

Flattening of the Phillips Curve with State-Dependent Prices and Wages

James Costain
Banco de España

Anton Nakov
ECB and CEPR

Borja Petit
CUNEF

June 2021

Abstract

We study monetary transmission in a model of state-dependent prices and wages based on “control costs”. Stickiness arises because precise choice is costly: decision-makers tolerate errors both in the timing of adjustments, and in the new level at which the price or wage is set. The model is calibrated to microdata on the size and frequency of price and wage changes. In our simulations, money shocks have less persistent real effects than in the Calvo framework; nonetheless, the model exhibits a substantial degree of non-neutrality, driven mainly by wage rigidity. State-dependent nominal stickiness implies a flatter Phillips curve as trend inflation declines, because price and wage adjustments become less frequent, making short-run inflation less reactive to shocks. Our model can explain almost half of the observed decline in the slope of the Phillips Curve since 2000.

Keywords: sticky prices, sticky wages, state-dependence, control costs

JEL Codes: E31, D81, C73

Thanks to Fernando Álvarez, Isaac Baley, Russell Cooper, Wouter den Haan, Jordi Galí, Erwan Gautier, Alok Johri, Francesco Lippi, Julián Messina, Alexander Meyer-Gohde, Ricardo Reis, Michael Reiter, Oreste Tristani, Ernesto Villanueva, and seminar participants at CEF (2016 and 2017), EEA-ESEM (2016 and 2017), DYNARE 2016, T2M (2017 and 2019), the Catalan Economic Society (2017 and 2019), the 2017 Inflation Targeting Seminar of the Banco Central do Brasil, the 2017 ESCB Monetary Economics Cluster Workshop (Bank of Spain), the 2018 Symposium of the Spanish Economic Association (SAEe), the 2018 Workshop on Theoretical and Experimental Macroeconomics (TU Berlin), the 6th Workshop in Macro, Banking, and Finance (U. Sassari), the 8th Bundesbank-CFS-ECB Workshop on Macro and Finance, the 6th Banco de España–CEMFI Research Workshop, and at De Nederlandsche Bank, Danmarks Nationalbanken, the ECB, EUI and CUNEF for helpful comments. Petit kindly acknowledges funding from Spain’s Ministerio de Economía y Competitividad (Grants ECO2016-79848-P, BES- 2017-082181 and María de Maeztu Programme for Units of Excellence in R&D, MDM-2016-0684). Views expressed here are those of the authors and do not necessarily coincide with those of the Bank of Spain, the Eurosystem, the ECB, or the CEPR. This paper was previously circulating under the title “Monetary policy implications of state-dependent prices and wages”.

1 Introduction

The nature of the nominal frictions that cause monetary policy to transmit to the real economy remains a matter of fundamental debate. Most applied work on monetary policy imposes reduced-form time-dependent frictions, such as the [Calvo \(1983\)](#) or [Taylor \(1979\)](#) frameworks. But influential papers have argued that if nominal stickiness is instead derived from rational choice—for example, from the menu cost framework, which makes the degree of stickiness “state-dependent”—then the real effects of monetary policy are negligible (see for example [Caplin and Spulber 1987](#), and [Golosov and Lucas Jr 2007](#)). This claim has motivated an extensive wave of recent research that investigates the impact of monetary policy in microfounded state-dependent models that are closely calibrated to retail price microdata. Much of this literature concludes that “selection effects” appear weak in microdata, and that macro models consistent with this micro evidence deliver relatively strong real effects of monetary policy. Hence, getting the specification of frictions right matters for understanding monetary transmission, and likewise how transmission varies with the trend inflation rate.¹

A key question regarding monetary transmission today is why the inflationary impact of increased economic activity—the slope of the Phillips curve—appears to be much lower now than it was some decades ago. But understanding the Phillips curve suggests that we should also consider the role of wage stickiness, instead of focusing on price stickiness only, as most literature on state dependence has done thus far. With this issue in view, this paper constructs and calibrates a macroeconomic model in which both price and wage adjustment are state-dependent. A natural point of departure for our work is [Erceg et al. \(2000\)](#), who study the interaction of monopolistic price- and wage-setters, both operating under Calvo frictions. Following [Erceg et al.](#), we set up the wage setters’ problem so that it closely parallels price setting, but we impose a microfounded model of state-dependent choice for both decisions. In this context, we find that wage rigidities account for most of the real effects of monetary policy shocks, consistent with the findings of [Huang and Liu \(2002\)](#) and [Christiano et al. \(2005\)](#) under the Calvo framework. Thus, enhancing a model of state-dependent prices by allowing for state-dependent wage rigidity too substantially strengthens the predicted real effects of monetary policy.

As the name indicates, “menu cost” models were originally motivated by the fixed costs associated with posting new prices. Given the greater complexity of labour contracting, compared with retail price posting, the menu cost paradigm has been applied less often to studies of wage dynamics, where time-dependent frameworks have been more common (*e.g.* [Taylor 1979](#), [Gertler and Trigari 2009](#)). Our model of nominal adjustment departs from both these alternatives, instead adopting the “control cost” approach from game theory, which is motivated by the costs

1. Surveys include [Klenow and Kryvtsov \(2008\)](#), [Nakamura and Steinsson \(2008\)](#), and [Klenow and Malin \(2010\)](#). For models with weaker selection effects resulting in greater non-neutrality, see [Midrigan \(2011\)](#); [Eichenbaum et al. \(2011\)](#); [Kehoe and Midrigan \(2015\)](#); [Dotsey and Wolman \(2020\)](#); [Alvarez et al. \(2015\)](#); [Alvarez and Lippi \(2020\)](#); [Costain and Nakov \(2011, 2019\)](#). On trend inflation, see [Gagnon \(2009\)](#) and [Alvarez et al. \(2019\)](#).

of making accurate decisions.² Firms (workers) in our model devote resources to choosing when to reset the price (wage), and which new value to set; both types of agents may make errors in their decisions, and a similar modelling structure can be applied to both.³ We know of only one previous study of state-dependent prices and wages in a DSGE model (Takahashi 2017). That paper differs from ours in that it abstracts from idiosyncratic shocks; hence it cannot be closely assessed relative to microdata on price and wage changes. Here, we calibrate our simulations to match the distribution of individual wage adjustments, as documented by the International Wage Flexibility Project (Dickens et al. 2007).

Our approach, motivated by decision costs, has the advantage that it endogenously delivers a combination of state-dependent and time-dependent features. On one hand, the cost of accurately choosing the price (wage) acts somewhat like a menu cost, since it is often better to simply leave the price or wage unchanged, so this cost component gives rise to state dependence. Errors in the chosen price or wage help our model fit the observed distribution of adjustments in microdata. Fitting these distributions is relevant since state-dependent models are often calibrated to match higher moments; in particular, some models make predictions about skewness (Luo and Villar 2021) or kurtosis (Alvarez et al. 2016). On the other hand, the costs associated with accurately choosing the right moment to adjust inject a flavor of time dependence; these costs reduce the selection effect because they cause the probability of adjustment to rise smoothly with the value of adjustment, instead of jumping from zero to one at the “S,s band”. By weakening the selection effect, this cost component increases the real impact of monetary policy, improving the model’s fit to macrodata.

Since our framework abstracts from any frictions in labour mobility, it does not directly address issues studied by the search and matching literature. However, since it endogenizes the degree of real versus nominal impact from monetary stimulus, it can be usefully applied to other macro-labour questions such as the dynamics of real wages, or the causes of the vanishing slope of the Phillips curve. Akerlof et al. (1996), Fahr and Smets (2010), Benigno and Ricci (2011), and Lindé and Trabandt (2018) have argued that asymmetric frictions, such as downward nominal wage rigidity, can make the Phillips curve flatter when inflation is low.⁴ However, the cause of the asymmetry in adjustment frictions that drives these results is left unexplained in these papers. In contrast, we show that asymmetric frictions are in fact not necessary to obtain a decrease in the slope of the Phillips curve as trend inflation falls. We argue that one plausible explanation for the recent flattening of the curve is the fact that *trend* inflation has declined, driving down the frequencies of adjustment of prices and wages, thus making *short-run* inflation

2. Control costs are a modelling device from game theory intended to capture the idea that the costs of precise decision-making sometimes lead players to make some mistakes. See Stahl (1990), Mattsson and Weibull (2002), or Van Damme (2002), Ch. 4.

3. This extends our earlier work (Costain and Nakov 2019, henceforth CN19), which applied the control cost framework to retail prices.

4. In a Calvo model without idiosyncratic shocks, Ascari (2004) shows analytically that *higher* trend inflation implies a flatter Phillips curve. This is because firms put more weight on future relative to current marginal costs, compared to the case with zero inflation.

less responsive to shocks.⁵ We evaluate this hypothesis using our calibrated model, and contrast our results with US data. Our model explains almost half of the observed flattening of the Phillips curve since 2000, as a consequence of the decline in trend inflation.

Several studies have documented changes in inflation persistence and in the size of the inflation-unemployment trade-off. For example, [King and Watson \(1994\)](#) found inflation to be close to $I(1)$ in the 1970s, whereas more recent studies such as [Ball and Mazumder \(2011\)](#) and [Blanchard \(2016\)](#) have found inflation to be stationary and the Phillips curve to be substantially flatter since 1990. [De Veirman \(2009\)](#) makes the point that the Phillips curve is non-linear due to endogenous price duration and therefore its slope depends on trend inflation. His approach however is not explicitly derived from micro-foundations. Our paper is more similar in spirit to [Alvarez and Lippi \(2020\)](#) who show how trend inflation affects the cumulative output response of a money shock in a menu cost model (see their Figure 7). [Ascari and Haber \(2021\)](#) offer empirical evidence that trend inflation affects the inflation-output tradeoff conditional on a monetary policy shock. Other plausible mechanisms that result in a flattening of the Phillips curve include better anchoring of inflation expectations ([Jorgensen and Lansing 2019](#); [Barnichon and Mesters 2021](#)) improved monetary policy ([Roberts 2006](#); [McLeay and Tenreyro 2020](#)); or increased competitive pressure due to structural change in industry ([Lombardi et al. 2020](#); [Basso and Rachedi 2021](#); [Andrés et al. 2021](#)) or globalization ([Forbes 2019](#); [Heise et al. 2020](#)).

2 Model

We embed the near-rational nominal adjustment model of [Costain and Nakov \(2019\)](#) in a discrete-time New Keynesian general equilibrium framework that combines elements of [Erceg et al. \(2000\)](#) and of [Golosov and Lucas Jr \(2007\)](#). There is a continuum of heterogeneous retail firms and a continuum of heterogeneous workers; retail goods markets and labour markets are both monopolistically competitive. Each firm is the unique seller of a differentiated retail good, and resets its nominal price intermittently. Each worker is the unique seller of a differentiated type of labour, and resets her nominal wage intermittently. Price and wage adjustments are driven by idiosyncratic as well as aggregate shocks. Workers belong to a representative household; the budget constraint is defined at the household level. In addition, there is also a monetary authority that sets an exogenous growth process for the nominal money supply.

Our approach to nominal rigidities is based on the assumption that cognitive costs cause people to make mistakes in their choices. To model errors, we treat decision outcomes as random variables, and we impose a cost function on choices with the property that reducing errors requires greater expenditures on decision-making. Given our functional form assumptions,

5. While more frequent adjustment does not necessarily reduce the real impact of monetary shocks ([Caplin and Spulber 1987](#)), when selection effects are relatively weak, more frequent adjustment is associated with decreased non-neutrality. Several papers have developed analytical approaches to calculating aggregate dynamics in heterogeneous agent models with intermittent adjustments, including [Baley and Blanco \(2021\)](#) and [Alvarez and Lippi \(2020\)](#). The latter paper shows that, fixing the total frequency of adjustments, weaker selection implies larger real effects.

the probabilities of different actions take the familiar form of a weighted logit, placing greater probability on more desirable actions. These assumptions apply both to the price or wage that is chosen (affecting the distribution of adjustments) and to the *timing* of the adjustment. The latter property drives nominal rigidity: firms and workers may fail to adjust to shocks that would make a price or wage change objectively desirable, so the “selection effect” is reduced, enhancing monetary non-neutrality.

2.1 Monopolistic firms

Each firm j produces output $Y_{j,t}$ under a constant returns technology $Y_{j,t} = A_{j,t}N_{j,t}$. Efficiency units of labour, denoted $N_{j,t}$, are the only input, and $A_{j,t}$ represents an idiosyncratic productivity process that follows a time-invariant Markov process on a bounded set, $A_{j,t} \in \Gamma^A \subseteq [\underline{A}, \bar{A}]$. Productivity innovations are *iid* across firms. Thus, $A_{j,t}$ is correlated with $A_{j,t-1}$, but it is uncorrelated with other firms’ shocks. Firm j is a monopolistic competitor that sets a price $P_{j,t}$, facing the demand curve $Y_{j,t} = C_t P_t^\epsilon P_{j,t}^{-\epsilon}$, where C_t is aggregate consumption and P_t is the price index.⁶ We assume each firm must fulfill all demand at its chosen price. Since firms are infinitesimal, each firm j ignores the effect of its own price $P_{j,t}$ on the aggregate price level P_t ,

$$P_t \equiv \left\{ \int_0^1 P_{j,t}^{1-\epsilon} dj \right\}^{\frac{1}{1-\epsilon}}. \quad (1)$$

It hires labour at wage rate W_t , generating profits

$$P_{j,t}Y_{j,t} - W_tN_{j,t} = \left(P_{j,t} - \frac{W_t}{A_{j,t}} \right) C_t P_t^\epsilon P_{j,t}^{-\epsilon} \quad (2)$$

per period. Firms are owned by the household, so they discount nominal income between times t and $t+1$ at the household’s stochastic discount rate $\Lambda_{t,t+1}$, defined below.

To clarify the structure of decision-making, it helps to distinguish value functions at several different points in time. First, let $V_t(P, A)$ be the value of a firm that begins period t with nominal price P and productivity A , prior to any time- t decisions, and prior to time- t output. Figure 1 presents the firms’ timeline. We assume that choices take time, so if the firm decides in period t to adjust its price, the new price only becomes effective at time $t+1$.⁷ Next, let $O_t(P, A)$ be the firm’s continuation value, net of current profits, when it still has the option to

6. We use succinct notation where time subscripts represent dependence on the economy’s aggregate state, Ω_t . Time subscripts on aggregate variables represent functions of Ω_t : $P_t \equiv P(\Omega_t)$ is the aggregate price level, $W_t \equiv W(\Omega_t)$ is the aggregate wage, and $C_t \equiv C(\Omega_t)$ is aggregate consumption demand. Likewise, time subscripts on value and policy functions represent dependence on Ω_t : $V_t(P, A) \equiv V(P, A, \Omega_t)$, $O_t(P, A) \equiv O(P, A, \Omega_t)$, and so forth.

7. A one-period lag corresponds to the timing of the model in its continuous-time limit. In continuous time, current prices are state variables, and decisions alter the state at an infinitesimally later moment. Imposing this structure on a discrete time version of the model simplifies the numerical solution (see footnote 17). This one-period delay may slightly distort our quantitative analysis of the most flexible calibrations of our model, but since we set the time period to one month, a one-period lag is not very restrictive in practice.

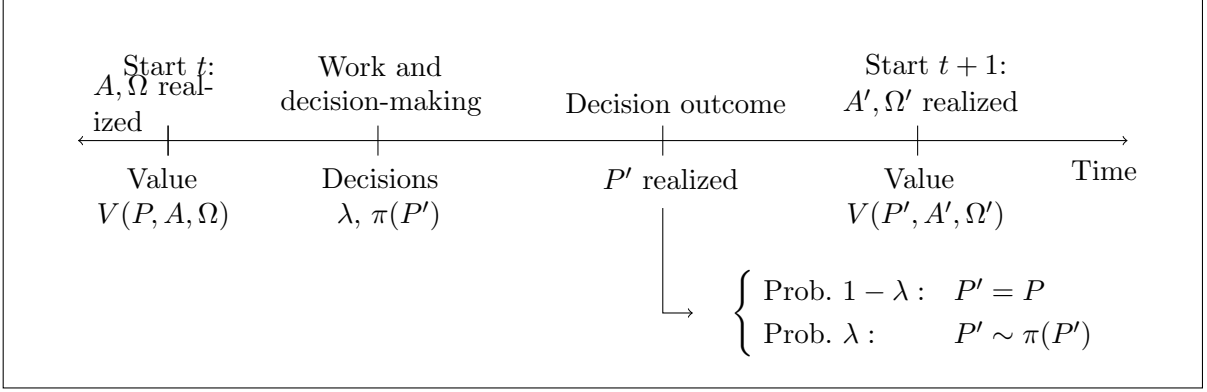


Figure 1: Firms' timeline.

adjust prices. That is,

$$V_t(P, A) = \left(P - \frac{W_t}{A} \right) C_t P_t^\epsilon P^{-\epsilon} + O_t(P, A). \quad (3)$$

The function $O_t(P, A)$ incorporates the value of the firm's two possible time- t decisions: whether to adjust its price, and if so, which new price P' to set for time $t + 1$. The firm may make errors in either of these choices. We discuss these two decisions in turn, beginning with the latter.

2.1.1 Choosing a new price

Our model formalizes the idea that nominal rigidities may derive primarily from the costs of decision-making. While one might assume that by paying a fixed cost, the firm can make the optimal choice, this would implicitly impose a corner solution with perfect precision. We find it more appealing and realistic to assume that firms can devote more or less time and resources to decision-making, thus choosing more or less precisely. In equilibrium in our framework firms will typically prefer to make choices with an interior degree of precision. Thus their chosen action will not always be the one that would have been optimal in the absence of decision costs; instead, most choices will involve some degree of "error".

Consistent with this general description, we adopt the "control cost" approach from game theory (see [Van Damme 2002](#), Chapter 4). A key feature of this approach is that we model the price decision indirectly: the firm's problem is written "as if" it chooses a probability distribution over prices, rather than choosing the price *per se*.⁸ The problem incorporates a cost function that increases with precision: concentrating greater probability on a smaller range of prices increases costs. Many measures of precision could be used to define decision costs; we choose a definition based on relative entropy, also known as Kullback-Leibler divergence, which is a measure of the difference between two probability distributions. For two possible distributions $\pi_1(x)$ and $\pi_2(x)$

8. [Luce \(1959\)](#) and [Machina \(1985\)](#) are early advocates of analyzing decisions in terms of a probability distribution over alternatives; this approach is also adopted by [Sims \(2003\)](#). See Chapter 2 of [Anderson et al. \(1992\)](#) for discussion.

of some random variable x with support on set \mathcal{X} , the Kullback-Leibler divergence $\mathcal{D}(\pi_1||\pi_2)$ of π_1 relative to π_2 is defined by⁹

$$\mathcal{D}(\pi_1||\pi_2) = \int_{\mathcal{X}} \pi_1(x) \ln \left(\frac{\pi_1(x)}{\pi_2(x)} \right) dx. \quad (4)$$

Following [Stahl \(1990\)](#) and [Mattsson and Weibull \(2002\)](#), we assume that the decision cost is proportional to the Kullback-Leibler divergence of the chosen distribution ($\pi_1(x)$ above) relative to an exogenous benchmark distribution ($\pi_2(x)$ above). Thus, if no decision costs are paid, the action x is distributed according to the benchmark distribution, $\pi_2(x)$. But by putting more effort into the decision process, the decision-maker can shrink the distribution of the action towards the most desirable alternatives.

We assume that decision costs are denominated in units of time. The only control variable that the firm must manage is its nominal price. We regard each adjustment of the nominal price as a costly decision, so when the firm sets a new nominal price \tilde{P} , this remains constant in nominal terms until the firm again chooses to adjust. We define the cost of the decision process relative to an exogenous benchmark distribution $\eta_t(\tilde{P})$ with support Γ_t^P .¹⁰ The time subscripts on η_t and Γ_t^P allow the benchmark price distribution to change over time, so the economy may have a nominal trend; later we detrend the model by restating it in real terms.

Assumption 1 *The time cost of choosing a distribution $\pi(\tilde{P})$ over nominal prices $\tilde{P} \in \Gamma_t^P$ is $\kappa_f \mathcal{D}(\pi||\eta_t)$, where $\kappa_f > 0$ is a constant, and $\eta_t(\tilde{P})$ is an exogenously-given benchmark distribution with support Γ_t^P .*

Here κ_f represents the marginal cost of entropy reduction, in units of labour time. The cost function described in Assumption 1 is nonnegative and convex.¹¹ The upper bound on the cost function is associated with a distribution that places all probability on a single price \tilde{P} (costs are maximized by placing all probability on one price that minimizes the benchmark probability $\eta_t(\tilde{P})$). The lower bound on this cost function is zero, associated with choosing the distribution $\pi(\tilde{P})$ equal to the benchmark distribution $\eta_t(\tilde{P})$.

Now consider the pricing decision under this cost function. If the firm sets a new nominal price \tilde{P} at time t , this new price only becomes effective at $t + 1$, and its value is:

$$V_t^e(\tilde{P}, A) \equiv E_t \left[\Lambda_{t,t+1} V_{t+1}(\tilde{P}, A') \mid A \right], \quad (5)$$

9. While we write (4) with an integral, we can be agnostic at this point about whether \mathcal{X} is a discrete or continuous set. If it is a continuous set, then π_1 and π_2 should be interpreted as density functions. If it is a discrete set, then π_1 and π_2 should be interpreted as vectors of probabilities, and the integral in (4) should be interpreted as a sum.

10. Our setup imposes a control cost function with an exogenous benchmark distribution. [Steiner et al. \(2017\)](#) show that a general dynamic rational inattention problem is equivalent to a control cost problem with an optimally-chosen benchmark distribution. Fixing distribution $\eta_t(\tilde{P})$ exogenously enhances the numerical tractability of our framework, but still yields a form of stickiness similar to that obtained from rational inattention.

11. See [Cover and Thomas \(2006\)](#), Theorem 2.7.2.

where $\Lambda_{t,t+1}$ is the household stochastic discount factor, and $E_t[\bullet|A]$ represents an expectation over the time $t+1$ variables $\Omega' \equiv \Omega_{t+1}$ and $A' \equiv A_{j,t+1}$ conditional on the time t aggregate state Ω_t and firm j 's time t productivity $A_{j,t} = A$. Following the control costs methodology, we assume the firm maximizes its value by allocating probability across possible nominal prices \tilde{P} , taking account of decision costs, as follows:

$$\tilde{V}_t(A) = \max_{\pi(\tilde{P})} \int \pi(\tilde{P}) V_t^e(\tilde{P}, A) d\tilde{P} - W_t \kappa_f \mathcal{D}(\pi || \eta_t) \quad (6)$$

$$\text{s.t.} \quad \int \pi(\tilde{P}) d\tilde{P} = 1. \quad (7)$$

Note that the decision costs in (6) are converted to nominal units by multiplying by the wage rate. We write the nominal value of the pricing decision as $\tilde{V}_t(A)$, where $A \equiv A_{j,t}$ is the firm's current productivity.

The first-order condition for $\pi(P)$ in problem (6) is¹²

$$V_t^e(P, A) - \kappa_f W_t \left[1 + \ln \left(\frac{\pi(P)}{\eta_t(P)} \right) \right] - \mu = 0,$$

where μ is the multiplier on the constraint (7). Some rearrangement yields a weighted multinomial logit formula:

$$\pi_t(P|A) \equiv \frac{\eta_t(P) \exp \left(\frac{V_t^e(P, A)}{\kappa_f W_t} \right)}{\int \eta_t(\tilde{P}) \exp \left(\frac{V_t^e(\tilde{P}, A)}{\kappa_f W_t} \right) d\tilde{P}}. \quad (8)$$

The parameter κ_f in the logit function can be interpreted as the degree of noise in the decision process; in the limit as $\kappa_f \rightarrow 0$, equation (8) converges to the policy function under full rationality, so that the optimal price is chosen with probability one. Plugging the logarithm of π_t into the objective, we can also derive an analytical formula for the value function:

$$\tilde{V}_t(A) = \kappa_f W_t \ln \left(\int \eta_t(\tilde{P}) \exp \left(\frac{V_t^e(\tilde{P}, A)}{\kappa_f W_t} \right) d\tilde{P} \right). \quad (9)$$

This formula gives the firm's nominal value when adjusting its current price, net of decision costs.

Some interpretive comments may be helpful at this point. First, while we write the firm's problem "as if" it chooses a probability distribution over prices, this should not be taken literally—actually choosing a distribution would be a complex, costly diversion from the true task of choosing the price itself. Rather, we model the decision as a choice of a mixed strategy because this is a way to model errors. And we write it as an optimization problem because this disciplines the errors; it means that the firm devotes time and effort to avoiding especially costly

12. Note that if we take future values $V_t^e(\tilde{P}, A)$, as given, problem (6) maximizes a concave objective subject to a linear constraint. Therefore a unique maximum exists for any given backwards induction step.

mistakes. Aspects of the model that we do take seriously include (a) making decisions is costly in terms of time and other resources; (b) therefore decision-makers do not always take the action that would otherwise be optimal; (c) *ceteris paribus*, more valuable actions are more probable; (d) in a retail pricing context, these considerations apply to the timing of price changes, as well as the actual price chosen, as we will see in the next subsection.

Second, the problem is written conditional on the true expected discounted values $V_t^e(\tilde{P}, A)$ of the possible nominal prices \tilde{P} , instead of conditioning on a prior, as a “rational inattention” model would. This reflects the fact that we are *not* assuming imperfect information. But this is different from saying that the firm “knows” the true values $V_t^e(\tilde{P}, A)$. Instead, our interpretation is that the firm has sufficient information to calculate $V_t^e(\tilde{P}, A)$. Even so, drawing correct conclusions from that information, and acting accordingly, may be costly.¹³

2.1.2 Choosing the timing of price adjustment

We next analyse, in an analogous manner, the decision whether or not to adjust at time t . As in section 2.1.1, we define costs relative to a benchmark probability distribution over possible actions. But for this choice, at any t , there are only two options: adjust now, or not. Since the probabilities of these two alternatives must sum to one, effectively the relevant benchmark is just a single number, which we can interpret as an exogenous benchmark adjustment hazard.

We suppose the time period is sufficiently short so that we can ignore multiple adjustments within a single period. If the firm chooses not to adjust its current price P , then its nominal price next period will be unchanged: $P' = P$; the expected value of this unchanged price, from the point of view of period t , is $V_t^e(P, A)$, given by (5). If instead the firm adjusts its price at t , then its expected value is $\tilde{V}_t(A)$, given by (9). Now suppose it adjusts its price with probability λ . We measure the cost of this adjustment probability in terms of Kullback-Leibler divergence, relative to some arbitrary Poisson process with arrival rate $\bar{\lambda}$:¹⁴

Assumption 2 *The time cost incurred in period t by setting the price adjustment probability $\lambda \in [0, 1]$ in period t is $\kappa_f \mathcal{D}(\lambda || \bar{\lambda})$, where $\kappa_f > 0$ and $\bar{\lambda} \in [0, 1]$ are exogenous parameters.*

The optimal adjustment probability at time t solves the following Bellman equation:

$$O_t(P, A) = \max_{\lambda \in [0, 1]} (1 - \lambda)V_t^e(P, A) + \lambda\tilde{V}_t(A) - W_t\kappa_f\mathcal{D}(\lambda || \bar{\lambda}). \quad (10)$$

Recall that $O_t(P, A)$ represents the continuation value of the firm, net of decision costs, when it

13. Since economists are accustomed to models of perfect rationality, they often equate observing a given information set with knowing all quantities that can be calculated from that information set. But when rationality is less than perfect, we cannot equate these two situations. Here, we assume firms can observe all relevant shocks and state variables, but we do not equate this with actually knowing $V_t^e(\tilde{P}, A)$ or knowing the optimal action, and therefore we do not equate it with implementing the optimal action with probability one.

14. In Assumption 2, we write the arguments of the Kullback-Leibler divergence as the adjustment probabilities λ and $\bar{\lambda}$. Under the standard notation used in eq. (4), the arguments should actually be the Bernoulli distributions $(\lambda, 1 - \lambda)$ and $(\bar{\lambda}, 1 - \bar{\lambda})$. We opt for the abbreviated notation since its meaning is clear.

still has the option to adjust, or not adjust. The first order condition from (10) is

$$V_t^e(P, A) - \tilde{V}_t(A) = \kappa_f W_t [\ln \lambda - \ln \bar{\lambda} - \ln(1 - \lambda) + \ln(1 - \bar{\lambda})]. \quad (11)$$

Rearranging, we obtain

$$\lambda_t(P, A) = \frac{\bar{\lambda}}{\bar{\lambda} + (1 - \bar{\lambda}) \exp\left(\frac{-D_t(P, A)}{\kappa_f W_t}\right)}, \quad (12)$$

where $D_t(P, A)$ is the expected gain from adjustment:

$$D_t(P, A) \equiv \tilde{V}_t(A) - V_t^e(P, A).$$

The weighted binary logit probability function (12) was also derived by Woodford (2009) from a model with a Shannon constraint. The free parameter $\bar{\lambda}$ measures the rate of decision making; concretely, the probability of adjustment in one discrete time period is $\bar{\lambda}$ when the firm is indifferent between adjusting and not adjusting (*i.e.* when $D_t(P, A) = 0$).

2.1.3 Deriving the Bellman equation

Next, to calculate the value function $V_t(P, A)$, we concatenate the two decision steps described in sections 2.1.1 and 2.1.2. If the firm starts period t with nominal price P , then its value $V_t(P, A) \equiv V_t(P, A, \Omega_t)$ at the beginning of t satisfies:

$$\begin{aligned} V_t(P, A) = \max_{\lambda, \pi(\tilde{P})} & \left(P - \frac{W_t}{A} \right) C_t P_t^\epsilon P^{-\epsilon} + (1 - \lambda) V_t^e(P, A) - W_t \kappa_f \mathcal{D}(\lambda | \bar{\lambda}) + \\ & + \lambda \left[\int \pi(\tilde{P}) V_t^e(\tilde{P}, A) d\tilde{P} - W_t \kappa_f \mathcal{D}(\pi | \eta) \right] \\ \text{s.t.} \quad & \int \pi(\tilde{P}) d\tilde{P} = 1. \end{aligned} \quad (13)$$

This Bellman equation subtracts off the two cost functions seen in the previous subsections.¹⁵ There is a time cost associated with monitoring whether or not a price adjustment is required, which we will call

$$\mu_t(P, A) \equiv \kappa_f \left[\lambda \ln \left(\frac{\lambda}{\bar{\lambda}} \right) + (1 - \lambda) \ln \left(\frac{1 - \lambda}{1 - \bar{\lambda}} \right) \right]. \quad (14)$$

15. For expositional transparency, we described pricing and timing above as two separate decisions, each with its own entropy costs. However, these two steps can equivalently be rewritten as a single decision, subject to a single entropy-based cost function, encompassing the alternatives of non-adjustment or of adjustment to any $\tilde{P} \in \Gamma_t^P$. For details, see CN19, Sec. 2.2. We will see below that the worker's problem must generally be written as a single combined decision, except in the special case of linear labour disutility.

The time cost of choosing which new price to set is

$$\tau_t(P, A) \equiv \lambda \kappa_f \int \pi(\tilde{P}) \ln \left(\frac{\pi(\tilde{P})}{\eta_t(\tilde{P})} \right) d\tilde{P}. \quad (15)$$

Finally, the time devoted to the actual production of goods will be written as

$$N_t(P, A) \equiv \frac{C_t}{A} \left(\frac{P_t}{P} \right)^\epsilon. \quad (16)$$

Hence, the firm's total demand for labour hours is

$$N_t^{tot}(P, A) = N_t(P, A) + \mu_t(P, A) + \tau_t(P, A). \quad (17)$$

2.2 Labour market

We next construct a model of nominal wage rigidity analogous to our treatment of nominal price rigidity. We suppose each worker i is the monopolistic supplier of a specific type of labour $H_{i,t}$, sold at wage $W_{i,t}$ per unit of time. The productivity of worker i 's labour $H_{i,t}$ is shifted by a shock process $Z_{i,t}$, which follows a time-invariant Markov process on a bounded set, $Z_{i,t} \in \Gamma^Z \subset [\underline{Z}, \bar{Z}]$. We will define $N_{i,t} = Z_{i,t}H_{i,t}$ as the ‘‘effective labour’’ of worker i . By this definition, we can say that the price of effective labour is $\frac{W_{i,t}}{Z_{i,t}}$. The idiosyncratic shock process $Z_{i,t}$ represents worker-specific productivity dynamics, which may include various forms of human capital accumulation.

Firm j 's labour input into goods production, $N_{j,t}$, is defined as a CES aggregate across varieties of effective labour i , with elasticity of substitution ϵ_n . That is,

$$N_{j,t} = \left\{ \int_0^1 N_{ijt}^{\frac{\epsilon_n-1}{\epsilon_n}} di \right\}^{\frac{\epsilon_n}{\epsilon_n-1}}. \quad (18)$$

We assume that firms use the same CES mix of labour for decision making that they use for goods production. This implies that the firm's optimal hiring for all purposes satisfies

$$H_{ijt} \equiv \frac{N_{ijt}}{Z_{i,t}} = Z_{i,t}^{\epsilon_n-1} \left(\frac{W_{i,t}}{W_t} \right)^{-\epsilon_n} N_{j,t}^{tot}, \quad (19)$$

where $N_{j,t}^{tot}$ is given by (17), and W_t is the following wage index:

$$W_t \equiv \left\{ \int_0^1 \left(\frac{W_{i,t}}{Z_{i,t}} \right)^{1-\epsilon_n} di \right\}^{\frac{1}{1-\epsilon_n}}. \quad (20)$$

Firm j 's nominal wage bill for all purposes is then

$$\int_0^1 W_{i,t} H_{ijt} di = W_t N_{j,t}^{tot}. \quad (21)$$

Total demand for worker i 's time is $H_{i,t} = H_t(W_{i,t}, Z_{i,t})$, defined by

$$H_{i,t} = Z_{i,t}^{\epsilon_n - 1} \left(\frac{W_{i,t}}{W_t} \right)^{-\epsilon_n} N_t^{tot} \equiv H_t(W_{i,t}, Z_{i,t}), \quad (22)$$

where N_t^{tot} is aggregate effective labour demand, given by integrating (17) across all firms.

The worker adjusts her nominal wage $W_{i,t}$ intermittently to maximize the value of labour income net of labour disutility. She faces control costs, both on her timing decision, and on the choice of which wage to set. We assume workers act in the interest of the households of which they form part, and that their consumption is fully insured by the household; hence they discount future income at the same rate $\Lambda_{t,t+1}$ that applies to the household and firms. Now let $L_t(W, Z)$ be the nominal value of a worker with wage W and productivity Z at the beginning of period t , before supplying labour, and before making any decisions. As in the case of price decisions, we assume that a wage adjustment in period t becomes effective in period $t + 1$. Therefore the value of setting the nominal wage to an arbitrary new value \widetilde{W} is

$$L_t^e(\widetilde{W}, Z) \equiv E_t \left[\Lambda_{t,t+1} L_{t+1}(\widetilde{W}, Z') \mid Z \right].$$

We make two assumptions about workers' decision costs that are analogous to our assumptions about firms.

Assumption 3 *The time cost of choosing a distribution $\pi^w(\widetilde{W})$ over nominal wages $\widetilde{W} \in \Gamma_t^w$ is $\kappa_w \mathcal{D}(\pi^w \parallel \eta_t^w)$, where $\kappa_w > 0$ is a constant, and $\eta_t^w(\widetilde{W})$ is an exogenously-given benchmark distribution with support Γ_t^w .*

Assumption 4 *The time cost incurred in period t by setting the wage adjustment probability $\rho \in [0, 1]$ in period t is $\kappa_w \mathcal{D}(\rho \parallel \bar{\rho})$, where $\kappa_w > 0$ and $\bar{\rho} \in [0, 1]$ are exogenous parameters.*

Now, let the disutility of labour be given by

$$X(H) = \chi \frac{H^{1+\zeta}}{1+\zeta}, \quad (23)$$

where H includes both the time a worker provides to the firms and the time devoted to the worker's own decisions, and $\zeta > 0$ is the inverse of the Frisch elasticity of labour supply. Adopting this convex disutility function implies that we will not be able to separate the choice of the wage from the choice of the *timing* of wage adjustment, as we did when we described the price-setting process. Nonetheless, the wage-setting problem takes a form closely analogous to the pricing

problem (13):

$$\begin{aligned}
L_t(W, Z) &= \max_{\tau^w, \mu^w, \rho, \pi^w(W)} WH_t(W, Z) - \frac{P_t}{u'(C_t)} \cdot X(H_t(W, Z) + \mu^w + \tau^w) + \\
&\quad + (1 - \rho)L_t^e(W, Z) + \rho \int \pi^w(\widetilde{W})L_t^e(\widetilde{W}, Z)d\widetilde{W} \\
\text{s.t.} \quad &\int \pi^w(\widetilde{W})d\widetilde{W} = 1, \\
\mu^w &= \kappa_w \left[\rho \ln \left(\frac{\rho}{\bar{\rho}} \right) + (1 - \rho) \ln \left(\frac{1 - \rho}{1 - \bar{\rho}} \right) \right], \\
\tau^w &= \rho \kappa_w \int \pi^w(\widetilde{W}) \ln \left(\frac{\pi^w(\widetilde{W})}{\eta_t^w(\widetilde{W})} \right) d\widetilde{W}. \tag{24}
\end{aligned}$$

Note that we scale labour disutility $X(H)$ by the factor $P_t/u'(C_t)$, to express the whole Bellman equation in nominal units. In (24), $\mu^w \equiv \kappa_w \mathcal{D}(\rho||\bar{\rho})$ represents time dedicated to monitoring whether or not it is a good moment to reset the wage. Since the probability of resetting the wage at a given time t is ρ , Assumption 4 implies that the (expected) time devoted to choosing a new wage in period t is $\tau^w \equiv \rho \kappa_w \mathcal{D}(\pi^w||\eta_t^w)$.

To clarify, recall that we stated the firm's decision in two separate steps, (10) and (6), representing the choice whether or not to adjust prices, and the choice of what price to set conditional on adjustment, respectively. This decomposition was possible because we assumed the firm could hire any quantity of labour at the (aggregate) wage rate W_t , making its labour costs linear in its labour demand. But imposing a linear cost function for a worker's time use would be highly restrictive, implying an infinite elasticity of labour supply. In the absence of decision costs, a worker would set the real wage as a function of the marginal utility of consumption only; productivity shocks would cause variation in hours worked without any change in the wage.¹⁶ To avoid this restrictive calibration, we allow for convex disutility. But therefore we cannot just condition on a given, constant marginal cost of labour; time supplied to firms affects the cost of time on each decision margin, so both margins are analysed simultaneously in problem (24).¹⁷

Nonetheless, the policy functions for wage setting and wage adjustment timing resemble the policy functions from the firm's problem. Following our previous calculations, we find that if

16. A specification with linear labour disutility is analysed in our working paper, [Costain et al. \(2019\)](#).

17. Inspecting problem (24) clarifies how a one-period decision lag simplifies the numerics. If the wage for time- t work were chosen at t , then the marginal cost of time- t choice would depend on an expectation over labour demand at each possible time- t wage. This would severely expand the dimension of the numerical calculation. This technicality does not arise in the firm's problem, where the marginal cost of time used is constant.

the worker adjusts, she chooses the following density over next period's nominal wages:¹⁸

$$\pi_t^w(W'|W, Z) \equiv \frac{\eta_t^w(W') \exp\left(\frac{L_t^e(W', Z)}{\kappa_w x_t(W, Z)}\right)}{\int \eta_t^w(\tilde{W}) \exp\left(\frac{L_t^e(\tilde{W}, Z)}{\kappa_w x_t(W, Z)}\right) d\tilde{W}}, \quad (25)$$

where $x_t(W, Z)$ denotes the marginal disutility of time in period t :

$$x_t(W, Z) \equiv \frac{P_t}{u'(C_t)} X'(H_t(W, Z) + \tau^w + \mu^w). \quad (26)$$

Likewise, if the worker's beginning-of-period wage and productivity are W and Z , her optimal adjustment probability must satisfy:

$$\rho_t(W, Z) = \frac{\bar{\rho}}{\bar{\rho} + (1 - \bar{\rho}) \exp\left(\frac{-D_t^w(W, Z)}{\kappa_w x_t(W, Z)}\right)}, \quad (27)$$

where

$$D_t^w(W, Z) \equiv \tilde{L}_t(W, Z) - L_t^e(W, Z)$$

represents the gain in value from adjusting rather than leaving the nominal wage unchanged.

The key to solving the worker's equations is to calculate the marginal disutility of time, $x_t(W, Z)$. If we know the aggregate variables P_t , W_t , C_t , and N_t , then labour demand $H_t(W, Z)$ is known from (22). In a context of backwards induction, where function $L_t^e(W, Z)$ is known, we can then solve a fixed-point problem to find $x_t(W, Z)$. By guessing $x_t(W, Z)$ at a given pair (W, Z) , we can construct the probabilities and the probability rate from (25) and (27), and then calculate the decision time costs $\tau_t^w(W, Z)$ and $\mu_t^w(W, Z)$ from the constraints on (24). Summing $H_t(W, Z) + \tau_t^w(W, Z) + \mu_t^w(W, Z)$, we can then update $x_t(W, Z)$ using (26).¹⁹

2.3 Household

The household consists of a continuum of heterogeneous workers of unit mass, who aggregate their resources to choose household consumption C_t , and bond and money holdings B_t and M_t . Utility is discounted by factor $\beta \equiv \beta_I \beta_S$ per period, where β_I represents the effect of pure impatience, and β_S is the probability of survival (each individual worker dies and is replaced by a new individual with probability $1 - \beta_S$ per period). Household consumption C_t is a CES

18. There is also an analytical formula for the worker's value function, analogous to (9); see the working paper Costain et al. (2019).

19. While problem (24) has two separate entropy cost terms, actually these two can be rewritten as a single entropy cost term (see CN19, Sec. 2.2, for discussion). Since labour disutility is convex, a unique well-defined solution exists for the maximization problem at any backwards induction step. Hence we conclude that the algorithm described here to calculate $x_t(W, Z)$ has a unique fixed point, which characterizes the marginal value of time in problem (24).

aggregate of differentiated products $C_{j,t}$:

$$C_t = \left\{ \int_0^1 C_{j,t}^{\frac{\epsilon-1}{\epsilon}} dj \right\}^{\frac{\epsilon}{\epsilon-1}}, \quad (28)$$

where ϵ is the elasticity of substitution across varieties.

Besides the wage setting decisions already discussed, the household must choose $\{C_{j,t}, B_t, M_t\}_{t=0}^{\infty}$ to maximize expected discounted utility:

$$\mathbb{E}_t \left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} \left(\frac{C_{\tau}^{1-\gamma} - 1}{1-\gamma} - \int_0^1 X(H_{i,t}^{tot}) di + \nu \ln \left(\frac{M_{\tau}}{P_{\tau}} \right) \right) \right], \quad (29)$$

subject to a per-period budget constraint:

$$\int_0^1 P_{j,t} C_{j,t} dj + M_t + R_t^{-1} B_t = \int_0^1 W_{i,t} H_{i,t} di + M_{t-1} + B_{t-1} + T_t^M + T_t^D. \quad (30)$$

Here $\int_0^1 P_{j,t} C_{j,t} dj$ is total nominal consumption, T_t^M is a lump sum transfer from the central bank, and T_t^D is a dividend payment from the firms. $\int_0^1 W_{i,t} H_{i,t} di$ is total labour compensation received from supplying the differentiated labour varieties $H_{i,t}$, and $H_{i,t}^{tot} = H_{i,t} + \tau_{i,t}^w + \mu_{i,t}^w$ is the total labour effort, including decision-making, of worker i . Each worker's labour and decision-making will vary with their current state (W, Z) as discussed previously.

Optimal consumption across the differentiated goods implies

$$C_{j,t} = (P_{j,t}/P_t)^{-\epsilon} C_t, \quad (31)$$

so nominal spending can be written as $P_t C_t = \int_0^1 P_{j,t} C_{j,t} dj$ under the price index P_t , given by equation (1). The first-order conditions for total consumption and for money use are

$$R_t^{-1} = 1 - \frac{v'(M_t/P_t)}{u'(C_t)} = E_t[\Lambda_{t,t+1}], \quad (32)$$

where the household's stochastic discount factor is given by:

$$\Lambda_{t,t+1} \equiv \frac{P_t u'(C_{t+1})}{P_{t+1} u'(C_t)}. \quad (33)$$

2.4 Monetary policy

We consider a monetary authority that generates an exogenous process for the money growth rate. We assume the nominal money supply is affected by an AR(1) shock g ,²⁰

$$g_t = \phi_g g_{t-1} + \epsilon_t^g, \quad (34)$$

where $0 \leq \phi_g < 1$ and $\epsilon_t^g \sim i.i.d.N(0, \sigma_g^2)$. Here g_t represents the time t rate of money growth:

$$M_t/M_{t-1} \equiv \mu_t = \mu^* \exp(g_t). \quad (35)$$

Seigniorage revenues are paid to the household as a lump sum transfer T_t^M , and the government budget is balanced each period, so that $M_t = M_{t-1} + T_t^M$.

In our money-in-the-utility model, an exogenous rise in nominal money growth causes households to demand more consumption, as the marginal utility of consumption rises, relative to that of money. Higher consumption represents a rise in output demand for firms which, given sticky prices, increase their production and labour demand. The aggregate price level rises too, yet it does so more slowly than nominal money. The increased labour demand leads some workers to set higher wages, but in the short run most continue supplying the labour demanded at their current sticky nominal wage.²¹

2.5 Aggregation

Summing across all goods, labour supply and goods demand must satisfy

$$N_t = \int_0^1 C_{j,t} A_{j,t}^{-1} dj = \int_0^1 C_t \left(\frac{P_{j,t}}{P_t} \right)^{-\epsilon} A_{j,t}^{-1} dj \equiv \bar{A}_t^{-1} C_t, \quad (36)$$

where N_t is the component of effective labour dedicated to goods production, and

$$\bar{A}_t \equiv \left(\int_0^1 \left(\frac{P_{j,t}}{P_t} \right)^{-\epsilon} A_{j,t}^{-1} dj \right)^{-1} \quad (37)$$

is a measure of aggregate productivity related to the degree of inefficient price dispersion, discussed by Yun (2005) and Nakamura et al. (2018).

Likewise, summing across goods and workers, the relation between labour hours and effective labour is

$$H_t = N_t^{tot} \left[\int Z_{i,t}^{\epsilon_n - 1} \left(\frac{W_{i,t}}{W_t} \right)^{-\epsilon_n} di \right] = N_t^{tot} \bar{Z}_t^{-1}, \quad (38)$$

20. In related work (Costain and Nakov 2011) we have studied state-dependent pricing when the monetary authority follows a Taylor rule. Our conclusions about the degree of state-dependence, microeconomic stylized facts, and the real effects of monetary policy were not greatly affected by the type of monetary policy rule considered. Therefore we focus here on the simple, transparent case of a money growth rule.

21. The same intuition holds also in the Calvo setup, *e.g.* in Galí (2008), Ch.6.

where H_t is the component of labour time sold to firms, and

$$\bar{Z}_t \equiv \left(\int_0^1 \left(\frac{W_{i,t}}{Z_{i,t} W_t} \right)^{-\epsilon_n} Z_{i,t}^{-1} di \right)^{-1} \quad (39)$$

measures the aggregate productivity of labour time in producing effective labour.

To compute a recursive equilibrium of this economy, we must identify its aggregate state, which will include aggregate shocks and the distribution of idiosyncratic states. Since nominal prices are predetermined under our assumed timing, it is natural to conjecture that the nominal state of the economy is summarized by:

$$\Omega_t \equiv (M_t, g_t, \Phi_t, \Phi_t^w), \quad (40)$$

where Φ_t and Φ_t^w are the distributions of nominal prices and productivities across firms, and nominal wages and productivities across workers. Online appendix B shows that the model is homogeneous of degree one in nominal variables, so the corresponding real state variable would be:

$$\Xi_t \equiv (g_t, \Psi_t, \Psi_t^w), \quad (41)$$

where Ψ_t and Ψ_t^w are the distributions of real prices and productivities across firms, and real wages and productivities across workers. It can be shown that this is a valid state variable for the economy by constructing an equilibrium in terms of Ξ .

Online Appendix B will show how to detrend the economy, so that it depends only on the real state Ψ_t instead of nominal Ω_t . Online Appendix C derives the distributional dynamics.²²

3 Calibration

We simulate the model at monthly frequency on a discrete grid, and discipline its key parameters using microdata on price and wage adjustments.

Data sources

As in CN19, our pricing data come from the Dominick’s supermarket dataset documented by [Midrigan \(2011\)](#).²³ These data represent weekly regular price changes, excluding temporary sales, and are displayed (in logs) as a blue-shaded histogram in the left panel of [Figure 2](#). We aggregate weekly adjustment rates to monthly rates for comparability with most related studies. We exclude sales because recent literature has argued that monetary nonneutrality depends primarily on the frequency of “regular” or “non-sale” price changes (*e.g.* [Eichenbaum](#)

22. After detrending, we discretize the two real distributions and solve for the model’s dynamics by the [Reiter \(2009\)](#) method. Replication codes are available at <https://github.com/borjapetit/costainnakovpetit2021>.

23. We are grateful to Virgiliu Midrigan for making his price data available to us, and to the James M. Kiltz Center at the Univ. of Chicago GSB, which is the original source of those data.

et al. 2011; Guimaraes and Sheedy 2011; or Kehoe and Midrigan 2015).²⁴.

Our wage change data are from the International Wage Flexibility Project (IWFP), seen as a blue histogram in the right panel of Figure 2; these data are taken from Figure 2a of Dickens et al. (2007). The histogram aggregates data on wage adjustments across multiple countries. While most of the underlying national data are drawn from surveys of firms, they refer to annual nominal wage changes of individual workers who remain employed by the same firm. The IWFP focused on annual changes because it observed a widespread tendency for wages to change once a year in many countries, which in turn means that most available surveys address annual changes. Clearly this makes our data on wage changes less than perfect for comparison with our price change evidence, which is at weekly frequency. Nonetheless, to get an empirical counterpart for our theoretical model, we will take the IWFP data at face value.²⁵ Therefore, we assume that the monthly frequency of nominal wage adjustment is $1/12=0.083$, and calculate nominal wage change statistics directly from the IWFP histogram.

Exogenous parameters

Some parameters are taken either from related papers or from standard values in the literature. We set the (inverse) Frisch elasticity to $\zeta = 0.5$, and $\nu = 1$. Following Golosov and Lucas Jr (2007), we set $\gamma = 2$, $\chi = 6$, and $\epsilon = 7$, and we set the same elasticity of substitution across varieties of labour as that across goods: $\epsilon_n = 7$. The discount factor is set to $\beta = 0.9967$, which corresponds to four percent annual discounting. The monthly survival probability is $\beta_S = 0.9979$, implying an expected working life of forty years. The log productivity of newborn workers is set to $z_0 = -0.6$, so that workers expect a 60 log points (82%) productivity gain over their life cycles. We assume two percent annual money growth in steady state, consistent with our retail pricing data. Table 1 collects these parameters.

Calibrated parameters

We next calibrate internally the decision cost parameters underlying the price- and wage-setting problems, as well as the productivity processes affecting firms and workers. These processes are assumed to follow discretized approximations of the following AR(1) dynamics:

$$\begin{aligned} a_{j,t} &= \rho_a a_{j,t-1} + \sigma_a \epsilon_t^a, \\ z_{i,t} &= \rho_z z_{i,t-1} + \sigma_z \epsilon_t^z, \end{aligned}$$

24. However, some authors dispute this conclusion; see Kryvtsov and Vincent (2014), Nevo and Wong (2019) and Alvarez and Lippi (2020)

25. Recently, Grigsby et al. (2021) study wage adjustment using higher-frequency data more comparable to retail price microdata. In U.S. data from a large payroll data processing firm, they find a wage adjustment probability of 26.0% quarterly and 72.7% annually; the mean absolute wage change, conditional on adjustment, is 10.7%. Considering job stayers only, they find a 66.3% annual wage change probability, with a mean absolute change of 6.34%. While their data imply somewhat larger wage changes than the IWFP data in our graphs, nonetheless, for job stayers, the order of magnitude is similar.

Table 1: Exogenous parameters

Parameter	Description	Value	Source
β	Discount factor (monthly)	0.9967	Annual real rate of 4%
β_S	Survival probability (monthly)	0.9979	Economic life span of 40 years
z_0	Log productivity at birth	-0.6	Normalization
ζ	Inverse Frisch elasticity	0.5	Standard value
ν	Coefficient on utility of money	1	Standard value
γ	Intertemporal elasticity of subs.	2	Golosov and Lucas Jr (2007)
χ	Coefficient on disutility of labor	6	Golosov and Lucas Jr (2007)
ϵ, ϵ_n	Elasticities of subs. across varieties	7	Golosov and Lucas Jr (2007)
μ^*	Long-run gross money growth	1.0017	Annual inflation of 2% (Dominicks')

where ϵ_t^a and ϵ_t^z are *i.i.d.* normal shocks with mean zero and unit variance.

Overall, we calibrate 8 parameters, collected in the vector $\mathcal{P} = (\bar{\lambda}, \kappa_f, \rho_a, \sigma_a, \bar{\rho}, \kappa_w, \rho_z, \sigma_z)$. We select \mathcal{P} by minimizing:

$$\mathcal{F}(\mathcal{P}) = (\mathcal{M}(\mathcal{P}) - \bar{\mathcal{M}})' \mathcal{W} (\mathcal{M}(\mathcal{P}) - \bar{\mathcal{M}}),$$

where $\mathcal{M}(\mathcal{P})$ is a vector of model-generated moments when the parameter vector is \mathcal{P} , $\bar{\mathcal{M}}$ is a vector of the corresponding moments computed from our price and wage data, and \mathcal{W} is a weighting matrix.²⁶ The vector $\mathcal{M}(\mathcal{P})$ contains the adjustment probabilities for prices and wages, as well as the histogram of (discretized) nonzero log price and log wage changes.

Table 2: Calibrated parameters

Parameter		Value	
Firms	Benchmark hazard (monthly)	$\bar{\lambda}$	0.2707
	Adjustment cost	κ_f	0.0177
	Productivity persistence	ρ_a	0.6441
	Standard deviation productivity shocks	σ_a	0.0703
Workers	Benchmark hazard (monthly)	$\bar{\rho}$	0.2317
	Adjustment cost	κ_w	0.0275
	Productivity persistence	ρ_z	0.9700
	Standard deviation productivity shocks	σ_z	0.0574

Table 2 collects the calibrated adjustment function and productivity process parameters: since wages are more rigid than prices in the data, the adjustment cost for wages is estimated to be substantially higher than that of prices, and the benchmark adjustment probability is lower. At the same time, workers' productivity is estimated to be more persistent but less volatile than firms' productivity. This makes sense, since worker productivity reflects persistent processes

26. Matrix \mathcal{W} weighs the adjustment probability for prices (wages) with the square root of the number of histogram bins for price (wage) changes, ensuring sufficient weight on the frequency component so that we match the empirical adjustment frequencies well.




Figure 2: Distribution of nonzero price and wage changes.

such as skill accumulation and life-cycle effects.²⁷

Model fit to aggregate steady-state

Figure 2 presents the model-generated histograms of non-zero price and wage changes (solid blue lines), as well as those in the data (grey area). Our simulation perfectly matches the adjustment probabilities of prices and wages (10.2% and 8.3% monthly, respectively). These “zero changes” are excluded from the graphs, which would otherwise need rescaling, making the histograms of non-zero changes harder to see. Fitting the complex shapes of the non-zero price and wage change histograms is more difficult, especially given our symmetric adjustment costs and technologies; the model-generated distributions are notably smoother than the data. But our model reproduces several features of the data. As in the empirical histograms, much of the mass is concentrated on small positive adjustments, but there is a fat right tail and a long, thinner left tail, and there is “missing mass” in the region of small negative wage adjustments.

Such a pattern of “missing mass”, compared with a normal distribution for example, is often taken to indicate downward nominal wage rigidity. It is interesting that our model, in which rigidities are entirely symmetric, also shows a dip in the mass just below zero, although this effect is weaker than it is in the data. In particular, wage changes between -2.5 and zero log points are 14% of all wage changes in the model, but only 4% in the data.²⁸ In our model, despite downward and upward wage adjustments being equally costly, workers rarely choose to make small negative wage changes because they expect their productivity to grow as they age, and because the nominal price level has a positive trend. While workers have an incentive to set a higher wage when they become more productive, they can react to small negative productivity shocks by waiting for price inflation to reduce their real wage, or for productivity growth to

27. Our estimation imposes an upper bound on productivity persistence, in order to limit the size of the required simulation grid. Workers’ estimated productivity persistence, 0.97, hits the upper bound.

28. Price and wage changes are reported in log points, where +1 log point corresponds approximately to a 1% increase (and +70 log points correspond approximately to multiplying by 2).

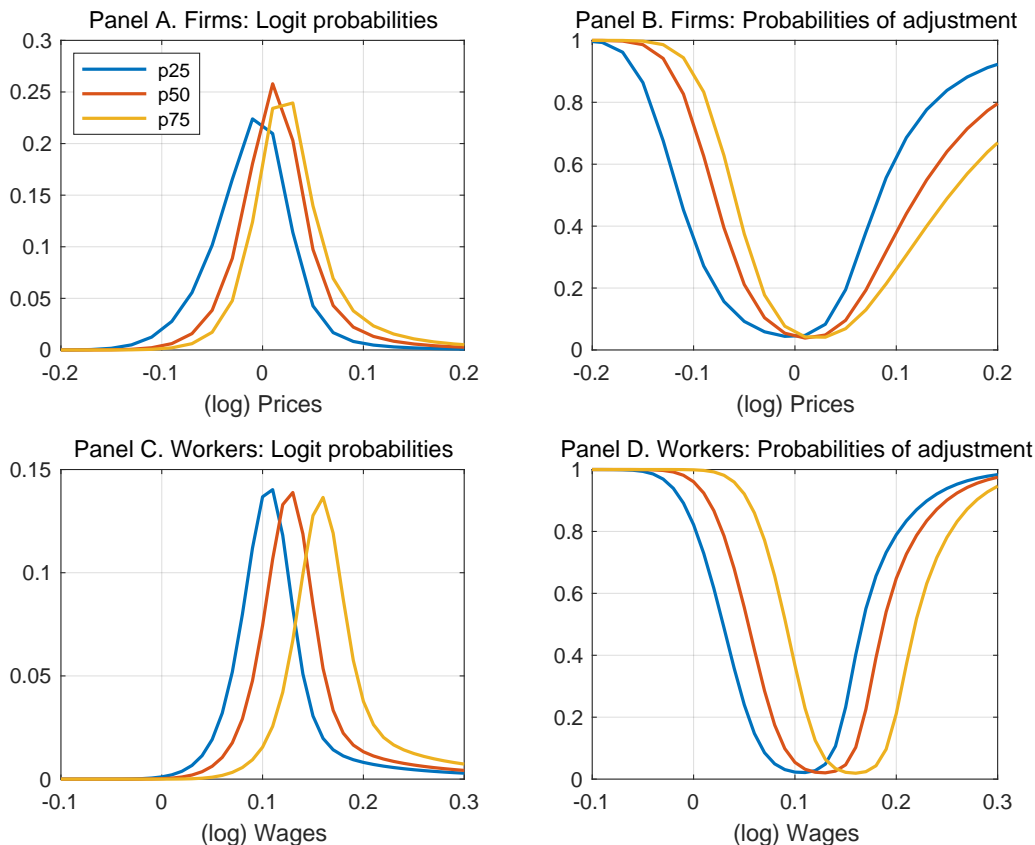


Figure 3: Adjustment behavior in the baseline model.

Notes: Distribution of price and wage changes (panels A and C) and price and wage adjustment probabilities (panels B and D); x -axis shows deviation of the log price or wage from its unconditional mean. The blue lines refer to a firm (worker) with cost (productivity) level in the 25th percentile of the cost (productivity) distribution; red lines refer to the median; yellow lines refer to the 75th percentile.

cancel out the shock.

Figure 3 shows the logit probabilities governing price resets and wage resets (left panels) and firms' and workers' adjustment probabilities (right panels) for low-, medium-, and high productivity firms/workers. For firms (top row), the probabilities are shown as functions of the lagged price, conditional on three possible realizations of marginal cost (inverse productivity); for workers, the probabilities are shown as functions of the lagged wage, conditional on three productivity levels. Firms prefer higher prices when costs are higher, and the probability of adjustment rises smoothly as firms deviate from the prices they prefer (conditional on costs).²⁹ As a consequence of our nonlinear disutility specification, workers set substantially higher wages as their productivity rises.³⁰ The preferred wage varies by roughly 6 log points as worker productivity varies from the 25th to the 75th percentile of the workers' productivity distribution.

29. Such evidence of state-dependent pricing in a retail context is presented in Eichenbaum et. al. (2011) using a weekly scanner dataset, see their Fig.8.

30. Our working paper, Costain et al. (2019), also presents a simulation of the linear labour disutility case. Linear disutility simplifies the numerical solution, but implies that the optimal wage choice does not vary with idiosyncratic productivity. Unsurprisingly, the resulting wage adjustment behaviour is highly counterfactual.

4 Results

We first study the steady-state and dynamic effects of nominal rigidity by comparing four calibrations of the model that vary the degree of noise in price- and wage-setting. We then turn to our main application: exploiting the state-dependence of price and wage setting to study nonlinearities in the effects of monetary policy, including how long-run inflation affects the slope of the Phillips curve.

4.1 Effects of noise in price and wage setting

We begin by analyzing how decision costs affect the frequency and the distribution of price and wage adjustments, asking which noise margin contributes most to the non-neutrality of monetary shocks. To do so, we compare our calibrated model to three counterfactual alternatives, presented in table 3. First, we simulate an economy in which we make prices flexible, by dividing κ_f by 100, while keeping wages sticky as in the baseline economy. Then, we keep prices sticky while making wages flexible, dividing κ_w by 100. Finally, we simulate an economy in which both prices and wages are flexible. To ease the exposition, we label the counterfactual versions as FP (flexible prices), FW (flexible wages), and FPFW (both flexible), respectively.

Table 3: Adjustment parameters for counterfactual exercises

	Baseline	FP	FW	FPFW
Firms (κ_f)	$\kappa_f^0 = 0.0177$	$\kappa_f^0/100$	κ_f^0	$\kappa_f^0/100$
Workers (κ_w)	$\kappa_w^0 = 0.0275$	κ_w^0	$\kappa_w^0/100$	$\kappa_w^0/100$

4.1.1 Steady-state adjustment behaviour

Table 4 reports the effects of reducing stickiness on steady-state adjustment statistics, comparing the four configurations of noise parameters. Decreased noise in price setting or wage setting makes adjustment more frequent: when prices are flexible, the frequency of price adjustment increases from 10% to almost 60%. Similarly, when wages are flexible, the frequency of wage changes rises from 8.34% in the baseline economy to 30.8%. When lower noise makes adjustment more frequent, prices and wages tend to deviate less from their desired levels than they do in the baseline economy. Consequently, decreasing the noise in price and wage setting implies smaller absolute changes on average. This effect is particularly strong for wages: in the baseline economy one fourth of wage changes were smaller than 2.5 log points, while when wages are flexible this increases to 80% of all wage changes. This large decrease in the size of wage changes occurs both because the baseline estimated noise parameter is higher for wages, and because convex labour disutility makes wage deviations especially costly.

This change in firms' and workers' adjustment strategies is illustrated in Figure 4, which compares the histograms of nonzero price and wage changes in the baseline economy with those

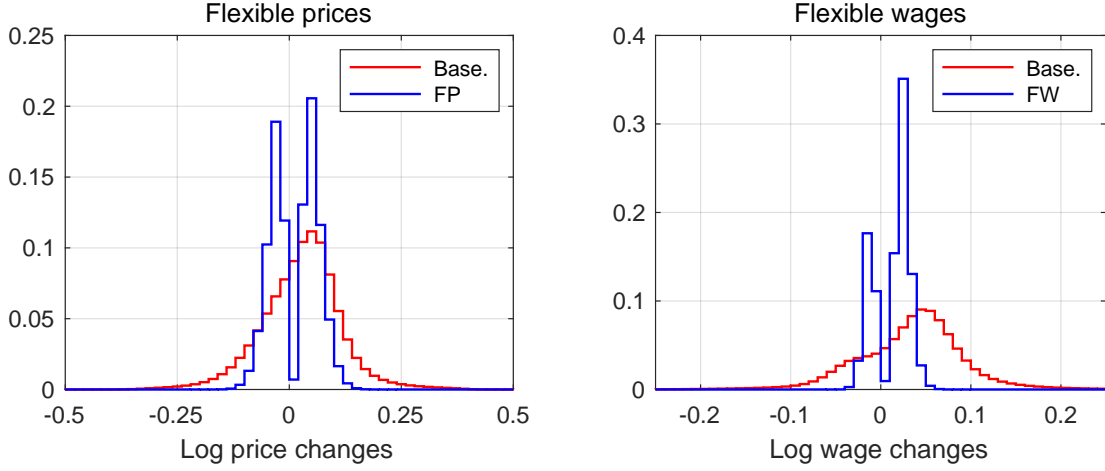


Figure 4: Distribution of nonzero price and wage changes: varying κ_f and κ_w .
Notes: left panel shows the effect of decreasing price stickiness on the distribution of nonzero price adjustments, keeping wages sticky. Right panel shows the effects of decreasing wage stickiness on the distribution of nonzero wage adjustments, keeping prices sticky.

of versions FP (left panel) and FW (right panel). When prices and wages are sticky (red lines), both histograms are smooth and display rather fat tails. As prices (wages) become more flexible, the price (wage) adjustment histogram becomes sharply bimodal, increasingly resembling the familiar (S, s) behaviour of a menu cost model. Errors in pricing and timing smooth out the distribution of changes under the baseline calibration, but as noise is reduced, stronger selection effects imply that the preponderance of price changes occur around two upper and lower thresholds. With sufficiently low noise, the firm (worker) rarely permits any large deviations from the optimal price (wage). Still, when the price or wage is already close to its target, it is optimal to avoid adjusting, thereby avoiding decision costs and the risk of errors. Therefore, the adjustment distributions in scenarios FP and FW are much spikier than their empirical counterparts (which were seen in Fig. 2), with typical changes that are less than half as large as those in the data (a mean absolute price change of 4.53 log points in specification FP, and a mean absolute wage change of 1.95 log points in specification FW).

4.1.2 Costs of frictions

Lower noise permits firms and workers to spend less on decision-making while setting prices and wages closer to their optimal values, which in effect increases the aggregate productivity \bar{A}_t of effective labour, and likewise implies higher \bar{Z}_t : more effective labour per unit of labour time. The top panel of Table 4 shows the steady state output losses Θ_t^p due to pricing frictions, expressed as a percentage of aggregate consumption C_t , which are 2.78% under the baseline parameters. The losses compared with aggregate consumption C_t^{fp} under flexible prices can be decomposed as follows (see Appendix A):

$$\Theta_t^p = \frac{C_t^{fp} - C_t}{C_t} = \left(\frac{\bar{A}^{fp} - \bar{A}_t}{\bar{A}_t} \right) \frac{N^{tot}}{N_t} + \frac{\mu_t}{N_t} + \frac{\tau_t}{N_t}. \quad (42)$$

Table 4: Evaluating the model with different values of κ_f and κ_w

		Data	Base.	FP	FW	FPFW	
<i>Prices</i>	Frequency of price change (%)	10.20	10.21	59.51	10.21	59.65	
	Mean absolute price change (%)	9.90	6.94	4.53	6.92	4.52	
	Skewness of price changes	-0.42	-0.12	-0.06	-0.12	-0.06	
	Kurtosis of price changes	4.81	4.60	2.01	4.60	2.01	
	Standard deviation of (log) prices ($\times 100$)	–	3.73	4.57	3.72	4.57	
	% of price changes > 0	65.10	56.47	52.37	56.49	52.37	
	% of abs price changes < 0.025	12.00	27.26	25.69	27.27	25.84	
	Output losses due to price stickiness (%) ^a , Θ^p	–	2.78	1.16	2.77	1.16	
	<i>Cost θ^{p*} of decision errors (%)^b</i>	–	1.51	0.91	1.51	0.91	
	<i>Cost of price setting (%)^b</i>	–	0.47	0.03	0.47	0.03	
	<i>Cost of timing choice (%)^b</i>	–	0.49	0.07	0.48	0.07	
	<i>Wages</i>	Frequency of wage change (%)	8.30	8.34	8.33	30.81	30.68
		Mean absolute wage change (%)	6.47	5.50	5.50	1.95	1.96
		Skewness of wage changes	0.35	0.17	0.17	-0.46	-0.46
Kurtosis of wage changes		4.39	11.94	11.70	2.00	2.00	
Standard deviation of (log) wages ($\times 100$)		–	3.52	3.38	3.60	3.45	
% of wage changes > 0		86.50	70.62	70.60	66.75	66.77	
% of abs wage changes < 0.025		11.80	25.17	25.17	80.21	80.02	
Output losses due to wage stickiness (%) ^a , Θ^w		–	1.98	2.00	0.08	0.08	
<i>Cost θ^{w*} of decision errors (%)^b</i>		–	0.74	0.76	0.18	0.18	
<i>Cost of wage setting (%)^b</i>		–	0.94	0.95	0.03	0.03	
<i>Cost of timing choice (%)^b</i>		–	1.09	1.10	0.08	0.08	
Total output losses relative to frictionless economy ^c , Θ		–	2.35	1.74	1.45	0.83	

^aOutput losses (Θ^p and Θ^w) differ from the sum of the following cost terms because output loss includes general equilibrium effects.

^bThe sum of the three cost terms in prices (wages) panel represents the gain accruing to a single firm (worker) unconstrained by decision costs ($\kappa_f = 0$ or $\kappa_w = 0$).

^cTotal output losses Θ differ from $\Theta^p + \Theta^w$ because Θ includes change in labor effort.

The time firms spend on decisions, $\mu_t + \tau_t$, costs 0.96% of output under the baseline parameters, falling to 0.1% of output in the flexible-price cases. The term $(\bar{A}^{fp} - \bar{A}_t)/\bar{A}_t$ measures how inefficient price dispersion changes, as aggregate productivity rises from \bar{A}_t (with frictions) to \bar{A}^{fp} (without). This misallocation term combines partial equilibrium gains from reduced price errors with the general equilibrium effect of changing the overall price level.³¹ But to isolate the partial equilibrium effects of price flexibility we instead report θ_t^{p*} , the potential increase in the representative firm's profits from eliminating errors, holding the rest of the economy fixed. Like the decision costs, these potential gains are expressed as a fraction of aggregate output (equivalently, aggregate consumption or revenues). The gains from eliminating errors are 1.51% of output in the baseline case, falling to 0.91% in cases FP and FPFW. The sum of the three rows in italics represents a single firm's gains from eliminating its price frictions (holding the

31. The misallocation term (not shown in the table) can be inferred by subtracting pricing and timing costs from output losses Θ_t^p .

rest of the economy fixed), which is 2.47% of output in the baseline specification.

Analogous cost terms apply for workers, whose losses (as a fraction of output, holding fixed labour effort) can be decomposed as

$$\Theta_t^w = \frac{C_t^{fw} - C_t}{C_t} = \frac{\bar{Z}_t^{fw} - \bar{Z}_t}{\bar{Z}_t} + \frac{\mu_t^w}{\underline{H}_t} + \frac{\tau_t^w}{\underline{H}_t}, \quad (43)$$

where $\underline{H}_t \equiv H_t - \bar{Z}_t^{-1}(\mu_t + \tau_t)$ is labour time devoted to goods production. Output costs derive in part from workers' time spent on decisions, $\mu_t^w + \tau_t^w$; these costs decrease from 2.03% of consumption (twice as large as the decision costs of firms) in the baseline case, to 0.11% of consumption in cases FW and FPFW. Misallocation effects derive from the difference in units of effective labour per unit of labour time when wages are flexible (\bar{Z}_t^{fw}) or sticky (\bar{Z}_t). We also report the gains θ^{w*} to a single worker who makes error-free decisions, holding fixed the rest of the economy. These gains fall from 0.74% of consumption in the baseline economy to 0.18% in the flexible wage simulations.

Finally, the last line of Table 4 reports the total consumption loss caused by price and wage frictions, compared to a frictionless economy. Appendix A shows that these losses equal

$$\Theta = \underbrace{\frac{\bar{A}^{fl} \bar{Z}_t^{fl}}{\bar{A}_t \bar{Z}_t} - 1}_{\text{Misallocation}} + \underbrace{\frac{\mu_t^w + \tau_t^w}{\underline{H}_t}}_{\text{Wage setting}} + \underbrace{\frac{\mu_t + \tau_t}{N_t}}_{\text{Price setting}} + \underbrace{\frac{\bar{A}_t^{fl} \bar{Z}_t^{fl} (H_t^{fl} - \underline{H}) - \bar{A}_t \bar{Z}_t (H_t^{tot} - \underline{H})}{\bar{A}_t \bar{Z}_t \underline{H}_t}}_{\text{Labour effort}} \quad (44)$$

The misallocation term incorporates the misallocation components of Θ^p and Θ^w (plus their interactions); therefore it combines partial equilibrium errors and general equilibrium effects due to changes in aggregate prices and wages. The overall loss is 2.34% of aggregate consumption in the baseline model, falling to 0.83% in the nearly frictionless model FPFW. This overall loss is less than the sum Θ^p and Θ^w because the negative wealth effect from imposing frictions motivates workers to supply more labour effort. This effect is the fourth term in (44); it is not included in Θ^p , which is computed fixing N_t^{tot} , or in Θ^w , which holds H_t^{tot} fixed.

4.1.3 Effects of monetary policy shocks

We next analyse how price- and wage-setting frictions affect the non-neutrality of monetary shocks. Figure 5 shows the effects of an autocorrelated money growth shock with monthly persistence 0.8. The figure compares the responses of the aggregate nominal price and wage levels, consumption, hours and the real wage as price and wage stickiness vary, across the different combinations of noise parameters. On impact, all the specifications exhibit similar effects, in particular consumption rises by 2.5%. The effects are similar on impact because we assume firms and workers cannot react to the shocks contemporaneously, but thereafter the four specifications behave differently. The baseline case (blue line), with both sticky prices and wages, implies substantial real effects of a monetary shock: the consumption and labour responses exhibit a half-lives of 7 and 8 months, respectively. This amounts to an 1.4% increase in

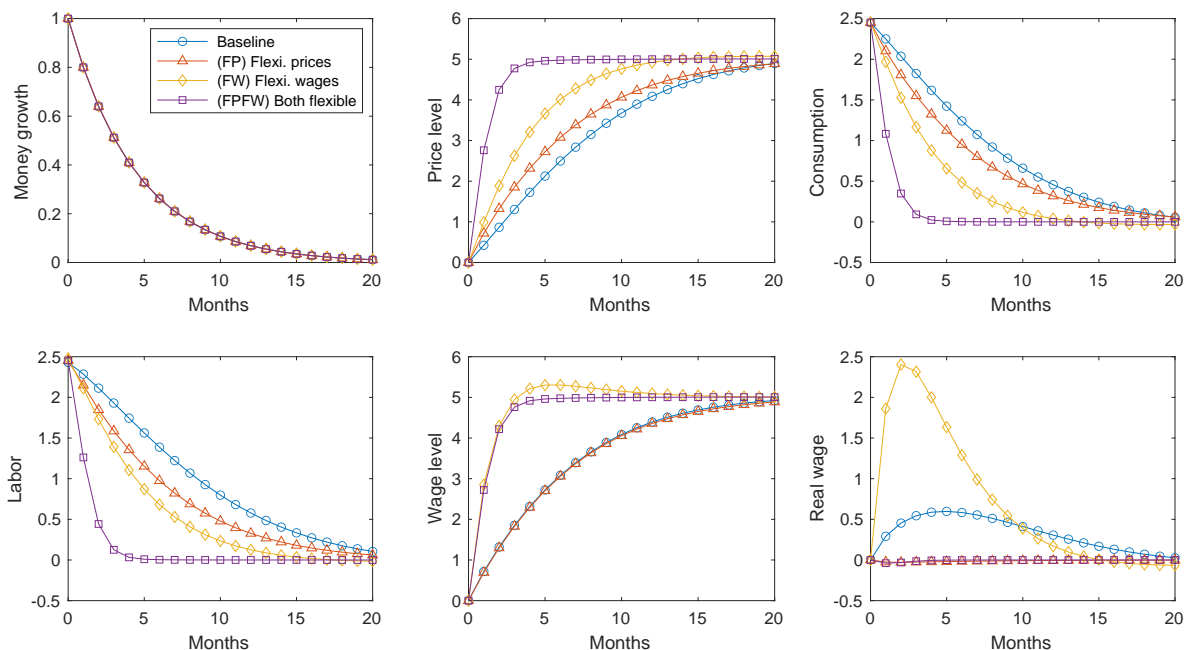


Figure 5: Money growth shock: effects of nominal rigidity.

Notes: Impulse responses to a 1 percent money supply shock (autocorrelation 0.8). Blue: Baseline model, with both prices and wages sticky. Red: (FP), flexible prices and sticky wages. Yellow: (FW), sticky prices and flexible wages. Purple: (FPFW): both prices and wages flexible.

consumption over the year following the shock, and an additional 0.17% increase in the following year. When prices become flexible, these real effects decrease to a consumption half-life of 6 months. The decrease in real effects is much more significant when wages are flexible: the half-life of the consumption IRF falls to only 4 months. As expected, the smallest real effect is found in the FPFW specification, which has very low real persistence, as in the [Golosov and Lucas Jr \(2007\)](#) menu cost model. In this case, with both price and wage flexibility, the half-life of the consumption and labour responses is only 2 months, with just a 0.33% increase in consumption over the first year.

Overall, these results show that the real effects of money shocks are large as long as wages are sticky. Version FP (sticky wages and flexible prices) has almost the same consumption response as the baseline economy, and lies substantially above FW (flexible wages and sticky prices). Intuitively, the reason why wage stickiness is crucial for non-neutrality is that it prevents rapid adjustment of firms' marginal costs, so even though prices are much more flexible in version FP than under the baseline calibration, the impulse responses of the price level in the two cases are quite similar.³² Both wages and prices adjust gradually in version FP, giving a real effect on consumption and output that is almost as large and persistent as that seen in the baseline case.

The key takeaway is that wage rigidity matters more than price rigidity for the overall degree of monetary non-neutrality in this model, consistent with the findings of [Christiano et al.](#)

32. In contrast, as Table 4 showed, when firms' marginal costs fluctuate for exogenous reasons, such as the productivity shocks $A_{j,t}$, prices do change more in version FP than in the baseline version.

(2005) and [Huang and Liu \(2002\)](#) for the Calvo model. The importance of wage rigidity for propagation of nominal shocks to real variables provides support for New Keynesian mechanisms in the light of empirical evidence of increased markups of price over marginal cost conditional on positive demand shocks ([Nekarda and Ramey 2013](#)). On the other hand, our findings do not offer any strong macroeconomic reason to favor the baseline specification with both rigidities versus version FP, where only wages are rigid. Empirical studies rarely find a significantly nonzero response of the real wage to monetary policy shocks (see for example [Christiano et al. 2005](#); [McCallum and Smets 2007](#); [Olivei and Tenreyro 2007](#); [Christiano et al. 2016](#)). This would suggest rejecting specification FW, but may not suffice to distinguish between parameterizations FP and FPFW on macroeconomic grounds alone.

4.2 Trend inflation and the slope of the Phillips curve

A crucial ongoing debate related to the effectiveness of monetary policy is how to understand the slope of the Phillips curve. We will see that state-dependent adjustment helps explain the curve’s apparent flattening, as a response to changes in trend inflation. But before we look directly at this issue, we first examine the effects of trend inflation in our model more generally.

4.2.1 Microeconomic effects of trend inflation

Considering the current prolonged “lowflation” episode, it is important to ask how trend inflation affects state-dependent adjustment behaviour. To this end, we simulate five economies with inflation trends varying from -2% to 8%, under the baseline noise parameters. We then repeat the simulations, changing the noise parameters, to shed light on the role of price and wage stickiness in accounting for the real effects of monetary policy as trend inflation varies.

Table 5 reports steady state statistics, under the baseline noise parameters, for trend inflation rates from -2% to 8%. The frequency of price and wage adjustment increases as the trend inflation rate deviates from zero. At exactly 0% inflation, firms and workers only need to adjust their prices and wages in response to productivity shocks. When trend inflation differs from 0%, though, they must also make adjustments to keep their prices and wages at their desired real levels. Therefore, firms and workers raise their monthly adjustment probabilities by 4.2 and 3 percentage points, respectively, as we move from 0% to a 4% inflation trend.

Likewise, the same increase in trend inflation causes the share of positive price and wage changes to rise too, going from 44% to 63% in the case of prices, and from 54% to 78% for wages³³. In contrast to the effects of a reduction in adjustment costs, the higher adjustment probability is not accompanied by smaller price and wage changes. As trend inflation rises from 2 to 8%, the average absolute sizes of price and wage changes both increase by 1.45 log points. This increase is somewhat at odds with the evidence of [Nakamura et al. \(2018\)](#), who showed

33. Figure D1 in Online Appendix D shows the histograms of price and wage changes at different trend inflation rates.

Table 5: Evaluating the baseline model at different trend inflation rates

		Trend inflation rate				
		-2%	0%	2%	4%	8%
<i>Prices</i>	Frequency of price change (%)	10.06	7.65	10.21	11.87	14.37
	Mean absolute price change (%)	6.78	6.21	6.94	7.49	8.39
	Skewness of price changes	0.33	0.09	-0.12	-0.25	-0.41
	Kurtosis of price changes	4.76	5.09	4.60	4.41	4.28
	Standard deviation of (log) prices (x100)	3.69	2.97	3.73	4.18	4.80
	% of price changes > 0	34.35	44.21	56.47	62.65	70.03
	% of abs price changes < 0.025	27.66	35.26	27.26	23.25	18.79
	Output losses due to price stickiness (%) ^a , Θ^p	2.77	2.69	2.78	2.85	2.99
	Cost θ^{p*} of decision errors (%) ^b	1.51	1.46	1.51	1.55	1.60
	Cost of price setting (%) ^b	0.47	0.54	0.47	0.44	0.42
Cost of timing choice (%) ^b	0.49	0.38	0.49	0.55	0.64	
<i>Wages</i>	Frequency of wage change (%)	7.81	6.98	8.34	9.99	13.05
	Mean absolute wage change (%)	4.98	4.94	5.50	6.04	6.95
	Skewness of wage changes	1.02	0.57	0.17	-0.04	-0.24
	Kurtosis of wage changes	12.48	12.22	11.94	12.02	11.83
	Standard deviation of (log) wages (x100)	3.55	3.52	3.52	3.53	3.55
	% of wage changes > 0	34.35	54.43	70.62	78.52	85.48
	% of abs wage changes < 0.025	27.12	30.36	25.17	20.73	15.39
	Output losses due to wage stickiness (%) ^a , Θ^w	1.96	1.82	1.98	2.17	2.52
	Cost θ^{w*} of decision errors (%) ^b	0.60	0.59	0.74	0.88	1.16
	Cost of wage setting (%) ^b	0.94	0.99	0.94	0.89	0.83
Cost of timing choice (%) ^b	1.08	0.95	1.09	1.26	1.55	
Total output losses relative to frictionless economy ^c , Θ	2.30	2.21	2.35	2.50	2.77	

^aOutput losses (Θ^p and Θ^w) differ from the sum of the following cost terms because output loss includes general equilibrium effects.

^bThe sum of the three cost terms in prices (wages) panel represents the gain accruing to a single firm (worker) unconstrained by decision costs ($\kappa_f = 0$ or $\kappa_w = 0$).

^cTotal output losses Θ differ from $\Theta^p + \Theta^w$ because Θ includes change in labor effort.

that the absolute size of price adjustments has been quite stable in US CPI microdata since 1975, including the Great Inflation episode. Yet, due to state-dependence, the increase in the size of price adjustments that our model generates is substantially smaller than what a Calvo model would predict.

Nakamura et al. (2018) argued that the absolute size of price changes could be an observable proxy for *inefficient* price dispersion, and hence for the costs of nominal rigidity and inflation. They also showed that inefficient price dispersion is inversely related to the productivity measure \bar{A}_t defined in (37). Therefore, the misallocation term that appears in our decomposition (42) measures the same concept—losses due to inefficient price dispersion—as a fraction of aggregate consumption. Using the fact that $\frac{P_{j,t}^{fp}}{P_t^{fp}} = \frac{\bar{A}_{j,t}^{fp}}{\bar{A}_{j,t}}$ and that $\int_0^1 \left(\frac{A_{j,t}}{\bar{A}^{fp}}\right)^{\epsilon-1} dj = 1$, the misallocation

term can be rewritten as:

$$\frac{\bar{A}^{fl} - \bar{A}_t}{\bar{A}_t} = \int \left(\frac{A_{j,t}}{\bar{A}^{fl}} \right)^{\epsilon-1} \left[\left(\frac{P_{j,t} P_t^{fl}}{P_{j,t}^{fl} P_t} \right)^{-\epsilon} - 1 \right]. \quad (45)$$

Hence misallocation is zero when all firms are at their optimal prices ($P_{j,t} = P_{j,t}^{fl}$) and the aggregate price is at its flexible value ($P_t = P_t^{fl}$); it deviates from zero as individual or aggregate prices deviate from their flexible levels.

In our baseline parameterization, misallocation relative to fully flexible prices is 1.82% (in the table, this is the difference between Θ^p and the costs of price setting and timing). Misallocation is minimized at zero inflation, falling to 1.77%, and it rises very gradually with inflation, reaching 1.93% at 8% annual trend inflation. Thus, while absolute price and wage changes get substantially larger with inflation in our simulations, the actual change in inefficient dispersion is small. This points to the importance of state dependence: firms tend to adjust prices when they get too far out of line, which is why the welfare costs of inflation are much more stable in state-dependent models than they are in the Calvo framework. On the other hand, it also indicates the relevance of substitution on the consumers' side, which reduces the welfare impact of inefficient price dispersion. The misallocation term is a welfare-based measure, taking account of consumers' ability to substitute across goods, with elasticity $\epsilon = 7$, when some goods are inefficiently priced.³⁴

4.2.2 Trend inflation and monetary nonneutrality

The effects of trend inflation on the economy go beyond the steady state, also affecting dynamic responses to shocks. Figure 6 compares the impulse responses of our estimated baseline model to a 1% money supply shock (with monthly autocorrelation 0.8, as before) as trend inflation varies. The largest real effects are obtained at 0% trend inflation, and these effects decrease monotonically as trend inflation deviates from zero. Concretely, the half-life of the labour response falls from 10 months at 0% trend inflation, to 8 months at the baseline 2%, and only 5 months at 8% trend inflation. The reduction in the real effects of monetary policy is due to the lower persistence induced by more frequent price and wage adjustments, as shown in Table 5. This is consistent with evidence of [Ascari and Haber \(2021\)](#) showing a quicker and less persistent reaction of prices to a monetary shock when trend inflation is higher.

34. Likewise, the corresponding term in (43),

$$\frac{\bar{Z}^{fw} - \bar{Z}_t}{\bar{Z}_t} = \bar{Z}^{fw} [\bar{Z}^{-1} - (\bar{Z}^{fw})^{-1}] = \bar{Z}^{fw} \int Z_{i,t}^{-1} \left[\left(\frac{W_{i,t}}{Z_{i,t} W_t} \right)^{-\epsilon_n} - \left(\frac{W_{i,t}^{fl}}{Z_{i,t} W_t^{fl}} \right)^{-\epsilon_n} \right] di,$$

measures losses in aggregate consumption due to wage dispersion. These losses are small, and hardly change with trend inflation, from a minimum of -0.12% at zero inflation to a maximum of 0.14% at 8% inflation. They are small because general equilibrium effects largely cancel out the partial equilibrium losses due to errors. In general this term has an ambiguous sign, because consumption gains (losses) may be compensated by leisure losses (gains).

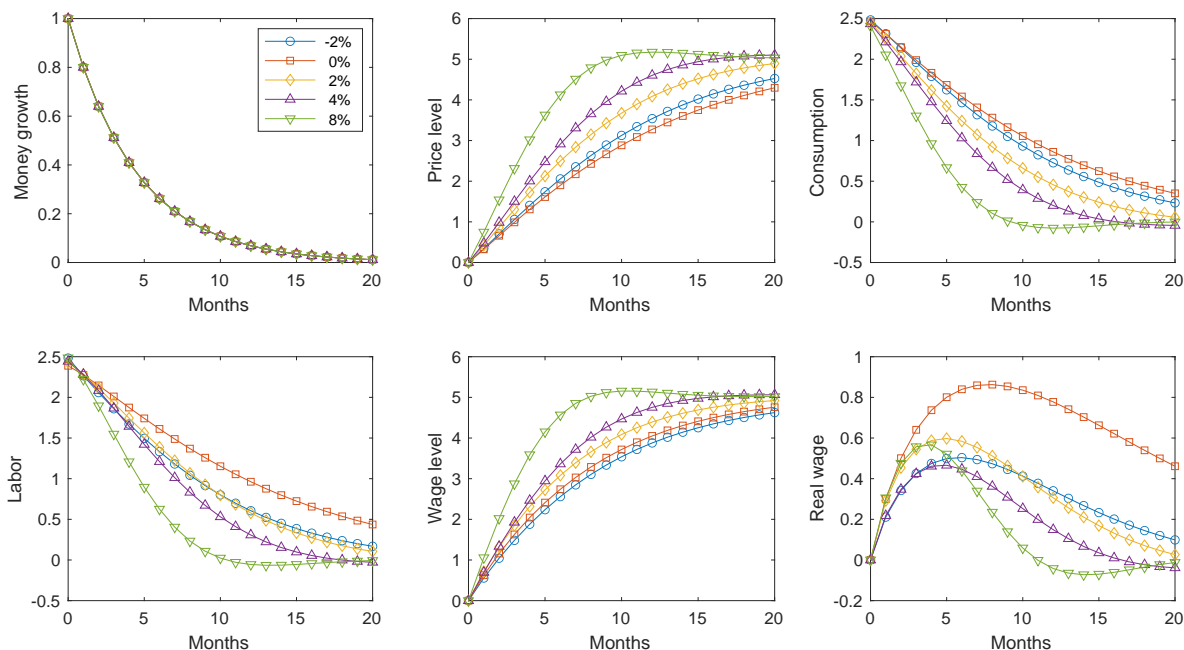


Figure 6: Impulse responses at different trend inflation rates in the baseline model. *Notes:* Impulse responses to a 1 percent money supply shock (autocorrelation 0.8), starting from annual trend inflation rates of -2% (blue), 0% (red), 2% (baseline case, yellow), 4% (purple) and 8% (green).

To further quantify how monetary transmission varies with trend inflation, we can compare our model to the Calvo framework by fitting Calvo specifications for price and wage adjustments that reproduce, as closely as possible, the state-dependent model’s output and inflation impulse responses.³⁵ The estimates are summarized in Table 6. The estimated Calvo probabilities for the baseline case with 2% inflation are considerably higher than the corresponding average frequencies of price and wage setting in our data, especially for wages. This is because, absent the selection effect, the only way for the Calvo model to replicate the more rapid nominal responses of a state-dependent model is by adjusting upward the exogenous probabilities of price and wage resetting. As trend inflation increases from 0 to 8%, the IRF-matching Calvo monthly adjustment probabilities increase from 11% to 26% for prices and from 22% to 48% for wages. This slope (roughly 2 for prices, or 3 for wages) can also be viewed as an elasticity of the hazard rate to the rate of price growth (one plus the inflation rate) of 11.2 for prices, and 10.1 for wages. These large elasticities highlight the relevance of the extensive margin of adjustment in response to positive trend inflation, and hence the importance of state dependence versus the “straitjacket” of the Calvo model.

Another way to quantify the effects of trend inflation on the model’s dynamics is to look at the Phillips multiplier, defined by [Barnichon and Mesters \(2021\)](#) as the ratio of the cumulative

35. To do so, we search with `csmine1` for a combination of Calvo parameters that minimizes the squared distance between the IRFs from the model with state dependent prices and wages, and the ones of the corresponding Calvo model. We truncate the IRFs at 36 months, by which time both the inflation and consumption responses are back to steady-state.

Table 6: Best fitting Calvo models

Inflation rate	Calvo Parameter	
	Prices	Wages
0%	0.1110	0.2219
2%	0.1530	0.2972
4%	0.1833	0.3487
8%	0.2645	0.4752

inflation response to the cumulative hours response over a sufficiently long horizon, *e.g.* 48 months. We display the multiplier in Table 7 for different trend inflation rates and adjustment costs scenarios. The “Baseline” column shows the Phillips multiplier corresponding to the scenario presented in Figure 6. The multiplier is lowest at 0% trend inflation, meaning that real effects are largest, as we discussed before.

Table 7: Phillips multipliers at different trend inflation rates and noise parameters

Trend inflation	Baseline	Flexible Prices	Flexible Wages	Flexible Prices and Wages
-2%	0.229	0.225	0.572	1.071
0%	0.167	0.212	0.267	1.080
2%	0.239	0.309	0.414	1.156
4%	0.297	0.404	0.502	1.230
8%	0.446	0.665	0.614	1.335

When we make prices flexible, firms increase their adjustment frequency. Hence we see in the “Flexible prices” column that the Phillips multiplier increases, except for the case of -2% trend inflation, where it is roughly unchanged. This non-monotonicity in the multiplier is not observed for wages; imposing wage flexibility increases the multiplier substantially at all the trend inflation rates we consider. Again, the numbers suggest that wage rigidities matter more than price rigidities for monetary non-neutrality; the increase in the multipliers (the reduction in the real effects) is much greater when we make wages flexible than when we make prices flexible, except at the highest inflation rate in the table. However, taken together with the Calvo estimates from Table 6, these results suggest that wage rigidity matters not because frictions are quantitatively larger for wages than for prices, but instead because wage stickiness feeds through to price stickiness, by way of firms’ cost structure.

4.2.3 Flattening of the Phillips curve

The paper of [Barnichon and Mesters \(2021\)](#) is one of many that have explored the apparent flattening of the Phillips curve in recent decades (see also [Ball and Mazumder 2011](#) and [Coibion and Gorodnichenko 2015](#), among others). Papers by [Benigno and Ricci \(2011\)](#) and [Lindé and Trabandt \(2018\)](#) have shown that asymmetric nominal rigidities can make the Phillips curve

flatter at low inflation rates. While the adjustment costs in our model are entirely symmetric, our model likewise implies a decrease in the slope of the Phillips curve as trend inflation falls. Therefore, in this section, we quantify how much of the observed change in the curve’s slope can be explained on the basis of state-dependent nominal adjustment.

Following [Jorgensen and Lansing \(2019\)](#), we take a traditional approach to evaluating this question, by estimating reduced-form Phillips curves both in the data and in our model, driven by money supply shocks.³⁶ We use US data from the period 1980-2020, and estimate the reduced-form relationship between the change in price inflation and the output gap. We split the sample in two sub-periods: 1980:1-1999:4 and 2000:1-2019:4. The first period is characterized by a higher inflation rate, with an average of 4.6%, while in the second period average inflation decreases to 2%. The top panels of [Figure 7](#) show the relationships between the change in inflation and the output gap for the two periods. As previous literature has found, the observed slope of the Phillips curve is higher in the period of higher inflation, between 1980 and 2000. To quantify how much of this flattening of the Phillips curve can be generated by the changing non-neutrality of our state-dependent model, we next simulate two scenarios with the observed trend inflation rates. In particular, we solve the model at 4.6% and at 2% trend inflation, subject to monetary policy shocks and then simulate 1000 monthly observations. Since the actual series are at quarterly frequency, we transform our model-generated data which is at monthly frequency. We then define the (one year) change in inflation as $\hat{i}_t = i_t - i_{t-4}$ where i_t is the (quarterly) inflation rate. We define the output gap as $\hat{y}_t = y_t - \bar{y}$ where \bar{y} is the steady-state level of output (in logs). We keep the last 80 quarterly observations from each simulated sample.³⁷ Using these transformed model-generated series, we run the same regression that we ran on the actual data. The bottom panels of [figure 7](#) show the outcome of this exercise.

Table 8: Slope of the Phillips curve. Data and Model

	1980:1-1999:4	2000:1-2019:4	Change	% Change
Data	0.3835	0.0114	-0.3721	-97.03
Model	0.3515	0.2050	-0.1464	-41.66

As shown in [figure 7](#), the Phillips curve derived from simulated data is substantially flatter in the second period, when annual trend inflation is 2.6 percentage points lower. [Table 8](#) collects the slopes of the Phillips curves for the two periods in the data and in the model. The slope of the simulated Phillips curve falls by 42%, while in the data the corresponding decline is of 97%. Thus, our state-dependent model explains almost half of the observed decline in the slope of

36. Our approach is limited to estimating conditional correlations between the change in inflation and the output gap, treating the economy as if it were subject to money shocks only. However, the logic of our mechanism—that price adjustment is sluggish when trend inflation is low—should be applicable to demand shocks in general, implying a smaller inflationary response together with a larger real response.

37. We calibrate the size of the money growth shocks in our model such that we match the standard deviation of inflation in each sample period.

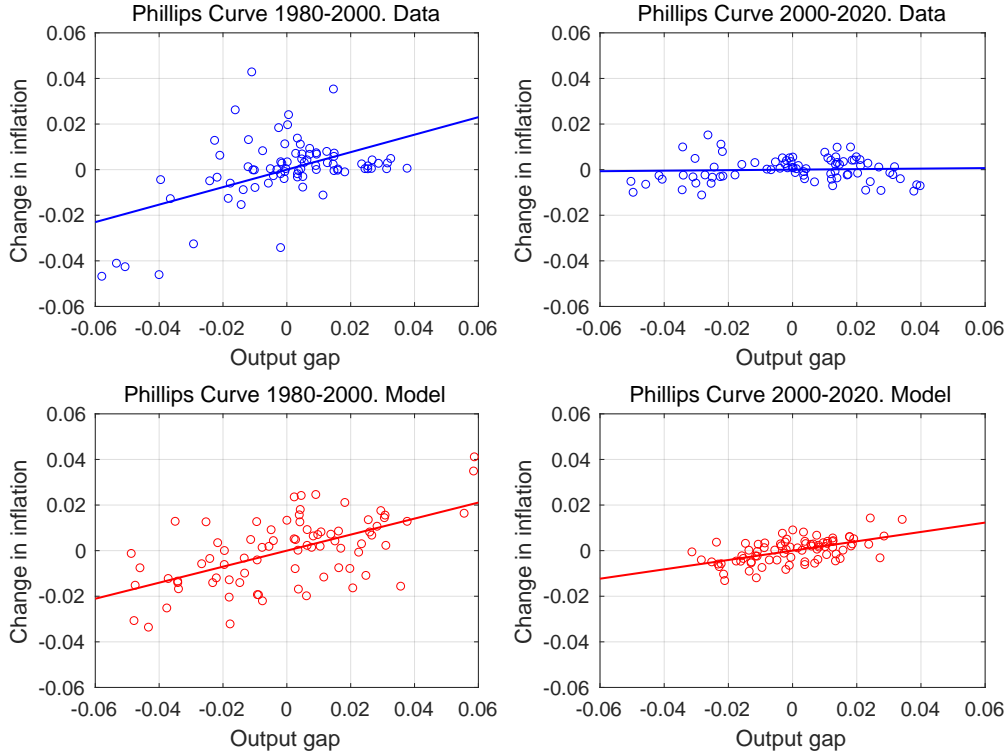


Figure 7: Phillips curves in the data and the model.

Source: FRED II (series CPILFESL, GDPPOT and GDPC1). The change in inflation is computed as $\hat{i}_t = i_t - i_{t-4}$, where $i_t = d \log(\text{CPILFESL}_t)$ and the output gap is $\hat{y}_t = \log(\text{GDPC1}_t / \text{GDPPOT}_t)$. See the main text for more details on how to construct the model-generated series.

the Phillips curve as a response to this relatively modest change in trend inflation. Our model, therefore, leaves room for other, complementary, explanations for a real or apparent flattening of the Phillips curve. Prominent hypotheses include improved monetary policy (Roberts 2006; McLeay and Tenreyro 2020), better anchoring of inflation expectations (Jorgensen and Lansing 2019; Barnichon and Mesters 2021), and increased competitive pressure related to technology or globalization (Forbes 2019).

4.3 Limits to monetary stimulus

The variation in impulse responses as trend inflation changes is not the only non-linearity displayed by our model. As stressed by Ascari and Haber (2021) and Alvarez and Lippi (2014), in models of state dependent pricing, the effects of monetary policy also vary dramatically with the size of the shock.

Figure 8 shows that as money supply shocks become larger, their impact falls proportionally less on consumption (and other real variables). The figure compares the cumulative responses of consumption to one-time, permanent, uncorrelated shocks to the money supply varying from zero to 10 log points. Under the baseline specification a small jump in the money supply causes a persistent increase in consumption; the cumulative effect on consumption is 2.7% for a jump of two log points in the money supply. However, the persistence of the real effects drops rapidly

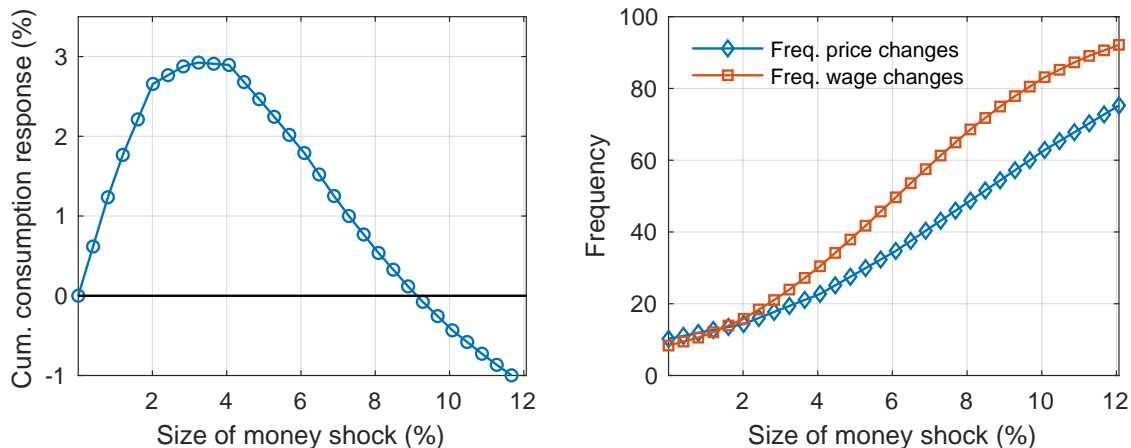


Figure 8: Comparing small and large money supply shocks.
Left: cumulative impulse responses of consumption to one-time increase in the money supply.
Right: change in adjustment frequency, on impact, for wages and prices.

with the size of the shock, so the cumulative real change is actually smaller for a money shock of six log points than it is for a shock of four log points. The reason is that larger shocks give firms and workers ever stronger incentives to adjust prices and wages immediately (stronger selection and extensive margin effects). Thus, most of the nominal reaction occurs immediately, making the real effects smaller. Indeed, for money supply shocks larger than 9 log points, the real stimulus on impact shrinks, and the brief initial rise of the impulse-response is followed by a prolonged slump in consumption and labour due to inflationary distortions, making the cumulative impact on consumption negative.³⁸

We can also decompose the total effect by looking at the FP and FW specifications.³⁹ In particular, it turns out that sticky wages are more important for non-neutrality for relatively smaller money shocks, while sticky prices are more important for larger shocks. This is linked to the behaviour of adjustment probabilities, which rise more steeply for wages than for prices, and is ultimately related to the relative absolute size of wage and price changes.⁴⁰

5 Conclusions

We have developed a DSGE model with state-dependent price and wage rigidity, combining monopolistic competition in goods and labour (as in [Erceg et al. 2000](#)) with nominal rigidity due to costly decision-making (as in [Costain and Nakov 2019](#)). Our heterogeneous-agents approach,

38. [Alvarez and Lippi \(2014\)](#) show that there are decreasing returns to monetary stimulus in state-dependent pricing models, and that the peak effect occurs for a money shock that is roughly half the size of firms' idiosyncratic shocks. This is consistent with our findings here: the peak effect in our baseline model is achieved by a 3.2 log points money shock, while the standard deviation of firms' idiosyncratic shocks is 5.7 log points.

39. Figure not shown but available upon request.

40. Recall that wage changes are estimated to be smaller than price changes in absolute value, consistent with the targeted empirical histograms.

with idiosyncratic shocks both to firms and to workers, allows us to fit our model to microdata on price and wage adjustments, but also permits us to calculate the dynamic effects of monetary policy shocks. Our model assumes that labour can be costlessly reallocated across firms at any time, so our study should be understood as documenting the interactions of nominal price stickiness with nominal wage stickiness, abstracting from matching frictions or any other forms of labour specificity. Fitting the data requires convex disutility of labour, so that workers prefer to vary their wages, as well as their hours, in response to shocks.

At a microeconomic level, we compare different calibrations to see how nominal rigidities affect price and wage adjustment behaviour. We estimate the decision cost parameters and productivity processes to match price and wage adjustment microdata; our estimates match the empirical frequency of adjustment, and produce a histogram somewhat smoother than the one observed in the data. If we instead reduce the level of decision costs, the adjustment histograms take on a bimodal shape with strong, counterfactual spikes resembling an (S,s) model. Allowing for a trend in idiosyncratic productivity over the life cycle implies that small negative wage changes are relatively infrequent; this helps explain a pattern which is often interpreted as evidence of downward nominal wage rigidity, in spite of the fact that there is no inherent downward rigidity in our setup.

We quantify the overall output losses due to frictions in our model. In the baseline specification they amount to about 2.35% of the output in a frictionless economy. We further decompose the losses to highlight misallocation components related to inefficient price and wage dispersion, and components that represent the costs of decision-making. Firms in our estimated model spend less than one percent of revenues on decisions related to price setting, while workers devote approximately two percent of their time to decisions about wage setting. Firms' decision costs are comparable to direct evidence, such as that of [Zbaracki et al. \(2004\)](#). The costs of misallocation are non-trivial but are much smaller than a Calvo model would predict. Our results suggest that the level of inefficient dispersion remains quite stable as trend inflation varies.

Regarding macroeconomic dynamics, a key conclusion from our framework is that wage stickiness is a stronger source of monetary non-neutrality than price stickiness. The version of our model with wage stickiness only produces almost as much non-neutrality as the version with wage and price frictions together. In contrast, without wage stickiness, our model has greatly reduced real effects of money shocks, and implies a strong, counterfactual rise in the real wage after a monetary stimulus. The reason for the importance of wages is not so much that they are inherently stickier, but that wage stickiness feeds through to marginal costs and thereby reduces the incentive for price changes after a monetary policy shock.

We then use our calibrated model to explore how trend inflation affects the slope of the Phillips curve. We find that decreased trend inflation makes nominal adjustment and short-run inflation less reactive to monetary shocks, enhancing their real effects and thus flattening the Phillips curve. Quantitatively, our model is able to explain almost half of the observed drop in the slope of the US Phillips curve between 1980-2000 and 2000-2020. In contrast to previous literature, the relationship between the Phillips curve slope and the trend inflation rate

is unrelated to any downward asymmetry in adjustment costs. Instead, in our context it arises naturally as a consequence of state-dependent changes in adjustment frequencies.

Finally, we find that monetary policy has a number of other highly nonlinear effects in our framework. Larger money shocks cause adjustment probabilities to rise, so inflation responds more quickly and the real stimulus is proportionally smaller. Indeed, the absolute size of the cumulative real impact is maximized by a rise of roughly 3.2 log points in the money supply, with a decreasing real stimulus thereafter. Money shocks larger than 9 log points instead have a mostly contractionary real impact in the model.

References

- Akerlof, G. A., W. T. Dickens, G. L. Perry, R. J. Gordon, and N. G. Mankiw (1996). The macroeconomics of low inflation. *Brookings Papers on Economic Activity* 1, 1–76.
- Alvarez, F., M. Gonzalez-Rozada, A. Neumeyer, and M. Beraja (2019). From hyperinflation to stable prices: Argentina’s evidence on menu cost models. *The Quarterly Journal of Economics* 134(1), 451–505.
- Alvarez, F. and F. Lippi (2014). Price setting with menu cost for multiproduct firms. *Econometrica* 82(1), 89–135.
- Alvarez, F. and F. Lippi (2020). Temporary price changes, inflation regimes, and the propagation of monetary shocks. *American Economic Journal: Macroeconomics* 12(1), 104–52.
- Alvarez, F. E., H. Le Bihan, and F. Lippi (2016). The real effects of monetary shocks in sticky price models: A sufficient statistics approach. *American Economic Review* 106(10), 2817–2851.
- Alvarez, F. E., F. Lippi, and L. Paciello (2015). Monetary shocks in models with inattentive producers. *The Review of Economic Studies* 83(2), 421–459.
- Anderson, S. P., A. de Palma, and J.-F. Thisse (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press.
- Andrés, J., Ó. Arce, and P. Burriel (2021). Market polarization and the Phillips curve. Working Paper 2106, Banco de España.
- Ascari, G. (2004). Staggered prices and trend inflation: Some nuisances. *Review of Economic Dynamics* 7(3), 642–667.
- Ascari, G. and T. Haber (2021). Non-linearities, state-dependent prices and the transmission mechanism of monetary policy. *The Economic Journal* (Forthcoming).
- Baley, I. and A. Blanco (2021). Aggregate dynamics in lumpy economies. *Econometrica* (Forthcoming).

- Ball, L. M. and S. Mazumder (2011). Inflation dynamics and the great recession. Working Paper 17044, National Bureau of Economic Research.
- Barnichon, R. and G. Mesters (2021). The Phillips multiplier. *Journal of Monetary Economics* 117, 689–705.
- Basso, H. and O. Rachedi (2021). Robot adoption and inflation dynamics. Manuscript.
- Benigno, P. and L. A. Ricci (2011). The inflation-output trade-off with downward wage rigidities. *American Economic Review* 101(4), 1436–66.
- Blanchard, O. (2016). The Phillips curve: Back to the 60s? *American Economic Review* 106(5), 31–34.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12(3), 383–398.
- Caplin, A. S. and D. F. Spulber (1987). Menu costs and the neutrality of money. *The Quarterly Journal of Economics* 102(4), 703–725.
- Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy* 113(1), 1–45.
- Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2016). Unemployment and business cycles. *Econometrica* 84(4), 1523–1569.
- Coibion, O. and Y. Gorodnichenko (2015). Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics* 7(1), 197–232.
- Costain, J. and A. Nakov (2011). Distributional dynamics under smoothly state-dependent pricing. *Journal of Monetary Economics* 58(6-8), 646–665.
- Costain, J. and A. Nakov (2019). Logit price dynamics. *Journal of Money, Credit and Banking* 51(1), 43–78.
- Costain, J., A. Nakov, and B. Petit (2019). Monetary policy implications of state-dependent prices and wages. Working Paper 2272, European Central Bank.
- Cover, T. and J. Thomas (2006). *Elements of Information Theory, Second edition*. Wiley Inter-science.
- De Veirman, E. (2009). What makes the output–inflation trade-off change? The absence of accelerating deflation in Japan. *Journal of Money, Credit and Banking* 41(6), 1117–1140.
- Dickens, W. T., L. Goette, E. L. Groshen, S. Holden, J. Messina, M. E. Schweitzer, J. Turunen, and M. E. Ward (2007). How wages change: micro evidence from the International Wage Flexibility Project. *Journal of Economic Perspectives* 21(2), 195–214.

- Dotsey, M. and A. Wolman (2020). Investigating nonneutrality in a state-dependent pricing model with firm-level productivity shocks. *International Economic Review* 61(1), 159–188.
- Eichenbaum, M., N. Jaimovich, and S. Rebelo (2011). Reference prices, costs, and nominal rigidities. *American Economic Review* 101(1), 234–62.
- Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of Monetary Economics* 46(2), 281–313.
- Fahr, S. and F. Smets (2010). Downward wage rigidities and optimal monetary policy in a monetary union. *Scandinavian Journal of Economics* 112(4), 812–840.
- Forbes, K. J. (2019). Inflation dynamics: Dead, dormant, or determined abroad? *Brookings Papers on Economic Activity* 2019(Fall), 257–319.
- Gagnon, E. (2009). Price setting under low and high inflation: evidence from Mexico. *The Quarterly Journal of Economics* 124, 1221–1263.
- Galí, J. (2008). *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press.
- Gertler, M. and A. Trigari (2009). Unemployment fluctuations with staggered Nash wage bargaining. *Journal of Political Economy* 117(1), 38–86.
- Golosov, M. and R. E. Lucas Jr (2007). Menu costs and Phillips curves. *Journal of Political Economy* 115(2), 171–199.
- Grigsby, J., E. Hurst, and A. Yildirmaz (2021). Aggregate nominal wage adjustments: New evidence from administrative payroll data. *American Economic Review* 111(2), 428–471.
- Guimaraes, B. and K. D. Sheedy (2011). Sales and monetary policy. *American Economic Review* 101(2), 844–876.
- Heise, S., F. Karahan, and A. Sahin (2020). The missing inflation puzzle: The role of wage-price pass-through. Working Paper 27663, National Bureau of Economic Research.
- Huang, K. X. and Z. Liu (2002). Staggered price-setting, staggered wage-setting, and business cycle persistence. *Journal of Monetary Economics* 49(2), 405–433.
- Jorgensen, P. L. and K. J. Lansing (2019). Anchored inflation expectations and the flatter Phillips curve. Working Paper 2019-27, Federal Reserve Bank of San Francisco.
- Kehoe, P. and V. Midrigan (2015). Prices are sticky after all. *Journal of Monetary Economics* 75, 35–53.
- King, R. G. and M. W. Watson (1994). The post-war U.S. Phillips curve: a revisionist econometric history. *Carnegie-Rochester Conference Series on Public Policy* 41, 157–219.

- Klenow, P. J. and O. Kryvtsov (2008). State-dependent or time-dependent pricing: Does it matter for recent US inflation? *The Quarterly Journal of Economics* 123(3), 863–904.
- Klenow, P. J. and B. A. Malin (2010). Microeconomic evidence on price-setting. In B. M. Friedman and M. Woodford (Eds.), *Handbook of Monetary Economics*, Volume 3, Chapter 6, pp. 231–284. North Holland.
- Kryvtsov, O. and N. Vincent (2014). On the importance of sales for aggregate price flexibility. Working Paper 2014-45, Bank of Canada.
- Lindé, J. and M. Trabandt (2018). Resolving the missing deflation puzzle. Manuscript, Sveriges Riksbank.
- Lombardi, M. J., M. Riggi, and E. Viviano (2020). Bargaining power and the Phillips curve: a micro-macro analysis. Working Paper 1302, Banca d’Italia.
- Luce, R. D. (1959). *Individual Choice Behavior a Theoretical Analysis*. John Wiley and sons.
- Luo, S. and D. Villar (2021). The skewness of the price change distribution: A new touchstone for sticky price models. *Journal of Money, Credit and Banking* 53(1), 41–72.
- Machina, M. (1985). Stochastic choice functions generated from deterministic preferences over lotteries. *The Economic Journal* 95(379), 575–594.
- Mattsson, L.-G. and J. W. Weibull (2002). Probabilistic choice and procedurally bounded rationality. *Games and Economic Behavior* 41(1), 61–78.
- McCallum, A. and F. Smets (2007). Real wages and monetary policy transmission in the euro area. Working Paper 1360, Kiel.
- McLeay, M. and S. Tenreyro (2020). Optimal inflation and the identification of the Phillips curve. *NBER Macroeconomics Annual* 34(1), 199–255.
- Midrigan, V. (2011). Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica* 79(4), 1139–1180.
- Nakamura, E. and J. Steinsson (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics* 123(4), 1415–1464.
- Nakamura, E., J. Steinsson, P. Sun, and D. Villar (2018). The elusive costs of inflation: Price dispersion during the U.S. Great Inflation. *The Quarterly Journal of Economics* 133(4), 1933–1980.
- Nekarda, C. J. and V. A. Ramey (2013). The cyclical behavior of the price-cost markup. Working Paper 19099, National Bureau of Economic Research.
- Nevo, A. and A. Wong (2019). The elasticity of substitution between time and market goods: Evidence from the Great Recession. *International Economic Review* 60(1), 25–51.

- Olivei, G. and S. Tenreyro (2007). The timing of monetary policy shocks. *American Economic Review* 97(3), 636–663.
- Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control* 33(3), 649–665.
- Roberts, J. (2006). Monetary policy and inflation dynamics. *International Journal of Central Banking* 2(3), 193–230.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50(3), 665–690.
- Stahl, D. O. (1990). Entropy control costs and entropic equilibria. *International Journal of Game Theory* 19(2), 129–138.
- Steiner, J., C. Stewart, and F. Matejka (2017). Rational inattention dynamics: Inertia and delay in decision-making. *Econometrica* 85(2), 521–553.
- Takahashi, S. (2017). State dependency in price and wage setting. *International Journal of Central Banking* 13(1), 151–189.
- Taylor, J. B. (1979). Staggered wage setting in a macro model. *American Economic Review, Papers and Proceedings* 69(2), 108–113.
- Van Damme, E. (2002). *Stability and perfection of Nash equilibria, Second edition*. Springer.
- Woodford, M. (2009). Information-constrained state-dependent pricing. *Journal of Monetary Economics* 56, S100–S124.
- Yun, T. (2005). Optimal monetary policy with relative price distortions. *American Economic Review* 95(1), 89–109.
- Zbaracki, M., M. Ritson, D. Levy, S. Dutta, and M. Bergen (2004). Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *The Review of Economics and Statistics* 86(2), 514–533.

Appendix

A Decomposing frictions

As (36) shows, aggregate consumption in our economy is $C_t = \bar{A}_t N_t$, where $N_t = N_t^{tot} - \mu - \tau$ is total effective labour used for production, and \bar{A}_t is the labour productivity measure defined in (37). This general definition of \bar{A}_t is valid for economies with flexible prices and/or flexible wages, as special cases.

A.1 Flexible prices

The first-order condition for a firm without price frictions implies that it will set the price $P_{j,t}^{fp} = \left(\frac{\epsilon}{\epsilon-1} \frac{W_t}{A_{j,t}} \right)$, for any aggregate wage level W_t , given its productivity $A_{j,t}$. So if all firms can flexibly set prices, the aggregate price level will be

$$P_t^{fp} = \frac{\epsilon}{\epsilon-1} \frac{W_t^{fp}}{\bar{A}^{fp}}. \quad (46)$$

Labour productivity, which we will call \bar{A}^{fp} , must satisfy

$$(\bar{A}^{fp})^{-1} = \int \left(\frac{P_{j,t}^{fl}}{P^{fl}} \right)^{-\epsilon} A_{j,t}^{-1} dj = \int \left(\frac{\bar{A}^{fl}}{A_{j,t}} \right)^{-\epsilon} A_{j,t}^{\epsilon-1} dj$$

Rearranging, we see that \bar{A}^{fp} is constant if the productivity distribution is time-invariant:

$$\bar{A}^{fp} = \left(\int A_{j,t}^{\epsilon-1} dj \right)^{1/(\epsilon-1)}, \quad (47)$$

This result holds regardless of the degree of wage-setting frictions.

Given effective labour N_t^{tot} , we conclude that under flexible prices, aggregate consumption will be $C_t^{fp} = \bar{A}^{fp} N_t^{tot}$. So given N_t^{tot} , the losses in aggregate consumption attributable to price stickiness alone are given by (42):

$$\Theta_t^p = \frac{C_t^{fp} - C_t}{C_t} = \left(\frac{\bar{A}^{fp} - \bar{A}_t}{\bar{A}_t} \right) \frac{N_t^{tot}}{N_t} + \frac{\mu_t}{N_t} + \frac{\tau_t}{N_t}$$

The first term in (42) captures the consumption losses due to inefficient price dispersion. The second and third terms capture the output losses from the fact that part of the labour input in the baseline economy must be devoted to decision-making rather than production.

The misallocation term $\left(\frac{\bar{A}^{fp} - \bar{A}_t}{\bar{A}_t} \right) \frac{N_t^{tot}}{N_t}$ combines partial equilibrium losses due to errors with general equilibrium effects. It is therefore also interesting to separate out the partial equilibrium component. In particular, in the table we report the representative firm's losses θ^{p*} from pricing

errors, as a fraction of aggregate consumption. These are:

$$\theta^{p*} = \int_0^1 \left\{ \left(\frac{P_{j,t}^{fp}}{P_t} \right)^{1-\epsilon} - \left(\frac{P_{j,t}}{P_t} \right)^{1-\epsilon} - A_{j,t}^{-1} \frac{W_t}{P_t} \left[\left(\frac{P_{j,t}^{fp}}{P_t} \right)^{-\epsilon} - \left(\frac{P_{j,t}}{P_t} \right)^{-\epsilon} \right] \right\} dj. \quad (48)$$

The error costs θ^{p*} , and the decision costs $N_t^{-1}\mu_t$ and $N_t^{-1}\tau_t$, are all reported in Tables 4 and 5. The sum of these three components is the firm's total loss caused by frictions. In other words, this sum represents the (per capita) gains that would accrue to a single firm that is capable of making decisions without control costs, holding fixed the rest of the economy.

A.2 Flexible wages

In general, in our frictional model, the aggregate relation between raw labour time and effective labour is given by (38). If instead there are no wage setting frictions, then for the same level of labour time H^{tot} we would obtain N_t^{fw} units of effective labour, defined by:

$$H_t^{tot} = N_t^{fw} \left[\int Z_{i,t}^{\epsilon_n - 1} \left(\frac{W_{i,t}^{fw}}{W_t^{fw}} \right)^{-\epsilon_n} di \right] = N_t^{fw} (\bar{Z}_t^{fw})^{-1}. \quad (49)$$

Hence the ratio of effective labour to labour time changes from \bar{Z}_t to \bar{Z}_t^{fw} when wages are flexible. For a given level of H_t^{tot} , the change in effective labour supply directly attributable to wage stickiness is given by (43):

$$\Theta_t^w = \frac{C_t^{fw} - C_t}{C_t} = \frac{N_t^{fw} - N_t^{tot}}{N_t^{tot}} = \frac{\bar{Z}_t^{fw} - \bar{Z}_t}{\bar{Z}_t} + \frac{\mu_t^w}{\underline{H}_t} + \frac{\tau_t^w}{\underline{H}_t}$$

To interpret this decomposition, consider an individual worker with productivity $Z_{i,t}$ who can set the wage flexibly, taking as given any arbitrary aggregate conditions. That worker will choose:

$$W_{i,t}^* = \left[\chi \left(\frac{\epsilon_n}{\epsilon_n - 1} \right) \left(\frac{P_t}{u'(C_t)} \right) \right]^{\frac{1}{1+\zeta\epsilon_n}} \left[W_t^{\epsilon_n} N_t^{tot} Z_{i,t}^{\epsilon_n - 1} \right]^{\frac{\zeta}{1+\zeta\epsilon_n}} \quad (50)$$

It can then be shown that, regardless of aggregate conditions (P_t, N_t^{tot}, C_t) , if no workers are subject to wage frictions, then the individual-to-aggregate wage ratio will be:

$$\frac{W_{i,t}^{fw}}{W_t^{fw}} = Z_{i,t}^{\frac{\zeta(\epsilon_n - 1)}{1+\zeta\epsilon_n}} (Z^\dagger)^{\frac{1}{1+\zeta\epsilon_n}}, \quad (51)$$

where

$$Z^\dagger \equiv \left(\int Z_{i,t}^{(1+\zeta)\left(\frac{\epsilon_n - 1}{1+\zeta\epsilon_n}\right)} \right)^{\frac{1+\zeta\epsilon_n}{\epsilon_n - 1}}. \quad (52)$$

Interestingly, this result holds independently of whether prices are flexible or sticky. Notice that

it implies that \bar{Z}^{fw} is a constant if the productivity distribution is time invariant:

$$\bar{Z}^{fw} = (Z^\dagger)^{\frac{\epsilon_n}{1+\zeta\epsilon_n}} \left(\int Z_{i,t}^{\frac{\epsilon_n-1}{1+\zeta\epsilon_n}} di \right)^{-1}. \quad (53)$$

As in the case of firms, the misallocation term $\frac{\bar{Z}^{fw}-\bar{Z}_t}{\bar{Z}_t}$ combines partial equilibrium and general equilibrium effects. As before, it is interesting to calculate the partial equilibrium component that represents the cost of errors in wage setting (converted from utility units to consumption units). We call these error costs θ^{w*} , and report them in Tables 4 and 5. Summing these error costs with the decision costs $\underline{H}_t^{-1}\mu_t^w$ and $\underline{H}_t^{-1}\tau_t^w$, also reported in the tables, gives the utility gains (per capita, evaluated in consumption units) that would accrue to a single worker capable of making decisions without control costs, holding fixed the rest of the economy.

A.3 Flexible general equilibrium

If neither workers nor firms face any nominal frictions, the steady-state general equilibrium of our economy can be computed analytically. The nominal price level is irrelevant; it can be normalized to $P_t^{fl} = 1$. For any aggregate price level, the nominal wage W_t^{fl} is given by (46). Aggregate consumption is $C_t^{fl} = \bar{A}^{fl}N_t^{fl}$, where \bar{A}^{fl} is given by (47). Integrating $W_{i,t}^{fl}/Z_{i,t}$ across i , where $W_{i,t}^{fl}$ is given by (50), we obtain an equation for the aggregate wage in a flexible economy. Substituting in for C_t^{fl} and W_t^{fl} , the only unknown is the level of effective labour N_t^{fl} in frictionless general equilibrium:

$$\left(\frac{\epsilon-1}{\epsilon} \right) \bar{A}^{fl} P_t^{fl} = \chi \left(\frac{\epsilon_n}{\epsilon_n-1} \right) \frac{P_t^{fl} (N_t^{fl})^\zeta}{(\bar{A}^{fl} N_t^{fl})^{-\gamma}} (Z^\dagger)^{-1}, \quad (54)$$

Inverting (54) we obtain N_t^{fl} , so we can calculate H_t^{fl} from (53):

$$N_t^{fl} = \left[\left(\frac{1}{\chi} \right) \left(\frac{\epsilon-1}{\epsilon} \right) \left(\frac{\epsilon_n-1}{\epsilon_n} \right) (\bar{A}^{fl})^{1-\gamma} Z^\dagger \right]^{\frac{1}{\gamma+\zeta}}, \quad (55)$$

Overall, aggregate consumption in the baseline economy is

$$C_t = \bar{A}_t N_t = \bar{A}_t (\bar{Z}_t H_t - \mu_t - \tau_t) = \bar{A}_t \bar{Z}_t \underline{H}_t = \bar{A}_t \bar{Z}_t H_t^{tot} - \bar{A}_t \bar{Z}_t (\mu_t^w + \tau_t^w) - \bar{A}_t (\mu_t + \tau_t),$$

where $\underline{H}_t \equiv H_t - \bar{Z}_t^{-1}(\mu_t + \tau_t)$. In the frictionless economy,

$$C_t^{fl} = \bar{A}^{fl} N_t^{fl} = \bar{A}^{fl} \bar{Z}_t^{fl} H_t^{fl}.$$

Hence the the overall loss relative to baseline consumption is given by (44):

$$\Theta = \frac{C_t^{fl} - C_t}{C_t} = \frac{\bar{A}_t^{fl} \bar{Z}_t^{fl} H_t^{fl} - \bar{A}_t [\bar{Z}_t (H_t^{tot} - \mu_t^w - \tau_t^w) - \mu_t - \tau_t]}{\bar{A}_t^{fl} \bar{Z}_t^{fl} \underline{H}_t} \quad (56)$$

$$= \frac{\bar{A}_t^{fl} \bar{Z}_t^{fl} - \bar{A}_t \bar{Z}_t}{\bar{A}_t \bar{Z}_t} + \frac{\mu_t^w + \tau_t^w}{\underline{H}_t} + \frac{\mu_t + \tau_t}{N_t} + \frac{\bar{A}_t^{fl} \bar{Z}_t^{fl} (H_t^{fl} - \underline{H}) - \bar{A}_t \bar{Z}_t (H_t^{tot} - \underline{H})}{\bar{A}_t \bar{Z}_t \underline{H}_t} \quad (57)$$

The first term incorporates the misallocation terms seen earlier in Θ^p and Θ^w (plus their interactions), combining partial and general equilibrium effects. The final term includes general equilibrium labour supply effects.

B Detrending

To describe the dynamics of the distributions of firms and workers, it helps to first remove the model's nominal trend. We assume benchmark distributions for nominal prices and wages, $\eta_t^P(\tilde{P})$ and $\eta_t^W(\tilde{W})$, that take the form of time-invariant distributions $\eta^p(\tilde{p})$ and $\eta^w(\tilde{w})$ over *real* prices and wages. This assumption implies that the firms' and workers' decision problems are homogeneous of degree one in nominal prices, so their Bellman equations can be stated in real rather than nominal terms.

Let Ω_t be a nominal aggregate state variable for this economy at time t . This means that there are functions P and W which define the nominal price and wage levels in terms of Ω_t :

$$P_t = P(\Omega_t), \quad (58)$$

$$W_t = W(\Omega_t). \quad (59)$$

We will define real variables by dividing by the aggregate price level, and we will treat all idiosyncratic real variables in logs. The real model requires notation for several real idiosyncratic quantities: $p_{j,t} \equiv \ln P_{j,t} - \ln P(\Omega_t)$, $\tilde{p}_{j,t} \equiv \ln \tilde{P}_{j,t} - \ln P(\Omega_t)$, $a_{j,t} \equiv \ln A_{j,t}$, $w_{i,t} \equiv \ln W_{i,t} - \ln P(\Omega_t)$, $\tilde{w}_{i,t} \equiv \ln \tilde{W}_{i,t} - \ln P(\Omega_t)$, $z_{i,t} \equiv \ln Z_{i,t}$, and $\xi_{i,t} \equiv x(W_{i,t}, Z_{i,t}, \Omega_t)/P(\Omega_t)$.

Assuming time-invariant benchmark distributions of *real* prices and wages places restrictions on the benchmark distributions of nominal variables. For any $\tilde{P} \equiv P(\Omega_t)e^{\tilde{p}}$, we must have $\eta_t^P(\tilde{P}) = \tilde{P}^{-1}\eta^p(\tilde{p})$. Likewise, given $\tilde{W} \equiv P(\Omega_t)e^{\tilde{w}}$, we must have $\eta_t^W(\tilde{W}) = \tilde{W}^{-1}\eta^w(\tilde{w})$.⁴¹

Now let Ξ_t be the real variable constructed by replacing each nominal state variable in Ω_t by its log real counterpart, and likewise replacing any distribution of nominal idiosyncratic state variables in Ω_t by the corresponding distribution of log real states. Then we may conjecture that Ξ_t is a valid real aggregate state variable for the model at time t . If so, there must exist functions m , w , and i that determine the real money supply, the real aggregate wage, and the

41. To see this, when we say that there is an unchanging distribution of \tilde{p} , we mean that $cdf_t^P(\tilde{P}) = cdf^p(\tilde{p})$, evaluated at the point $\tilde{P} = P_t e^{\tilde{p}}$. Using the chain rule, this implies $\frac{\partial cdf_t^P}{\partial P}(\tilde{P})P_t e^{\tilde{p}} = \frac{\partial cdf^p}{\partial p}(\tilde{p})$. Then since $\eta_t^P(\tilde{P}) \equiv \frac{\partial cdf_t^P}{\partial P}(\tilde{P})$ and $\eta^p(\tilde{p}) \equiv \frac{\partial cdf^p}{\partial p}(\tilde{p})$ we obtain $\eta_t^P(\tilde{P}) = \tilde{P}^{-1}\eta^p(\tilde{p})$.

inflation rate in terms of Ξ :

$$m_t \equiv M_t/P(\Omega_t) = m(\Xi_t), \quad (60)$$

$$w_t \equiv W(\Omega_t)/P(\Omega_t) = w(\Xi_t), \quad (61)$$

$$i_t \equiv \ln P(\Omega_t) - \ln P(\Omega_{t-1}) = i(\Xi_t, \Xi_{t-1}). \quad (62)$$

Likewise, aggregate consumption and labor must be functions of the real state, so $c(\Xi_t) = C_t \equiv C(\Omega_t)$ and $n(\Xi_t) = N_t \equiv N(\Omega_t)$, and firm-specific labor demand can be written as

$$h(w, z, \Xi_t) \equiv H(P(\Omega_t)e^w, e^z, \Omega_t) = e^{z(\epsilon_n-1)}n(\Xi_t)w(\Xi_t)^{\epsilon_n}e^{-\epsilon_n w}. \quad (63)$$

Now, given the real state Ξ , the firms' Bellman equations can be expressed in terms of real value functions v and v^e that satisfy the identities

$$v(p, a, \Xi) \equiv \frac{V(P(\Omega)e^p, e^a, \Omega)}{P(\Omega)}, \quad (64)$$

$$v^e(p, a, \Xi) \equiv \frac{V^e(P(\Omega)e^p, e^a, \Omega)}{P(\Omega)} = \beta E \left\{ \frac{u'(c(\Xi_{t+1}))}{u'(c(\Xi_t))} v(p - i_{t+1}, a', \Xi_{t+1}) \middle| a, \Xi_t \right\}. \quad (65)$$

We see in (65) that, absent any nominal price adjustment, a log real price p at time t becomes $p - i_{t+1}$ at time $t + 1$. The real version of Bellman equation (13) is:

$$\begin{aligned} v(p, a, \Xi_t) &= \max_{\lambda, \pi^p(\tilde{p})} \left(e^p - \frac{w(\Xi_t)}{e^a} \right) c(\Xi_t) e^{-\epsilon_p} + (1 - \lambda) v^e(p, a, \Xi_t) + \\ &+ \lambda \int \pi^p(\tilde{p}) v^e(\tilde{p}, a, \Xi_t) d\tilde{p} - \lambda \kappa_f w(\Xi_t) \int \pi^p(\tilde{p}) \ln \left(\frac{\pi^p(\tilde{p})}{\eta^p(\tilde{p})} \right) d\tilde{p} - \\ &- \kappa_f w(\Xi_t) \left[\lambda \ln \left(\frac{\lambda}{\bar{\lambda}} \right) + (1 - \lambda) \ln \left(\frac{1 - \lambda}{1 - \bar{\lambda}} \right) \right] \\ \text{s.t. } &\int \pi^p(\tilde{p}) d\tilde{p} = 1. \end{aligned} \quad (66)$$

The worker's Bellman equation (24) can be detrended in the same manner. To do so, we postulate real value functions l and l^e that satisfy the identities

$$l(w, z, \Xi) \equiv \frac{L(P(\Omega)e^w, e^z, \Omega)}{P(\Omega)}, \quad (67)$$

$$l^e(w, z, \Xi) \equiv \frac{L^e(P(\Omega)e^w, e^z, \Omega)}{P(\Omega)} = \beta E \left\{ \frac{u'(c(\Xi_{t+1}))}{u'(c(\Xi_t))} l(w - i_{t+1}, z', \Xi_{t+1}) \middle| z, \Xi_t \right\}. \quad (68)$$

The worker's real Bellman equation can now be written as follows:

$$l(w, z, \Xi_t) = \max_{\tau^w, \mu^w, \rho, \pi^w(\tilde{w})} e^w h(w, z, \Xi_t) - \frac{X(h(w, z, \Xi_t) + \tau^w + \mu^w)}{u'(c(\Xi_t))} \\ + (1 - \rho) l_t^e(w, z, \Xi_t) + \rho \int \pi^w(\tilde{w}) l^e(\tilde{w}, z, \Xi_t) d\tilde{w} \quad (69)$$

$$\text{s.t.} \quad \int \pi^w(\tilde{w}) d\tilde{w} = 1, \\ \rho \kappa_w \int \pi^w(\tilde{w}) \ln \left(\frac{\pi^w(\tilde{w})}{\eta^w(\tilde{w})} \right) d\tilde{w} = \tau^w, \\ \kappa_w \left[\rho \ln \left(\frac{\rho}{\bar{\rho}} \right) + (1 - \rho) \ln \left(\frac{1 - \rho}{1 - \bar{\rho}} \right) \right] = \mu^w. \quad (70)$$

Equation (70) implies the following distribution of wages:

$$\pi_t^w(\tilde{w}|w, z) \equiv \frac{\eta^w(\tilde{w}) \exp \left(\frac{l_t^e(\tilde{w}, w)}{\kappa_w \xi_t(w, z)} \right)}{\int \eta^w(w') \exp \left(\frac{l_t^e(w', z)}{\kappa_w \xi_t(w, z)} \right) dw'} \quad (71)$$

where

$$\xi_t(w, z) \equiv \frac{X'(h_t(w, z) + \tau_t^w(w, z) + \mu_t^w(w, z))}{u'(C_t)} \quad (72)$$

is the worker's marginal disutility of time spent working, expressed in units of consumption goods. Similarly, the first-order condition for ρ implies the following adjustment probability:

$$\rho_t(w, z) = \frac{\bar{\rho} \exp \left(\frac{\tilde{l}_t(w, z)}{\kappa_w \xi_t(w, z)} \right)}{\bar{\rho} \exp \left(\frac{\tilde{l}_t(w, z)}{\kappa_w \xi_t(w, z)} \right) + (1 - \bar{\rho}) \exp \left(\frac{l_t^e(w, z)}{\kappa_w \xi_t(w, z)} \right)}. \quad (73)$$

Thus, the noise in both the timing and wage-setting decisions is proportional to the worker's marginal disutility of labor.

For purposes of backwards induction, to calculate the worker's decision in any state (w, z, Ξ) , it suffices to find the unique value of $\xi_t(w, z)$ that solves (72). The worker's decision time costs $\mu_t^w(w, z)$ and $\tau_t^w(w, z)$ can be calculated using (71) and (73); their sum is strictly decreasing in ξ . Since marginal disutility increases strictly with total time use (and since $h_t(w, z)$ does not depend on ξ), the right-hand side of (72) is a strictly decreasing function of ξ . Therefore (72) can be solved by bisection to give a unique solution $\xi_t(w, z) \geq 0$ in any state (w, z, Ξ_t) .

C Distributional dynamics

The distributions of firms' and workers' state variables evolve over time as firms and workers respond to idiosyncratic and aggregate shocks. We begin by describing firms' dynamics.

$P_{j,t}$ is the nominal price at which firm j produces in period t , prior to adjustment; this may differ from its price $\tilde{P}_{j,t}$ at the end of t , when price adjustments are realized. But instead of tracking nominal prices, it is simpler to focus on log real prices $p_{j,t}$. Therefore, we define

$\Psi_t(p_{j,t}, a_{j,t})$ as the real distribution at the beginning of t , when production takes place, and $\tilde{\Psi}_t(\tilde{p}_{j,t}, a_{j,t})$ as the real distribution at the end of t . Also, we use lower-case letters to represent the joint densities associated with these distributions, namely $\psi_t(p_{j,t}, a_{j,t})$ and $\tilde{\psi}_t(\tilde{p}_{j,t}, a_{j,t})$, respectively.⁴²

Two stochastic processes drive the dynamics of the distribution. First, there is the Markov process for firm-specific log productivity, which we can write in terms of the following *c.d.f.*:

$$S(a'|a) = \text{prob}(a_{j,t} \leq a' | a_{j,t-1} = a), \quad (74)$$

or in terms of the corresponding density function:

$$s(a'|a) = \frac{\partial}{\partial a'} S(a'|a). \quad (75)$$

Thus, suppose that the density of real prices and log productivities at the end of period $t-1$ is $\tilde{\psi}_{t-1}(\tilde{p}, a)$. Holding fixed a firm's nominal price, its real log price is changed by inflation, from $\tilde{p}_{i,t-1}$ to $p_{i,t} \equiv \tilde{p}_{i,t-1} - i_t$ at the beginning of t . At the same time, productivities a will be shifted by the Markov transitions s . Therefore the density of real log prices and log productivities at the beginning of t is given by

$$\psi_t(\tilde{p} - i_t, a') = \int s(a'|a) \tilde{\psi}_{t-1}(\tilde{p}, a) da, \quad (76)$$

and hence the cumulative distribution at the beginning of t , in real terms, is

$$\Psi_t(p, a') = \int^p \int^{a'} \left(\int s(b|a) \tilde{\psi}_{t-1}(q + i_t, a) da \right) db dq. \quad (77)$$

The definition of the aggregate price level (1) implies that the following identity must hold for distribution $\psi_t(p, a)$ at all times:

$$\int \int e^{(1-\epsilon)p} \psi_t(p, a) da dp = 1. \quad (78)$$

The second stochastic process that determines the dynamics is the process of real price updates, which we have defined in terms of a conditional density of logit form in (8). A firm with real log price p and log productivity a at the beginning of period t adjusts its price with probability $\lambda \left(\frac{d_t(p, a)}{\kappa_f w_t} \right)$, where

$$d_t(p, a) \equiv \tilde{v}_t(a) - v_t^e(p, a).$$

Upon adjustment, its new real log price is distributed according to $\pi_t(\tilde{p}|a)$. Therefore, if the

42. Our notation here assumes that all densities are well-defined on a continuous support, but we do not actually impose this assumption on the model. With slightly more sophisticated notation we could allow explicitly for distributions with mass points, or with discrete support.

density of firms at the beginning of t is $\psi_t(p, a)$, the density at the end of t is given by

$$\tilde{\psi}_t(\tilde{p}, a) = \left(1 - \lambda \left(\frac{d_t(\tilde{p}, a)}{\kappa_f w_t}\right)\right) \psi_t(\tilde{p}, a) + \int \lambda \left(\frac{d_t(p, a)}{\kappa_f w_t}\right) \pi_t(\tilde{p}|a) \psi_t(p, a) dp.$$

The cumulative distribution at the end of t is simply given by integrating up this density:

$$\tilde{\Psi}_t(p, a) = \int^{\tilde{p}} \int^a \tilde{\psi}_t(q, b) db dq.$$

The dynamics of wages and worker productivities are analogous, except that an individual worker may die and be replaced by a new worker with probability $1 - \beta_S$ per period. Let $\Psi_t^w(w_{i,t}, z_{i,t})$ be the distribution of real log prices and log worker productivities at the beginning of the period, when production takes place, and let $\tilde{\Psi}_t^s(\tilde{w}_{i,t}, z_{i,t})$ be the corresponding distribution of *surviving* workers at the end of the period. We write the corresponding densities as $\psi_t^w(w_{i,t}, z_{i,t})$ and $\tilde{\psi}_t^s(\tilde{w}_{i,t}, z_{i,t})$, respectively.

Now, consider a worker with real log wage w and log productivity z at the beginning of period t . She adjusts her wage with probability $\rho \left(\frac{d_t^w(w, z)}{\kappa_w \xi_t(w, z)}\right)$, where

$$d_t^w(w, z) \equiv \tilde{l}_t(w, z) - l_t^e(w, z).$$

Upon adjustment, her new real log wage is distributed according to $\pi_t^w(\tilde{w}|w, z)$. Therefore, if the density of workers at the beginning of t is $\psi_t^w(w, z)$, the density at the end of t is:

$$\tilde{\psi}_t^w(\tilde{w}, z) = \left(1 - \rho \left(\frac{d_t^w(\tilde{w}, z)}{\kappa_w \xi_t(\tilde{w}, z)}\right)\right) \psi_t^w(\tilde{w}, z) + \int \rho \left(\frac{d_t^w(w, z)}{\kappa_w \xi_t(w, z)}\right) \pi_t^w(\tilde{w}|w, z) \psi_t^w(w, z) dw.$$

The cumulative distribution at the end of t integrates up this density:

$$\tilde{\Psi}_t^w(\tilde{w}, z) = \int^{\tilde{w}} \int^z \psi_t^w(q, b) db dq.$$

The aggregate wage definition (20) implies that $\psi_t^w(w, z)$ always satisfies the following identity:

$$\int \int e^{(1-\epsilon_n)(w-z)} \psi_t^w(w, z) dz dw = w_t^{1-\epsilon_n}. \quad (79)$$

A worker alive in period t survives to period $t + 1$ with probability β_S . Her productivity, conditional on survival, is driven by the Markov process S^z :

$$S^z(z'|z) = \text{prob}(z_{i,t+1} \leq z' | z_{i,t} = z), \quad (80)$$

with the following density function:

$$s^z(z'|z) = \frac{\partial}{\partial z'} S(z'|z).$$

Meanwhile, holding fixed a worker's nominal wage, her real log wage is changed by inflation,

from $\tilde{w}_{i,t}$ at the end of t , to $w_{i,t+1} \equiv \tilde{w}_{i,t} - i_{t+1}$. Therefore the density of real log wages and log worker productivities among surviving workers at the beginning of $t + 1$ is:

$$\psi_{t+1}^s(\tilde{w} - i_{t+1}, z') = \int s^z(z'|z) \tilde{\psi}_t^w(\tilde{w}, z) dz. \quad (81)$$

The cumulative distribution at the beginning of t integrates up (81) and adds on the component of new-born workers, with distribution Ψ_t^0 :

$$\Psi_{t+1}^w(w, z) = \beta_S \int^w \int^z \left(\int s^z(b|y) \tilde{\psi}_t^w(q + i_{t+1}, y) dy \right) db dq + (1 - \beta_S) \Psi_{t+1}^0(w, z).$$

Considering birth and death matters here because it permits us to impose an upward trend in productivity over the course of an individual's working life: a worker typically ends her career at a wage higher than the one she started with. This trend is important for matching the distribution of wage adjustments. We denote the distribution of wages and productivity for newborn workers at time t by Ψ_t^0 . For simplicity, we assume that a newborn worker's wage is the one she would choose, conditional on her productivity, if wage setting were frictionless. With this simplifying assumption, we avoid modelling an initial decision-making state at birth. Since our analysis focuses on wage changes, ignoring the level of the initial wage, this assumption has a negligible impact on the results.

D Additional figures

Figure D.1: Histograms of price and wage changes at different trend inflation rates

