Uncovering the costs of high inflation *

Ken Miyahara ⓡ **Alberto Cavallo** ⓡ **Francesco Lippi**

March 20, 2025

Abstract

Inflation gives rise to inefficient price dispersion in New Keynesian models. Yet empirical analyses suggest that such costs are small (see e.g., Nakamura et al. (2018)). We study price dispersion using a new model that enriches the canonical sticky-price model with an information-gathering activity that describes the firm's effort to identify the best price to be chosen upon adjustment (akin to e.g., Reis (2006); Caballero (1989)). We refer to this activity as "price-research", involving for instance a review of the firm's balance sheets to measure production costs.

Firms optimally manage prices allocating resources to price-changing and price-research activities. When inflation is low firms invest heavily in price-research and the reset price is accurately chosen. Instead, high inflation shrinks the firm's markup rapidly and the firm engages in price changes with less precise information on its own idiosyncratic costs. High inflation leads firms to shift resources from pricing-research towards price-adjustment, generating a less efficient distribution of relative prices.

We use two granular datasets, from Turkey and Argentina during 2019-2025, to study the size and frequency of price changes. In each country there is a period of "moderate" two-digit inflation which subsequently doubles. We use the micro pricing patterns to parametrize our structural model and its fundamental frictions. The calibrated model fits several observed pricing patterns better than a simple one where the only friction is due to "sticky prices". The model suggests that the cost of a doubling of inflation, from 30% to 60% per year as observed in Turkey, increases the costs of inflation by about 80 basis points of GDP, which is about 8 times larger than the costs predicted by a canonical sticky-price model with only price-adjustment frictions.

JEL Classification Numbers: E5

Key Words: sticky prices, flexible prices, high inflation, generalized hazard function.

*Cavallo: Harvard University (email: acavallo@hbs.edu); Lippi: Luiss University (email: flippi@luiss.it); Miyahara: University of Chicago (email: kmiyahara@uchicago.edu). We thank, without implicating, several colleagues and coauthors who provided feedback on a preliminary draft of the paper. We benefited from discussions with Fernando Alvarez. Lippi acknowledges financial support from the ERC grant 101054421-DCS.

1 Introduction

Inflation gives rise to inefficient price dispersion in New Keynesian models (see Gali (2008); Woodford (2003)). Yet empirical analyses suggest that the misallocation costs implied by price dispersion are elusive (e.g., Nakamura, Steinsson, Sun, and Villar (2018)). We present a new model to study how inflation affects price dispersion. The model enriches the canonical sticky-price model with a "price-research" activity, an information-gathering process that allows the firm to discover its "optimal" price.

The model thus features two frictions: a canonical price-adjustment friction (as in a menu-cost model) and an "information gathering" friction. We refer to the latter activity as price-research, an activity that is aimed at measuring e.g., the firm's idiosyncratic production costs, in the spirit of Reis (2006); Caballero (1989). The traditional sticky price model is nested as a special case where information is freely available to the firm.

Firms optimally allocate resources between price-adjustment and price-research activities. The latter is essential for identifying the firm's "optimal" price in presence of idiosyncratic shocks, while the former is required for actually changing the posted price. In the absence of aggregate shocks and inflation, the firm never adjust prices without gathering new information (as in Alvarez, Lippi, and Paciello (2011)). Under low inflation, firms invest heavily in price-research and the reset price is accurately chosen.

When inflation is high the firm's markup shrink rapidly, prompting firms shift resources away from price-research into price-adjustments. Since the main driver of price-research investment is the magnitude of idiosyncratic shocks, its level remains relatively stable across inflation regimes. In contrast, the frequency of price adjustment increases markedly. As a result, high-inflation environments lead to more price adjustments being made with relatively less precise information, generating a less efficient price system. The model thus features a new channel through which high inflation gives rise to welfare costs: by inducing firms to gather less information, inflation leads to a less efficient price system.

To quantify these effects we use two granular datasets from Turkey and Argentina during

2019-2025 to study price-setting behavior before and after inflation doubles from moderate two-digit levels. We document key moments such as the variance of "inflation-centered" price changes and the elasticity of the adjustment frequency to inflation, which we use to discipline our model parameters. A key new result, generalizing a well-known formula in simpler models, shows that the variance of "inflation-centered" price changes is revealing of both the size of idiosyncratic shocks and the magnitude of information frictions in the economy.

Results. We use the calibrated model to quantify the misallocation costs of inflation. In a model experiment, we estimate the welfare costs of an inflation increase from 30% to 60% in Turkey. Our model predicts a misallocation cost of inflation of 81 basis points of GDP, compared to just 12 basis points in a standard sticky-price model. These results highlight the significant role of information frictions in the welfare costs of inflation.

2 Model

This section describes the model for preferences, production technology and price setting behavior.

Households. The household preferences are given by

$$\int_0^\infty e^{-\rho t} U\left(C(t), H(t), M(t)\right) dt \quad \text{where} \quad U(C, H, M) \equiv \frac{C^{1-\epsilon}}{1-\epsilon} - \alpha H + \log\left(\frac{M}{P}\right), \tag{1}$$

where ρ is the discount rate, ϵ the reciprocal elasticity of intertemporal substitution, α the disutility of labor, $C(t)$ is a consumption composite, $H(t)$ the labor supply and $M(t)/P(t)$ real money balances at time t. Let $c_i(t)$ denote the consumption of good i. The composite $C(t)$ is

$$C(t) \equiv \left[\int_0^1 \left(A_i(t) c_i(t)\right)^{\frac{\eta-1}{\eta}} di\right]^{\frac{\eta}{\eta-1}} \tag{2}$$

where $A_i(t)$ denote preference shocks and η is a parameter determining the price elasticity of demand. Labor supply equals the aggregate labor hired by firms, $H(t) = \int_0^1 h_i(t)di$. Households maximize utility subject to the budget constraint, choosing consumption for each good i, labor, and real balances for each t (see Appendix **??**).

Firms. The economy is populated by a unit mass of firms, indexed by $i \in [0, 1]$. Each firm produces one good, $c_i(t)$, using labor (h_i), according to

$$c_i(t) = \frac{h_i(t)}{Z_i(t)} \tag{3}$$

The marginal cost of production of firm i is: $\text{mc}_i(t) = W(t)Z_i(t)$ where $W(t)$ is the nominal wage and $Z_i(t)$ is a firm-specific productivity level. We assume that $Z_i(t) = \exp(\sigma\, z_i(t))$, where $\{z_i\}$ are standard Brownian motions, independent across i, with standard deviation parameter σ. We assume that $A_i(t) = Z_i(t)$ so the (log of) marginal cost and the preference shock are perfectly correlated.[1]

The *target price* is the price that maximizes current profits and is given by the marginal cost times a constant markup:

$$p_i^*(t) = \frac{\eta}{\eta - 1}\text{mc}_i(t) \ .$$

In presence of price adjustment frictions, the firm keeps track of its price gap $x_i(t)$, the wedge between the actual price $p_i(t)$ and the target price $p_i^*(t)$:

$$x_i(t) \equiv \log p_i(t) - \log p_i^*(t) \ . \tag{4}$$

Absent frictions the gap is identically zero, i.e. each firm charges the target price $p_i(t) = p_i^*(t)$. If the price is not adjusted, the price gap evolves due to trend inflation (increasing nominal

[1]This assumption, also used in Woodford (2009); Bonomo et al. (2010); Midrigan (2011); Alvarez and Lippi (2022), allows the problem to be described by a scalar stationary state variable, the price gap x. This is used to write the dynamic programming problem of the firm as well as to keep the expenditure shares stationary across goods in the presence of permanent idiosyncratic shocks.

costs) and idiosyncratic productivity shocks. As long as the nominal price is kept constant the evolution of the price gap obeys

$$dx_i(t) = -\mu\, dt + \sigma\, dz_i(t), \tag{5}$$

where z_i is a standard Brownian motion describing productivity shocks and μ is the growth rate of nominal wages per unit of time, i.e. the inflation rate.[2]

Adjustment costs. In every moment, the firm chooses an effort rate α for price resetting, so that with probability α per unit of time, the nominal price $p_i(t)$ is reset. The optimal choice for α balances the benefit of controlling its price gap with the cost of effort, which is given by a convex function.

Information costs. A firm does not know its price gap, $x(t)$, at every moment because it does not observe $p_i^*(t)$ or the corresponding evolution of the idiosyncratic productivity shock $z_i(t)$ at each t. At every instant, the firm chooses an effort rate of pricing-research ω, such that with probability ω per unit of time, it discovers the target price $p_i^*(t)$ and, consequently, its price gap $x(t)$. The optimal choice for pricing-research effort ω balances the benefits of eliminating uncertainty about the price gap against the cost of effort, which is given by a convex function.

Beliefs. A firm that last discovered a price gap of x_0 forms (rational) beliefs about its price gap after τ time units have elapsed without a new pricing-research or adjustment. Specifically, it believes that $x(\tau) \sim \Phi(x_0 - \mu\,\tau, \sigma^2\,\tau)$. where Φ is the cumulative distribution function (CDF) of the Normal distribution. The state of the firm is summarized by the (mean) *predicted price gap* $\bar{x}$ and the time τ elapsed since the last discovery of its price gap, which we refer to as the *research age*. The state $(\bar{x}, \tau)$ allows the firm to form beliefs about

[2]This is rationalized by assuming the money supply grows at an exogenous rate μ.

its price gap at each instant such that

$$x \sim \Phi(\bar{x}, \sigma^2 \tau).$$

Absent adjustment or pricing-research opportunities, the state evolves according to

$$\begin{aligned} d\bar{x}(t) &= -\mu\, dt, \\ d\tau(t) &= 1\, dt. \end{aligned}$$

Thus, the predicted gap $\bar{x}$ drifts downward at the inflation rate while the research age τ increases one-to-one with calendar time.

The price-setting problem. The firm's value function $v(\bar{x}, \tau)$ solves the following problem:

$$\begin{aligned} \rho\, v(\bar{x}, \tau) = \mathbb{E}\,[F(x)|(\bar{x}, \tau)] - \mu \partial_{\bar{x}} v(\bar{x}, \tau) + \partial_\tau v(\bar{x}, \tau) \\ + \min_{\alpha \geq 0} \{\alpha \cdot [v(\bar{x}^*(\tau), \tau) - v(\bar{x}, \tau)] + (\kappa_a\, \alpha)^{\gamma_a}\} \\ + \min_{\omega \geq 0} \{\omega \cdot [\mathbb{E}\,[v(x, 0)|(\bar{x}, \tau)] - v(\bar{x}, \tau)] + (\kappa_o\, \alpha)^{\gamma_o}\} \end{aligned} \tag{6}$$

where the flow cost function $F(x) = 1 - \eta \left(e^x - \frac{\eta - 1}{\eta}\right) e^{-\eta x}$ represents forgone profits measured as a fraction of maximum (static) profits. The expectation $\mathbb{E}$ is taken with respect to firm's beliefs, i.e. $x \sim \Phi(\bar{x}, \sigma^2 \tau)$. The variable $\bar{x}^*(\tau)$ is the optimal predicted price gap that is chosen upon adjustment, satisfying $\partial_{\bar{x}} v(\bar{x}^*(\tau), \tau) = 0$. Because of the different nature of the problem for firms with old or recent pricing-research, the reset price choice will in general depend on the research age τ.

From the recursive formulation we can see that the benefit of adjusting prices, given by $v(\bar{x}^*(\tau), \tau) - v(\bar{x}, \tau)$, is balanced with the effort cost of price adjustment. Correspondingly, the benefit of eliminating uncertainty around the price gap prediction, given by $\mathbb{E}\,[v(x, 0)|(\bar{x}, \tau)] -$

$v(\bar{x}, \tau)$, is balanced with the effort costs of pricing-research.

Similar to the setting in Cavallo et al. (2024), $F(x)$ is asymmetric around $x = 0$, which implies that a firm's incentive to adjust prices is larger when its price is "low", a negative x, than when it is "high". This is a standard feature of the CES demand system.

The price-setting behavior implied by equation (6) is summarized by two *generalized hazard functions*, $\alpha, \omega : \mathbb{R} \times [0, \infty) \to \mathbb{R}_+$. The function α gives the probability (per unit of time) that a firm with price gap and observation age pair $(\bar{x}, \tau)$ will change its price. The function ω gives the probability that a firm with state $(\bar{x}, \tau)$ will discover its price gap and reset their research age to zero.

The first order conditions for α and ω give the generalized hazard functions $\alpha(\bar{x}, \tau)$, $\omega(\bar{x}, \tau)$:

$$\alpha(\bar{x}, \tau) = \frac{1}{\kappa_a} \left(\frac{v(\bar{x}^*(\tau), \tau) - v(\bar{x}, \tau)}{\kappa_a \, \gamma_a} \right)^{\frac{1}{\gamma_a - 1}}, \quad \text{for each } (\bar{x}, \tau) \in \mathbb{R} \times [0, \infty) \quad , \tag{7}$$

$$\omega(\bar{x}, \tau) = \frac{1}{\kappa_o} \left(\frac{\mathbb{E}\left[v(x, 0) | (\bar{x}, \tau) \right] - v(\bar{x}, \tau)}{\kappa_o \, \gamma_o} \right)^{\frac{1}{\gamma_o - 1}}, \quad \text{for each } (\bar{x}, \tau) \in \mathbb{R} \times [0, \infty) \quad . \tag{8}$$

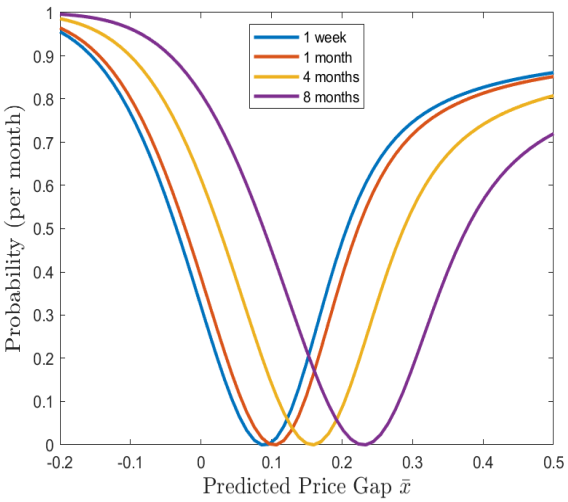
These policy rules imply that price changes and discoveries occur probabilistically, with an intensity that depends on the benefits of each activity and the parameters of the effort cost functions. Hence these probabilities depend on the predicted gap $\bar{x}$ and the research age τ.

Policy functions. Figure 1 depicts the policy function of the firm for a 30% inflation rate (see calibration inTable 2. The adjustment hazard in panel (a) depicts state-dependent behavior. For a given the research age, the firm increases its hazard rate of adjustment when the gap between its predicted price gap is far from its optimal value $\bar{x}^*(\tau)$ and more so for negative gaps. Both Figure 1 and Figure 2 show that firms with older research have higher optimal markups $\bar{x}^*$. This is a tell-tale of firm's uncertainty. These firms are aware of the possibility of having large negative gaps and set a higher markup to protect against this possibility.

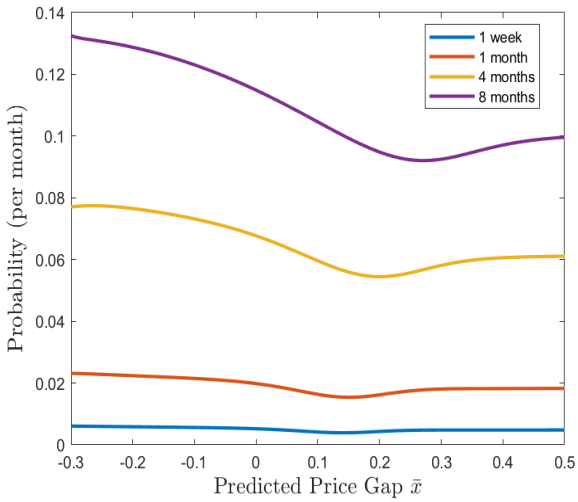
The discovery hazard as a function of predicted price gaps in panel (b) Figure 1 illustrates

Figure 1: Policy Functions for $\mu = 0.3$

(a) Hazard of adjustment α

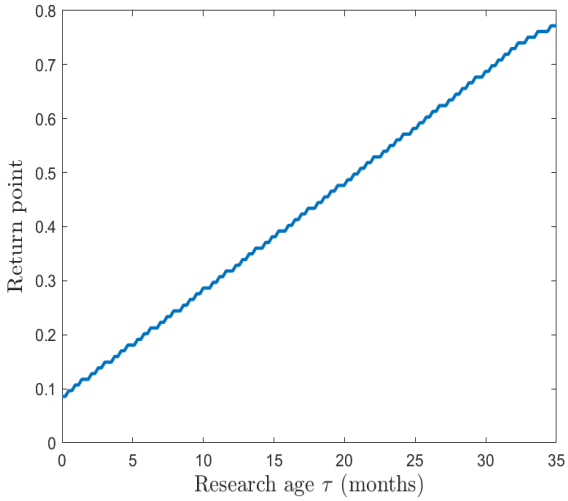


(b) Hazard of target-price discovery ω



Note: Adjustment and discovery hazards for different research ages. The y-axis depicts the probability per month computed as $1 - \exp(\alpha(\bar{x}, \tau)/12)$ for the adjustment hazard and similarly for the discovery hazard. The calibration corresponds to $\mu = .3$, $\sigma = 20\%$, an average frequency of price changes of 3.7 price changes per year and 1 target-price discovery per year (see Table 2). The curvature parameters of the cost functions are quadratic $\gamma_a = \gamma_o = 2$.

Figure 2: Return Point $\bar{x}^*(\tau)$



Note: The calibration corresponds to $\mu = .3$, $\sigma = 20\%$, an average frequency of price changes of 3.7 price changes per year and 1 target-price discovery per year (see Table 2). The curvature parameters of the cost functions are quadratic $\gamma_a = \gamma_o = 2$.

the benefits of observation and eliminating uncertainty about the price gap. For a given the predicted price gap, the discovery hazard is higher the older the research, revealing an increased benefit of reducing uncertainty the longer the time elapsed since the last target-price discovery. For a given research age τ, the hazard of discovery increases with the size of the predicted gap but a flatter slope than the hazard of adjustment. This is clear when noticing that the pricing-research benefit, $\mathbb{E}\left[v(x,0)|(\bar{x},\tau)\right] - v(\bar{x},\tau)$, does not alter the expected value of the predicted gap. For example, suppose that v was linear in its first argument, then the benefit of pricing-research will always be zero because $\mathbb{E}\left[x|(\bar{x},\tau)\right] = \bar{x}$. In general, the value of pricing-research is less sensitive to $\bar{x}$ than the value of price-adjustment.

Aggregation. The hazards α and ω together with the law of motion of $(\bar{x},\tau)$ determine the steady-state distribution of the firms' predicted price gaps and research ages, with density function $\bar{m}:\mathbb{R}\to\mathbb{R}_+$ solving the Kolmogorov forward equation:

$$\left[\alpha(\bar{x},\tau)+\omega(\bar{x},\tau)\right]\cdot\bar{m}(\bar{x},\tau) = \mu\,\partial_{\bar{x}}\bar{m}(\bar{x},\tau) - \partial_{\tau}\bar{m}(\bar{x},\tau), \quad \text{for each } (\bar{x},\tau)\notin\mathcal{R}. \tag{9}$$

where $\mathcal{R} = \{(\bar{x},\tau) : \bar{x}\in \text{Ran }\bar{x}^* \text{ or } \tau = 0\}$ denotes the set of re-injection states. The left hand side counts outflows from the state and the right hand side, inflows. Equation (9) thus indicates that the steady-state distribution satisfies an outflows-inflows balance condition. The density satisfies boundary conditions $\lim_{x\downarrow x^*} m(x) = \lim_{x\uparrow x^*} m(x)$, $1 = \int_{-\infty}^{\infty} m(x)dx$, and $\lim_{x\to\pm\infty} m(x) = 0$.

The steady-state density of (actual) price gaps and research ages, m, the marginal density of price gaps, m_x, and the marginal density of research ages m_τ are obtained using firm's beliefs

$$m(x,\tau) = \int_{-\infty}^{\infty} \phi\left(x|\bar{x},\sigma^2\tau\right)\, \bar{m}(\bar{x},\tau)\, d\bar{x}, \quad \text{all } (x,\tau), \tag{10}$$

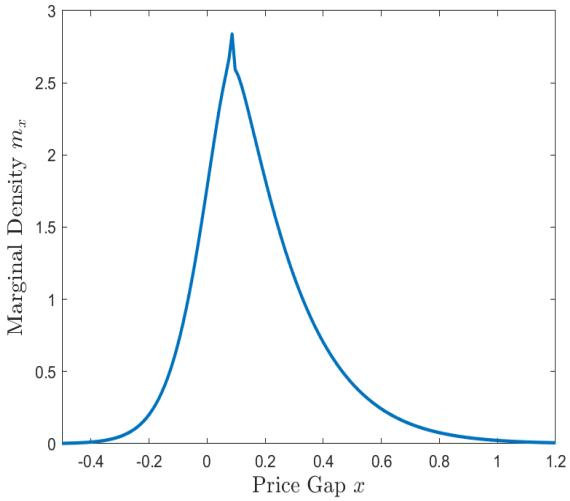
$$m_x(x) = \int_{0}^{\infty} m(x,\tau)\, d\tau, \quad \text{all } x. \tag{11}$$

$$m_\tau(\tau) = \int_{-\infty}^{\infty} m(x,\tau)\, dx, \quad \text{all } \tau. \tag{12}$$

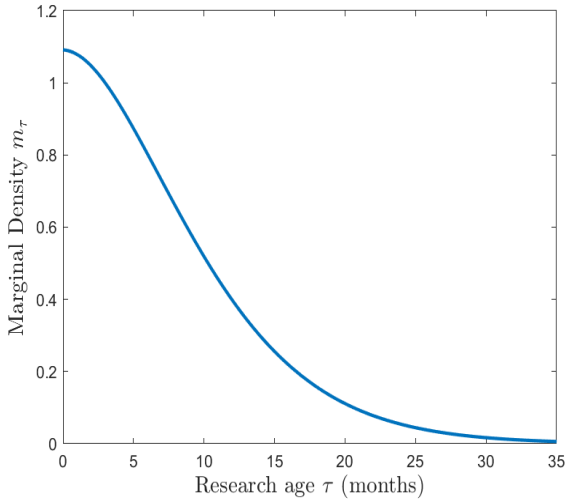
The steady-state marginal density functions are pictured in Figure 3.

Figure 3: Steady-state densities with $\mu = 0.30$

(a) Distribution of price gaps m_x



(b) Distribution of research ages m_τ



Note: The calibration corresponds to $\mu = .3$, $\sigma = 20\%$, an average frequency of price changes of 3.7 price changes per year and 1 target-price discovery per year (see Table 2). The curvature parameters of the cost functions are quadratic $\gamma_a = \gamma_o = 2$.

3 Mapping the model to the data

This section develops the model's prediction for moments of the distribution of the size of price adjustments and the time elapsed between adjustments. These mappings between the data and the fundamental parameters of the model will be used in Section 4 to calibrate the model and analyze the implications of a shift to a high steady state inflation.

Observables and mappings. The frequency of price adjustments N_a is

$$N_a = \int_0^\infty \int_{-\infty}^\infty \alpha(\bar{x}, \tau)\, \bar{m}(\bar{x}, \tau)\, d\bar{x}\, d\tau. \tag{13}$$

The formula counts the number of adjustments per unit of time by averaging the rate at which a firm with state $(\bar{x}, \tau)$ adjusts with the density of firms across predicted gaps and research ages. Analogously, the frequency of discoveries N_o is given by

$$N_o = \int_0^\infty \int_{-\infty}^\infty \omega(\bar{x}, \tau)\, \bar{m}(\bar{x}, \tau)\, d\bar{x}\, d\tau. \tag{14}$$

A desired price change becomes a price change $\Delta p \equiv \bar{x}^*(\tau) - \bar{x}$ at rate $\alpha(\bar{x}, \tau)$. Therefore, the distribution of price changes has the following density $q : \mathbb{R} \to \mathbb{R}_+$

$$q(\Delta p) = \int_{\Psi(\Delta p)} \alpha(\bar{x}, \tau)\, \bar{m}(\bar{x}, \tau)\, d\bar{x}\, d\tau. \tag{15}$$

with $\Psi(\Delta p) = \{(\bar{x}, \tau) \,:\, \bar{x}^*(\tau) - \bar{x} = \Delta p\}$.

Let t^a denote the time elapsed between adjustments. To obtain the joint distribution of the sizes of price changes and time elapsed between adjustments, $\{\Delta p, t^a\}$, notice that given an initial state $(\bar{x}^*(\tau_0), \tau_0)$, immediately after repricing, one can compute the conditional density, $P_0(\Delta p, \tau^a, t^a)$, of size of price changes Δp, research age at adjustment dates τ^a, and time elapsed between adjustments t^a, where the subscript in P denotes the dependence on the initial condition. In particular, the size of the price change Δp will be

$$\Delta\, p = \bar{x}^*(\tau^a) - \bar{x}^*(\tau_0) + \mu\, t^a - \sigma\, Z(\tau_0 + t^a - \tau^a). \tag{16}$$

where τ_0 is given by the initial condition and the $\sigma\, Z$ term is a Brownian increment with variance $\sigma^2 \cdot (\tau_0 + t^a - \tau^a)$. The right hand side contains three terms, let us explain each in order. The first term $\bar{x}^*(\tau^a) - \bar{x}^*(\tau_0)$ accounts for the change in optimal markup relative to the last price change. This term arises because the observation ages at dates of adjustments

can be different for consecutive price spells. The second term corrects the markup erosion due to accumulated inflation between adjustments. The third term adjusts any idiosyncratic shocks that were discovered between adjustments. For example, if a discovery happens at the same time of adjustment, $\tau^a = 0$, then the firm learns all idiosyncratic shocks that transpired between the last discovery and the last adjustment, τ_0, plus all shocks that happened between adjustments, t^a. Instead, if the research age at the new adjustment date is $\tau^a = \tau_0 + t^a$ then no information was learned since the last adjustment and the idiosyncratic adjustment is zero.

Now, we will use moments of the joint distribution $\{\Delta p, t^a\}$ to learn about the parameters of the model through the lens of equation (16).

PROPOSITION 1. Let $\tilde{\mathrm{E}}$ denote expectations with respect to P_0 where the initial condition is the steady-state distribution of states upon adjustment , $m_0(\bar{x}, \tau) = \frac{1}{N_a} \int_{-\infty}^{\infty} \alpha(\bar{x}', \tau)\, \bar{m}(\bar{x}', \tau)\, d\bar{x}'$ for $\bar{x} = \bar{x}^*(\tau)$ and zero otherwise. Then the following formulas hold

$$N_a = \frac{1}{\tilde{\mathrm{E}}\, t^a} \tag{17}$$

$$N_a\, \tilde{\mathrm{E}}\, \Delta p = \mu, \tag{18}$$

$$N_a\, \tilde{\mathrm{E}}\, (\Delta p - \mu\, t^a)^2 = N_a\, \tilde{\mathrm{E}}\, (\bar{x}^*(\tau^a) - \bar{x}^*(\tau_0))^2 + \sigma^2. \tag{19}$$

where Δp is the size of price changes, τ^a, the observation age at adjustment dates, t^a, the time elapsed between adjustments and $\tau_0 \sim m_0$.

Equation (17) expresses the frequency of price changes as the reciprocal of the expected time between adjustments. Equation (18) provides a formula for the drift in price gaps, incorporating both the frequency of adjustments and the average price change. Equation (19) presents a new result relating the squared size of price changes–beyond keeping up with inflation–to the magnitude of idiosyncratic innovations σ^2.

Our new Equation (19) provides a generalization to a well known formula in the price-setting literature. In the case without drift and no information costs, this standard result

states that $N_a \operatorname{Var}(\Delta x) = \sigma^2$.[3] In contrast, our new result introduces two key deviations. First, the variance term is replaced by a moment of the joint distribution of price changes and time between adjustments, $\tilde{\mathrm{E}}\,(\Delta p - \mu\, t^a)^2$. This term incorporates t^a, ensuring that only information about markup differences and idiosyncratic shocks is retained on the right-hand side (see also equation (16)). Second, the σ^2 term is augmented by the squared difference in optimal markups, accounting for variations in research ages at consecutive price changes.[4] As a result, the left-hand side of equation (19) no longer estimates the size of shocks σ^2 but rather an upper bound on it. The extent to which this upper bound exceeds σ^2 depends on the dispersion of consecutive markups, $\bar{x}^*$, and the magnitude of the adjustment frequency. In the special case of zero drift and no information frictions, one obtains constant return points $\bar{x}^*$ and the standard formula holds.

Another key observable that provides insight about fundamental parameters is the elasticity of the frequency of adjustments to inflation. To understand why, consider the case of a firm without observation costs and fixed adjustment costs.[5] Alvarez et al. (2019) show that when the drift component μ is large relative to the shock size σ^2 ($\mu/\sigma^2 \to \infty$), that model converges to the framework of Sheshinski and Weiss (1977), where the elasticity of the adjustment frequency to inflation is 2/3. Conversely, when σ^2 is large relative to μ ($\mu/\sigma^2 \to 0$), this elasticity approaches zero. Our analysis focuses on intermediate cases, where both the inflation rate and idiosyncratic shocks are sizable. In these cases, the elasticity of frequency to inflation is positively related to the ratio μ/σ^2, which we leverage to guide our calibration.

Unobservables. Methods to recover certain model moments from data, including the frequency of observations and the joint distribution of consecutive research ages, are not available. However, equation (19) and the elasticity of frequency to inflation provide two key relationships that link the ratio μ/σ^2 to observable variables. We leverage this logic in Sec-

[3] In the case with drift but no information cost the first term in the right hand side vanishes (see RefBaileyBlanco).

[4] Since τ^a and τ_0 are potentially correlated, the expression does not further simplify.

[5] And a symmetric flow return function.

tion 4 to calibrate the extent of information frictions in the economy.

	2019-2021			2022-2024			
	μ	N_a	$N_a \cdot E(\Delta p - \mu\, t^a)^2$	μ	N_a	$N_a \cdot E(\Delta p - \mu\, t^a)^2$	Elasticity
Turkey	0.30	3.7	0.10	0.60	5.3	0.15	0.43
Argentina	0.46	7.8	0.25	0.108	10.3	0.43	0.32

Table 1: Price Setting Behavior in Turkey and Argentina

Note: Frequency units are number of price changes per year. Average moments for Food and Beverages sectors in Argentina and Turkey after removing COICOP-3 sectors one standard deviation away from the mean. In Turkey 113, 116 and 117 were removed while in Argentina 113 and 117.

The Dataset.

Calibration. While identification results exist for the standard "sticky-price" model (see Cavallo et al. (2024)), there are not off-the-shelf, one-to-one mappings between observables and model parameters in our framework. To address this, we pursue the following strategy.

We impose a standard quadratic effort cost for adjustment and pricing-research activities, $\gamma_a = \gamma_o = 2$. We then jointly use equation (17), equation (18), and equation (19), along with the elasticity of frequency to inflation between our moderate and high inflation episodes. This allows us to estimate N_a and μ and calibrate the level of effort costs of adjustment and observation, κ_a, κ_o, as well as the magnitude of idiosyncratic shocks, σ^2.

We present the steady-state calibration for the moderate inflation regime in Turkey corresponding to the 2019-2021 period

	Matched moments			Parameters			Others	
	μ	N_a	$N_a \cdot E(\Delta p - \mu\, t^a)^2$	σ	κ_a	κ_o	N_o	Elasticity
Baseline Model	0.30	3.7	0.10	0.20	0.03	0.25	1	0.21
Only adj. cost	0.30	3.7	0.10	0.31	0.04	-	-	0.07

Table 2: Matched moments, calibrated parameters, and unobservables for **Turkey**

Note: As described in the text the curvature parameters for the effort cost functions are quadratic.

For the calibration, we match three key moments: the frequency of price changes, the inflation rate, and the variance of "inflation-centered" price changes. Notably, our model better captures the sizable elasticity of price change frequency to inflation, yielding an elasticity of 0.21 compared to 0.07 in the standard "sticky-price" model. This difference arises from the calibration of the idiosyncratic shock parameter σ. In the standard model, all observed variance is attributed to idiosyncratic shocks, and larger idiosyncratic shocks reduce the elasticity of price adjustments to inflation (recall Alvarez et al. (2019)). In our model, information frictions introduce a wedge between the variance term in equation (19) and σ^2 driven by heterogeneity in consecutive reset prices $\bar{x}^*(\tau)$.

While we could calibrate the model to match the empirical elasticity of price change frequency to inflation (0.43, as reported in Table 1), doing so would require a much higher ratio of μ/σ^2. This, in turn, would require a lower variance of the idiosyncratic shocks, σ^2, and significantly larger information frictions. However, such a calibration would imply an extremely low rate of price-research (about one "review" every 2 years [TO BE CHECKED], a prediction that is inconsistent with the empirical evidence on the frequency of pricing-research, see e.g., Alvarez et al. (2011).

4 Welfare Costs of High Inflation

We consider an economy in steady state with an inflation rate μ and perform a comparative statics exercise in which the inflation rate is doubled to $2\,\mu$. We then compare the resulting steady-state outcomes. In this experiment, inflation is the only fundamental parameter that changes, ensuring that any differences in behavior and economic efficiency can be attributed to the increase in inflation. We use the calibration to Turkish firm's price setting behavior during the moderate 30% inflation period and only increase the inflation rate tp 60%.

In order to study the welfare costs of inflation we must define what we mean by welfare costs. We focus on two inefficiencies that arise in the new Keynesian framework due to the

sticky-price friction. The first is due to price dispersion: sticky prices introduce a wedge between the marginal rate of substitution and the marginal rate of transformation for consumers and workers. These wedges give rise to an inefficient allocation of resources. We call this the welfare cost of "misallocation" and denote it by χ. The second one arises when firms use resources in price adjustments and pricing-research. These correspond to a deadweight loss, akin to the shoe-leather cost of inflation in textbook models of money demand. We call this the welfare cost of "price management" and denote it by ϕ. It could be further decomposed in costs of adjustment and research, ϕ_{adj} and ϕ_{res}, correspondingly.

As explained above, the first component of the welfare costs is due to wedges introduced by price gaps. Namely, it is related to price gap dispersion. These welfare costs also scale with the parameter of elasticity of demand. Intuitively, this is due to a larger resource-shifting reaction caused by a more sensitive demand. It can be shown that the welfare cost of misallocation denoted by χ is given by

$$\chi(\mu) \equiv \frac{U - \bar{U}}{U_c \cdot \bar{C}} \approx \frac{\eta}{2}\text{Var}(x),$$

The cost χ is expressed in terms of GDP percentage loss relative to the efficient GDP level. The μ argument emphasizes that the variable generally depends on inflation. Notice that the expression above, up to first order, is the one used in traditional monetary policy textbook analyses.

The cost of price-management ϕ is given by

$$\phi(\mu) \equiv \mathbb{E}\left[\,(\kappa_o\omega)^{\gamma_o} + (\kappa_a\alpha)^{\gamma_a}\right],$$

where ϕ_{adj} and ϕ_{res} are defined in the obvious ways. Again, the argument μ emphasizes that the variable generally depends on inflation.

Experiment. We estimate the welfare costs of an inflation increase from 30% to 60% in Turkey. Our model predicts a misallocation cost of inflation of 81 basis points of GDP, compared to just 12 basis points in a standard sticky-price model. The measures are obtained as level differences between the steady state with 60% and with 30% inflation (e.g. $\Delta\chi = \chi(0.60) - \chi(0.30)$)

In contrast, the price-management costs of inflation are equal in both models. In a further decomposition, the response of the costs of price adjustments are similar in the two models, whereas the decrease in pricing-research investments signifies a small welfare gain.

	$\Delta \operatorname{Var} x$	$\Delta\chi$	$\Delta\,\phi_{\text{adj}}$	$\Delta\,\phi_{\text{res}}$
Baseline Model	0.0027	0.0081	0.0028	-0.0002
Only adj. cost	0.0004	0.0012	0.0026	-

Table 3: Welfare Costs of 30% to 60% inflation

Note: The second column is obtained from the first one multiplying times $\eta/2$ where η is the elasticity of substitution. For the calibration see Table 2.

References

Alvarez, F., M. Beraja, M. Gonzalez-Rozada, and P. A. Neumeyer (2019). From hyperinflation to stable prices: Argentinas evidence on menu cost modles. *The Quarterly Journal of Economics 143*(1), 451–505.

Alvarez, F. and F. Lippi (2022). The analytic theory of a monetary shock. *Econometrica 90*(4), 1655–1680.

Alvarez, F. E., F. Lippi, and L. Paciello (2011). Optimal price setting with observation and menu costs. *The Quarterly Journal of Economics 126*(4), 1909–1960.

Bonomo, M., C. Carvalho, and R. Garcia (2010). State-dependent pricing under infrequent information: a unified framework. Staff Reports 455, Federal Reserve Bank of New York.

Caballero, R. J. (1989). Time dependent rules, aggregate stickiness and information externalities. Discussion Papers 198911, Columbia University.

Cavallo, A., F. Lippi, and K. Miyahara (2024, December). Large shocks travel fast. *American Economic Review: Insights 6*(4), 558–74.

Gali, J. (2008). *Monetary policy, inflation, and the business cycle : an introduction to the New Keynesian framework*. Princeton University Press.

Midrigan, V. (2011, July). Menu costs, multi-product firms, and aggregate fluctuations. *Econometrica, 79*(4), 1139–1180.

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.

Reis, R. (2006). Inattentive producers. *Review of Economic Studies 73*(3), 793–821.

Sheshinski, E. and Y. Weiss (1977, June). Inflation and costs of price adjustment. *Review of Economic Studies 44*(2), 287–303.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Woodford, M. (2009, October). Information-constrained state-dependent pricing. *Journal of Monetary Economics 56*, s100–s124.