partial\mathcal{O}}{\partial y}, \tag{56}$$

and the semi-elasticity of LM tightness with respect to ι as

$$\varepsilon_{\theta,\iota} \equiv \frac{1}{\theta}\frac{\partial\theta}{\partial \iota} = \frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{O}}{\mathcal{O}-b}\varepsilon_{\mathcal{O},\iota}, \qquad \varepsilon_{\mathcal{O},\iota} \equiv \frac{1}{\mathcal{O}}\frac{\partial\mathcal{O}}{\partial \iota}, \tag{57}$$

where $\epsilon_{\Upsilon,\theta} = -\theta\Upsilon'(\theta)/\Upsilon(\theta)$, with Υ defined in (54). We adopt the convention of writing $\varepsilon_{\theta,\iota}$ and $\varepsilon_{\mathcal{O},\iota}$ as semi-elasticities since ι is typically written in percentage terms. Once the elasticities $\varepsilon_{\theta,y}$ and $\varepsilon_{\theta,\iota}$ are computed, the corresponding elasticities of employment can be obtained through $\varepsilon_{n,y} = \varepsilon_{n,\theta}\varepsilon_{\theta,y}$ and $\varepsilon_{n,\iota} = \varepsilon_{n,\theta}\varepsilon_{\theta,\iota}$, where $\varepsilon_{n,\theta}$ is computed from (55) as

$$\varepsilon_{n,\theta} = (1-n)\frac{\theta f'(\theta)}{f(\theta)}. \tag{58}$$

Elasticities of unemployment $u = 1-n$ with respect to either y or ι are computed similarly.

We next turn to the role of the decentralized goods market (DM). Match output itself depends on agents' money demand and DM frictions, as captured by the terms $\varepsilon_{\mathcal{O},y}$ and $\varepsilon_{\mathcal{O},\iota}$. Note that the steady-state version of (40) gives

$$\begin{aligned} \mathcal{O} =& y + \frac{\alpha(n)}{n}(1-\varphi)(\mathrm{u}(x) - \mathrm{c}(x)) \\ \equiv& y + \mathcal{P}, \end{aligned} \tag{59}$$

where x is the steady-state DM quantity traded, and $\mathcal{P}$ denotes the DM output in units of the numeraire. In turn, x is determined by the optimality condition (38) for real balances in steady state, which can be rearranged as

$$\varphi\frac{\mathrm{c}'(x)}{\mathrm{u}'(x)} = \frac{\alpha(n)}{\iota + \alpha(n)} - (1-\varphi). \tag{60}$$

Differentiating with respect to y, we find (see Appendix C.1 for details) that $\varepsilon_{\mathcal{O},y}$ is given by

$$\varepsilon_{\mathcal{O},y} = \underbrace{\frac{y}{\mathcal{O}}}_{\text{direct effect}} + \underbrace{\frac{\mathcal{P}}{\mathcal{O}}\varepsilon_{\mathcal{P},n}\varepsilon_{n,\theta}\varepsilon_{\theta,y}}_{\text{GE effect}}, \tag{61}$$

where

$$\varepsilon_{\mathcal{P},n} = \underbrace{\frac{1}{\sigma_{u,x}+\sigma_{c,x}}\frac{\iota\alpha(n)}{\iota+\alpha(n)}\frac{\epsilon_{\alpha,n}}{\varphi\alpha(n)-(1-\varphi)\iota}\frac{x[u'(x)-c'(x)]}{u(x)-c(x)}}_{\text{money-demand effect}} - \underbrace{(1-\epsilon_{\alpha,n})}_{\text{congestion effect}} \tag{62}$$

and $\epsilon_{\alpha,n} = n\alpha'(n)/\alpha(n)$ is the elasticity of the goods market matching function. Productivity y affects match output not only directly, but also through a general-equilibrium effect, via its effect on decentralized-market output $\mathcal{P}$. The latter effect, captured by (62), has two components. First, higher employment increases a buyer's probability of meeting a seller, and thereby the buyer's demand for real money balances. Second, there is also a congestion effect: higher employment reduces a firm's probability of meeting a buyer, and thereby its profits for any given amount of money balances. Depending on which effect dominates, $\varepsilon_{\mathcal{P},n}$ may be positive or negative.

In a similar vein, we find that the elasticity of $\mathcal{O}$ with respect to ι is given by

$$\varepsilon_{\mathcal{O},\iota} = \underbrace{\frac{\mathcal{P}}{\mathcal{O}}\varepsilon_{\mathcal{P},\iota}}_{\text{direct effect}} + \underbrace{\frac{\mathcal{P}}{\mathcal{O}}\varepsilon_{\mathcal{P},n}\varepsilon_{n,\theta}\varepsilon_{\theta,\iota}}_{\text{GE effect}}, \tag{63}$$

where

$$\varepsilon_{\mathcal{P},\iota} = -\frac{1}{\sigma_{u,x}+\sigma_{c,x}}\frac{\alpha(n)}{\iota+\alpha(n)}\frac{1}{\varphi\alpha(n)-(1-\varphi)\iota}\frac{x[u'(x)-c'(x)]}{u(x)-c(x)}. \tag{64}$$

In other words, inflation ι reduces match output both directly – by lowering the demand for real balances – and through the aforementioned general-equilibrium effect, via the effect of n on $\mathcal{P}$.

Substituting (61) and (63) into (56) and (57) and rearranging, we arrive at the following expressions for the elasticities:

$$\varepsilon_{\theta,y} = \underbrace{\frac{1}{\epsilon_{\Upsilon,\theta}}}_{\text{LMM}} \times \underbrace{\frac{\mathcal{O}}{\mathcal{O}-b}}_{\text{FS}} \times \underbrace{\left(1-\frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{P}}{\mathcal{O}-b}\varepsilon_{\mathcal{P},n}\varepsilon_{n,\theta}\right)^{-1}}_{\text{GEM}} \times \underbrace{\frac{y}{\mathcal{O}}}_{\text{direct}_{\mathcal{O},y}}, \tag{65}$$

$$\varepsilon_{\theta,\iota} = \underbrace{\frac{1}{\epsilon_{\Upsilon,\theta}}}_{\text{LMM}} \times \underbrace{\frac{\mathcal{O}}{\mathcal{O}-b}}_{\text{FS}} \times \underbrace{\left(1-\frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{P}}{\mathcal{O}-b}\varepsilon_{\mathcal{P},n}\varepsilon_{n,\theta}\right)^{-1}}_{\text{GEM}} \times \underbrace{\varepsilon_{\mathcal{P},\iota}\frac{\mathcal{P}}{\mathcal{O}}}_{\text{direct}_{\mathcal{O},\iota}}. \tag{66}$$

The expressions in (65) and (66) illuminate precisely the mechanisms active in our framework, and how exactly they differ from the benchmark DMP framework. As an illustration, consider the decomposition of the elasticity $\varepsilon_{\theta,y}$ in equation (65). The effect of a shock to productivity y on market tightness can be decomposed into four components. First, the term labeled "direct$_{\mathcal{O},y}$," captures the direct effect of a shock to y on *match output*. Second, the fundamental surplus effect, labeled "FS," captures how much an increase in *match output* raises *match surplus*. This is the effect highlighted by, e.g. Hagedorn and Manovskii (2008) and Ljungqvist and Sargent (2017). Third, the labor market matching effect, labeled "LMM," captures how much an increase in match surplus raises the equilibrium market tightness through the free entry condition. Finally, the general-equilibrium multiplier, labeled "GEM," measures how the induced increase in employment further affects match output through the goods market. This feedback effect may be either greater or smaller than one, depending on whether $\varepsilon_{\mathcal{P},n}$ is positive or negative, that is, depending on whether the money-demand effect or the congestion effect dominates.[14] In the benchmark DMP model, the direct$_{\mathcal{O},y}$ and GEM terms would both be identically equal to one, so that only the labor market channels, LMM and FS, are operational. The decomposition of $\varepsilon_{\theta,\iota}$ is identical, except for the direct effect of inflation on match output, labeled "direct$_{\mathcal{O},\iota}$". Finally, once the elasticity of θ with respect to each shock is computed, the elasticity of either employment or unemployment can be calculated from the steady-state employment equation (55), and therefore obeys a similar decomposition.

7.1 Decomposing the mechanisms: numerical results

We use our calibrated model to compute the elasticities in (65) and (66), as well as their individual components. Figure 3 plots the overall elasticities $\varepsilon_{\theta,y}$ and $\varepsilon_{\theta,\iota}$, as well as the corresponding unemployment elasticities $\varepsilon_{u,y}$ and $\varepsilon_{u,\iota}$, on a grid for steady state ι and y. To be precise, for each value of ι and y on the grid, we solve numerically for the steady state of our calibrated model and then, based on the resulting steady-state values for the endogenous variables, calculate the values of the analytically-derived elasticities and their respective components. On the grid, both $\varepsilon_{\theta,y}$ and $\varepsilon_{\theta,\iota}$ are monotonically increasing (in absolute value) in the interest rate and monotonically decreasing in productivity, confirming the main claim of this paper. The rest of this sub-section is devoted to pinpointing precisely which components of the elasticities drive the result.

To do so, we first plot, in Figure 4, the individual components: LMM, FS, GEM, direct$_{\mathcal{O},y}$, and direct$_{\mathcal{O},\iota}$, on a grid for steady state ι and y.[15] Several messages emerge from these figures. First, the two channels pertaining to the labor market side, LMM and FS, are the dominant ones in absolute value: ι and y do induce some variation in the goods-market effects, as shown in panels 4c-4e, but this variation is much smaller in magnitude than that in panels 4a-4b. We will further confirm

[14] In Appendix C.1, we argue that, regardless of the sign of $\varepsilon_{\mathcal{P},n}$ the GEM effect should be non-negative in a locally stable equilibrium.

[15] We perform the decomposition for the elasticities of market tightness, rather than the elasticities of employment or unemployment, but the results and main message are similar if, e.g., we instead decompose $\varepsilon_{u,y}$ and $\varepsilon_{u,\iota}$.

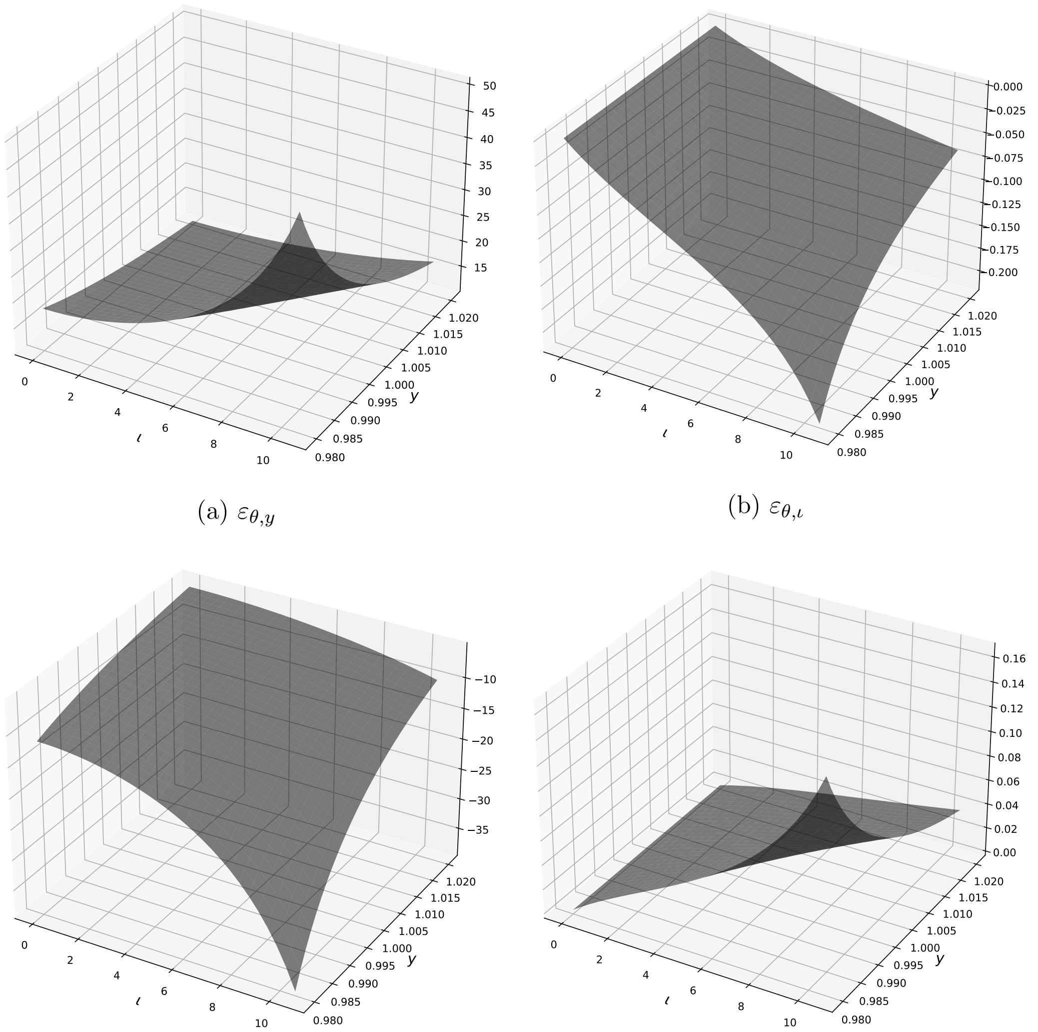

(a) $\varepsilon_{\theta,y}$

(b) $\varepsilon_{\theta,\iota}$

(c) $\varepsilon_{u,y}$

(d) $\varepsilon_{u,\iota}$

Figure 3: Overall elasticities on a grid for y and ι (annualized).

this below in a decomposition exercise. Second, and relatedly, not only the FS effect (emphasized in the literature) but also the LMM effect is important for the model's amplification mechanism. The FS effect captures how an increase in match output translates, in percentage terms, into a large increase in match *surplus*. This increase is larger when the match surplus is small. The LMM effect, in turn, captures how an increase in the match surplus translates into an increase in LM tightness. The numerical analysis shows that both effects exhibit substantial variation with respect to productivity and inflation. The LMM effect varies with economic conditions primarily because of the nonlinearity of the matching function.

Third, as shown in panel 4c, the GEM effect is less than one under our parameterization; in other words, the feedback effect from employment to the goods market acts as a dampening mechanism. This occurs because the congestion effect of higher employment dominates its effect on buyers' money demand. Moreover, the GEM effect is smaller (i.e. dampens the overall effect more) for high ι and low y.[16] The intuition for this result is as follows. The GEM effect is dampening because the congestion effect in (62) dominates the money-demand effect: an increase in firms' entry and employment lowers the firm's probability of finding a buyer, hence its profits. But the effect of shocks on employment is higher for low y and high ι, precisely because of the small-surplus effects outlined above; hence the congestion effect is also stronger for low y and high ι. We also observe, however, that the GEM effect is close to one regardless of y and ι. The reason for this is that the money demand and congestion effects are roughly of the same magnitude in absolute value, combined with the fact that the fraction of the DM in total trade – which matters for the importance of the goods market friction – is small. These results are discussed in detail in Appendix C.2 and illustrated in Figures 11 and 12. The findings show that both the size and the overall variation in the GEM effect are small compared to the LM effects, as pointed out above, and hence is not the crucial driving factor behind the main quantitative results.

Fourth, we inspect the direct effects in panels 4d and 4e. The direct effect of y on $\mathcal{O}$ is simply given by $y/\mathcal{O} = y/\left(y + \mathcal{P}\right)$, hence mechanically increasing in y and increasing in ι. The direct effect of ι on $\mathcal{O}$ is more complex, as it is given by the product $\varepsilon_{\mathcal{P},\iota}\frac{\mathcal{P}}{\mathcal{O}}$. Figure 13 in Appendix C.2 shows the decomposition of this product into the two individual effects. The elasticity $\varepsilon_{\mathcal{P},\iota}$ is negative and decreasing in ι: an increase in inflation lowers the DM quantity traded, x, and more so when x is already low. On the other hand, the quantity $\frac{\mathcal{P}}{\mathcal{O}}$ is positive and decreasing in ι: an increase in inflation lowers DM trade as a proportion of total trade, thereby making further increases in inflation less consequential for total output. The overall effect of ι on $\varepsilon_{\mathcal{P},\iota}\frac{\mathcal{P}}{\mathcal{O}}$ is therefore ambiguous theoretically. As panel 4e together with Figure 13 show, the effect of ι on $\varepsilon_{\mathcal{P},\iota}$ dominates, leading $\text{direct}_{\mathcal{O},\iota}$ to be amplified at higher ι. The primary reason for this is that $\frac{\mathcal{P}}{\mathcal{O}}$ is small to start with, so the effect of ι on $\frac{\mathcal{P}}{\mathcal{O}}$, while qualitatively acting in the opposite direction, is not quantitatively significant.

[16]For $y < b$, the GEM effect is non-monotone in ι. In Appendix C.2, we discuss why.

	$\iota = 2.00\%$ $y = 1.00$				$\iota = 8.00\%$ $y = 1.00$
$\text{direct}_{\mathcal{O},y}$	0.9110	0.9269	0.9269	0.9269	0.9269
GEM	0.9827	–	0.9841	0.9841	0.9841
FS	10.1558	–	–	12.0806	12.0806
LMM	1.5278	–	–	–	1.7591
$\varepsilon_{\theta,y}$	13.8911	14.1328	14.1539	16.8364	19.3843

Table 5: Change in $\varepsilon_{\theta,y}$ when ι increases from 2% to 8%.

	$\iota = 2.00\%$ $y = 1.00$				$\iota = 8.00\%$ $y = 1.00$
$\text{direct}_{\mathcal{O},\iota}$	-0.0017	-0.0036	-0.0036	-0.0036	-0.0036
GEM	0.9827	–	0.9841	0.9841	0.9841
FS	10.1558	–	–	12.0806	12.0806
LMM	1.5278	–	–	–	1.7591
$\varepsilon_{\theta,\iota}$	-0.0253	-0.0556	-0.0557	-0.0662	-0.0762

Table 6: Change in $\varepsilon_{\theta,\iota}$ when ι increases from 2% to 8%.

7.2 Comparative statics of the elasticities: an example

We next provide a numerical example for how a discrete change in the steady state ι affects the elasticities $\varepsilon_{\theta,y}$ and $\varepsilon_{\theta,\iota}$. Tables 5 and 6 display the values of $\varepsilon_{\theta,y}$ and $\varepsilon_{\theta,\iota}$, respectively, for $\iota = 2\%$ and $\iota = 8\%$. In each case, when moving from $\iota = 2\%$ to $\iota = 8\%$, we change the magnitudes of the previously described individual channels one by one in order to decompose the overall change into the individual components. Table 5 indicates that the change in $\varepsilon_{\theta,y}$ is mainly driven by the FS and LMM channels. On the other hand, as shown in Table 6, the direct effect of ι on DM trade is important for the change in $\varepsilon_{\theta,\iota}$ but the FS and LMM channels also play a significant role. The general-equilibrium multiplier, operating through the feedback effect between labor and goods markets, does not play much of a role.

To summarize our findings so far, the model-implied steady-state elasticities are consistent with our main claim that low productivity and high inflation amplify the effects of further shocks. The primary driving forces appear to be the labor market effects, manifested in the small match surplus and nonlinear matching function. In addition, the direct negative effect of inflation on traded quantities amplifies further effects of inflation. Feedback effects coming through goods market frictions are less significant, largely because the fraction of total trade taking place in the decentralized goods market is small to start with.

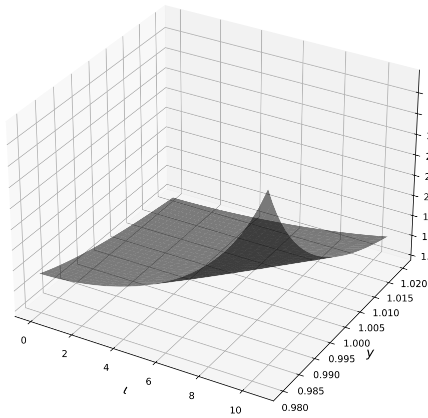

(a) LMM

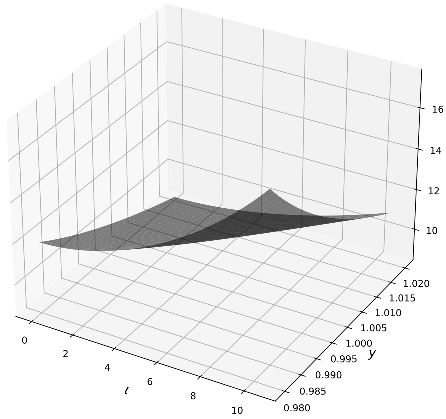

(b) FS

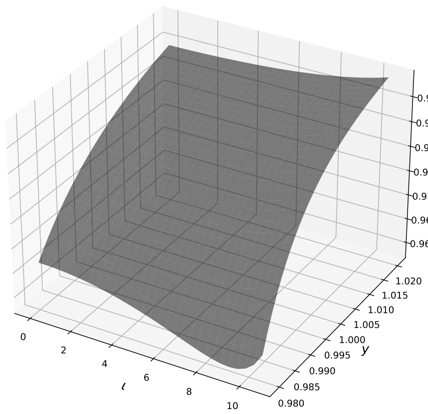

(c) GEM

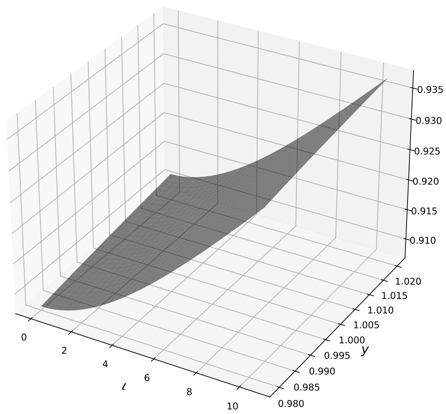

(d) $\text{direct}_{\mathcal{O},y}$

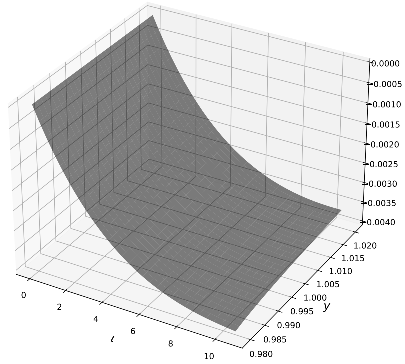

(e) $\text{direct}_{\mathcal{O},\iota}$

Figure 4: Overall elasticities on a grid for y and ι (annualized).

8 Quantitative results

We now confirm, by simulating the fully dynamic calibrated model, that it can generate the correlation patterns we illustrated in the cross-country data. We then use the model to compute the welfare cost of inflation and conduct counterfactual experiments.

8.1 Nonlinear inflation-unemployment correlations

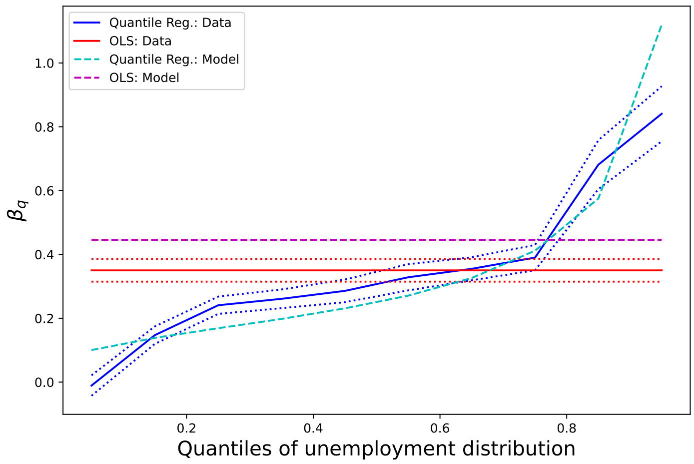

Figure 5: OLS and quantile regression coefficients of $\bar{u}$ on $\bar{\imath}$ for various quantiles of $\bar{u}$: model vs. data.

Notes: The dotted lines represent the 95% confidence intervals. Data are from the OECD. Both unemployment and long-term nominal interest rate series for each country are filtered out of high frequency variations using the HP filter with $\lambda = 1,600$. Model-based regressions are computed using 1,000 simulations of 1,000 months each. For each simulation, we burn the first 136 periods to match the length of the data series and then take quarterly averages.

In order to assess the ability of the model to quantitatively replicate the stylized facts discussed in the empirical section, we simulate the model under both productivity and nominal interest rate shocks.[17] First, we examine the inflation-unemployment relationship in the calibrated model. Figure 5 shows the coefficients obtained by running OLS and quantile regressions of HP-filtered trend unemployment on trend interest rates using the simulated quarterly data. For comparison purposes, we plot along the regression coefficients obtained from the OECD panel data seen in Figure 1. The slope of the OLS regression stands at 0.45. This is close to the value of 0.35

[17] We run 1,000 simulations of the model each extending for 1,000 months and drop the first 136 months to match the length of the US data series used in the calibration. We then aggregate the simulated data by taking quarterly averages.

Table 7: Regression of unemployment volatility on $\bar{\iota}$ using simulated data

	log unemployment volatility
	(1)
Constant	0.030***
	(0.000)
Trend long-term rate	0.014***
	(0.000)
Observations	269,000
R^2	0.184
	*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parenthesis. Data are based on model simulations. Unemployment volatility is measured as the 5-year rolling window standard deviation of HP-detrended log unemployment. Long-term nominal interest rate series is filtered out of high frequency variations using the HP filter with $\lambda = 1600$.

that we obtained using a pooled OLS regression on OECD data. The coefficients of the quantile regressions computed using the simulated data exhibit a very similar shape to the empirical ones. As the level of trend unemployment increases, the effect of an increase in trend interest rate is strengthened. At the 5th percentile of the distribution of unemployment, a 1pp increase in $\bar{\iota}$ leads to a 0.10pp increase in $\bar{u}$. In contrast, at the 95th percentile the same increase in $\bar{\iota}$ leads to a 1.12pp increase in $\bar{u}$ – about a tenfold increase in the effect. In the model, this nonlinear effect is caused by the higher elasticity of job creation when the surplus of the firm-worker match is smaller. The latter occurs in states where unemployment is high, either due to low productivity or high inflation or both.

Next, we regress unemployment volatility on the trend nominal interest rate using simulated data. Table 7 presents the results. In particular, a 1pp increase in $\bar{\iota}$ leads to an increase of 0.013 in unemployment volatility. This is very close to the value of 0.01 obtained from the fixed effects panel regression in Table 2 using OECD data.

8.2 Correlation between inflation and labor productivity

Our model implies that – similarly to Berentsen et al. (2011) – the effect of anticipated inflation on unemployment is transmitted through lower measured labor productivity. In fact, as illustrated in Section 7, the effect of inflation on the labor market surplus through reduced output per worker is the primary force behind its nonlinear effects, and is more important quantitatively than the

Table 8: Regression of labor productivity on $\bar{\iota}$ using simulated data

	Labor productivity (in log)
	(1)
Constant	0.103***
	(0.000)
Trend long-term rate	-0.003***
	(0.000)
Observations	288,000
R^2	0.123
	*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parentheses. Data are based on model simulations. Labor productivity is measured in logs. Long-term nominal interest rate series is filtered out of high frequency variations using the HP filter with $\lambda = 1600$.

subsequent feedback effect from labor to goods markets. A testable prediction of our model, therefore, is that the correlation between anticipated inflation and measured output per worker is negative as well. We illustrate this in Table 8, which displays the results of an OLS regression of output per worker ($\mathcal{O}$) on trend nominal interest rate ($\bar{\iota}$) in model-simulated data.

To investigate the relationship between anticipated inflation and labor productivity in OECD data, we regress real output per worker (i.e. GDP per person employed) on trend nominal interest rates. Column (1) of Table 9 presents the results of the pooled OLS regression

$$\ln \mathcal{O}_{jt} = \alpha + \beta\bar{\iota}_{jt} + \varepsilon_{jt}, \tag{67}$$

where $\ln \mathcal{O}_{jt}$ is the natural logarithm of real output per worker and $\bar{\iota}_{jt}$ is the trend component of the long-term nominal interest rate in country j at quarter t. It indicates that a 1pp increase in $\bar{\iota}$ is associated with a 4% decrease in output per worker. Columns (2), (3) and (4) present results from fixed-effects panel regressions of the type

$$\ln \mathcal{O}_{jt} = \alpha + \beta\bar{\iota}_{jt} + \gamma_j + \delta_t + \varepsilon_{jt}, \tag{68}$$

where γ_j and δ_t represent country and time fixed effects. All three panel regressions confirm the negative relationship between output per worker and long-term interest rates. In particular, the specification depicted in Column (4) includes both country and time fixed effects. The latter

Table 9: Regression of labor productivity on $\bar{\iota}$

	Labor productivity (in log)			
	(1)	(2)	(3)	(4)
Constant	4.470***	4.639***	4.498***	4.550***
	(0.020)	(0.019)	(0.026)	(0.022)
Trend long-term rate	-0.040***	-0.038***	-0.010*	-0.020***
	(0.004)	(0.004)	(0.006)	(0.004)
Observations	3,658	3,658	3,658	3,658
R^2	0.327	0.337	0.024	0.075
F-Statistic	1779.30***	1844.00***	84.22***	271.14***
Country fixed effects	No	Yes	No	Yes
Time fixed effects	No	No	Yes	Yes
Clustered errors (country level)	Yes	Yes	Yes	Yes

*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parentheses. Data are from the OECD. Labor productivity is measured as the natural logarithm of real GDP per person employed. Long-term nominal interest rate series for each country is filtered out of high frequency variations using the HP filter with $\lambda = 1600$.

control for any time-trends in labor productivity that are common across countries. In this specification, a 1pp increase in long-term interest rates is associated with a 2% decrease in labor productivity. Table 16 in the appendix presents the same regressions using the 5-year moving average instead of the HP filter to compute $\bar{\iota}$. The results are quantitatively similar.

Comparison of Tables 8 and 9 shows that the model's prediction is qualitatively consistent with the data. In fact, the estimated negative correlation between $\bar{\iota}$ and output per worker is even stronger in the data than it is in model-simulated data. This may be potentially rationalized by the fact that our model abstracts from other sources of such a negative correlation. In particular, we assume that the exogenous stochastic processes for ι_t and y_t (i.e. the exogenous productivity shock) are uncorrelated, so that our model accounts for only the endogenous relationship between trend inflation and output per worker.

8.3 Generalized impulse response functions

To further examine the dynamics of the model, and in particular the state-dependent reaction of unemployment and other endogenous variables to shocks, we follow Gallant et al. (1993) and

Koop et al. (1996) in computing the generalized, nonlinear, impulse response function

$$GIRF_Y(k, \varepsilon_t, \Omega_t) = \mathbb{E}[Y_{t+k}|\varepsilon_t, \Omega_t = \omega_t] - \mathbb{E}[Y_{t+k}|\Omega_t = \omega_t], \tag{69}$$

where $\Omega_t = \omega_t$ is the state of the economy at the beginning of period t and ε_t is an innovation to the exogenous variable at time t. This function measures the update in the conditional expectation of the variable Y_{t+k} implied by a shock ε_t given the state of the economy $\Omega_t = \omega_t$. The shock ε_t could refer to either the productivity process y_t or the cyclical interest rate process $\hat{\imath}_t$. The GIRF in equation (69) is a random variable and its shape will in general be a function of $\Omega_t = \{\bar{\imath}_t, \hat{\imath}_t, y_t, u_t\}$, the state of the economy at the moment of the shock. The top row of Figure 6 plots the distribution of the GIRFs (light blue) for key model variables in reaction to a shock to the productivity process y_t of size $\varepsilon_t = \sigma_y$. This is done by evaluating the GIRF above at different initial states randomly drawn from the ergodic distribution of Ω_t.[18] By averaging across the initial states, one can obtain the mean GIRF given by

$$\mathbb{E}[GIRF_Y(k, \varepsilon_t, \Omega_t)] = \mathbb{E}[Y_{t+k}|\varepsilon_t] - \mathbb{E}[Y_{t+k}], \tag{70}$$

where the expectation operator is taken over the ergodic distribution of the state Ω_t. These are depicted in dark blue in the top row panels of Figure 6. Notice that depending on the initial state of the economy the reaction to productivity shocks can be dramatically different from the average. While on average unemployment increases by 0.26pp on impact, that reaction can go up to 0.67pp at the 95th percentile. The reaction of DM consumption is strongly hump-shaped, in particular at higher levels of trend inflation. This is a result of the hump-shaped reaction of unemployment which gradually amplifies goods market matching frictions and reduces money demand.

The bottom row of Figure 6 depicts the mean GIRFs conditional on the level of trend inflation. In states where trend inflation is high (in red), the reaction of the economy to shocks is stronger. For example, the average reaction of unemployment on impact is 1.5 times (2.1 times) stronger when trend inflation is about 8% compared to 3% (0%). Labor market tightness reacts on average 1.5 times (1.9 times) stronger under 8% trend inflation compared to 3% (0%).

[18] The ergodic distribution of the model's state variables is obtained by running 10,000 simulations of 1,000 months each and burning the first 136 months to obtain the same length as the data. We then take 1,000 random draws from the distribution. For each initial state drawn, we obtain the GIRF by computing the difference between the conditional expectations with and without the shock, each averaged over 10,000 simulations of the model, of 100 periods' length, starting from that initial state. We use percentage point deviations for unemployment. We use percentage deviations for the other variables.

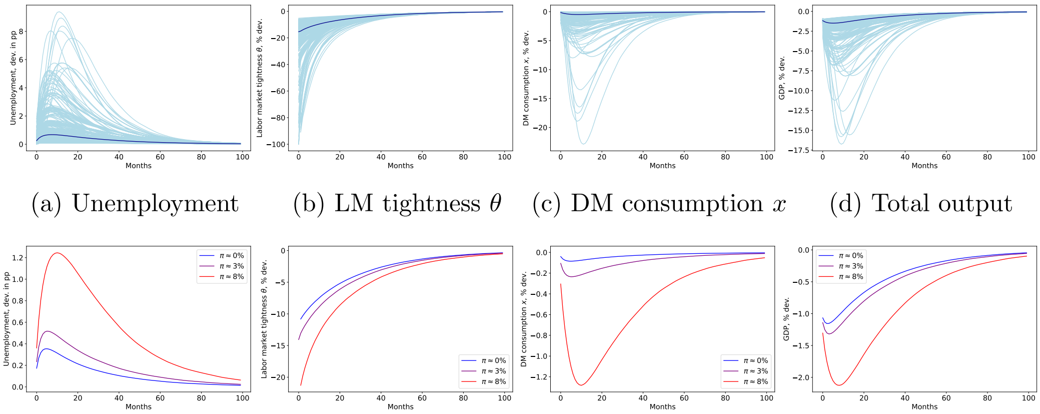

(a) Unemployment (b) LM tightness θ (c) DM consumption x (d) Total output

(e) Unemployment (f) LM tightness θ (g) DM consumption x (h) Total output

Figure 6: Reaction to a negative productivity shock: unconditional and conditional GIRFs

Notes: GIRFs for u, θ, x and total output following a 1 standard deviation negative productivity shock. The top row depicts the distribution of GIRFs (light blue) and their mean (dark blue). Each curve represents the GIRF evaluated at a particular initial state drawn randomly from the ergodic distribution of the calibrated model. The bottom row depicts the mean GIRFs conditional on the level of trend inflation being at about 0%, 3% and 8%.

8.4 Robustness checks

In this section we explore how, and to what extent, the nonlinear behavior predicted by the model depends on specific parameter choices. Throughout, we focus on the three statistics emphasized above: the average relationship between $\bar{\iota}$ and $\bar{u}$, the state-dependent nature of that relationship and the relationship between $\bar{\iota}$ and the volatility of u. The first is measured using the coefficient of the OLS regression $\bar{u}$ on $\bar{\iota}$. We measure the second using the coefficient of the quantile regression of $\bar{u}$ on $\bar{\iota}$ at the 5th and 95th percentiles of the distribution of $\bar{u}$. The third is measured using the coefficient of the OLS regression of the volatility of $\log u$ on $\bar{\iota}$.

Unemployment flow value. As is well-known in the literature on unemployment volatility and confirmed in section 7, the value of non-market activity, captured by b, is crucial in determining the strength of the fundamental surplus (FS) effect. We therefore assess the sensitivity of our quantitative results with respect to this parameter. Our baseline calibration puts b at 0.913 of average output, which amounts to $b = 0.99$. We experiment with both higher and lower values of non-market activity. Since changing b also changes the average unemployment rate, we re-calibrate the cost of vacancy posting k in each experiment so as to keep the average unemployment rate fixed. Table 10 reports the resulting calibrated parameters, as well as the corresponding statistics of interest. As expected from the analysis in Section 7, lower values of b dampen the effect of ι on unemployment as well as its nonlinearity. In particular, lower values of b reduce the

Table 10: Sensitivity analysis – Unemployment flow value

Statistic	**Parameter change**				
	$\boldsymbol{b = 1.029}$ $k = 0.63$	$b = 0.99$ $k = 1.472$	$\boldsymbol{b = 0.921}$ $k = 2.790$	$\boldsymbol{b = 0.867}$ $k = 3.776$	$\boldsymbol{b = 0.759}$ $k = 5.715$
$b/\mathcal{O}$	0.950	0.913	0.850	0.800	0.700
standard deviation of log u	0.226	0.138	0.073	0.053	0.035
$\bar{u}$ on $\bar{\iota}$: OLS coefficient	0.926	0.446	0.191	0.136	0.087
$\bar{u}$ on $\bar{\iota}$: QReg. coef. at 5th perc.	0.050	0.100	0.099	0.085	0.065
$\bar{u}$ on $\bar{\iota}$: QReg. coef. at 95th perc.	2.854	1.120	0.337	0.209	0.117
log u vol. on $\bar{\iota}$: OLS coefficient	0.030	0.014	0.004	0.002	0.001

Notes: Data are based on monthly model simulations aggregated on a quarterly basis. Nominal interest rate and unemployment series are filtered out of high frequency variations using the HP filter with $\lambda = 1600$. Unemployment volatility is measured as the standard deviation of HP-detrended log unemployment. The volatility regression uses a 5-year rolling window standard deviation.

implied regression coefficient of $\bar{u}$ on $\bar{\iota}$, but also shrink the gap in the coefficients at 95th and 5th percentiles. Lower values of b also dramatically reduce the effect of ι on unemployment volatility. This is not surprising, as lower values of b also drastically reduce unemployment volatility itself, as expected from the previous literature (Shimer (2005), Hagedorn and Manovskii (2008)) and confirmed in our results.

Elasticity of the labor market matching function. As we highlight in Section 7, not only the fundamental surplus (FS) effect but also the labor market matching (LMM) effect is quantitatively important for the model's amplification mechanism. The former captures by how much an increase in match output affects match surplus. The latter captures how much an increase in the match surplus in turn affects labor market tightness. The strength of this channel depends on the elasticity of the matching function, which in turn depends on the parameter χ.

In order to illustrate this, we simulate the model under different values of χ. For each value of χ, we recalibrate the cost of entry k to keep unemployment at its level in the baseline calibration. Table 11 reports the parameter values as well as the simulation results. Compared to our baseline calibration, a lower χ leads to a weaker transmission from $\bar{\iota}$ to $\bar{u}$ both on average and at the 95th and 5th percentiles of the distribution of $\bar{u}$. In addition, the nonlinear effects become weaker, as evidenced by the ratio of the quantile regression coefficients at 95th and 5th percentiles. The coefficient of the volatility regression likewise drops. Intuitively, a lower χ implies a lower elasticity of labor market tightness to match surplus, and furthermore implies that this elasticity is less sensitive to changes in labor market tightness.

Table 11: Sensitivity analysis – labor market matching function

Statistic	**Parameter change**			
	$\boldsymbol{\chi = 1.4}$ $k = 1.531$	$\chi = 1.269$ $k = 1.472$	$\boldsymbol{\chi = 1}$ $k = 1.249$	$\boldsymbol{\chi = 0.7}$ $k = 0.753$
$\bar{u}$ on $\bar{\iota}$: OLS coefficient	0.519	0.446	0.302	0.182
$\bar{u}$ on $\bar{\iota}$: QReg. coef. at 5th perc.	0.094	0.100	0.106	0.094
$\bar{u}$ on $\bar{\iota}$: QReg. coef. at 95th perc.	1.414	1.120	0.638	0.318
log u vol. on $\bar{\iota}$: OLS coefficient	0.016	0.014	0.009	0.004

Notes: Data are based on monthly model simulations aggregated on a quarterly basis. Nominal interest rate and unemployment series are filtered out of high frequency variations using the HP filter with $\lambda = 1600$. Unemployment volatility is measured as the 5-year rolling window standard deviation of HP-detrended log unemployment.

Relative size of the decentralized market. Finally, we conduct robustness checks with respect to the relative size of the decentralized market, measured below by $\mathcal{P}/\mathcal{O}$. Intuitively, this magnitude is important because the direct effect of inflation is the distortion of x, the quantity traded in the decentralized market. We manipulate the size of this market by changing the value of A, which controls the average money demand. For each experiment, we recalibrate the average productivity of the centralized market, $\overline{y}$, so as to keep the average *total* output per worker, $\mathcal{O}$, constant. Table 12 reports the results of changes to A. As expected, larger values of A lead to stronger effects of inflation and stronger nonlinearities; however, the results are qualitatively and quantitatively robust, even to large changes (e.g. a doubling) in the relative size of the monetary sector.

To summarize the analysis of this section, our results are broadly robust qualitatively to changes in the key parameters. Quantitatively, the main results are quite sensitive to changes in the key labor market parameters, most notably the value of unemployment consumption and the elasticity of the matching technology. On the other hand, goods market parameters do affect the model's amplification mechanism, but even a substantial increase in the size of the decentralized goods market does not have a drastic effect. This confirms the discussion of Section 7, which argued that labor market parameters that have conventionally guided the unemployment volatility literature continue to be of primary importance here.

Table 12: Sensitivity analysis – relative DM size

Statistic	Parameter change			
	$\boldsymbol{A = 1.294}$ $\bar{y} = 1.029$	$A = 1.421$ $\bar{y} = 1$	$\boldsymbol{A = 1.505}$ $\bar{y} = 0.975$	$\boldsymbol{A = 1.644}$ $\bar{y} = 0.921$
$\mathcal{P}/\mathcal{O}$	0.050	0.077	0.100	0.150
$\bar{u}$ on $\bar{\iota}$: OLS coefficient	0.277	0.446	0.630	1.273
$\bar{u}$ on $\bar{\iota}$: QReg. coef. at 5th perc.	0.064	0.100	0.133	0.202
$\bar{u}$ on $\bar{\iota}$: QReg. coef. at 95th perc.	0.775	1.120	1.504	2.892
log u vol. on $\bar{\iota}$: OLS coefficient	0.009	0.013	0.018	0.028

Notes: Data are based on monthly model simulations aggregated on a quarterly basis. Nominal interest rate and unemployment series are filtered out of high frequency variations using the HP filter with $\lambda = 1600$. Unemployment volatility is measured as the 5-year rolling window standard deviation of HP-detrended log unemployment.

9 Welfare cost of inflation

We compute the welfare cost of trend inflation in consumption equivalent terms as follows. Given the stochastic processes $\{\bar{\iota}_t, \hat{\iota}_t, y_t\}_{t=0}^{\infty}$, equilibrium welfare is given by

$$\mathcal{W}(\{\bar{\iota}_t, \hat{\iota}_t, y_t\}_{t=0}^{\infty}) = \mathbb{E}\sum_{t=0}^{\infty} \beta^t \left[\alpha(n_t)[\mathrm{u}(x_t) - \mathrm{c}(x_t)] + n_t y_t + (1-n_t)b - \kappa v_{t+1}\right]. \tag{71}$$

We can write

$$Y_t = n_t y_t - \alpha(n_t)\mathrm{c}(x_t) + (1-n_t)b - \kappa v_{t+1} \tag{72}$$

for the net consumption of the CM good at time t, whereas x_t is the consumption of DM good at time t, so that

$$\mathcal{W}(\{\bar{\iota}_t, \hat{\iota}_t, y_t\}_{t=0}^{\infty}) = \mathbb{E}\sum_{t=0}^{\infty} \beta^t \left[\alpha(n_t)\mathrm{u}(x_t) + Y_t\right]. \tag{73}$$

We can also compute welfare when the trend component of inflation is fixed at $\bar{\iota}_t = 0$, $\forall t$. Letting $\{n_t^*, x_t^*, Y_t^*\}_{t=0}^{\infty}$ be the equilibrium allocation induced by the stochastic process $\{0, \hat{\iota}_t, y_t\}_{t=0}^{\infty}$, we can write

$$\mathcal{W}(\{0, \hat{\iota}_t, y_t\}_{t=0}^{\infty}) = \mathbb{E}\sum_{t=0}^{\infty} \beta^t \left[\alpha(n_t^*)\mathrm{u}(x_t^*) + Y_t^*\right]. \tag{74}$$

We then define the welfare cost of trend inflation as the fraction $1-\Delta$ of consumption agents are willing to give up in both markets in order to move from the trend inflation process $\bar{\iota}_t$ to a zero

Table 13: Welfare cost of inflation in the baseline economy

Annual trend inflation	**Implied trend interest rate,** $\bar{\iota}$	**Welfare cost,** $(1-\Delta(\bar{\iota}))\%$
-2.75%	0.00%	-
0.00%	2.82%	0.48%
2.50%	5.40%	1.40%
5.00%	7.97%	2.54%
7.50%	10.54%	3.81%
10.00%	13.11%	5.28%

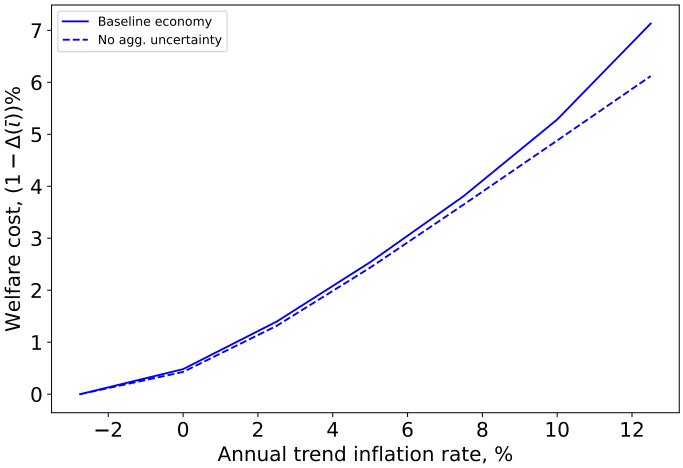

Figure 7: Welfare cost of trend inflation with and without aggregate uncertainty.

trend nominal interest rate, all else equal. In other words, the quantity Δ solves

$$\mathcal{W}(\{\bar{\iota}_t, \hat{\iota}_t, y_t\}_{t=0}^{\infty}) = \mathbb{E}\sum_{t=0}^{\infty}\beta^t\left[\alpha(n_t^*)\mathrm{u}(\Delta x_t^*) + \Delta Y_t^*\right]. \tag{75}$$

To measure this welfare cost of inflation, we simulate the model with cyclical interest rate and productivity shocks under various levels of trend inflation. We then calculate the welfare cost for each trend inflation level by averaging across the simulations. Table 13 presents the results for various levels of trend inflation. In particular, increasing the trend inflation rate from -2.75% (corresponding to $\bar{\iota} = 0$) to 10% leads to a welfare loss of 5.28%.[19]

[19]This figure is comparable in magnitude to the range of estimates of welfare cost of inflation from search-based models surveyed in Lagos et al. (2017). It should be noted that the estimates of the welfare cost of inflation are not immediately comparable across models; in particular, the estimates surveyed in Lagos et al. (2017) come from models without frictional labor markets.

Nonlinear cost of inflation. The solid blue line in Figure 7 plots the welfare cost of inflation as a function of the annual trend inflation rate in the baseline economy. One can clearly see that the slope is increasing in $\bar{\iota}$. This means that the effect of a marginal increase in trend inflation becomes stronger as the latter increases. In particular, increasing the trend nominal interest rate from -2.75% to 5% reduces welfare by 2.54%; increasing it from 5% to 10% leads to an additional 2.75pp of welfare loss. This is a 8% stronger effect. In summary, the model predicts that the welfare loss from inflation is significantly nonlinear as well.

The role of aggregate uncertainty. Compared to the previous literature, our model has an additional element that amplifies the welfare cost of inflation: aggregate uncertainty. Its role is related to the observation that higher levels of trend inflation lead to a smaller match surplus, resulting in higher volatility of unemployment following aggregate shocks. Because the reaction of unemployment to shocks is asymmetric, average unemployment under aggregate uncertainty is higher (see e.g. Hairault et al. (2010), Jung and Kuester (2011), and Petrosky-Nadeau and Zhang (2020)). This leads to a lower level of welfare compared to an economy without aggregate shocks. To highlight the effects of this amplification channel on welfare, Figure 7 plots in dashed blue the average welfare as a function of trend inflation for an economy where we shut down aggregate uncertainty. The economy without aggregate shocks implies a welfare cost of inflation of 4.88%, as opposed to 5.28% in the baseline economy. That represents an additional 0.40pp or 7.64% of the total welfare cost of inflation in the baseline economy. Furthermore, Figure 7 shows that the aggregate uncertainty channel seems to matter substantially only at high levels of inflation (above 7%).

Decomposing the welfare cost of inflation. In addition to the direct distortionary effect of inflation on DM trade, our model environment is plagued by congestion externalities standard in frictional labor and product markets. On the labor market side, firms do not internalize the effect of their vacancy posting decisions on the matching probability of workers and other firms, as explained in the seminal work of Hosios (1990). Furthermore, on the product market side, firms do not internalize the effect of aggregate employment on the matching probabilities of buyers and sellers. As pointed out by Rocheteau and Wright (2005), the latter effect implies that there may be excessive entry of firms, which may be mitigated by higher inflation.

Motivated by this observation, we seek to understand to what extent our measured welfare cost of inflation is driven by its direct effect as a tax on DM trade as opposed to congestion externalities. While there are many ways to conduct such a decomposition, we choose one that is simple to implement and to interpret. Consider an equilibrium allocation of employment and DM quantities $\{n_t, x_t\}_{t=0}^{\infty}$ induced by a given path of inflation. We can then consider the problem of a fictitious social planner who takes the DM traded quantities x_t as given, but can dictate firms' vacancy posting and thereby choose n_t. The gap in welfare between the solution to this problem and the true equilibrium provides a measure of congestion inefficiencies, for a fixed (intensive-margin) amount of DM trade. We can then decompose the welfare change in response to an increase in nominal rates from ι to ι' into a component stemming from the change in x and a component stemming from the *change* in the magnitude of the congestion inefficiency. Appendix D reports the details of this decomposition exercise. Consistent (qualitatively) with the intuition of e.g. Rocheteau and Wright (2005), we find that the welfare-reducing effect of inflation is driven by its intensive-margin effect on x, whereas the congestion inefficiency is in fact mitigated by higher inflation.

10 Conclusion

In this paper we have argued that there is an important interaction between the inflation tax and business cycles. Such interaction arises naturally in a standard monetary search framework incorporating both labor market frictions and a role for money as a medium of exchange. The analysis illustrates that such interaction is potentially more important at high levels of inflation. We have also shown quantitatively that the small-surplus logic driving unemployment dynamics in the DMP model continues to be the dominant force here, even in the presence of confounding effects from goods-labor market interactions.

Our findings are of potential relevance given the recent increase in inflation in many countries: an immediate implication of our analysis is that unemployment volatility in response to various shocks is not invariant to this higher level of inflation. More generally, the mechanism we highlight has to do with the interaction between two forces driving unemployment: in our case, inflation lowers the match surplus, which in turn makes unemployment more sensitive to both non-monetary shocks and further increases in inflation. The insight is applicable much more broadly to studying unemployment dynamics in response to multiple shocks.

References

Mohammed Ait Lahcen. Informality, frictional markets, and monetary policy. 2020.

S Boragan Aruoba, Christopher J Waller, and Randall Wright. Money and capital. *Journal of Monetary Economics*, 58(2):98–116, 2011.

Guido Ascari and Argia M. Sbordone. The macroeconomics of trend inflation. *Journal of Economic Literature*, 52(3):679–739, 2014.

Guido Ascari, Paolo Bonomolo, and Qazi Haque. The Long-Run Phillips Curve is... a Curve. 2022.

Regis Barnichon. Building a composite help-wanted index. *Economics Letters*, 109(3):175–178, 2010.

Aleksander Berentsen, Guido Menzio, and Randall Wright. Inflation and unemployment in the long run. *American Economic Review*, 101(1):371–98, 2011.

Joshua Bernstein, Alexander W Richter, and Nathaniel A Throckmorton. Nonlinear search and matching explained. *Federal Reserve Bank of Dallas Working Paper*, 2106, 2021.

Zach Bethune and Guillaume Rocheteau. Unemployment, aggregate demand, and the distribution of liquidity. 2017.

Zachary Bethune, Guillaume Rocheteau, and Peter Rupert. Aggregate unemployment and household unsecured debt. *Review of Economic Dynamics*, 18(1):77–100, 2015.

Andreas Beyer and Roger EA Farmer. Natural rate doubts. *Journal of Economic Dynamics and Control*, 31(3):797–825, 2007.

Satyajit Chatterjee and Dean Corbae. On the aggregate welfare cost of great depression unemployment. *Journal of monetary Economics*, 54(6):1529–1544, 2007.

Thomas F Cooley and Gary D Hansen. The inflation tax in a real business cycle model. *The American Economic Review*, pages 733–748, 1989.

Jan De Loecker, Jan Eeckhout, and Gabriel Unger. The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644, 2020.

Wouter J Den Haan, Garey Ramey, and Joel Watson. Job destruction and propagation of shocks. *American economic review*, 90(3):482–498, 2000.

Mei Dong and Sylvia Xiaolin Xiao. Liquidity, monetary policy, and unemployment: a new monetarist approach. *International Economic Review*, 60(2):1005–1025, 2019.

A Ronald Gallant, Peter E Rossi, and George Tauchen. Nonlinear dynamic structures. *Econometrica: Journal of the Econometric Society*, pages 871–907, 1993.

Pedro Gomis-Porqueras, Benoit Julien, and Chengsi Wang. Optimal monetary and fiscal policies in a search-theoretic model of money and unemployment. *Macroeconomic Dynamics*, 17(6): 1330–1354, 2013.

Pedro Gomis-Porqueras, Stella Huangfu, and Hongfei Sun. The role of search frictions in the long-run relationships between inflation, unemployment and capital. *European Economic Review*, 123:103396, 2020.

Chao Gu, Janet Hua Jiang, and Liang Wang. Credit condition, inflation and unemployment. 2021.

Marcus Hagedorn and Iourii Manovskii. The cyclical behavior of equilibrium unemployment and vacancies revisited. *American Economic Review*, 98(4):1692–1706, 2008.

Jean-Olivier Hairault, François Langot, and Sophie Osotimehin. Matching frictions, unemployment dynamics and the cost of business cycles. *Review of Economic Dynamics*, 13(4):759–779, 2010.

Alfred A Haug and Ian King. In the long run, us unemployment follows inflation like a faithful dog. *Journal of Macroeconomics*, 41:42–52, 2014.

Chao He and Min Zhang. Money allocation, unemployment, and monetary policy. 2020.

Robert J Hodrick and Edward C Prescott. Postwar us business cycles: an empirical investigation. *Journal of Money, credit, and Banking*, pages 1–16, 1997.

Kuk Mo Jung and Ju Hyun Pyun. A long-run approach to money, unemployment and equity prices. 2020.

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

Gary Koop, M Hashem Pesaran, and Simon M Potter. Impulse response analysis in nonlinear multivariate models. *Journal of econometrics*, 74(1):119–147, 1996.

Ricardo Lagos and Randall Wright. A unified framework for monetary theory and policy analysis. *Journal of Political Economy*, 113(3):463–484, 2005.

Ricardo Lagos, Guillaume Rocheteau, and Randall Wright. Liquidity: A new monetarist perspective. *Journal of Economic Literature*, 55(2):371–440, 2017.

Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A llvm-based python jit compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, pages 1–6, 2015.

Lars Ljungqvist and Thomas J Sargent. The fundamental surplus. *American Economic Review*, 107(9):2630–65, 2017.

Robert E Lucas. Inflation and welfare. *Econometrica*, 68(2):247, 2000.

Robert E Lucas and Juan Pablo Nicolini. On the stability of money demand. *Journal of Monetary Economics*, 73:48–65, 2015.

Nicolas Petrosky-Nadeau and Etienne Wasmer. Macroeconomic dynamics in a model of goods, labor, and credit market frictions. *Journal of Monetary Economics*, 72:97–113, 2015.

Nicolas Petrosky-Nadeau and Lu Zhang. Solving the dmp model accurately. *Quantiative Economics*, 2017.

Nicolas Petrosky-Nadeau and Lu Zhang. Unemployment crises. *Journal of Monetary Economics*, 2020.

Nicolas Petrosky-Nadeau, Lu Zhang, and Lars-Alexander Kuehn. Endogenous disasters. *American Economic Review*, 108(8):2212–45, 2018.

Robert H Rasche. M1-velocity and money-demand functions: Do stable relationships exist? In *Carnegie-Rochester Conference Series on Public Policy*, volume 27, pages 9–88. Elsevier, 1987.

Guillaume Rocheteau and Antonio Rodriguez-Lopez. Liquidity provision, interest rates, and unemployment. *Journal of Monetary Economics*, 65:80–101, 2014.

Guillaume Rocheteau and Randall Wright. Money in search equilibrium, in competitive equilibrium, and in competitive search equilibrium. *Econometrica*, pages 175–202, 2005.

Geert K. Rouwenhorst. 'asset pricing implications of equilibrium business cycle models' chapter 10, frontiers of business cycle research, t. cooley, 1995.

Francisco Ruge-Murcia. Estimating nonlinear dsge models by the simulated method of moments: With an application to business cycles. *Journal of Economic Dynamics and Control*, 36(6): 914–938, 2012.

Thomas Sargent and John Stachurski. Quantitative economics. Technical report, Citeseer, 2014.

Robert Shimer. The cyclical behavior of equilibrium unemployment and vacancies. *American economic review*, 95(1):25–49, 2005.

Robert Shimer. Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15(2):127–148, 2012.

Appendix A Tables and figures for empirical results

Table 14: Regression of $\bar{u}$ on $\bar{\imath}$ (5-year moving averages)

	Trend unemployment			
	(1)	(2)	(3)	(4)
Constant	5.837***	6.066***	3.005**	2.837***
	(0.597)	(0.385)	(1.302)	(1.029)
Trend long-term rate	0.366***	0.324***	0.885***	0.915***
	(0.112)	(0.071)	(0.269)	(0.188)
Observations	3,262	3,262	3,262	3,262
R^2	0.083	0.142	0.167	0.200
F-Statistic	295.68***	532.55***	600.89***	744.80***
Country fixed effects	No	Yes	No	Yes
Time fixed effects	No	No	Yes	Yes
Clustered errors (country level)	Yes	Yes	Yes	Yes

*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parentheses. Data are from the OECD. Both unemployment and long-term nominal interest rate series for each country are filtered out of high frequency variations by taking 5-year moving averages.

Table 15: Regression of log $\bar{u}$ on $\bar{\iota}$ (HP filter)

	Trend log unemployment			
	(1)	(2)	(3)	(4)
Constant	1.708***	1.755***	1.556***	1.706***
	(0.076)	(0.051)	(0.175)	(0.114)
Trend long-term rate	0.039***	0.030***	0.067**	0.039*
	(0.011)	(0.009)	(0.033)	(0.021)
Observations	4,007	4,007	4,007	4,007
R^2	0.072	0.090	0.072	0.024
F-Statistic	311.62***	393.96***	289.76***	92.32***
Country fixed effects	No	Yes	No	Yes
Time fixed effects	No	No	Yes	Yes
Clustered errors (country level)	Yes	Yes	Yes	Yes

*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parenthesis. Data are from the OECD. Both the logarithm of unemployment and long-term nominal interest rate series for each country are filtered out of high frequency variations using the HP filter with $\lambda = 1600$.

Table 16: Regression of log $\bar{u}$ on $\bar{\imath}$ (5-year moving averages)

	Trend log unemployment			
	(1)	(2)	(3)	(4)
Constant	1.729***	1.776***	1.493***	1.663***
	(0.082)	(0.045)	(0.156)	(0.098)
Trend long-term rate	0.041***	0.032***	0.084***	0.053***
	(0.012)	(0.008)	(0.028)	(0.018)
Observations	3,262	3,262	3,262	3,262
R^2	0.075	0.101	0.110	0.052
F-Statistic	263.02***	364.35***	374.11***	164.11***
Country fixed effects	No	Yes	No	Yes
Time fixed effects	No	No	Yes	Yes
Clustered errors (country level)	Yes	Yes	Yes	Yes

*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parenthesis. Data are from the OECD. Both the logarithm of unemployment and long-term nominal interest rate series for each country are filtered out of high frequency variations using a 5-year moving average.

Table 17: Regression of (HP-detrended) unemployment volatility on $\bar{\iota}$

	Unemployment volatility			
	(1)	(2)	(3)	(4)
Constant	0.391***	0.354***	0.224*	-0.023
	(0.050)	(0.053)	(0.123)	(0.166)
Trend long-term rate	0.046***	0.053***	0.079***	0.127***
	(0.008)	(0.010)	(0.028)	(0.033)
Observations	3,616	3,616	3,616	3,616
R^2	0.077	0.132	0.090	0.135
F-Statistic	300.46***	544.05***	332.22***	519.68***
Country fixed effects	No	Yes	No	Yes
Time fixed effects	No	No	Yes	Yes
Clustered errors (country level)	Yes	Yes	Yes	Yes

*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parenthesis. Data are from the OECD. Unemployment volatility is measured as the 5-year rolling window standard deviation of HP-detrended unemployment. Long-term nominal interest rate series for each country is filtered out of high frequency variations using the HP filter with $\lambda = 1600$.

Table 18: Regression of (5-year moving average detrended) log unemployment volatility on $\bar{\iota}$

	log unemployment volatility			
	(1)	(2)	(3)	(4)
Constant	0.099***	0.085***	0.064***	0.007
	(0.014)	(0.019)	(0.023)	(0.049)
Trend long-term rate	0.009***	0.011***	0.015***	0.026***
	(0.003)	(0.004)	(0.005)	(0.009)
Observations	2,882	2,882	2,882	2,882
R^2	0.066	0.113	0.078	0.097
F-Statistic	201.77***	364.07***	224.29***	282.66***
Country fixed effects	No	Yes	No	Yes
Time fixed effects	No	No	Yes	Yes
Clustered errors (country level)	Yes	Yes	Yes	Yes

*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parenthesis. Data are from the OECD. Unemployment volatility is measured as the 5-year rolling window standard deviation of log unemployment detrended using a 5-year moving average. Long-term nominal interest rate series are filtered out of high frequency variations using a 5-year moving average.

Table 19: Regression of (5-year moving average detrended) unemployment volatility on $\bar{\iota}$

	Unemployment volatility			
	(1)	(2)	(3)	(4)
Constant	0.588***	0.523***	-0.234	-0.740**
	(0.143)	(0.129)	(0.288)	(0.357)
Trend long-term rate	0.098***	0.110***	0.256***	0.354***
	(0.031)	(0.025)	(0.065)	(0.069)
Observations	2,882	2,882	2,882	2,882
R^2	0.079	0.139	0.196	0.216
F-Statistic	248.16***	460.15***	650.05***	721.46***
Country fixed effects	No	Yes	No	Yes
Time fixed effects	No	No	Yes	Yes
Clustered errors (country level)	Yes	Yes	Yes	Yes

*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parenthesis. Data are from the OECD. Unemployment volatility is measured as the 5-year rolling window standard deviation of unemployment detrended using a 5-year moving average. Long-term nominal interest rate series are filtered out of high frequency variations using a 5-year moving average.

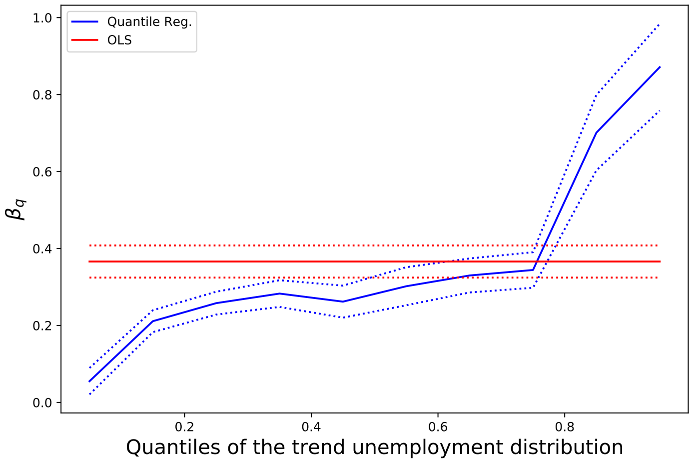

Figure 8: Quantile regression coefficients of $\bar{u}$ on $\bar{\iota}$ for various quantiles of $\bar{u}$ (5-year moving averages)

Notes: The dashed lines represent the 95% confidence intervals. Data are from the OECD. Both unemployment and long-term nominal interest rate series for each country are filtered out of high frequency variations using a 5-year moving average.

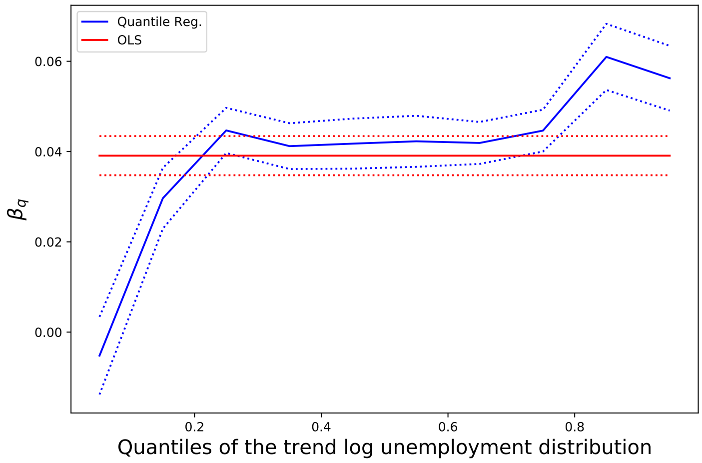

Figure 9: Quantile regression coefficients of trend log u on $\bar{\imath}$ for various quantiles of trend log u (HP filtered)

Notes: The dashed lines represent the 95% confidence intervals. Data are from the OECD. Both log unemployment and long-term nominal interest rate series for each country are filtered out of high frequency variations using the HP filter with $\lambda = 1600$.

Figure 10: Quantile regression coefficients of trend log u on $\bar{\imath}$ for various quantiles of trend log u (5-year moving averages)

Notes: The dashed lines represent the 95% confidence intervals. Data are from the OECD. Both log unemployment and long-term nominal interest rate series for each country are filtered out of high frequency variations using a 5-year moving average.

Table 20: Regression of labor productivity on $\bar{\iota}$ (5-year moving average)

	Labor productivity (in log)			
	(1)	(2)	(3)	(4)
Constant	4.670***	4.653***	4.496***	4.550***
	(0.024)	(0.021)	(0.020)	(0.034)
Trend long-term rate (5y moving average)	-0.037***	-0.034***	-0.006	-0.015**
	(0.004)	(0.004)	(0.004)	(0.006)
Observations	3,308	3,308	3,308	3,308
R^2	0.323	0.327	0.011	0.064
F-Statistic	1578.80***	1588.70***	35.18***	207.58***
Country fixed effects	No	Yes	No	Yes
Time fixed effects	No	No	Yes	Yes
Clustered errors (country level)	Yes	Yes	Yes	Yes

*p<0.1; **p<0.05; ***p<0.01

Notes: Standard errors are in parentheses. Data are from the OECD. Labor productivity is measured as the natural logarithm of real GDP per person employed. Long-term nominal interest rate series for each country is filtered out of high frequency variations using a 5-year moving average.

Appendix B Computations and calibration

Solution method. The model is solved numerically using value function iteration.[20] We discretize the state space $\Omega_t = \{\bar{\iota}_t, \hat{\iota}_t, y_t, u_{t-1}\}$ as follows: The continuous state stochastic processes for y_t and $\hat{\iota}_t$ are each approximated by a 30-state Markov chain using the Rouwenhorst (1995) procedure, which is shown by Petrosky-Nadeau and Zhang (2017), for the case of productivity shocks, to provide a better approximation when solving the DMP model nonlinearly.[21] The trend component $\bar{\iota}_t$ is modeled as a very persistent Markov chain with 5 states. The state values of the chain are computed by dividing the distribution of $\bar{\iota}_t$ into five quintiles and taking the average value of each quintile. This yields the values: $\{3.24\%, 4.37\%, 6.06\%, 7.74\%, 10.62\%\}$. We then estimate the transition probabilities by maximum likelihood as in Chatterjee and Corbae (2007). The maximum likelihood estimate of p_{jk}, the probability of transitioning from states j to k, is computed as the ratio of the number of times the economy transitions from state j to k to the number of times the economy is in state j. We obtain the following estimated transition probability matrix:

$$\begin{pmatrix} 0.994 & 0.006 & 0 & 0 & 0 \\ 0.006 & 0.988 & 0.006 & 0 & 0 \\ 0 & 0.006 & 0.988 & 0.006 & 0 \\ 0 & 0 & 0.006 & 0.988 & 0.006 \\ 0 & 0 & 0 & 0.006 & 0.994 \end{pmatrix}.$$

The space of the state variable u_{t-1} is discretized using a grid of 30 equidistant points. We linearly interpolate the value function between the grid points of u_t in order to improve accuracy.

Calibration procedure. We calibrate the internal parameters following a Simulated Method of Moments (SMM) procedure.[22] Let Θ be the vector containing the internal parameters, μ the vector of the targeted empirical moments and $\mu_s(\Theta)$ the vector of their model-based counterparts, obtained by simulating the model using a random draw s of productivity and interest rate shocks. The simulated moments are averaged over $S = 1,000$ simulations each of length $T = 1,000$.[23] We burn the first 136 observations to match the length of the empirical data series (864 monthly observations). The SMM procedure consists in solving for the vector $\hat{\Theta}$ that minimizes the distance $G(\Theta) = \mu - \frac{1}{S}\sum_{s=1}^{S} \mu_s(\Theta)$ such that

$$\hat{\Theta} = \arg\min_{\Theta} G(\Theta)^T W^{-1} G(\Theta) \tag{76}$$

[20] Our code is written in Python and uses the Numba library extensively for just-in-time compilation and parallelization (Lam et al., 2015).

[21] We use the Rouwenhorst routine from the QuantEcon Python library (Sargent and Stachurski, 2014).

[22] See for example Ruge-Murcia (2012) and references therein.

[23] To match empirical moments based on quarterly data, we aggregate our monthly simulations quarterly and compute the corresponding model-based moments.

where W is a semi-definite weighting matrix. Since our calibrated model is exactly identified, the choice of the weighting matrix is irrelevant. We use the percent difference to ensure that the distance is unit free and thus avoid unintended weighting.

Appendix C Additional details on model mechanism

This section contains additional details on both the derivations of the elasticities in Section 7 and the associated numerical decomposition.

C.1 Omitted derivations

This section presents the detailed derivations of the elasticities $\varepsilon_{\theta,y}$ and $\varepsilon_{\theta,\iota}$ in (65) and (66). We first derive the expressions for $\varepsilon_{\mathcal{P},n}$ and $\varepsilon_{\mathcal{P},\iota}$ in (62) and (64). Totally differentiating equation (60) with respect to x, y and ι, we find

$$\varphi\left(\frac{\mathrm{c}''(x)}{\mathrm{u}'(x)}-\frac{\mathrm{c}'(x)\mathrm{u}''(x)}{\mathrm{u}'(x)^2}\right)\mathrm{d}x=\left(\frac{\alpha'(n)}{\iota+\alpha(n)}-\frac{\alpha(n)\alpha'(n)}{[\iota+\alpha(n)]^2}\right)\mathrm{d}n-\frac{\alpha(n)}{[\iota+\alpha(n)]^2}\mathrm{d}\iota. \tag{77}$$

Rearranging terms yields

$$\varphi\frac{\mathrm{c}'(x)}{\mathrm{u}'(x)}\left(\frac{x\mathrm{c}''(x)}{\mathrm{c}'(x)}-\frac{x\mathrm{u}''(x)}{\mathrm{u}'(x)}\right)\frac{\mathrm{d}x}{x}=\frac{\iota\alpha(n)}{[\iota+\alpha(n)]^2}\left(\frac{n\alpha'(n)}{\alpha(n)}\frac{\mathrm{d}n}{n}-\frac{\mathrm{d}\iota}{\iota}\right). \tag{78}$$

Using equation (60) to substitute out $\varphi\mathrm{c}'(x)/\mathrm{u}'(x)$ together with $\sigma_{\mathrm{u},x}=-x\mathrm{u}''(x)/\mathrm{u}'(x)$, $\sigma_{\mathrm{c},x}=x\mathrm{c}''(x)/\mathrm{c}'(x)$ and $\epsilon_{\alpha,n}=n\alpha'(n)/\alpha(n)$, yields

$$\frac{\mathrm{d}x}{x}=\frac{1}{\sigma_{\mathrm{u},x}+\sigma_{\mathrm{c},x}}\frac{\iota\alpha(n)}{\iota+\alpha(n)}\frac{1}{\varphi\alpha(n)-(1-\varphi)\iota}\left(\epsilon_{\alpha,n}\frac{\mathrm{d}n}{n}-\frac{\mathrm{d}\iota}{\iota}\right). \tag{79}$$

Note that the term $\varphi\alpha(n)-(1-\varphi)\iota$ is strictly positive whenever we are in the relevant case in which x is strictly positive. Equation (60) namely implies that, keeping everything else constant, x is strictly decreasing in ι with $\lim_{\iota\to\frac{\varphi\alpha(n)}{1-\varphi}}=0$.

Totally differentiating $\mathcal{P}=\frac{\alpha(n)}{n}(1-\varphi)\left[\mathrm{u}(x)-\mathrm{c}(x)\right]$ with respect to x and n yields

$$\mathrm{d}\mathcal{P}=\left(\frac{\alpha'(n)}{n}-\frac{\alpha(n)}{n^2}\right)(1-\varphi)\left[\mathrm{u}(x)-\mathrm{c}(x)\right]\mathrm{d}n+\frac{\alpha(n)}{n}(1-\varphi)\left[\mathrm{u}'(x)-\mathrm{c}'(x)\right]\mathrm{d}x. \tag{80}$$

Rearranging terms yields

$$\frac{\mathrm{d}\mathcal{P}}{\alpha(n)(1-\varphi)\left[\mathrm{u}(x)-\mathrm{c}(x)\right]/n}=(\epsilon_{\alpha,n}-1)\frac{\mathrm{d}n}{n}+\frac{x\left[\mathrm{u}'(x)-\mathrm{c}'(x)\right]}{\mathrm{u}(x)-\mathrm{c}(x)}\frac{\mathrm{d}x}{x} \tag{81}$$

with $\epsilon_{\alpha,n}=n\alpha'(n)/n$. Substituting the expression for $\mathcal{P}$ into (81) and using (79) to substitute out $\mathrm{d}x/x$, we obtain

$$\frac{\mathrm{d}\mathcal{P}}{\mathcal{P}}=(\epsilon_{\alpha,n}-1)\frac{\mathrm{d}n}{n}+\frac{1}{\sigma_{\mathrm{u},x}+\sigma_{\mathrm{c},x}}\frac{\iota\alpha(n)}{\iota+\alpha(n)}\frac{1}{\varphi\alpha(n)-(1-\varphi)\iota}\frac{x\left[\mathrm{u}'(x)-\mathrm{c}'(x)\right]}{\mathrm{u}(x)-\mathrm{c}(x)}\left(\epsilon_{\alpha,n}\frac{\mathrm{d}n}{n}-\frac{\mathrm{d}\iota}{\iota}\right). \tag{82}$$

This immediately implies that

$$\frac{\mathrm{d}\mathcal{P}}{\mathcal{P}} = \varepsilon_{\mathcal{P},n}\frac{\mathrm{d}n}{n} + \varepsilon_{\mathcal{P},\iota}\mathrm{d}\iota. \tag{83}$$

with $\varepsilon_{\mathcal{P},n}$ defined by equation (62) and $\varepsilon_{\mathcal{P},\iota}$ defined by equation (64), as desired. Note that the sign of $\varepsilon_{\mathcal{P},n}$ is ambiguous while $\varepsilon_{\mathcal{P},\iota} = 0$ if $\iota = 0$ and $\varepsilon_{\mathcal{P},\iota} < 0$ if $\iota > 0$.

We proceed by totally differentiating the free entry condition (54) with respect to θ and $\mathcal{O}$:

$$0 = \Upsilon'(\theta)(\mathcal{O} - b)\mathrm{d}\theta + \Upsilon(\theta)\mathrm{d}\mathcal{O}. \tag{84}$$

Rearranging terms and defining $\epsilon_{\Upsilon,\theta} = -\theta\Upsilon'(\theta)/\Upsilon$, which is strictly positive, yields

$$\frac{\mathrm{d}\theta}{\theta} = \frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{O}}{\mathcal{O} - b}\frac{\mathrm{d}\mathcal{O}}{\mathcal{O}}. \tag{85}$$

Also notice that the output equation (59) can be totally differentiated to yield

$$\frac{\mathrm{d}\mathcal{O}}{\mathcal{O}} = \frac{y}{\mathcal{O}}\frac{\mathrm{d}y}{y} + \frac{\mathcal{P}}{\mathcal{O}}\frac{\mathrm{d}\mathcal{P}}{\mathcal{P}}. \tag{86}$$

Combining equations (83), (85) and (86) then gives us

$$\frac{\mathrm{d}\theta}{\theta} = \frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{O}}{\mathcal{O} - b}\left[\frac{y}{\mathcal{O}}\frac{\mathrm{d}y}{y} + \frac{\mathcal{P}}{\mathcal{O}}\left(\varepsilon_{\mathcal{P},n}\frac{\mathrm{d}n}{n} + \varepsilon_{\mathcal{P},\iota}\mathrm{d}\iota\right)\right]. \tag{87}$$

Differentiating the steady-state employment expression (55) with respect to θ shows that $\varepsilon_{n,\theta} \equiv \frac{\mathrm{d}n}{\mathrm{d}\theta}\frac{\theta}{n}$ satisfies (58). Substituting (58) into (87) and rearranging terms, we obtain

$$\frac{\mathrm{d}\theta}{\theta} = \left(1 - \frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{O}}{\mathcal{O} - b}\frac{\mathcal{P}}{\mathcal{O}}\varepsilon_{\mathcal{P},n}\varepsilon_{n,\theta}\right)^{-1}\frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{O}}{\mathcal{O} - b}\left(\frac{y}{\mathcal{O}}\frac{\mathrm{d}y}{y} + \frac{\mathcal{P}}{\mathcal{O}}\varepsilon_{\mathcal{P},\iota}\mathrm{d}\iota\right). \tag{88}$$

It then follows immediately that

$$\frac{\mathrm{d}\theta}{\theta} = \varepsilon_{\theta,y}\frac{\mathrm{d}y}{y} + \varepsilon_{\theta,\iota}\mathrm{d}\iota, \tag{89}$$

with $\varepsilon_{\theta,y}$ given by equation (65) and $\varepsilon_{\theta,\iota}$ given by equation (66), as desired.

Finally, we relate the sign of the GEM feedback effect

$$\left(1 - \frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{O}}{\mathcal{O} - b}\frac{\mathcal{P}}{\mathcal{O}}\varepsilon_{\mathcal{P},n}\varepsilon_{\theta,n}\right)^{-1} \tag{90}$$

in equations (65) and (65) to the dynamic stability of steady states. For this purpose, we reconsider the free-entry condition (54). Taking into account the equilibrium relationship between $\mathcal{O}$ and n, which operates through the dependency of $\mathcal{P}$ on n, as well as the equilibrium relationship between θ and n, which operates through equation (55), we find that the derivative of the right-hand side of equation (54) with respect to n is given by

$$\Upsilon'(\theta)(\mathcal{O}-b)\frac{\theta}{n}\frac{1}{\varepsilon_{n,\theta}}+\Upsilon(\theta)\frac{\mathcal{P}}{n}\varepsilon_{\mathcal{P},n}, \tag{91}$$

which has the opposite sign of

$$1-\frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{O}}{\mathcal{O}-b}\frac{\mathcal{P}}{\mathcal{O}}\varepsilon_{\mathcal{P},n}\varepsilon_{n,\theta}. \tag{92}$$

Whenever the feedback term is positive, a small increase in employment relative to its steady-state value therefore implies a violation of the free-entry condition because the value of posting a vacancy becomes too small compared to the cost of posting a vacancy. As a result, market tightness and the job finding rate go down, so that employment gradually converges back to the steady state. This points towards the local stability (instability) of the steady state whenever the feedback effect is positive (resp. negative). In addition, due to the feedback effect, multiple steady states can exist. The steady state achieving the highest rate of employment is locally stable, as the right-hand side of the free-entry condition approaches zero when θ approaches infinity.

C.2 Decomposing the mechanism

This section provides additional details on the decomposition of the elasticities in Section 7.1. We first decompose the GEM effect,

$$\left(1-\frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{P}}{\mathcal{O}-b}\varepsilon_{\mathcal{P},n}\varepsilon_{n,\theta}\right)^{-1} \tag{93}$$

We begin by inspecting the term $\varepsilon_{\mathcal{P},n}$, which, from (62), equals the sum of a money-demand effect and a congestion effect. The money-demand effect captures the fact that higher employment raises a buyer's probability of finding a seller, thereby raising buyers' real money balances and hence sellers' profits. The congestion effects captures the fact that higher employment lowers a seller's probability of finding a buyer, thereby lowering seller's profits. Because these two effects drive profits in opposite directions, the sign of $\varepsilon_{\mathcal{P},n}$ is theoretically ambiguous.

Figure 11 shows the decomposition of $\varepsilon_{\mathcal{P},n}$ into the money-demand and congestion effects, for a range of y and ι. There are two main lessons from the figure. First, the congestion effect dominates the money-demand effect on the entire grid – even for high interest rates and low productivity, where the money-demand effect is strongest and the congestion effect is weakest. As a result, $\varepsilon_{\mathcal{P},n}$ is negative for relevant values of ι and y. Second, however, because of the two offsetting effects, the overall magnitude of $\varepsilon_{\mathcal{P},n}$ is not very large in absolute value, as will be important below.

We next observe that, in the expression (93) for the GEM effect, the elasticity $\varepsilon_{\mathcal{P},n}$ is multiplied by the positive term $\frac{1}{\epsilon_{\Upsilon,\theta}}\frac{\mathcal{P}}{\mathcal{O}-b}\varepsilon_{n,\theta}$. This term captures how changes in $\mathcal{P}$ translate, in turn, into changes in θ and subsequently n, which then feeds back into $\mathcal{P}$ through $\varepsilon_{\mathcal{P},n}$. Figure 12 displays

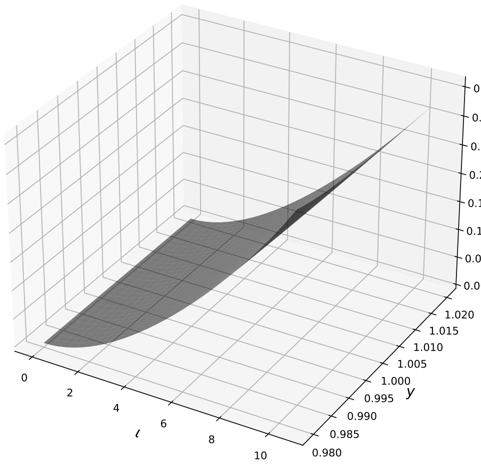

(a) The money-demand effect

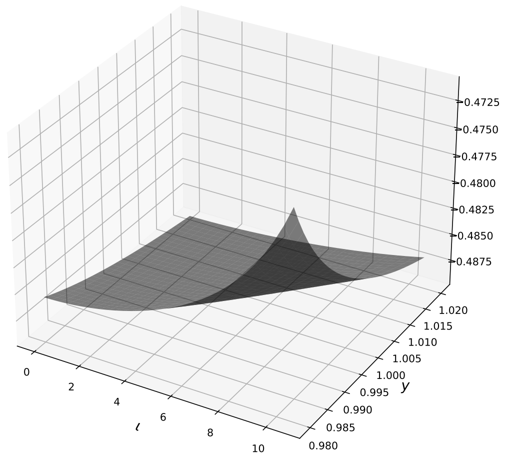

(b) The congestion effect.

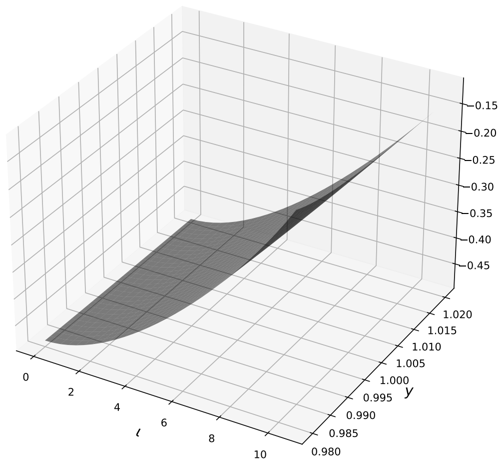

(c) The overall elasticity $\varepsilon_{\mathcal{P},n}$.

Figure 11: $\varepsilon_{\mathcal{P},n}$ and its decomposition into a money-demand effect and a congestion effect.

the individual components of this product, illustrating two results. First, the positive components, illustrated in panels 12a-12c, are amplified for large ι and small y. Though $\varepsilon_{\mathcal{P},n}$ (which is negative) becomes less negative due to a stronger money-demand effect and a weaker congestion effect, the fact that it is negative, and all the other (positive) terms tend to become stronger for large ι and small y, implies that the GEM effect becomes more dampening for large ι and small y. To put it differently, the GEM effect is composed, roughly speaking, of the product of two components: the feedback effect from employment to profits, and the feedback effect from profits back into employment. The comparative statics of the latter with respect to ι and y dominate the comparative statics of the former. Second (perhaps more importantly), because each individual component of Figure 12 is modest in absolute value, the product of the individual components is small, and hence the GEM effect is not too consequential for the model's behavior in response to changes in ι or y. Third, we note that for small y, the GEM effect is non-monotone in ι. This is purely driven by the fact that for $y < b$ ($y > b$), $\frac{\mathcal{P}}{\mathcal{O}-b}$ is decreasing (resp. increasing) in $\mathcal{P}$, with in turn $\mathcal{P}$ decreasing in ι. This is illustrated in panel 12c.

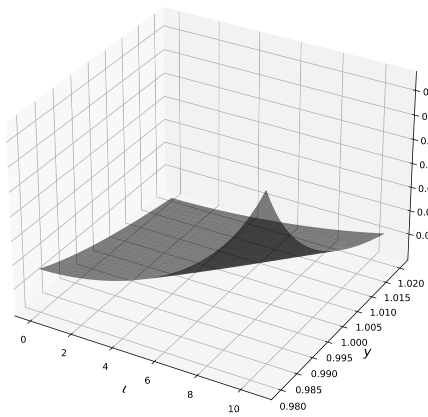

(a) $\varepsilon_{n,\theta}$

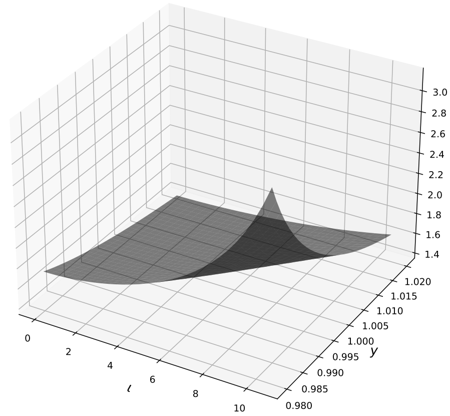

(b) $\frac{1}{\varepsilon_{\Upsilon,\theta}}$

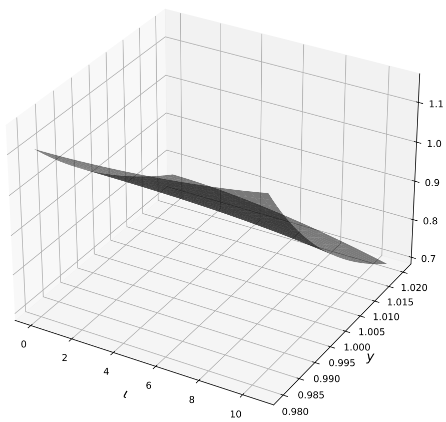

(c) $\frac{\mathcal{P}}{\mathcal{O}-b}$

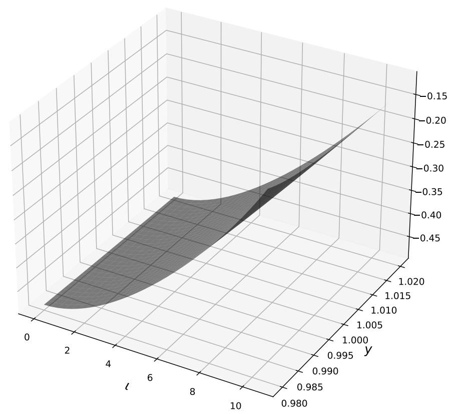

(d) $\varepsilon_{\mathcal{P},n}$

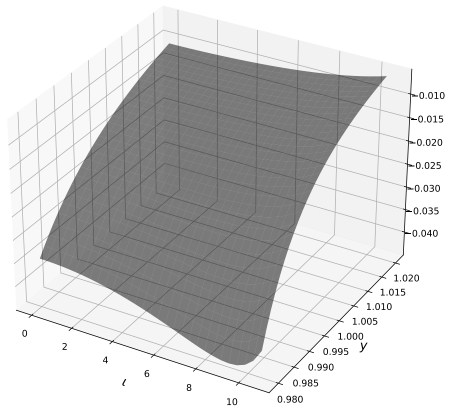

(e) $\frac{1}{\varepsilon_{\Upsilon,\theta}} \frac{\mathcal{P}}{\mathcal{O}-b} \varepsilon_{\mathcal{P},n} \varepsilon_{n,\theta}$

Figure 12: Decomposition of the general-equilibrium multiplier.

We next decompose the term $\text{direct}_{\mathcal{O},\iota}$, the direct effect of ι on $\mathcal{O}$. This is given by the product $\varepsilon_{\mathcal{P},\iota}\frac{\mathcal{P}}{\mathcal{O}}$. Figure 13 shows the decomposition of this product into its two individual components. The elasticity $\varepsilon_{\mathcal{P},\iota}$ is negative as long as $\iota > 0$, and its absolute value increases as ι increases. On the other hand, $\frac{\mathcal{P}}{\mathcal{O}}$ (positive) is small when ι is large, as a higher ι reduces the share of DM trade in total trade. When combined, we see that the effect of ι on $\varepsilon_{\mathcal{P},\iota}$ is the dominant force. This occurs because the quantity $\frac{\mathcal{P}}{\mathcal{O}}$ is small to begin with, and changes in this quantity do not significantly drive the changes in the product $\varepsilon_{\mathcal{P},\iota}\frac{\mathcal{P}}{\mathcal{O}}$ in percentage terms.

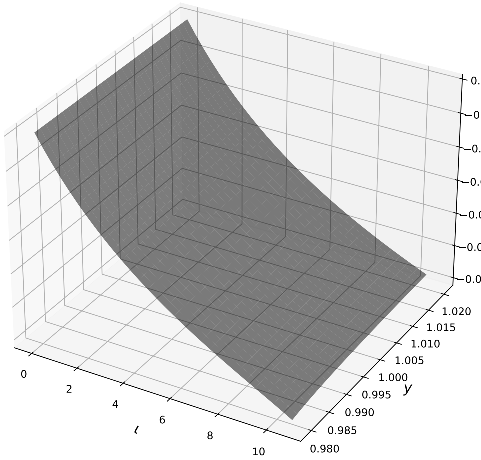

(a) $\varepsilon_{\mathcal{P},\iota}$.

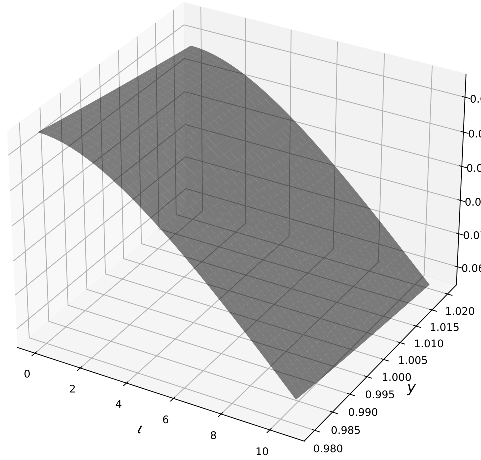

(b) $\frac{\mathcal{P}}{\mathcal{O}}$

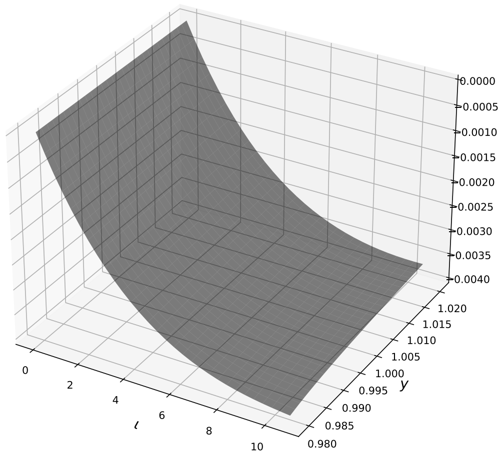

(c) $\varepsilon_{\mathcal{P},\iota}\frac{\mathcal{P}}{\mathcal{O}}$.

Figure 13: $\varepsilon_{\mathcal{P},\iota}\frac{\mathcal{P}}{\mathcal{O}}$ and its decomposition.

Appendix D Decomposing the welfare cost of inflation

This section provides the details on decomposing the welfare cost of inflation, as discussed at the end of Section 9. To better understand the contribution of frictions in different markets to the welfare cost of inflation, we attempt to disentangle two components: the direct effect of the inflation tax on DM consumption, and the congestion externalities in labor and goods markets.

We conduct the following analysis restricting attention to steady states, and can therefore refer to a single nominal interest rate. For a given nominal interest rate ι, we can compute the equilibrium allocation of market tightness θ_t, employment n_t, and DM consumption, x_t, yielding the corresponding level of welfare $\mathcal{W}(\iota)$. Next, consider a fictitious social planner's problem of choosing θ_t and n_t while taking the path of x_t as given. The solution to such a fictitious problem answers the question: what would be the efficient path of employment, *given* the anticipated amount of DM trade? Such a path of θ_t, n_t maximizes

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\alpha(n_t)(\mathrm{u}(x_t) - \mathrm{c}(x_t)) + n_t y_t + (1 - n_t)b - k\theta_{t+1}(1 - n_t)\right] \tag{94}$$

subject to

$$n_t = (1 - \delta)n_{t-1} + f(\theta_t)(1 - n_{t-1}) \tag{95}$$

taking the sequence x_t as given. Imposing steady state, the efficient θ, n allocation given x must satisfy the system of equations:

$$n = \frac{f(\theta)}{\delta + f(\theta)}$$

$$\frac{\alpha'(n)(\mathrm{u}(x) - \mathrm{c}(x)) + (y - b)}{k} = \frac{\beta^{-1} - 1 + \delta + f(\theta) - \theta f'(\theta)}{f'(\theta)}$$

Call the resulting allocation $(\tilde{\theta}(x), \tilde{n}(x))$ and the corresponding welfare level $\widetilde{\mathcal{W}}(x)$. For any given ι we can compute the true equilibrium $x(\iota)$ and $\mathcal{W}(\iota)$, as well as $\widetilde{\mathcal{W}}(\iota) = \widetilde{\mathcal{W}}(x(\iota))$, i.e. the counterfactual level of welfare that one would obtain if x is the same as the equilibrium $x(\iota)$ but n and θ are chosen efficiently given that x. By construction, for any ι, $\mathcal{W}(\iota)$ differs from $\widetilde{\mathcal{W}}(\iota)$ only by virtue of a different θ and n. Then we can decompose the welfare cost of inflation (when moving from the Friedman rule to some $\iota > 0$) as follows:

$$\mathcal{W}(\iota) - \mathcal{W}(0) = \underbrace{\left(\widetilde{\mathcal{W}}(\iota) - \widetilde{\mathcal{W}}(0)\right)}_{\text{inflation tax effect}} + \underbrace{\left(\mathcal{W}(\iota) - \widetilde{\mathcal{W}}(\iota)\right) - \left(\mathcal{W}(0) - \widetilde{\mathcal{W}}(0)\right)}_{\text{change in congestion inefficiency}} \tag{96}$$

The first term measures the change in welfare stemming from the inflation tax distortion on x. The second term measures the change in the congestion inefficiency when the nominal interest rate changes from 0 to ι. Note that this second term can be positive, negative, or zero. For example,

Table 21: Decomposition of the non-stochastic steady state consumption-equivalent welfare cost of inflation in the baseline economy

Annual inflation rate	**Implied interest rate** ι	**Welfare cost,** $(1-\Delta(\iota))\%$	Inflation tax effect	Change in congestion inefficiency
-2.75%	0.00%	-	-	-
0.00%	2.82%	0.43%	0.52%	-0.13pp
2.50%	5.40%	1.31%	1.59%	-0.34pp
5.00%	7.97%	2.43%	2.92%	-0.57pp
7.50%	10.54%	3.65%	4.35%	-0.80pp
10.00%	13.11%	4.88%	5.78%	-1.01pp

if this term is zero, this would indicate that, while there may be a congestion inefficiency, its magnitude remains the same as inflation changes. If the term is positive, this would indicate that a higher inflation mitigates the congestion inefficiency.

Figure 14 presents the net welfare cost and its decomposition. As is apparent from the Figure, the change in congestion inefficiency indeed *attenuates* the welfare cost of inflation, because inflation reduces excessive entry by firms. We also compute, similarly to our baseline welfare cost calculation, the consumption-equivalent for each component in the same manner as described in equations (73), (74), and (75). We thus solve for the consumption-equivalent compensation terms corresponding to the inflation tax and the congestion inefficiency terms in (96). Table 21 presents the decomposition of the consumption-equivalent welfare cost of inflation at the non-stochastic steady state for various levels of ι. In consumption-equivalent terms, the "inflation tax" cost of changing trend inflation from the Friedman rule to 10% is about 5.78%. However, this cost is attenuated by the welfare gain from lowering congestion inefficiencies, as inflation operates as a tax on firms' entry. In particular, the welfare cost resulting from congestion inefficiencies at the Friedman rule is about 1.31%. Moving to a 10% annual trend inflation reduces this inefficiency by about 1.01pp. This explains why the net welfare cost of moving to 10% inflation is only 4.88%. Figure 15 presents the net welfare cost and its decomposition in consumption-equivalent terms, likewise confirming the result in Figure 14.

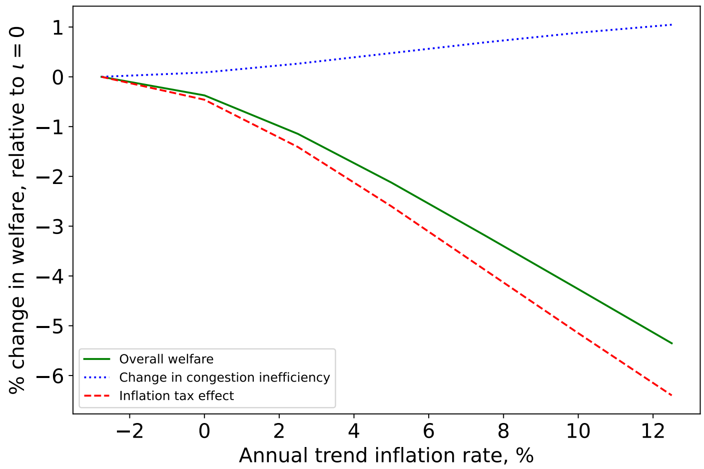

Figure 14: Decomposition of the welfare cost of inflation.

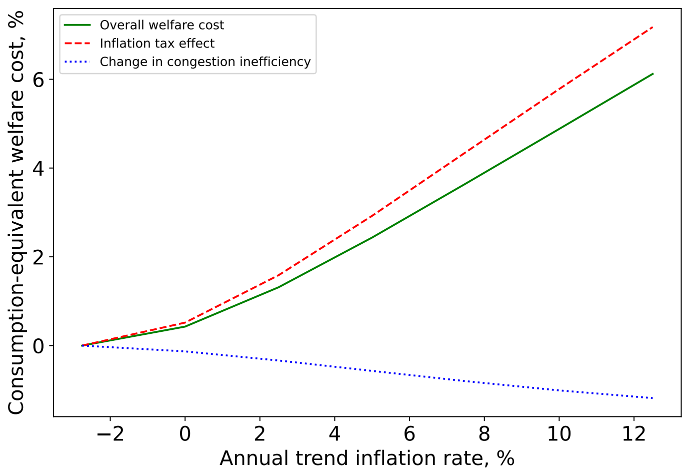

Figure 15: Decomposition of the consumption-equivalent welfare cost of inflation.