Discussion of Beaudry, Hou, Portier: The Dominant Role of Expectations and Broad-Based Supply Shocks in Driving Inflation

Gabriel Chodorow-Reich*
Harvard University

June 2024

*Contact: chodorowreich@fas.harvard.edu.

Beaudry, Hou, and Portier (hereafter BHP) have written an excellent and provocative paper. Rather than attempt to comment on each element of the analysis, I will focus my discussion on two main claims. First, the paper makes an empirical case that a linear and flat Phillips curve can explain inflation during 2021-2023. Second, it provides a theory through which broad-based supply shocks change inflation expectations and hence inflation itself and fits this theory to the U.S. experience.

I can state my bottom line succinctly up front. The paper makes a strong case that supply shocks can drive inflation expectations. I remain less convinced of the claim in the paper's title that these forces are dominant in driving inflation and by implication that tight labor markets play a subordinate role.

I structure my comment in the order of the paper. Section 1 reviews the New Keynesian Phillips Curve (NKPC). The theory motivates a few changes to the empirical exercise assessing non-linearity in the Phillips Curve slope, which I implement in section 2. While my perspective shares ingredients with existing work (Benigno and Eggertsson, 2023; Gagliardone and Gertler, 2023), this part of the comment contains some novel theoretical and empirical arguments. I conclude that the U.S. time series are perfectly consistent with a non-linear Phillips curve in which only very tight labor markets cause inflation, although I acknowledge that the time series lack the variation to make a stronger claim. Section 3 illustrates the main mechanics of the BHP model using a simplified version and offers some comments.

1 Review of NKPC

I start with a review of the textbook NKPC. The key decision concerns the reset price of a firm changing its price at date t and that expects to change its price with probability $1 - \alpha$ each period thereafter. This firm sets its price to minimize deviations from the expected desired markup over marginal cost over the course of the price spell. Letting $p_{t|t}$ denote the reset price, μ^*_{t+h} the desired markup at date $t+h$, p_{t+h} the nominal price level, and $c_{t+h|t}$ the firm's marginal cost all in log deviation from their steady state values, and letting β denote the firm's discount factor, this equation takes the form:

$$0 = \mathbb{E}_t \sum_{h=0}^{\infty} (\alpha\beta)^h \left[\underbrace{p_{t|t}}_{\substack{\text{Reset}\\ \text{price}}} - \big(\underbrace{\mu^*_{t+h}}_{\substack{\text{Desired}\\ \text{markup}}} + \underbrace{p_{t+h} + c_{t+h|t}}_{\substack{\text{Nominal}\\ \text{marginal cost}}} \big) \right]. \tag{1}$$

Equation (1) can be solved and written to express the period t optimal reset price as a weighted average of the price set in period t under fully flexible prices and the expected optimal reset price in $t+1$:

$$p_{t|t} = (1-\omega)\left(\mu^*_t + p_t + c_{t|t}\right) + \omega \mathbb{E}_t p_{t+1|t+1}, \tag{2}$$

where $\omega \equiv \frac{\alpha\beta + \alpha\beta\chi}{1+\alpha\beta\chi}$ and $\chi > 0$ denotes the (opposite of the) elasticity of a firm's relative marginal cost to its relative price. Using the definition of inflation and its decomposition into the extensive and intensive margin $\pi_t = p_t - p_{t-1} = (1-\alpha)\left(p_{t|t} - p_{t-1}\right)$, one can manipulate the reset price equation to arrive at the marginal cost Phillips curve:

$$\pi_t = \lambda\left(\mu^*_t + c_t\right) + \beta \mathbb{E}_t \pi_{t+1}, \tag{3}$$

where $\lambda \equiv \left(\frac{1-\alpha}{\alpha}\right)\left(\frac{1-\alpha\beta}{1+\chi}\right)$.

These steps illustrate three key points for measurement. First, expectations enter directly through price setters considering what price they would like in the future. Higher expected inflation implies a higher future desired price, which implies a higher price today and hence more inflation today. Importantly, price setters' expectations of their future costs and desired markups matter.[1] Second, a theory of price setters implies the appropri-

[1] This statement holds true even if expectations of other agents influence future costs, as discussed in Reis (2023). For example, consumer demand may depend on consumers' inflation expectations through the Euler equation. However, consumer demand affects prices through either changing marginal costs or desired markups. Likewise, workers' inflation expectations may affect wage demands, but these influence prices

ate measure of inflation is the GDP price index, not the CPI. Third, the theory predicts a linear relationship between inflation and log marginal cost. Estimations of the Phillips curve instead typically relate inflation to the output gap. This switch results from an unheralded divine coincidence between the central bank's dual mandate of inflation and output and what researchers can more easily measure in the data, since marginal cost typically is not directly observed.[2] Of course, this switch requires determining the relationship between aggregate output and marginal cost, which may introduce non-linearity.

2 Evidence of Non-linearity

I now make a simple case in the spirit of the BHP exercise for a steepening Phillips curve during 2021-2023. My analysis incorporates three theory-motivated changes to the data. First, I use the GDP price index to measure inflation. Second, I use professional forecasters' inflation expectations in place of the Michigan Survey. Third, I use either the vacancy-unemployment ratio or the expected vacancy duration as the measure of labor market tightness, rather than the unemployment gap. I have already discussed the rationale for using the GDP price index and this change matters relatively little to the results. I now motivate the alternative measures of expectations and the output gap before presenting the results.

2.1 Inflation Expectations

According to equation (2), the NKPC wants firms' expectations of their desired price growth. Coibion and Gorodnichenko (2015) argue that household expectations from the Michigan survey best fit a (linear) Phillips curve and suggest that in the absence of direct

through firms' costs.

[2]See Gagliardone, Gertler, et al. (2023) for a recent attempt at estimation of of a marginal cost Phillips Curve. Of course, potential output and hence the output gap also is not directly observed.

Figure 1: Measures of Expected Inflation

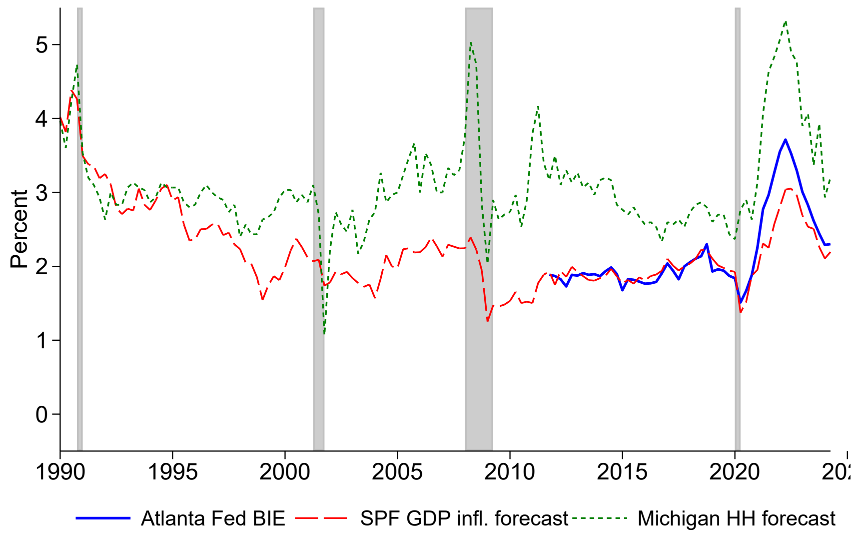

measurement of firm expectations there is no reason not to think they coincide with expectations of households. BHP follow this line. Figure 1 shows the Michigan expectations in green.

Following publication of Coibion and Gorodnichenko (2015), new measures of firm expectations have come into being. The blue line in figure 1 plots the Federal Reserve Bank of Atlanta firm expectations of their own cost growth, referred to as Business Inflation Expectations (BIE). The BIE arguably comes closest among existing survey measures to the proper theoretical construct, since what enters into the firm's pricing decision is its expectation of its own desired future price.[3] Meyer and Sheng (2021) offer an extensive empirical exploration of the BIE.

The main drawback of the BIE is that it starts only in 2011. However, the red line in Figure 1 shows that the median forecast of GDP price index inflation from the Survey of

[3]Of course, linking the desired future price to expected cost growth imposes a constant markup. Surprisingly, I am not aware of any survey that circumvents this assumption by directly asking firms about what price they would like to have at future dates. The Federal Reserve Bank of Cleveland reports an expectation of CPI inflation from a panel of CEOs, but this variable contains many differences from the firm's own desired price growth.

Professional Forecasters (SPF) closely tracks the BIE. Importantly, both measures differ notably from the Michigan survey; during the 20-teens both the BIE and SPF have roughly rational forecasts of inflation of 2%, while the Michigan Survey expectations are always too high and excessively volatile. In what follows I equate expected inflation in the Phillips Curve with the SPF forecast to obtain a longer sample. Given the short time series of the BIE, I understand that reasonable people may disagree with this choice, although I do not think that reasonable people can disagree that it is a plausible measure of price setters' expectations.

2.2 Relation of Marginal Cost to Labor Market Tightness

Figure 2 shows why the choice of the measure of the output gap might matter. The figure shows three series: the gap between the unemployment rate u and the CBO estimate of NAIRU u^*, the vacancy-unemployment ratio $\theta = V/U$ (setting V equal to JOLTS job openings and extended backward by Barnichon (2010)), and the Davis, Faberman, and Haltiwanger (2013) measure of expected vacancy duration (extended forward by me). The gap u^*/u, here shown in logs, is the paper's preferred measure. Since u is typically above u^*, $\log(u^*/u)$ is typically negative. One can immediately see why this measure does not generate non-linearity in the Phillips curve; given the downward trend in the NAIRU, it shows a tighter labor market in 2000 than in 2019 and about the same tightness in 2019 and 2022. Furthermore, since purely frictional unemployment appears to bound u from below at around 3.5%, the unemployment gap becomes less informative in very tight labor markets.

In contrast, both θ and vacancy duration are historically high in 2021-2022. The high sensitivity of θ and vacancy duration in tight labor markets reflects their unboundedness from above. I next argue that this feature also may capture an important economic force

Figure 2: Measures of the Business Cycle

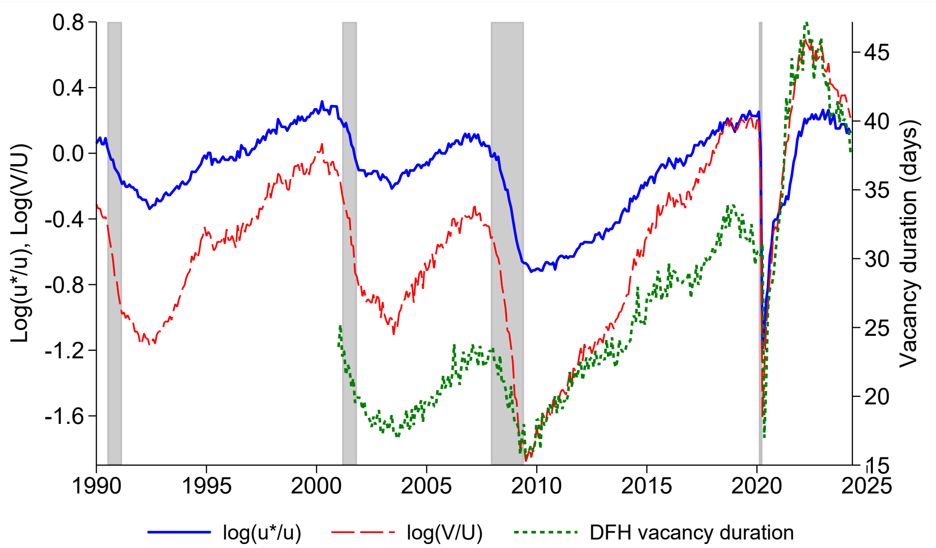

— a high level of vacancies signals that firms face capacity constraints.

In the remainder of this sub-section I provide a simple theory that puts θ into the Phillips curve and shows why non-linearity might arise. Broadly, there are two (non-mutually exclusive) ways to link θ to inflation, either through wages (Blanchard and Bernanke, 2023; Benigno and Eggertsson, 2023; Moscarini and Postel-Vinay, 2023) or through recruiting costs (Gagliardone and Gertler, 2023). I adopt the latter formulation. Relative to Gagliardone and Gertler (2023), by adopting a more general matching function and not linearizing, I illustrate two sources of non-linearity in the relationship between θ and marginal cost. This contribution is new.

The model consists of workers, wholesale firms, and retail firms. Competitive wholesalers transform labor into output with a linear technology:

$$Y_{i,t} = L_{i,t}, \tag{4}$$

where $Y_{i,t}$ and $L_{i,t}$ are output and labor input of wholesaler i at date t. Wholesalers take the price of their output C_t as given. They must hire each period, so that:

$$L_{i,t} = \bar{L}_i + H_{i,t}, \tag{5}$$

where $\bar{L}_i$ is the firm's fixed stock of labor, $H_{i,t}$ is the number of new hires, and I assume that $\bar{L}_i$ is small enough that firms always hire.[4] The number of hires equals the product of the equilibrium job-filling rate, $q(\theta_t)$, and the number of vacancies posted by the firm, $V_{i,t}$:

$$H_{i,t} = q(\theta_t) V_{i,t}, \tag{6}$$

where $\theta_t = V_t/U_t$, V_t is the aggregate quantity of vacancies, U_t the aggregate quantity of unemployed job-seekers, and $q(\theta_t)$ follows from the Haan, Ramey, and Watson (2000) matching function:

$$q(\theta_t) = \frac{1}{\left(\eta + (1-\eta)\,\theta_t^{\rho}\right)^{1/\rho}}, \tag{7}$$

where $\eta \subseteq [0,1]$ and $\rho \geq 0$ are parameters of the matching function.[5] Firms pay a flow cost k per vacancy and pay all workers a wage W_t.

The wholesaler's marginal cost, denoted $C_{i,t}$, equals the sum of the recruiting cost and the wage:

$$C_{i,t} = \frac{k}{q(\theta_t)} + W_t. \tag{8}$$

In the competitive equilibrium, the wholesaler's marginal cost equals the output price of the wholesaler sector, $C_{i,t} = C_t$. That is, taking the output price C_t as given, wholesalers

[4] This formulation simplifies the exposition by making the wholesaler's marginal cost a static function of its hiring. More complicated laws of motion are possible.

[5] If $\rho = 0$, the matching function is Cobb-Douglas and the vacancy-filling rate is $q(\theta_t) = \theta^{\eta-1}$.

hire until marginal cost equals C_t. In log deviation form, the wholesale price then satisfies:

$$c_t \equiv \ln\left(\frac{C_t}{C^*}\right) = -R(\theta_t)\ln\left(\frac{q(\theta_t)}{q(\theta^*)}\right) = R(\theta_t)\,h(\theta_t)\ln\left(\frac{\theta_t}{\theta^*}\right), \tag{9}$$

$$\text{where: } R(\theta_t) = \frac{k}{q(\theta_t)\,C_t}, \tag{10}$$

$$h(\theta_t) = \left(\frac{(1-\eta)\,\theta_t^{\rho}}{\eta + (1-\eta)\,\theta_t^{\rho}}\right), \tag{11}$$

and, importantly, the share of recruiting costs in total marginal cost $R(\theta_t)$ satisfies $R'(\theta_t) > 0$ and the inverse elasticity of the job-filling rate to tightness $h(\theta_t)$ satisfies $h'(\theta_t) > 0$ as long as $\rho > 0$.[6]

Retailers transform one unit of the wholesale good into one unit of a differentiated final good and reset their price with Calvo probability $1-\alpha$. Thus, the retail price Phillips Curve is given by equation (3) with a constant markup and marginal cost given by equation (9). One can close the model by positing a process for nominal demand.

To summarize, in this model the correct output gap measure in the Phillips curve is $\ln\theta_t$ because marginal cost depends on the hiring cost. Because the Beveridge curve is non-linear (unemployment is bounded from below), in general equilibrium $\ln\theta_t$ is a non-linear function of u and a Phillips curve in unemployment gap space inherits this non-linearity. Furthermore, the slope of the Phillips curve in θ space is non-linear because as tightness increases recruiting costs become a larger share of marginal cost and the elasticity of the hiring rate to θ_t diminishes.

[6]The condition $\rho > 0$ corresponds to gross complementarity between vacancies and searchers in the matching function. This property seems intuitively appealing. For example, consider adding either 1% more vacancies or 1% more searchers. With $\rho > 0$, the relative change in matches from 1% more vacancies diminishes as θ increases. Empirical evidence of ρ remains scarce. Haan, Ramey, and Watson (2000) internally calibrate a value of 1.27, while Lange and Papageorgiou (2020) find evidence consistent with $\rho < 0$.

Figure 3: Phillips Curve Non-Linearity

(a) Against $\ln \theta$ (b) Against $-\ln q(\theta)$

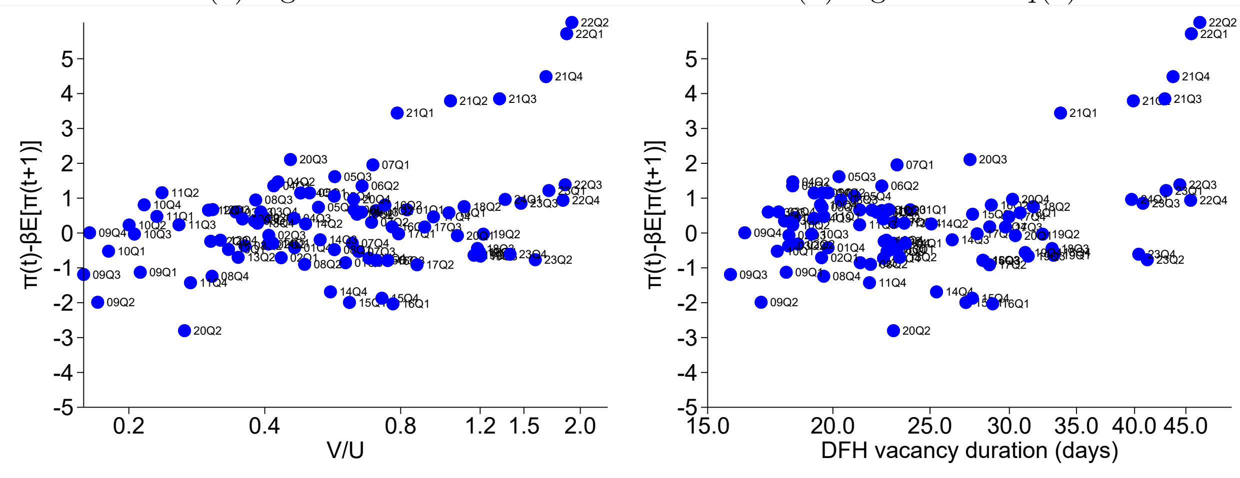

2.3 Evidence of Non-Linearity

I now undergo an exercise in the spirit of BHP figure 3 but with the three changes just discussed: (i) I use the GDP price index to measure inflation; (ii) I use the SPF to measure inflation expectations, and (iii) I use $\log \theta$ or log vacancy duration instead of the unemployment rate gap. I make one final change required by changing the output gap measure: instead of calibrating the slope of the Phillips curve, I impose a coefficient of 0.99 on expected inflation and plot $\pi_t - 0.99\mathbb{E}_t\pi_{t+1}$ against the output gap measure.

Panel (a) of Figure 3 shows the results with $\ln \theta$ as the output gap measure. Panel (b) shows the results using the Davis, Faberman, and Haltiwanger (2013) vacancy duration, which is the inverse of the job-filling rate and hence in logs equal to $-\ln q(\theta)$. The period from 2021 to 2022 stands out in both panels as having very high tightness and high excess inflation. The main difference in the tightness measures occurs in 2021Q1-2021Q3, which register as even tighter using vacancy duration than θ. I suspect these measures may even understate labor market tightness during this period, as some establishments remained closed because of the difficulty of finding workers.

In the context of this theory and calibration, the confluence of very high labor market tightness in 2021-2022 and high excess inflation has two possible interpretations, since $\pi_t - 0.99\mathbb{E}_t\pi_{t+1} = R(\theta_t)\, h(\theta_t) \ln(\theta_t/\theta^*) + \epsilon_t$. Either the Phillips curve slope $R(\theta_t)\, h(\theta_t)$ has a strong non-linearity, or large positive supply shocks ϵ_t happened to coincide with the period of maximum tightness in the sample.

The aggregate JOLTS-era post-2000 data cannot on their own distinguish these two explanations. Benigno and Eggertsson (2023) point to a similar episode of high inflation during the last period of high θ in the late 1960s, doubling the effective sample size to two. Clearly, the aggregate time series data leave room for disagreement. Hence my statement at the outset that the U.S. time series are perfectly consistent with a non-linear Phillips curve in which only very tight labor markets cause inflation, without rejecting the possibility of the alternative view.

3 The BHP Model

I now turn to the BHP model of expectations formation. Section 3.1 provides a simplified version of the model. Section 3.2 offers some comments.

3.1 A Simplified Version

My simple version has two ingredients, a flat Phillips curve and an equation for expectations formation:

$$\pi_t = \beta\mathbb{E}_t[\pi_{t+1}] + \epsilon_t, \tag{12}$$

$$\mathbb{E}_t[\pi_{t+1}] = a\pi_{t-1} + b\pi_t. \tag{13}$$

Solving this system:

$$\pi_t = \frac{\beta a \pi_{t-1} + \epsilon_t}{1 - \beta b}, \tag{14}$$

$$\mathbb{E}_t[\pi_{t+1}] = \frac{a}{1 - \beta b}\pi_{t-1} + \frac{b}{1 - \beta b}\epsilon_t. \tag{15}$$

The impulse response of inflation to a transitory supply shock is:

$$\frac{d\pi_{t+h}}{d\epsilon_t} = \frac{1}{1 - \beta b}\left(\frac{\beta a}{1 - \beta b}\right)^h. \tag{16}$$

Equations (14) to (16) contain the key amplification and persistence present in the full BHP model. The dependence of expected inflation on current inflation generates amplification of supply shocks ϵ_t through higher expected inflation. The backward-looking component of the expectations process, $a > 0$, generates persistent inflation in response to a transitory supply shock.

3.2 Comments

My "1960s version" of the BHP model yields the key result that transitory supply shocks can have amplified and persistent effects on overall inflation by raising inflation expectations. In fact, setting $b = \tilde{\rho}K$, the impact effect of inflation coincides with the impact in the BHP model. In this sense, the information frictions and bounded rationality in BHP serve the purpose of microfounding the contemporaneous and backward-looking components of expectations formation.

Information frictions in price-setting have a distinguished history that includes Lucas (1972) and Mankiw and Reis (2002). In the BHP microfoundation, the main friction concerns price-setters' belief formation taking account of an imprecise signal of contemporaneous inflation and neglecting that the error in their signal is correlated with actual inflation. Broadly, this seems plausible.[7] The details have interesting implications, including that inflation expectations exhibit less sensitivity to movements in oil prices than to other items because agents understand that oil is a volatile sector subject to large supply shocks. This implication runs counter to much conventional wisdom, but BHP provide some evidence in favor of it.

[7]A minor quibble with the details. Price-setters in the BHP model incorporate into their expectations the current inflation rate, which itself depends on the actions of those same price-setters. On the one hand, such criticism runs a little cheap; we are all guilty of writing down simultaneous systems of equations to determine equilibrium without fully specifying the adjustment process (Dogra, 2024). On the other, in a model of expectations formation such details have greater bite. In his exploration of the pass-through of expected to current inflation, Werning (2022) goes so far as to date the expectation in the Phillips curve as based on only $t-1$ information.

References

Barnichon, Regis (2010). "Building a composite Help-Wanted Index". *Economics Letters* 109.3, pp. 175–178. ISSN: 0165-1765.

Benigno, Pierpaolo and Gauti B Eggertsson (Apr. 2023). *It's Baaack: The Surge in Inflation in the 2020s and the Return of the Non-Linear Phillips Curve*. Working Paper 31197. National Bureau of Economic Research.

Blanchard, Olivier J and Ben S Bernanke (June 2023). *What Caused the US Pandemic-Era Inflation?* Working Paper 31417. National Bureau of Economic Research.

Coibion, Olivier and Yuriy Gorodnichenko (Jan. 2015). "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation". *American Economic Journal: Macroeconomics* 7.1, pp. 197–232.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger (2013). "The Establishment-Level Behavior of Vacancies and Hiring". *The Quarterly Journal of Economics* 128.2, pp. 581–622.

Dogra, Keshav (2024). "Paradoxes and Problems in the Causal Interpretation ofEquilibrium Economics".

Gagliardone, Luca and Mark Gertler (May 2023). *Oil Prices, Monetary Policy and Inflation Surges*. Working Paper 31263. National Bureau of Economic Research.

Gagliardone, Luca, Mark Gertler, et al. (June 2023). *Anatomy of the Phillips Curve: Micro Evidence and Macro Implications*. Working Paper 31382. National Bureau of Economic Research.

Haan, Wouter J. den, Garey Ramey, and Joel Watson (June 2000). "Job Destruction and Propagation of Shocks". *American Economic Review* 90.3, pp. 482–498.

Lange, Fabian and Theodore Papageorgiou (Apr. 2020). *Beyond Cobb-Douglas: Flexibly Estimating Matching Functions with Unobserved Matching Efficiency*. Working Paper 26972. National Bureau of Economic Research.

Lucas, Robert E (1972). "Expectations and the neutrality of money". *Journal of Economic Theory* 4.2, pp. 103–124. ISSN: 0022-0531.

Mankiw, N. Gregory and Ricardo Reis (Nov. 2002). "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve*". *The Quarterly Journal of Economics* 117.4, pp. 1295–1328. ISSN: 0033-5533.

Meyer, Brent H. and Xuguang Simon Sheng (Mar. 2021). *Unit Cost Expectations and Uncertainty: Firms' Perspectives on Inflation*. Working Paper 2021-12b. Working Paper.

Moscarini, Giuseppe and Fabien Postel-Vinay (July 2023). *The Job Ladder: Inflation vs. Reallocation*. Working Paper 31466. National Bureau of Economic Research.

Reis, Ricardo (May 2023). “Four Mistakes in the Use of Measures of Expected Inflation”. *AEA Papers and Proceedings* 113, pp. 47–51.

Werning, Iván (July 2022). *Expectations and the Rate of Inflation*. Working Paper 30260. National Bureau of Economic Research.