

A Theory of Supply Function Choice and Aggregate Supply

Joel P. Flynn*
Yale

George Nikolakoudis†
Princeton

Karthik A. Sastry‡
Princeton

October 4, 2023

Abstract

Many modern theories of the business cycle generate demand-driven fluctuations by assuming that monopolistic firms set a price in advance and commit to supplying the market-clearing quantity. In this paper, we enrich firms' supply decisions by allowing them to choose any *supply function*: a description of the price charged at each quantity of production. By changing only the strategy space of firms and leaving fixed all other microfoundations in a standard monetary business cycle model, we provide a novel theory in which the nature of uncertainty and market power endogenously determine the aggregate supply curve. We find that aggregate supply flattens when there is: (i) lower inflation uncertainty, (ii) greater demand uncertainty, and (iii) increased market power. Money is maximally non-neutral if firms' quantities are perfectly elastic to prices (price-setting) and neutral if only if firms' quantities are perfectly inelastic to prices (quantity-setting). When mapped to the data, our theory explains: (i) the long-run flattening of the aggregate supply curve as an outcome of more hawkish monetary policy and rising market power and (ii) the steepening of aggregate supply in times of high inflation uncertainty (*e.g.*, the 1970s and 2020s) but not times of high real uncertainty (*e.g.*, the Great Recession).

*Yale University Department of Economics, 30 Hillhouse Avenue, New Haven, CT, 06511. Email: joel.flynn@yale.edu.

†Princeton Department of Economics, Julis Romo Rabinowitz Building, Princeton, NJ 08544. Email: nikolakoudis@princeton.edu.

‡Princeton University Department of Economics, Julis Romo Rabinowitz Building, Princeton, NJ, 08544, Email: ksastry@princeton.edu.

This paper supersedes our earlier work, circulated as: “Prices *vs.* Quantities: A Macroeconomic Analysis.” We are grateful to Daron Acemoglu, Mark Aguiar, Narek Alexanian, George-Marios Angeletos, Markus Brunnermeier, Gabriel Chodorow-Reich, Philip Haile, Basil Halperin, Joe Hazell, Nobuhiro Kiyotaki, Ricardo Lagos, Jennifer La’O, Ernest Liu, Alessandro Pavan, Michael Peters, Ricardo Reis, Richard Rogerson, Kunal Sangani, Aleh Tsyvinski, Michael Woodford, Gianluca Violante, Pierre Yared and seminar participants at Columbia University, the University of Pennsylvania, the NBER Summer Institute, Yale University, Princeton University, and the University of Wisconsin for helpful comments. First version: June 24, 2023.

1 Introduction

At the heart of modern models of aggregate supply are monopolistic firms that make decisions under uncertainty. It is common to restrict these firms' supply decisions to an important but narrow class: setting a price and committing to produce enough to meet *ex post* demand. For example, price-setting is assumed in classic models of aggregate supply based on exogenous, infrequent adjustment (Taylor, 1980; Calvo, 1983), menu costs (Barro, 1972; Caplin and Spulber, 1987; Golosov and Lucas, 2007), and limited information (Mankiw and Reis, 2002; Woodford, 2003a; Maćkowiak and Wiederholt, 2009).

In this paper, we enrich the baseline model to allow firms to choose *any* supply function: a mapping that describes the price charged at each quantity of production.¹ Supply function choice is a standard approach in microeconomic theory (e.g., Grossman, 1981; Hart, 1985; Klemperer and Meyer, 1989; Vives, 2011, 2017; Pavan et al., 2022) to model firms' ability to adjust decisions to realized demand, while remaining consistent with a foundation of information, contracting, or organizational frictions. Seen through this lens, committing to a fixed price is not generally the optimal choice. Our goal is to understand how allowing for richer behaviors at the microeconomic level affects our understanding of the macroeconomy.

Methodologically, we show how to tractably model supply function choice in a macroeconomic setting. To understand the determinants of supply function choice, we first characterize how optimal supply functions depend on the elasticity of demand and firms' uncertainty about demand, costs, and competitors' prices in partial equilibrium. To understand the implications of supply function choices, we then characterize how firm-level supply curves aggregate in general equilibrium to determine the aggregate supply curve. Putting these together, we characterize the equilibrium fixed point by which firms' uncertainty, and hence supply function choice, depends on the supply function choices of others.

We find that allowing for supply functions in an otherwise standard monetary business cycle model yields a theory of the aggregate supply curve with novel implications. Concretely, we find that aggregate supply flattens under: (i) lower price-level uncertainty, (ii) greater demand uncertainty, (iii) increased market power, and (iv) more hawkish monetary policy. Applying our model to the data, we find that our model generates variation in the slope of aggregate supply that coincides with empirical evidence on how the slope of the aggregate supply curve (the Phillips curve) has changed over time. Taken together, our analysis shows how structural changes like the commitment to price stability in the 1980s or a secular rise in market power can cause a flattening of the aggregate supply curve.

¹This is different from *nonlinear pricing*, whereby firms transact different quantities at different prices. A supply function specifies the uniform price that everyone pays as a function of the total quantity sold.

Supply Function Choice of a Single Firm. We study the standard model of a monopolistically competitive firm that lies at the core of many modern business cycle models. The firm faces a constant-price-elasticity demand curve and operates a constant-returns-to-scale production function. It has log-normal uncertainty about its competitors’ prices, demand, productivity, input prices, and the stochastic discount factor.

Given its information and beliefs, the firm commits to price-quantity pairs described by a *supply function* $f : \mathbb{R}_{++}^2 \rightarrow \mathbb{R}$. The firm produces and prices where the market demand curve intersects the locus of points such that $f(p, q) = 0$. Internalizing this, the firm chooses its optimal, non-parametric supply function to maximize its expected real profits under the stochastic discount factor. In other words, we allow firms to implement the ECON 101 notion of a supply curve: a systematic relationship between the price that they charge and the quantity that they produce. The ubiquitous “price-setting” model is nested by functions of the form $f(p) = 0$, or perfectly elastic supply; the opposite “quantity-setting” model (*e.g.*, as in Jaimovich and Rebelo, 2009; Angeletos and La’O, 2010, 2013; Benhabib et al., 2015) is nested by functions of the form $f(q) = 0$, or perfectly inelastic supply. Thus, while we allow firms to set prices or quantities, we also allow them to choose richer and potentially more preferable strategies that are nonetheless consistent with their informational constraints.

We solve in closed-form for the optimal supply function and show that it is endogenously log-linear: $\log p = \alpha_0 + \alpha_1 \log q$. Thus, the firm’s behavior in response to shocks to market demand is described by its optimally chosen inverse supply elasticity, α_1 : the percentage by which the firm increases prices in response to a one percent increase in production. In turn, this elasticity depends on the firm’s market power and its relative uncertainty about demand, competitors’ prices, and real marginal costs. These relationships arise because uncertainty and market power shape firms’ relative desires to hedge against different types of shocks, and manipulating the supply function gives firms a natural tool to do so.

Three comparative statics are particularly important for our subsequent macroeconomic conclusions. First, higher uncertainty about firm-level demand pushes toward a lower α_1 , or firms behaving more like price-setters. Price-setting perfectly insulates firms against demand shocks as the optimal response of a firm to changing demand conditions is to simply set its relative price equal to a constant markup on its real marginal cost. Second, higher uncertainty about competitors’ prices pushes toward a higher α_1 , or firms behaving more like quantity-setters. Quantity-setting perfectly insulates firms against shocks to competitors’ prices as it allows the firm’s relative price to adjust perfectly in response to such changes. Third, an increase in market power (a lower elasticity of demand) pushes toward a lower α_1 , or firms behaving more like price-setters. More market power reduces the cost to the firm of setting the “wrong” price.

General Equilibrium: From Supply Functions to Aggregate Supply. To study the aggregate implications of supply-function choice, we embed our framework in a monetary business-cycle model with incomplete information, following [Woodford \(2003a\)](#) and [Hellwig and Venkateswaran \(2009\)](#). In addition to exogenous microeconomic and macroeconomic uncertainty, the model generates endogenous macroeconomic uncertainty about firms’ demand, aggregate prices, and real marginal costs. In particular, because of imperfect competition between firms, the model features aggregate demand externalities ([Blanchard and Kiyotaki, 1987](#)) whereby firms face greater demand when aggregate output is high. Moreover, because households demand money, both the level of the money supply and aggregate output jointly determine the aggregate price level. Finally, because of income effects in labor supply, real marginal costs are higher when aggregate output is higher.

We first characterize aggregate outcomes given a fixed supply function. We show that the price level and real output follow an *aggregate supply and aggregate demand* representation. In this representation, there is a well-defined “slope of aggregate supply,” which also corresponds to the relative response of the price level and real GDP to an aggregate demand (money supply) shock. In particular, aggregate supply is inelastic—or, money is neutral—if and only if firms are quantity-setters. Aggregate supply is maximally elastic—or, money is as non-neutral as possible—if firms are price-setters.² Finally, increased market power flattens the aggregate supply curve. This effect is present as long as firms are *not* pure price-setters. Intuitively, as a lower elasticity of demand decreases how much a given change in the price level moves any given firm’s demand curve, with greater market power firms adjust their prices *ex-post* by less. This strategic interaction is missing in pure price-setting models, which is why the elasticity of aggregate supply is independent of market power in these (more familiar) cases.

We next characterize how the slope of aggregate supply is *endogenously* determined, via the fixed point relating macroeconomic uncertainty with firms’ supply-function choice. This reveals feedback loops: uncertainty affects supply functions, which affects the slope of aggregate supply, and in turn shapes macroeconomic uncertainty.

Under an empirically reasonable parameter restriction that balances strategic complementarity (from aggregate demand externalities) with substitutability (from wage pressure), we can further derive the slope of aggregate supply in closed form. This slope smoothly decreases in firms’ *relative uncertainty* about idiosyncratic demand shocks *vs.* the money supply. For example, an economy with more hawkish monetary policy, defined as a less

²There are other theoretically interesting, but perhaps empirically irrelevant, cases in which firm-level supply curves are *downward-sloping* and money has even larger real effects than other the price-setting model.

volatile money supply, features a flatter aggregate supply curve and therefore an endogenously smaller effect of demand shocks on the price level. In that sense, this economy has “more stable prices” for two *reinforcing* reasons. First, there are fewer demand shocks that directly decrease variation in the price level. Second, firms respond to more stable prices by flattening their demand curves, which endogenously reduces the responsiveness of prices to money supply shocks. This observation would be consistent with the narrative that more hawkish monetary policy in the United States (e.g., during and after the tenure of Paul Volcker) achieved price stability by actively flattening the aggregate supply curve.

An economy with higher idiosyncratic demand variation also features a flatter aggregate supply curve. Combining this with Bloom et al.’s (2018) observation that firms’ idiosyncratic uncertainty rises substantially in recessions, our theory offers the following resolution to the puzzle of “missing disinflation” during the Great Recession: aggregate supply itself endogenously *flattened* in the face of a large jump in microeconomic uncertainty.

Aggregate Supply in the Model and the Data. In the final section, we study the model’s implications for the slope of aggregate supply in the United States. To do so, we obtain estimates of time-varying uncertainty for macroeconomic aggregates by estimating a Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model on aggregate time series for output, the price level, and real marginal costs. We then embed these estimates in the general equilibrium structure of our model to analyze how the elasticity of aggregate supply has changed throughout time.

First, we find that our model can explain a quantitatively significant portion of the flattening of aggregate supply from the 1970s to the Great Moderation. In particular, we find that the reduction in the slope of aggregate supply is primarily driven by two components of uncertainty: (i) a reduction in the volatility of the price level, and (ii) an increased covariance between real marginal costs and demand. Both of these forces contribute to steeper firm-level supply functions: pure price-setting becomes less favorable to firms both when relative price variation is large, and when marginal costs covary positively with demand (as firms commit to producing exactly when demand is high). This steepening at the micro level translates to a steeper aggregate supply curve. Notably, our model implies that aggregate supply remained flat during the Great Recession (a period characterized by a spike in real, rather than nominal, uncertainty) before steepening again in the post-Covid period.

Second, we show that a rise in market power can lead to a further flattening of aggregate supply by creating additional incentives for firms to set steeper supply functions in times of high price volatility. Taken together, our empirical results emphasize that both uncertainty and market power are crucial in shaping the relative responsiveness of prices and output to aggregate demand shocks.

Related Literature. To our knowledge, ours is the first study to derive aggregate supply in a business-cycle model from a foundation of supply function competition. In microeconomic theory, Grossman (1981) and Hart (1985) introduce supply function equilibrium in oligopoly models under certainty; Klemperer and Meyer (1989) study this model under uncertainty; and Vives (2011, 2017) and Pavan et al. (2022) study further implications in specializations with quadratic preferences, private information, and Gaussian uncertainty.³ These papers’ mathematical results cannot be applied “off the shelf” to our macroeconomic setting due to the combination of the following features: non-quadratic preferences; imperfect substitutability; multiple, correlated sources of uncertainty; and strategic interactions in both input and product markets.

The closest analysis in the literature on firms’ optimal supply decisions is performed by Reis (2006), who studies a binary comparison of price-setting and quantity-setting for a rationally inattentive firm in a canonical macroeconomic setting.⁴ Our analysis differs from and builds on Reis’ analysis by studying supply schedules beyond price-setting and quantity-setting, allowing for multiple, correlated shocks to the firm, and studying equilibrium implications. Of course, one could argue that price- and quantity-setting are preferable to firms because they are simple. For this reason, we perform our entire general equilibrium analysis under a binary choice of price-setting and quantity-setting in Appendix C and recover similar insights.

Our finding that uncertainty shapes the slope of aggregate supply is shared with two literatures. First, and most relatedly, the classic “islands model” analysis of Lucas (1972, 1973, 1975). Unlike Lucas’ model (1972), ours features *monopolist* producers engaging in supply-function competition instead of price-taking producers (competitive markets). The “inference problem” that links uncertainty to supply decisions in our model arises for a different reason, without reference to the migration or physically separated markets hypothesized by Phelps (1970). Our model moreover allows roles for market power and uncertainty about costs which, in turn, are important in our mapping to the data.

Second, menu cost models of price-setting allow the extent of uncertainty to matter for the extent of monetary non-neutrality (see *e.g.*, Barro, 1972; Sheshinski and Weiss, 1977; Golosov and Lucas, 2007). In menu cost models, it is the extent of *total* uncertainty about optimal reset prices that determines non-neutrality, with greater uncertainty making money more neutral in many models (see Vavra, 2014). By contrast, in our theory, we show *total* uncertainty has no effect on the slope of aggregate supply, while *relative* uncertainty does.

³In asset pricing, Kyle (1989) studies an analogous problem of demand-function competition, building on prior work in perfectly competitive (“price-taking”) markets (Grossman, 1976; Hellwig, 1980).

⁴This analysis of “prices *vs.* quantities choice” echoes classic works on choice of instruments (Weitzman, 1974; Poole, 1970; Klemperer and Meyer, 1986).

Our analysis contributes to an extensive literature studying how assumptions about market structure and/or uncertainty affect predictions for the aggregate supply curve. We highlight a few particularly relevant connections below. [Mongey \(2021\)](#) and [Wang and Werning \(2022\)](#) study oligopoly with staggered price-setting, and [Fujiwara and Matsuyama \(2022\)](#) study monopolistic competition with entry and non-CES demand. Some of our findings regarding the relationship between market power and the Phillips curve are qualitatively similar to these authors', but our microfoundations and economic arguments are drastically different. [Woodford \(2003a\)](#), [Maćkowiak and Wiederholt \(2009\)](#), [Afrouzi and Yang \(2021\)](#), and [Angeletos and Huo \(2021\)](#), *inter alios*, study how the slope of aggregate supply in price-setting models depends on the structure of incomplete information. One implication of our results is that such conclusions about the role of incomplete information depend critically on assumptions about the (endogenous or exogenous) shape of supply functions.

Finally, our analysis mapping theory to data provides a theoretical rationalization of the results from the empirical literature that documents a flattening of the US Phillips curve (*e.g.*, [Ball and Mazumder, 2011](#); [Blanchard, 2016](#); [Hazell et al., 2022](#)) from the 1970s through the Great Moderation, as well as more recent work that tries to reconcile this with high inflation after the Covid pandemic (*e.g.*, [Blanchard and Bernanke, 2023](#)).

Outline. Section 2 solves for the firm's optimal supply function in partial equilibrium. Section 3 introduces the monetary business cycle model in which we embed supply functions. Section 4 characterizes equilibrium with supply function choice and shows how supply function choices affect aggregate supply. Section 5 compares the model's predictions for the slope of aggregate supply to existing empirical evidence. Section 6 concludes.

2 Supply Function Choice in Partial Equilibrium

In this section, we introduce and solve the problem of a monopolist that chooses a supply function in the presence of uncertainty about demand, costs, aggregate prices, and the stochastic discount factor. First, to build intuition, we study the case in which firms choose between two supply functions: price-setting (horizontal supply) or quantity-setting (vertical supply). We derive a formula for the advantage of price-setting relative to quantity-setting in units of log expected profits and interpret its comparative statics. The logic and formalism are reminiscent of [Weitzman \(1974\)](#): firms choose the supply function that best hedges them against uncontrollable shocks. We next study the more general problem of supply-function choice. Flexible supply functions give firms more tools to hedge against demand shocks, but remain an imperfect hedge against cost and price-level shocks. We show that optimal supply functions are log-linear and characterize their slope and intercept in terms of primitives. In

particular, we show how changes in uncertainty and market power (*i.e.*, the elasticity of demand) endogenously affect the firm-level elasticity of supply.

2.1 The Firm's Problem of Choosing Supply Functions

Set-up. A firm produces output $q \in \mathbb{R}_+$ via a constant-returns-to-scale production technology using a single input $x \in \mathbb{R}_+$:

$$q = \Theta x \tag{1}$$

where $\Theta \in \mathbb{R}_{++}$ is the firm's Hicks-neutral productivity. The firm can purchase the input at price $p_x \in \mathbb{R}_{++}$. The firm faces a constant-elasticity-of-demand demand curve given by:

$$\frac{p}{P} = \left(\frac{q}{\Psi} \right)^{-\frac{1}{\eta}} \tag{2}$$

where $p \in \mathbb{R}_+$ is the market price, $\Psi \in \mathbb{R}_{++}$ is a demand shifter, $P \in \mathbb{R}_{++}$ is the aggregate price level, and $\eta > 1$ is the price elasticity of demand. We interpret the elasticity of demand as an (inverse) measure of market power: when η is high, firms cannot easily manipulate market prices by altering quantities. The firm's profits are priced according to a real stochastic discount factor $\Lambda \in \mathbb{R}_{++}$. For simplicity, we define the firm's real marginal cost as $\mathcal{M} = P^{-1}\Theta^{-1}p_x$.

At the beginning of the decision period, the firm is uncertain about demand, costs, others' prices, and the stochastic discount factor (SDF). Specifically, they believe that the state $(\Psi, \mathcal{M}, P, \Lambda)$ follows a log-normal distribution with mean μ and variance Σ .⁵ The firm's payoff is given by its expected real profits (revenue minus costs), as priced by the real SDF:

$$\mathbb{E} \left[\Lambda \left(\frac{p}{P} - \mathcal{M} \right) q \right] \tag{3}$$

where $\mathbb{E}[\cdot]$ is the firm's expectation given some joint beliefs about $(\Lambda, P, \mathcal{M}, p, q)$.

Supply Functions. The firm commits to implementing price-quantity pairs described by the implicit equation $f(p, q) = 0$ where $f : \mathbb{R}_{++}^2 \rightarrow \mathbb{R}$. We will refer to f as the supply function. Price-setting is nested as a case in which $f(p, q) \equiv f^P(p)$. Quantity-setting is nested as a case in which $f(p, q) \equiv f^Q(q)$. More generally, we allow plans to be given by any non-parametric function f , even allowing for non-monotonicity and discontinuities.

The idea behind a supply function is familiar from ECON 101: it is the price that the firm plans to charge given any level of production. Of course, even a supply function is a metaphor for a more complicated organizational process at the firm and market level.

⁵Of course, \mathcal{M} is log-normal so long as P , Θ , and p_x are log-normal.

Klemperer and Meyer (1989) provide a detailed justification of how internal firm processes often correspond to supply functions.⁶ To exemplify what setting a supply function means, we borrow and extend an analogy from Reis (2006). Consider a bakery that must decide upon how much bread to produce and how to price its bread. A price-setting bakery fixes its price and keeps selling bread until it exhausts customer demand at that price. A quantity-setting bakery produces bread and sells it at the greatest price such that all of the bread produced is sold. Extending this analogy, a bakery that chooses a supply function can be thought of as observing how much bread is being sold at a given price and then adjusting its price to maximize profits. Thus, if demand happens to be strong and they are selling a higher quantity than they expected, the baker may decide to raise the price of its bread. The responsiveness of the baker’s production to its price is precisely its inverse supply elasticity. The supply function f captures this inverse supply elasticity as the slope of the locus of price-quantity pairs that it implies. Thus, supply functions allow us to more flexibly model the short-run responsiveness of the firm to demand fluctuations. This allows us to study behavior beyond the potentially restrictive assumption of a price-setting bakery (nested as an inverse elasticity of zero) or a quantity-setting bakery (nested as an inverse elasticity of infinity).


More formally, after choosing a supply function f , and following the realization of Ψ and P , the firm produces at a point where f intersects the demand curve. That is, the market clears. Toward making this rigorous, we define the *nominal demand state* $z = \Psi P^\eta$ and rewrite the demand curve as $q = zp^{-\eta}$. Thus, having set f and following the realization of z , the firms’ price is given by some solution \hat{p} to the equation $f(\hat{p}, z\hat{p}^{-\eta}) = 0$ with the realized quantity being $\hat{q} = z\hat{p}^{-\eta}$.

Figure 1 illustrates *ex post* market clearing visually. The columns respectively correspond to a price-setting function, $\log p = 0$ or $f(p, q) = 1 - p$; a quantity-setting function, $\log q = 0$ or $f(p, q) = 1 - q$; and a flexible rule, $\log p = \log q$ or $f(p, q) = 1 - \frac{p}{q}$. In the first row, we

⁶For example, they argue that:

“If a consulting firm sticks to a fixed rate per hour, it is fixing a price (perhaps subject to a capacity constraint). In fact, however, even when firms quote fixed rates, the real price often varies. When business is slack, more hours are worked on projects than are reported, but when the office is busy, marginally related training, travel time, and the time spent originally may all be charged to the client. Top management in effect commits to a supply function by choosing the number of employees and the rules and organizational values that determine how both the real price and the number of hours supplied adjust to demand [...] As another example, airlines use computer systems to adjust the number of discount seats they offer according to current reservation levels. Here management chooses a particular supply function by its choice of computer program. Running the program finds the price-quantity pair on this supply function that lies on the actual market demand curve.” – Klemperer and Meyer (1989).

Figure 1: An Illustration of Supply-Function Choice



Note: The columns correspond to different supply functions. The top row illustrates *ex post* market clearing for two realizations of the demand curve. The bottom curve illustrates the induced joint distribution of quantities and prices given log-normal uncertainty about z . See the main text for more details.

show each supply function (solid black line) and two demand curves, corresponding to a large demand realization z_1 (red dashed line) and a low demand realization z_0 (blue dotted line). The dots indicate the respective intersections of supply and demand, or realized quantity-price pairs in these states. In the second row, we illustrate the induced joint distribution of quantity-price pairs in purple, when z is log-normally distributed. The price- and quantity-setting policies fix uncertainty about one dimension, but induce uncertainty in the other. The flexible rule induces joint uncertainty in both variables.

We assume that the firm chooses the profit-maximizing selection from the set of solutions if there are many. We assume that the firm does not produce if no solution exists. We let the solution function be $\hat{p} : \mathbb{R}_+ \rightarrow \mathbb{R}_{++}$, with $\hat{p}(z)$ being the price charged in state $z \in \mathbb{R}_+$. As we placed no restrictions on f , it is equivalent to think of the firm as choosing \hat{p} directly. For a given choice of \hat{p} , the firm's payoff is given by:

$$J(\hat{p}) = \int_{\mathbb{R}_{++}^4} \Lambda \left(\frac{\hat{p}(z)}{P} - \mathcal{M} \right) z \hat{p}(z)^{-\eta} dG(\Lambda, P, \mathcal{M}, z) \quad (4)$$

where G is the cumulative distribution function representing the firm's beliefs. We therefore

study the problem:

$$\sup_{\hat{p}: \mathbb{R}_+ \rightarrow \mathbb{R}_{++}} J(\hat{p}) \quad (5)$$

Given a solution \hat{p} for how firms optimally adapt their prices to demand, we will recover the optimal plan f for how firms optimally set a supply function.

2.2 Building Intuition: Prices *vs.* Quantities

To build intuition for how uncertainty and market power shape the optimal supply function, we first study the restricted problem of a firm that can either fix a price or a quantity.

Optimal Price-Setting. In this case, \hat{p} is a constant function and Problem 5 reduces to:

$$V^P = \max_{p \in \mathbb{R}_+} \mathbb{E} \left[\Lambda \left(\frac{p}{P} - \mathcal{M} \right) \Psi \left(\frac{p}{P} \right)^{-\eta} \right] \quad (6)$$

Taking first-order conditions, the optimal price is given by:

$$p^* = \frac{\eta}{\eta - 1} \frac{\mathbb{E}[\Lambda \mathcal{M} P^\eta \Psi]}{\mathbb{E}[\Lambda P^{\eta-1} \Psi]} \quad (7)$$

where the numerator is the expected marginal benefit of charging higher prices in reducing costs and the denominator is the expected marginal cost of charging higher prices in increasing revenue. In the absence of uncertainty, this reduces to the statement that the optimal relative price is a constant markup of $\frac{\eta}{\eta-1}$ on real marginal costs. Substituting the optimal price into the firm's payoff function and rearranging, we obtain that:

$$V^P = \frac{1}{\eta - 1} \left(\frac{\eta}{\eta - 1} \right)^{-\eta} \mathbb{E}[\Lambda \mathcal{M} P^\eta \Psi]^{1-\eta} \mathbb{E}[\Lambda P^{\eta-1} \Psi]^\eta \quad (8)$$

Optimal Quantity-Setting. In this case, in Problem 5, the implemented price function must be the demand curve evaluated at a fixed q : $\hat{p}(z) = z^{\frac{1}{\eta}} q^{-\frac{1}{\eta}}$. The quantity-setting problem is therefore:

$$V^Q = \max_{q \in \mathbb{R}_+} \mathbb{E} \left[\Lambda \left(\left(\frac{q}{\Psi} \right)^{-\frac{1}{\eta}} - \mathcal{M} \right) q \right] \quad (9)$$

and optimal quantity is given by:

$$q^* = \left(\frac{\eta}{\eta - 1} \frac{\mathbb{E}[\Lambda \mathcal{M}]}{\mathbb{E}[\Lambda \Psi^{\frac{1}{\eta}}]} \right)^{-\eta} \quad (10)$$

where the numerator is the expected marginal cost of expanding production and the denominator is the expected marginal revenue from expanding production. In the absence of uncertainty, this is the quantity that the firm sells by setting its relative price equal to a constant markup on its real marginal cost. Substituting the optimal quantity into the firm's payoff, we obtain:

$$V^Q = \frac{1}{\eta - 1} \left(\frac{\eta}{\eta - 1} \right)^{-\eta} \mathbb{E} [\Lambda \mathcal{M}]^{1-\eta} \mathbb{E} \left[\Lambda \Psi^{\frac{1}{\eta}} \right]^\eta \quad (11)$$

Result: When to Set Prices *vs.* Quantities. A cursory inspection of the values of price-setting and quantity-setting (Equations 8 and 11) reveals that they are not generally equal. Define the log-difference between the values of price-setting and quantity-setting as:

$$\Delta = \log V^P - \log V^Q \quad (12)$$

We obtain the following formula for the proportional benefit of prices over quantities:

Proposition 1 (Prices *vs.* Quantities). *The comparative advantage of prices over quantities is given by:*

$$\Delta = \frac{1}{2}(\eta - 1) \left(\frac{1}{\eta} \sigma_\Psi^2 - \eta \sigma_P^2 - 2\sigma_{\Psi, \mathcal{M}} - 2\eta \sigma_{P, \mathcal{M}} \right) \quad (13)$$

This is increasing in demand uncertainty σ_Ψ^2 . It is decreasing in price uncertainty σ_P^2 , the covariance between demand and real marginal costs $\sigma_{\Psi, \mathcal{M}}$, and the covariance between prices and real marginal costs $\sigma_{P, \mathcal{M}}$. Moreover, when $\sigma_{P, \mathcal{M}} \geq -\frac{1}{2}\sigma_P^2$, $\text{sgn}(\Delta)$ is decreasing in η .

Proof. See Appendix A.1. □

The key idea is that price- and quantity-setting hedge the firm against different shocks— in analogy to price *vs.* quantity regulation in Weitzman (1974) or interest-rate *vs.* money supply targeting in Poole (1970).

The Role of Uncertainty. To understand the intuition for the comparative statics, we go case by case. First, in the presence of demand shocks alone, setting relative prices equal to a constant markup on marginal costs coincides with the first-best. By contrast, fixing the quantity supplied induces losses. Thus, demand shocks favor price-setting. Second, in the face of aggregate price shocks, fixing an optimal quantity allows relative prices to adjust perfectly while fixing an optimal price leads the firm's price to diverge from the aggregate price and loses revenue. Thus, aggregate price shocks favor quantity-setting. Third and fourth, when demand and real marginal costs or aggregate prices and real marginal costs

negatively covary, price-setting causes the firm to produce a large amount exactly when costs are low, favoring price-setting.

The Role of Market Power. The extent to which the firm values (i)-(iv) is mediated by the price elasticity of demand (*i.e.*, the extent of market power), since this determines how rapidly prices respond to underlying changes. In particular, as long as the covariance between prices and real marginal costs is not sufficiently negative, lower market power (higher η) favors quantity-setting and greater market power (lower η) favors price-setting. Fixing marginal costs, as firms' demand curves become more flat, price-setting exposes the firm to significant risk of setting the "wrong" price and either making zero profit (if their price is higher than competitors') or a very negative profit (if their price is lower than competitors'). This is potentially counteracted by a strong negative correlation between costs and others' prices: if when others charge high prices your marginal cost is very low, it is better to do price-setting as you undercut your competition precisely when it is very profitable to do so. In practice, we will later estimate a positive correlation between the price level and real marginal costs, making this second potential mechanism more of a theoretical curiosity.

2.3 The Optimal Supply Function

We now study the globally optimal supply function, or the solution to Problem 5. The firm now has considerably more flexibility than it did in the last subsection's study of prices *vs.* quantities choice. Nonetheless, we will be able to characterize the firm's optimal policy in closed form and illustrate comparative statics in the extent of uncertainty and the price elasticity of demand.

Toward the Solution: A Necessary Condition. Recall that we reduced the problem of choosing an optimal function to the problem of choosing an optimal mapping \hat{p} from nominal demand $z = \Psi P^\eta$ to the price. In an optimal policy, for any given realization $z = t$, it is necessary that there is no local benefit to changing the price $\hat{p}(t)$. This yields the following condition:

$$\begin{aligned} 0 &= \mathbb{E} \left[\Lambda \frac{1}{P} z \hat{p}(z)^{-\eta} - \eta \Lambda \left(\frac{\hat{p}(z)}{P} - \mathcal{M} \right) z \hat{p}(z)^{-\eta-1} \mid z = t \right] \\ &= (1 - \eta) \mathbb{E} \left[\Lambda \frac{1}{P} \mid z = t \right] t p(t)^{-\eta} + \eta \mathbb{E} [\Lambda \mathcal{M} \mid z = t] t p(t)^{-\eta-1} \end{aligned} \quad (14)$$

Re-arranging, we obtain the following:

$$p(t) = \frac{\eta}{\eta - 1} \frac{\mathbb{E}[\Lambda \mathcal{M} \mid z = t]}{\mathbb{E}[\Lambda P^{-1} \mid z = t]}, \forall t \in \mathbb{R} \quad (15)$$

This condition resembles the optimal price under price-setting (Equation 7), with the key difference that it conditions on nominal demand z . Outcomes under optimal rules therefore differ from optimal outcomes under price-setting (or quantity-setting) due to the firm’s ability to make inferences about the stochastic discount factor, real marginal costs, and the price level (e.g., the relevant deflator for profits). We will be able to solve for optimal functions in closed form, despite the high dimensionality of Problem 5, because Equation 15 reduces to a log-linear relation between p and z given lognormality.

Optimal Supply Functions. We now derive the optimal plan in closed form using rigorous variational arguments. This reveals how the slope of the optimal supply function, the “smooth” analog to prices *vs.* quantities choice, depends on the firm’s uncertainty and elasticity of demand:

Theorem 1 (The Optimal Supply Function). *Any optimal supply curve is almost everywhere given by:*

$$f(p, q) = \log p - \alpha_0 - \alpha_1 \log q \quad (16)$$

where the slope of the optimal price-quantity locus, $\alpha_1 \in \overline{\mathbb{R}}$, is given by:

$$\alpha_1 = \frac{\eta\sigma_P^2 + \sigma_{\mathcal{M},\Psi} + \sigma_{P,\Psi} + \eta\sigma_{\mathcal{M},P}}{\sigma_\Psi^2 - \eta\sigma_{\mathcal{M},\Psi} + \eta\sigma_{P,\Psi} - \eta^2\sigma_{\mathcal{M},P}} \quad (17)$$

Proof. See Appendix A.2. □

To obtain more intuition for the result, we re-express α_1 as the relative rate at which the firm wants log prices and log quantities to increase with $\log z$, or

$$\alpha_1 = \frac{d \log p}{d \log z} \bigg/ \frac{d \log q}{d \log z} = \frac{\text{Cov}[\log z, \log p^{**}]}{\text{Cov}[\log z, \log q^{**}]} \quad (18)$$

where we define p^{**} and q^{**} as the optimal *ex post* choices,

$$p^{**} = \frac{\eta}{\eta - 1} \mathcal{M}P \quad \text{and} \quad q^{**} = \left(\frac{\eta}{\eta - 1} \right)^{-\eta} \frac{z}{(\mathcal{M}P)^\eta} \quad (19)$$

Under this interpretation, the firm’s optimal policy is equivalent to running the following two-stage least squares (2SLS) regression: the firm wants to work out how its optimal price should change with its optimal quantity and uses the nominal demand state z as an instrument for the optimal quantity. The supply function is steep ($|\alpha_1|$ is large) if nominal demand predicts large movements in the *ex post* optimal price. In the 2SLS metaphor, the supply function is steep if the coefficient in the “reduced form” regression of p^{**} on z is high. The supply

function is shallow ($|\alpha_1|$ is small) if nominal demand predicts large movements in the *ex post* optimal quantity. In the 2SLS metaphor, the supply function is flat if the coefficient in the “first stage” regression of q^{**} on z is high.

Turning to how the nature of uncertainty affects the firm’s optimal inverse supply elasticity, we begin by focusing on the case in which the firm’s supply schedule is upward-sloping. This occurs if $0 \leq \text{Cov}[\log z, \log(\mathcal{M}P)] \leq \frac{1}{\eta} \text{Var}[\log z]$: high demand predicts that nominal costs are higher, but not too much higher. In this case, greater price-level uncertainty increases the slope of the optimal supply schedule. The intuition is similar to the one we derived earlier: not knowing the prices of your competitors makes quantity-setting relatively more attractive relative to price-setting because quantity-setting allows your *relative* price to adjust *ex post*. On the other hand, greater demand uncertainty decreases the slope of the optimal supply schedule. Intuitively, as in the binary comparison of prices *vs.* quantities, demand uncertainty favors price-setting as it allows production to adjust to accommodate greater demand. Finally, greater covariances between real marginal costs and demand and real marginal costs and the price level increase the firm’s inverse supply elasticity. Intuitively, when these covariances increase, the firm expects to produce more exactly when it is more costly. Thus, the firm optimally sets a steeper supply schedule to avoid over-producing in response to changes in demand.

We finally observe that a positively sloped supply function is not guaranteed: if nominal costs move sufficiently with nominal demand, then a monopolist may prefer a *downward* sloping supply function in order to hedge against high costs in high-demand states. In practice, however, we will find no empirical evidence for this condition, and it remains a theoretical curiosity.

Market Power and Supply Functions. The elasticity of demand plays two roles in determining the optimal (inverse) elasticity of supply. The first relates to *payoffs*: when η is high, *ex post* optimal quantities depend more on nominal marginal costs (holding fixed nominal demand). This echoes our earlier discussion in the context of prices *vs.* quantities choice: when goods are more substitutable, the firm’s optimal policy depends dramatically on whether its marginal costs are above or below others’ prices. The second role relates to *information*: when η is high, nominal demand contains relatively more information about the price level P and less about real demand Ψ . When studied in our general-equilibrium environment (Sections 3 and 4), these forces will open up the possibility that the slope of aggregate supply systematically depends on the extent of competition in the macroeconomy.

In general, the interaction of these two forces can make the optimal supply function steeper or flatten when η increases. But, below, we describe a sufficient condition under which steeper demand induces steeper supply:

Corollary 1 (Market Power and Optimal Supply). *A sufficient condition for greater market power to lower the inverse supply elasticity $\frac{\partial \alpha_1}{\partial \eta} > 0$ is that each of the following three inequalities holds:*

$$\alpha_1 \geq 0, \sigma_{M,P} \geq 0, 2\eta\sigma_{M,P} + \sigma_{M,\Psi} \geq \sigma_{P,\Psi} \quad (20)$$

Proof. See Appendix A.3. □

The force of these conditions is to restrict the extent to which high nominal demand predicts low marginal costs. In this case, the dominant logic is the following. When there is fiercer competition, an upward-sloping aggregate supply function better allows a firm to index its prices relative to its nominal costs. As discussed earlier, this allows the firm to better hedge its risks from setting the “wrong” price when products are very substitutable.

Later, in our empirical analysis (Section 5), we find that the condition of Corollary 1 always holds in US data since 1960 as long as $\eta > 3$. Thus, the empirically relevant case appears to be that market power flattens the firms’ optimal supply function.

Price- and Quantity-Setting Revisited. The previous result and discussion make clear that pure price- and quantity-setting are “edge cases” amongst the larger space of supply functions. Nonetheless, we observe below that they are intuitively obtained in the limiting cases of extreme demand or price-level uncertainty:

Corollary 2 (A Foundation for Price-Setting and Quantity-Setting). *The following statements are true:*

1. As $\sigma_P^2 \rightarrow \infty$, $|\alpha_1| \rightarrow \infty$ and the optimal plan converges to quantity-setting.
2. As $\sigma_\Psi^2 \rightarrow \infty$, $\alpha_1 \rightarrow 0$ and the optimal plan converges to price-setting.

This result provides a foundation for focusing on price- and quantity-setting when only one source of risk is dominant. In a macroeconomic environment, however, we may expect all sources of risk to be present in comparable orders of magnitude. In such a scenario, the extreme policies may perform poorly, for both the firm and the economic analyst.

A Demand-Supply Representation. Finally, we observe that outcomes in this market for the monopolist’s good are outcome-equivalent to those that would have been generated by the following simple demand and supply model:

Corollary 3 (Demand-Supply Representation). *The outcomes generated in the market are equivalent to those generated by the following “Demand-Supply” Model:*

$$\log p = -\frac{1}{\eta} \log q + \frac{1}{\eta} \log z \quad (D)$$

$$\log p = \alpha_0 + \alpha_1 \log q \quad (S)$$

We will now study how the demand and supply curves in this individual market map into the aggregate supply and aggregate demand curves in the context of a general equilibrium macroeconomic model. In this context, we will primarily be interested in understanding how the microeconomic inverse supply elasticity maps into the elasticity of aggregate supply: the relative responsiveness of aggregate prices to aggregate output. Moreover, we will study how equilibrium macroeconomic dynamics endogenously influence the optimal microeconomic supply elasticity.

Extension: Multiple Inputs, Decreasing Returns to Scale, and Monopsony. In Appendix B, we generalize this analysis to allow for a Cobb-Douglas production technology with multiple inputs, decreasing returns to scale, and convex costs of hiring additional inputs (capturing monopsony). We show that all of these forces complicate the analysis solely by introducing a single composite parameter that aggregates the decreasing returns and monopsony forces across inputs. We show in Proposition 6 that the optimal supply function remains optimally log-linear. Moreover, as is perhaps intuitive, decreasing returns to scale and monopsony power both reduce the optimal supply elasticity of the firm.

3 A Monetary Macroeconomic Model

We now embed supply-function choice in a monetary macroeconomic model. We otherwise use intentionally standard microfoundations (see, *e.g.*, Woodford, 2003b; Hellwig and Venkateswaran, 2009; Drenik and Perez, 2020). We use this model to derive a fully micro-founded, general-equilibrium specialization of Theorem 1 and to study its implications for the aggregate supply curve.

3.1 Households

Time is discrete and infinite $t \in \mathbb{N}$. There is a continuum of differentiated goods indexed by $i \in [0, 1]$, each of which is produced by a different firm.

A representative household has standard (Hellwig and Veldkamp, 2009; Golosov and Lucas, 2007) expected discounted utility preferences with discount factor $\beta \in (0, 1)$ and per-period utility defined over consumption of each variety, C_{it} ; holdings of real money balances, $\frac{M_t}{P_t}$; and labor effort supplied to each firm, N_{it} :

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\gamma}}{1-\gamma} + \ln \frac{M_t}{P_t} - \int_{[0,1]} \phi_{it} N_{it} di \right) \right] \quad (21)$$

where $\gamma \geq 0$ indexes income effects in both money demand and labor supply and $\phi_{it} > 0$ is the marginal disutility of labor supplied to firm i at time t , which is an IID lognormal variable with time-dependent variance, or $\log \phi_{it} \sim N(\mu_\phi, \sigma_{\phi,t}^2)$. The consumption aggregate C_t is a constant-elasticity-of-substitution aggregate of the individual consumption varieties with elasticity of substitution $\eta > 1$:

$$C_t = \left(\int_{[0,1]} \vartheta_{it}^{\frac{1}{\eta}} c_{it}^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}} \quad (22)$$

where ϑ_{it} is an IID preference shock that is also lognormal with time-dependent variance, or $\log \vartheta_{it} \sim N(\mu_\vartheta, \sigma_{\vartheta,t}^2)$. We also define the corresponding ideal price index:

$$P_t = \left(\int_{[0,1]} \vartheta_{it} p_{it}^{1-\eta} \right)^{\frac{1}{1-\eta}} \quad (23)$$

Households can save in either money or risk-free one-period bonds B_t (in zero net supply) that pay an interest rate of $(1 + i_t)$. The household owns the firms in the economy, each of which has profits of Π_{it} . Thus, the household faces the following budget constraint at each time t :

$$M_t + B_t + \int_{[0,1]} p_{it} C_{it} di = M_{t-1} + (1 + i_{t-1})B_{t-1} + \int_{[0,1]} w_{it} N_{it} di + \int_{[0,1]} \Pi_{it} di \quad (24)$$

where p_{it} is the price of variety of variety i and w_{it} is a variety-specific nominal wage.

The aggregate money supply follows an exogenous random walk with drift μ_M and time-dependent volatility σ_t^M :

$$\log M_t = \log M_{t-1} + \mu_M + \sigma_t^M \varepsilon_t^M \quad (25)$$

where the money innovation is an IID random variable that follows $\varepsilon_t^M \sim N(0, 1)$. So that interest rates remain strictly positive, we assume that $\frac{1}{2}(\sigma_t^M)^2 \leq \mu_M$ for all $t \in \mathbb{N}$. We will say that monetary policy is more *hawkish* when monetary volatility σ_t^M is lower.

3.2 Firms

The production side of the model follows closely the model from Section 2. Each consumption variety is produced by a separate monopolist firm, also indexed by $i \in [0, 1]$. Each firm operates a production technology that is linear in labor:

$$q_{it} = \zeta_{it} A_t L_{it} \quad (26)$$

where L_{it} is the amount of labor employed, ζ_{it} is IID lognormal with time-dependent volatility $\sigma_{\zeta,t}$, or $\log \zeta_{it} \sim N(\mu_{\zeta}, \sigma_{\zeta,t}^2)$, and $\log A_t$ follows an AR(1) with time-varying volatility σ_t^A :

$$\log A_t = \rho \log A_{t-1} + \sigma_t^A \varepsilon_t^A \quad (27)$$

where the productivity innovations are IID and follow $\varepsilon_t^A \sim N(0, 1)$. When the firm sells output at price p_{it} and hires labor at wage w_{it} , its nominal profits are given by $\Pi_{it} = p_{it}q_{it} - w_{it}L_{it}$. Since firms are owned by the representative household, their objective is to maximize expectations of real profits, discounted by a real stochastic discount factor Λ_t . Thus, the firm's payoff is $\frac{\Lambda_t}{P_t} \Pi_{it}$.

At the beginning of time period t , firms first observe A_{t-1} and M_{t-1} . Firms also receive private signals about aggregate productivity s_{it}^A and the money supply s_{it}^M :

$$\begin{aligned} s_{it}^A &= \log A_t + \sigma_{A,s,t} \varepsilon_{it}^{s,A} \\ s_{it}^M &= \log M_t + \sigma_{M,s,t} \varepsilon_{it}^{s,M} \end{aligned} \quad (28)$$

where the signal noise is IID and follows $\varepsilon_{it}^{s,A}, \varepsilon_{it}^{s,M} \sim N(0, 1)$. Firms are uninformed about the idiosyncratic productivity shock z_{it} , demand shock ϑ_{it} , and labor supply shock ϕ_{it} .

3.3 Markets and Equilibrium

In each period, conditional on the aforementioned information set, firms choose a supply function. As in Section 2, firms make this decision under uncertainty about demand, costs, and the stochastic discount factor. But, as will become clear, this uncertainty is now partially about *endogenous* objects. After firms make their choices, the money supply, idiosyncratic demand shocks, and both aggregate and idiosyncratic productivity are realized. Finally, the household makes its consumption and savings decisions and any prices that were not fixed adjust to clear the market.

Formally, we define an equilibrium as follows:

Definition 1 (Supply-Function General Equilibrium). *An equilibrium is a collection of variables*

$$\left\{ \{p_{it}, q_{it}, C_{it}, N_{it}, L_{it}, w_{it}, \Pi_{it}\}_{i \in [0,1]}, C_t, P_t, M_t, A_t, B_t, N_t, \Lambda_t \right\}_{t \in \mathbb{N}}$$

and a sequence of supply functions $\{f_t : \mathbb{R}_{++}^2 \rightarrow \mathbb{R}\}_{t \in \mathbb{N}}$ such that, in all periods:

1. All firms choose their supply function f_t to maximize expected real profits under the household's stochastic discount factor.

2. The household chooses consumption C_{it} , labor supply N_{it} , money holdings M_t , and bond holdings B_t to maximize their expected utility subject to their lifetime budget constraint, while Λ_t is the household's marginal utility of consumption.
3. Money supply M_t and productivity A_t and evolve exogenously via Equations 25 and 27.
4. Firms' and consumers' expectations are consistent with the equilibrium law of motion.
5. The markets for the intermediate goods, final good, labor varieties, bonds, and money balances all clear.

Note that this definition restricts to symmetric supply functions for all firms. We argue that this is a natural restriction natural given firms' *ex ante* symmetry, but it would not otherwise be guaranteed in equilibrium.⁷ We will also often be interested in describing equilibrium dynamics conditional on a (potentially suboptimal) supply function for firms. Formally, these *temporary equilibria* are equilibria in which we do not require statement (1) of Definition 1.

4 Supply Function Choice and Aggregate Supply

We now study the equilibrium determination of the aggregate supply curve. We do this in four steps. First, we solve for all equilibrium conditions except for the firm's supply function decision. By doing this, we pin down the equilibrium interest rate as a function of exogenous parameters and show that the dynamic economy can be studied as a sequence of one-shot economies. Second, within such a one-shot economy, we hypothesize that firms use log-linear supply functions and show that equilibrium macroeconomic dynamics are endogenously log-linear. This allows us to derive a simple Aggregate Demand and Aggregate Supply (AD/AS) representation of the equilibrium dynamics. In this representation, the equilibrium aggregate supply elasticity governs the responsiveness of the economy to shocks. Third, we study how the aggregate supply elasticity depends on firms' microeconomic supply elasticities, the extent of information about the money supply, and market power. From this, we learn that price-setting maximizes the non-neutrality of monetary and allows for no role of market power in determining the aggregate supply curve while quantity-setting implies that money is neutral. Finally, we combine the previous steps, find that firms' supply functions are endogenously log-linear, and fully characterize equilibrium in terms of a single, scalar fixed-point equation for the firm-level supply elasticity. We study how strategic interactions, market power, and the combination of microeconomic demand uncertainty

⁷In particular, there will be an (essentially unique) supply function given that aggregates are lognormal, applying Theorem 1. But heterogeneous supply functions could support non-log-normal aggregates, conditional on which there is not a unique supply function.

alongside aggregate productivity and monetary uncertainty affect the equilibrium aggregate supply elasticity.

4.1 Firms' Uncertainty in Equilibrium

We begin by deriving the general-equilibrium analogs of the four objects that were central to the firm's problem in Section 2: firm-specific demand shocks, firm-specific marginal costs, the price level, and the stochastic discount factor.

From the intratemporal Euler equation for consumption demand *vs.* labor supply, the household equates the marginal benefit of supplying additional labor $w_{it}C_t^{-\gamma}P_t^{-1}$ with its marginal cost ϕ_{it} . Thus, variety-specific wages are given by

$$w_{it} = \phi_{it}P_tC_t^\gamma \quad (29)$$

From the intertemporal Euler equation between consumption and money today, the cost of holding an additional dollar today equals the benefit of holding an additional dollar today plus the value of an additional dollar tomorrow:

$$C_t^{-\gamma}\frac{1}{P_t} = \frac{1}{M_t} + \beta\mathbb{E}_t\left[C_{t+1}^{-\gamma}\frac{1}{P_{t+1}}\right] \quad (30)$$

Further, from the intertemporal choice between bonds, the cost of saving an additional dollar today equals the nominal interest rate $1+i_t$ times the value of an additional dollar tomorrow:

$$C_t^{-\gamma}\frac{1}{P_t} = \beta(1+i_t)\mathbb{E}_t\left[C_{t+1}^{-\gamma}\frac{1}{P_{t+1}}\right] \quad (31)$$

By combining Equations 30 and 31, we obtain that aggregate consumption follows:

$$C_t = \left(\frac{i_t}{1+i_t}\right)^{\frac{1}{\gamma}} \left(\frac{M_t}{P_t}\right)^{\frac{1}{\gamma}} \quad (32)$$

which implies that aggregate consumption is increasing in real money balances, with elasticity given by $\frac{1}{\gamma}$. Intuitively, when consumption utility has greater curvature, income effects in money demand are larger and money demand is more responsive to changes in consumption. Thus, consumption responds less to real money balances when γ is large. The level of real money balances naturally depends on the opportunity cost of holding money i_t , and so money demand is lower when interest rates are high, all else equal.

Moreover, by substituting Equation 32 back into Equation 31, we obtain a recursion that

interest rates must satisfy:

$$\frac{1+i_t}{i_t} = 1 + \beta \mathbb{E}_t \left[\frac{1+i_{t+1}}{i_{t+1}} \frac{M_t}{M_{t+1}} \right] \quad (33)$$

As money follows a random walk, solving this equation forward and employing the household's transversality condition, we obtain that:⁸

$$\frac{1+i_t}{i_t} = 1 + \beta \exp\left\{-\mu + \frac{1}{2}(\sigma_t^M)^2\right\} \sum_{i=1}^{\infty} \prod_{j=1}^i \beta \exp\left\{-\mu + \frac{1}{2}\sigma_{M,t+j}^2\right\} \quad (35)$$

which is deterministic, but depends on the full future path of monetary volatility.

From the household's choice among varieties, the demand curve for each variety i is

$$\frac{p_{it}}{P_t} = \left(\frac{c_{it}}{\vartheta_{it} C_t} \right)^{-\frac{1}{\eta}} \quad (36)$$

Firm i faces strong demand when aggregate consumption is high, its competitors' prices are low, or its idiosyncratic demand is high. Moreover, η is the price elasticity of demand.

Summarizing the above, we have derived the following equilibrium mapping from endogenous objects to the objects that are relevant to the firm in partial equilibrium.

Proposition 2 (Firm-Level Shocks in General Equilibrium). *In any temporary equilibrium, demand shocks, aggregate price shocks, stochastic discount factor shocks, and marginal cost shocks follow:*

$$\Psi_{it} = \vartheta_{it} C_t, \quad P_t = \frac{i_t}{1+i_t} C_t^{-\gamma} M_t, \quad \Lambda_t = C_t^{-\gamma}, \quad \mathcal{M}_{it} = \frac{\phi_{it} C_t^\gamma}{z_{it} A_t} \quad (37)$$

Proof. See Appendix A.4. □

The first expression says that demand shocks have two components: an idiosyncratic shock deriving from consumer preferences and an aggregate shock corresponding to the aggregate demand externality (Blanchard and Kiyotaki, 1987). The second expression says that the price level must increase in nominal money balances, increase in the nominal interest rate, and decrease in consumption to lie on this demand curve. The third expression says that marginal utility is higher when consumption is lower. The fourth expression says that costs

⁸Observe also that in the case of time-invariant money volatility, interest rates follow the familiar equation:

$$1+i^* = \beta^{-1} \exp\left\{\mu_M - \frac{1}{2}(\sigma^M)^2\right\} \quad (34)$$

increase when wages increase due to disutility shocks or income effects and when productivity decreases at the micro or macro level.

The uncertainty described in Proposition 2 concerns endogenous objects. This introduces strategic uncertainty (*i.e.*, meaningful uncertainty about other firms' choices), familiar from the analysis of Woodford (2003a) under price-setting but missing, for example, from the analysis of Lucas (1972). Moreover, firms' uncertainty is correlated across variables due to macroeconomic linkages in the product, money, and labor markets.

An important technical implication of Proposition 2 is that, if C_t is log-normal, then so too is $(\Psi_{it}, P_t, \Lambda_t, \mathcal{M}_{it})$. This follows from the fact that all four expressions are log-linear and all other fundamentals $(A_t, M_t, \vartheta_{it}, \phi_{it}, z_{it})$ are log-normal by assumption. Therefore, if we can find that C_t is log-normal in equilibrium, our Theorem 1 can be directly applied to determine the optimal supply function in general equilibrium in a fully non-linear setting. We will call an equilibrium in which $\log C_t$ is linear in $(\log A_t, \log M_t)$ a *log-linear equilibrium*.

4.2 From Supply to Aggregate Supply with Fixed Functions

We next characterize aggregate outcomes given a fixed supply function. This allows us to define the *aggregate supply elasticity* and describe how it depends on firms' supply functions.

In particular, we consider a class of economies in which firms' exogenously set log-linear supply functions of the form

$$\log p_{it} = \alpha_{0t,i}^*(\alpha_{1,t}) + \alpha_{1,t} \log q_{it} \quad (38)$$

where $\alpha_{1,t} \in \mathbb{R}$ is a fixed parameter and $\alpha_{0t,i}^*(\alpha_{1,t})$ is the profit-maximizing "intercept" conditional on this slope. This optimal intercept depends on the slope $\alpha_{1,t}$, the firm's beliefs, and realized demand, but not (independently) on the realized quantity.

Conditional on these supply function choices, we guess and verify that there exists a unique equilibrium allocation in which aggregate consumption and the price level are log-linear in aggregate shocks:

$$\begin{aligned} \log P_t &= \chi_{0,t}(\alpha_{1,t}) + \chi_{A,t}(\alpha_{1,t}) \log A_t + \chi_{M,t}(\alpha_{1,t}) \log M_t \\ \log C_t &= \tilde{\chi}_{0,t}(\alpha_{1,t}) + \tilde{\chi}_{A,t}(\alpha_{1,t}) \log A_t + \tilde{\chi}_{M,t}(\alpha_{1,t}) \log M_t \end{aligned} \quad (39)$$

where, in each equation, the "intercepts" and "slopes" depend on $\alpha_{1,t}$.

Define the posterior weight on the firms' signals of productivity and the aggregate money

supply as, respectively,

$$\kappa_t^A = \frac{1}{1 + \left(\frac{\sigma_{A,s,t}}{\sigma_t^A}\right)^2}, \quad \kappa_t^M = \frac{1}{1 + \left(\frac{\sigma_{M,s,t}}{\sigma_t^M}\right)^2} \quad (40)$$

Moreover, define the slope of supply functions in $\log z_{it} = \eta \log P_t + \log \Psi_{it}$ as

$$\omega_{1,t} = \frac{\alpha_{1,t}}{1 + \eta \alpha_{1,t}} \quad (41)$$

We now characterize equilibrium macroeconomic dynamics with fixed supply functions:

Proposition 3 (Macroeconomic Dynamics with Supply Functions). *If all firms use log-linear supply functions of the form 38, output in the unique log-linear temporary equilibrium follows:*

$$\log C_t = \tilde{\chi}_{0,t} + \frac{1}{\gamma} \frac{\kappa_t^A}{1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma}\right) (1 - \kappa_t^A)} \log A_t + \frac{1}{\gamma} \frac{(1 - \kappa_t^M)(1 - \eta \omega_{1,t})}{1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma}\right) (1 - \kappa_t^M)} \log M_t \quad (42)$$

and the aggregate price in the unique log-linear temporary equilibrium is given by:

$$\log P_t = \chi_{0,t} - \frac{\kappa_t^A}{1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma}\right) (1 - \kappa_t^A)} \log A_t + \frac{\kappa_t^M + \frac{\omega_{1,t}}{\gamma} (1 - \kappa_t^M)}{1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma}\right) (1 - \kappa_t^M)} \log M_t \quad (43)$$

where $\chi_{0,t}$ and $\tilde{\chi}_{0,t}$ are constants that depend only on parameters (including $\alpha_{1,t}$) and past shocks to the economy.

Proof. See Appendix A.5 □

To build intuition for this result, it is helpful to think of these dynamics as being generated by a simple AD/AS model, in which productivity shocks shift the AS curve and money shocks shift the AD curve. Indeed, the dynamics derived in Proposition 3 are formally equivalent to those that we would have derived if we had written down the following AD/AS model.

Corollary 4 (AD/AS Representation). *The dynamics generated in the unique log-linear equilibrium are equivalent to those of the following “Aggregate Demand and Aggregate Supply” model:*

$$\log P_t = \log \left(\frac{i_t}{1 + i_t} \right) - \epsilon_t^D \log Y_t + \log M_t \quad (\text{AD})$$

$$\log P_t = (\chi_{0,t} - \epsilon_t^S \tilde{\chi}_{0,t}) + \epsilon_t^S \log Y_t + (\chi_{A,t} - \epsilon_t^S \tilde{\chi}_{A,t}) \log A_t \quad (\text{AS})$$

where the inverse supply and demand elasticities are given by:

$$\epsilon_t^S = \gamma \frac{\kappa_t^M + \frac{\omega_{1,t}}{\gamma}(1 - \kappa_t^M)}{(1 - \omega_{1,t}\eta)(1 - \kappa_t^M)} \quad \text{and} \quad \epsilon_t^D = \gamma \quad (44)$$

Proof. See Appendix A.6. □

In this representation, the aggregate demand curve is Equation 32, which combines the Euler equations for money and bonds with the transversality condition and implies that: i) the interest rate is a function of exogenous parameters and ii) aggregate consumption demand has an elasticity of $1/\gamma$ to changes in real money balances. Thus, the “inverse elasticity of aggregate demand” is fixed in our model as γ . The aggregate supply curve describes the equilibrium relationship between aggregate output and aggregate prices by aggregating firms’ microeconomic pricing and production decisions conditional on a fixed inverse supply elasticity (which we will soon endogenize). Note that the elasticity of the aggregate supply curve intuitively captures the relative response of prices and output to a monetary shock: $\epsilon_t^S = \frac{\chi_{M,t}}{\bar{\chi}_{M,t}}$.

The Propagation of Demand Shocks. To illustrate this representation, we use it to understand the propagation of a “demand shock,” an increase of the money supply by $\log M_1 = \log M_0 = \Delta \log M > 0$, in Figure 2. The shock shifts up the aggregate demand curve. The effects of this shock on output and prices depend on the relative size of the inverse supply and demand elasticities. In particular, prices move more and quantities move less if ϵ^S is large. Thus, the state-dependent effects of demand (and supply) shocks hinge on the slope of the aggregate supply curve.


We can gain some intuition for this formula by expanding out the total effect into a partial equilibrium effect and a series of higher-order general equilibrium effects:⁹

$$\frac{\Delta \log P}{\Delta \log M} = \frac{\epsilon_t^S}{\epsilon_t^D + \epsilon_t^S} = \underbrace{\left(\kappa_t^M + \frac{\omega_{1,t}}{\gamma}(1 - \kappa_t^M) \right)}_{\text{PE}} \times \underbrace{\sum_{j=0}^{\infty} \left(\omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right)^j}_{\text{GE}} \quad (45)$$

To understand the partial equilibrium (PE) effect, observe that when M goes up by 1%, all else equal, real money balances increase by 1%. From the household’s optimality conditions, this increases their consumption demand by $\epsilon_t^{D-1}\% = \frac{1}{\gamma}\%$. This has two effects. First, the firm experiences a $\frac{1}{\gamma}\%$ demand shock. As the firm has inverse supply elasticity of

⁹This summation only converges when $\left| \omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right| < 1$. Our fixed point arguments establish that the claimed formulae hold more generally whenever $\left| \omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right| \neq 1$.

Figure 2: An Aggregate Supply and Demand Representation



Note: An aggregate supply and demand illustration of dynamics after a shock of size $\Delta \log M$ to the money supply (see Corollary 4).

$\omega_{1,t}$, this leads the firm to increase its prices by $\frac{\omega_{1,t}}{\gamma}\%$. Second, from the household's labor supply condition, real marginal costs increase by $\gamma \times \frac{1}{\gamma}\% = 1\%$. As the firm wishes to set its relative price equal to a constant markup on its real marginal costs, this makes the firm want to increase prices by 1%. As they have already increased their prices by $\frac{\omega_{1,t}}{\gamma}\%$, they would achieve this 1% total price increase by increasing prices by $1 - \frac{\omega_{1,t}}{\gamma}\%$ in response to the 1% increase in real marginal costs. However, as firms receive imperfect signals of the money supply, their posterior means after the 1% shock increase in money and real marginal costs increase by only $\kappa_t^M\%$. Thus, on average, they respond to the increase in marginal costs by raising prices by $\kappa_t^M \times (1 - \frac{\omega_{1,t}}{\gamma})\%$. Thus, in partial equilibrium, the firms increase their prices on average by $\frac{\omega_{1,t}}{\gamma} + \kappa_t^M \times (1 - \frac{\omega_{1,t}}{\gamma})\% = \kappa_t^M + \frac{\omega_{1,t}}{\gamma}(1 - \kappa_t^M)\%$, which leads to an equal-sized effect on the aggregate price level.

To understand the general equilibrium effects, consider a 1% increase in the aggregate price level. This has three effects. First, as others' prices have risen by 1%, the firm experiences an $\eta\%$ demand shock. Second, as the price level has risen by 1%, real money balances fall and consumption demand falls by $\frac{1}{\gamma}\%$. Together, given the inverse supply elasticity of $\omega_{1,t}$, these effects lead firms to increase their prices by $\omega_{1,t} \times (\eta - \frac{1}{\gamma})\%$. Third, from the households' labor supply condition, the fall in real consumption demand induces a reduction in real marginal costs by $\frac{1}{\gamma} \times \gamma\% = 1\%$. With perfect information of the 1% increase in the price level, the firm would wish to reduce its price by $\omega_{1,t} \times (\eta - \frac{1}{\gamma})\%$, since this would imply that its *relative* price (which would fall by 1%) is maintained as a constant mark-up

over real marginal costs (which has fallen by 1%). However, as firms are imperfectly informed of the monetary shock, if a money shock induced a 1% increase in prices, then they would only on average perceive a $\kappa_t^M\%$ increase in the price level. Thus, they reduce their prices by $\kappa_t^M \times \omega_{1,t} \times (\eta - \frac{1}{\gamma})\%$. In total, out of a monetarily induced 1% increase in the aggregate price level, the average increase in firms' prices is therefore $\omega_{1,t} \times (\eta - \frac{1}{\gamma}) \times (1 - \kappa_t^M)\%$. Applying this logic to the initial $\kappa_t^M + \frac{\omega_{1,t}}{\gamma}(1 - \kappa_t^M)\%$ increase in prices from the PE effect and iterating it to all subsequent price increases in GE yields Equation 45.

This increase in the aggregate price level is necessarily less than one-for-one to the monetary shock when $\epsilon_t^S < \infty$ as $\frac{\epsilon_t^S}{\epsilon_t^D + \epsilon_t^S} < 1$. Thus, from the household's consumption demand, we have that output in the economy increases according to:

$$\frac{\Delta \log Y}{\Delta \log M} = \frac{1}{\epsilon_t^D + \epsilon_t^S} > 0 \quad (46)$$

The rich economics of this situation would be absent if price-setting ($\omega_{1,t} = 0$) were exogenously assumed: the PE effect would be $\kappa_t^M\%$ and the GE effect would be 0%. More generally, a novel implication of our model is that general-equilibrium strategic interaction, known to depend on uncertainty via higher-order beliefs in price-setting models (*e.g.*, Woodford, 2003a) and quantity-setting models (*e.g.*, Angeletos and La'O, 2010), hinges critically on the slope of the supply function. We will explore in more detail how the inverse supply elasticity affects aggregate supply in the following analysis.

The Propagation of Supply Shocks. While our study is primarily focused on predictions for the aggregate supply curve and transmission of demand shocks, our model also makes predictions for the aggregate demand curve and the transmission of supply shocks. These predictions are mediated by the slope of supply functions.

The response of the economy to a supply shock follows a very similar logic:

$$\frac{\Delta \log P}{\Delta \log A} = \underbrace{\kappa_t^A}_{\text{PE}} \times \underbrace{\sum_{j=0}^{\infty} \left(\omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A) \right)^j}_{\text{GE}} \quad (47)$$

The PE effect is immediate: firms perceive a $\kappa_t^A\%$ decrease in their real marginal costs and adjust their prices by an equal percentage. The GE effects of the change in the price level are identical other than that productivity uncertainty may differ from monetary uncertainty. Thus, strategic interactions are attenuated by a factor of $1 - \kappa_t^A$ rather than $1 - \kappa_t^M$. A key takeaway from our analysis is that the general equilibrium *transmission* of shocks crucially depends on the slope of microeconomic supply curves in the economy.

4.3 The Elasticity of Aggregate Supply in Temporary Equilibrium

Having demonstrated that the response of the economy to demand and supply shocks can be understood solely in terms of the slope of the aggregate supply curve, we now formally investigate how various microeconomic forces affect it.

Corollary 5 (How Microeconomic Forces Affect Aggregate Supply). *If firms' supply curves are upward-sloping (i.e., $\omega_{1,t} \in [0, 1/\eta]$), then the following statements are true:*


1. Steeper microeconomic supply steepens the AS curve: $\partial \epsilon_t^S / \partial \omega_{1,t} \geq 0$
2. Monetary uncertainty flattens the AS curve: $\partial \epsilon_t^S / \partial \kappa_t^M \geq 0$.
3. Income effects steepen the AS curve: $\partial \epsilon_t^S / \partial \gamma \geq 0$.
4. Market power flattens the AS curve: $\partial \epsilon_t^S / \partial \eta \geq 0$.

The proof is immediate from differentiation of Equation 44. To understand the first statement, observe that a steeper microeconomic supply function makes prices more responsive to realized quantity *ex-post*. At the aggregate level, this implies that the price level is also more responsive to changes in output. Second, less monetary uncertainty steepens the AS curve because firms respond to the perceived increase in the money supply by increasing *average* prices (as modulated through the intercept $\alpha_{0,t}^*$). This reduces variation in *real* money balances, thereby attenuating the effect of demand shocks on aggregate output. Third, output responds less to money balances the higher is γ (cf. Proposition 2). Consequently, a higher γ steepens the AS curve.

Finally, market power flattens the AS curve. Crucially, this effect is non-zero if and only if $\omega_{1,t} \neq 0$, i.e., firms do not undertake pure *price-setting*. This flattening operates through the general equilibrium transmission mechanisms of the model. When other firms raise their prices in response to a money supply shock, firm-level demand increases because the firm's *relative* price is now lower. The magnitude of this demand change is exactly parameterized by the elasticity of substitution η . If the responsiveness of prices to quantities at the firm level is non-zero, this demand increase generates an additional price level response. Consequently, higher market power flattens the AS curve by lowering the responsiveness of *firm-level* prices to *relative* price changes. Our analysis highlights that strategic considerations that arise from market structure disappear under the well-studied case of pure price-setting (and only in this case). Thus, we emphasize that monetary models can allow for a non-trivial role for market power without resorting to a change in their microfoundations, such as the addition of oligopoly (Mongey, 2021; Wang and Werning, 2022) or more complicated demand curves (Fujiwara and Matsuyama, 2022).

The precision of private information about the money supply, κ_t^M , interacts with the slope of supply functions. An important implication is that existing quantitative results about the

Figure 3: Aggregate Supply Under Price-Setting and Quantity-Setting



Note: An aggregate supply and demand illustration of dynamics after a shock of size $\Delta \log M$ to the money supply (see Corollary 4) under price-setting (panel a) and quantity-setting (panel b).

influence of uncertainty on the slope of the Phillips curve (*e.g.*, in Woodford, 2003a) depend on the assumed slope of firms' supply functions;

Corollary 6 (Aggregate Supply Under Price- and Quantity-Setting). *If firms engage in price-setting ($\omega_{1,t} = 0$), then:*

$$\epsilon_t^S = \gamma \frac{\kappa_t^M}{1 - \kappa_t^M} \quad (48)$$

If firms engage in quantity-setting ($\omega_{1,t} = \frac{1}{\eta}$), then:

$$\epsilon_t^S = \infty \quad (49)$$

In particular, the AS curve is vertical under quantity-setting and money cannot have any real effects. This is not a foregone conclusion, but an equilibrium result. Indeed, quantity-setting firms could condition their production on their monetary signal and money would have real effects if they did so.¹⁰ However, if firms set quantities, there is no equilibrium in which firms' quantities depend on the monetary signal as this is suboptimal. Under price-setting, money has real effects but market power has no role in shaping the slope of the aggregate supply curve. Moreover, as ϵ_t^S is increasing in $\omega_{1,t}$, price-setting provides a *lower bound* on the inverse elasticity of the aggregate supply curve and therefore maximizes the real effects of demand shocks. We illustrate these two “extreme” predictions for aggregate supply and demand in Figure 3.

¹⁰As a simple example, setting $\log q_{it} = s_{it}^M$ is feasible for firms and this would allow money to have real effects: $\log C_t = \text{const}_t + \log M_t$.

4.4 The Equilibrium Elasticity of Supply

The final step in our analysis is to endogenize the firm-level inverse supply elasticity as a best response to equilibrium macroeconomic dynamics. Observe that we have verified that if firms use log-linear supply functions, then aggregate dynamics are endogenously log-linear (by Proposition 3). Moreover, we have verified that if aggregate dynamics are log-linear, then firms' uncertainty is endogenously log-normal (by Proposition 2). Thus, we have shown that firms' supply curves are endogenously log-linear in a log-linear equilibrium (by Theorem 1). By combining all of these results, we reduce the determination of log-linear equilibrium in the full dynamic economy with functional supply decisions by firms to a single, scalar fixed-point equation for firms' transformed inverse supply elasticities:

Theorem 2 (Equilibrium Supply Elasticity Characterization). *All (and only all) solutions $\omega_{1,t} \in (-\infty, \infty)$ of the following equation correspond to transformed inverse supply elasticities in log-linear equilibrium:*

$$\omega_{1,t} = T_t(\omega_{1,t}) \equiv \frac{\frac{(\eta - \frac{1}{\gamma})\kappa_t^A}{1 - \omega_{1,t}(\eta - \frac{1}{\gamma})(1 - \kappa_t^A)}(\sigma_{t|s}^A)^2 + \frac{\frac{1}{\gamma} + (\eta - \frac{1}{\gamma})\kappa_t^M}{1 - \omega_{1,t}(\eta - \frac{1}{\gamma})(1 - \kappa_t^M)}(\sigma_{t|s}^M)^2}{\sigma_{\vartheta,t}^2 + \left(\frac{(\eta - \frac{1}{\gamma})\kappa_t^A}{1 - \omega_{1,t}(\eta - \frac{1}{\gamma})(1 - \kappa_t^A)}\right)^2(\sigma_{t|s}^A)^2 + \left(\frac{\frac{1}{\gamma} + (\eta - \frac{1}{\gamma})\kappa_t^M}{1 - \omega_{1,t}(\eta - \frac{1}{\gamma})(1 - \kappa_t^M)}\right)^2(\sigma_{t|s}^M)^2} \quad (50)$$

Proof. See Appendix A.7. □

This fixed-point equation incorporates the variances and covariances that enter the optimal supply function as a function of equilibrium macroeconomic dynamics when firms use supply functions with transformed inverse supply elasticities $\omega_{1,t}$. This depends on the responsiveness of aggregate prices and output to aggregate productivity and monetary shocks as well as the conditional uncertainty about these shocks when firms set their supply functions. Firms' idiosyncratic uncertainty about demand matters, but firms' uncertainty about idiosyncratic productivity and factor prices do not as the variance of marginal costs *per se* does not matter for the choice of an optimal supply function.

In the remainder of this section, we will study this equation to understand equilibrium dynamics. First, we can use this result to establish log-linear equilibrium existence and provide a bound on the number of equilibria by rewriting the fixed-point equation as a quintic polynomial in $\omega_{1,t}$:

Proposition 4 (Existence, Number, and Uniqueness of Equilibria). *There exists a log-linear equilibrium. There exist at most five log-linear equilibria.*

Proof. See Appendix A.8 □

We now study how uncertainty, strategic interactions, and market power shape the aggregate supply elasticity in equilibrium.

A Simple Characterization Under Balanced Strategic Interactions. We first characterize the elasticity of aggregate supply under the parametric condition $\eta\gamma = 1$. Recall from our discussion in Section 4.2 that η parameterizes the strength of *strategic complementarities*: the additional increase in demand a firm faces from an increase in the aggregate price level due to a change in *relative* prices. In contrast, $1/\gamma$ parameterizes the strength of *strategic substitutabilities*: the reduction in demand a firm faces from an increase in the aggregate price level due to a reduction in *aggregate consumption* (that results from the reduction in real money balances). Hence, $\eta\gamma = 1$ considers the case when these two strategic interactions exactly balance. This allows us to simplify the fixed point in Equation 50 considerably.

Corollary 7 (Idiosyncratic vs. Aggregate Demand Uncertainty). *When $\eta\gamma = 1$, the unique inverse elasticity of aggregate supply is*

$$\epsilon_t^S = \gamma \frac{\kappa_t^M}{1 - \kappa_t^M} \left(1 + \frac{1}{\gamma^2 \rho_t^2 \kappa_t^M} \right) \quad (51)$$

where $\rho_t = \frac{\sigma_{\theta,t}}{\sigma_{i|s}^M}$ is the relative uncertainty about demand vs. the money supply.

Proof. See Appendix A.9 □

First, observe that uncertainty about *aggregate* productivity does not enter the elasticity of aggregate supply as $\eta\gamma = 1$. This is because a perceived increase in aggregate productivity induces all firms to decrease their prices. In the absence of additional strategic interactions, firms will not respond to other firms' price reductions. Hence, the demand state z (Equation 15) is not useful for inference on *nominal* marginal costs and κ_t^A does not enter the fixed point. The same is not true for uncertainty about the money supply, as it induces *direct* variation on the demand state z by changing aggregate consumption through real money balances. Consequently, firms can condition on the demand state z to learn about their real marginal costs when the money supply is uncertain.

Second, as $\rho_t \rightarrow \infty$, the inverse elasticity of aggregate supply approaches $\frac{1}{\gamma} \frac{\kappa_t^M}{1 - \kappa_t^M}$. It is easily seen that this is the “classic” AS curve that arises from *pure* price-setting, in which $\omega_{1,t} = 0$. Intuitively, *idiosyncratic* demand conditions do not affect a given firm's marginal cost. Hence, as idiosyncratic demand becomes relatively more volatile, the firm optimally sets a constant price to keep its markup over marginal cost constant. Had the firm chosen $\omega_{1,t} \neq 0$, the firm would induce unprofitable variation in its price by responding to idiosyncratic demand conditions.

Third, as $\rho_t \rightarrow 0$, the inverse elasticity of aggregate supply becomes vertical. Consequently, money has no real effects on output. It is easily seen that this is the AS curve that arises from *pure* quantity-setting ($\omega_{1,t} = \frac{1}{\eta}$), which generates monetary neutrality. Intuitively, as uncertainty about the money supply – and therefore the aggregate price level – increases, firms find it optimal to keep their quantities constant and let their *relative* price adjust to demand.

This discussion highlights that *relative* uncertainty about idiosyncratic *vs.* aggregate demand shocks is a crucial determinant of the responsiveness of output to monetary shocks. Moreover, this feature only becomes relevant once firms are allowed to optimally choose their supply functions. As Corollary 6 demonstrates, if one were to exogenously impose price-setting or quantity-setting, the slope of aggregate supply is independent of any feature of idiosyncratic or aggregate uncertainty other than the signal-to-noise ratio for the money supply.

Equilibrium Under Dominant-Uncertainty Limits. To better understand how each source of uncertainty matters, we next characterize how equilibria behave as each source of uncertainty becomes dominant:¹¹

Corollary 8 (Dominant-Shock Limits). *The following statements are true:*

1. As $\sigma_{\vartheta,t} \rightarrow \infty$, in any equilibrium $\omega_{1,t} \rightarrow 0$ (price-setting)
2. As $\sigma_{t|s}^M \rightarrow \infty$, in any equilibrium $\omega_{1,t} \rightarrow \frac{1}{\eta}$ (quantity-setting)
3. As $\sigma_{t|s}^A \rightarrow \infty$ and $\eta\gamma \neq 1$, in any equilibrium $\omega_{1,t} \rightarrow \frac{1}{\eta - \frac{1}{\gamma}}$

Proof. See Appendix A.10 □

The intuition for this result mirrors that of Corollary 7. As idiosyncratic uncertainty about demand becomes dominant, firms find it optimal to set prices to keep their markup over real marginal costs constant. As uncertainty about the money supply becomes dominant, firms become more uncertain about the aggregate price level. Consequently, firms find it optimal to set quantities and let their relative prices adjust to meet demand. Finally, as uncertainty about aggregate productivity becomes dominant, firms use the demand state z to make inferences solely about the realization of aggregate productivity. Under perfect information, a 1% decrease in productivity would imply that firms raise their prices by 1%. This translates to an $(\eta - \frac{1}{\gamma})\%$ increase in demand for a given firm. Since firms would optimally like to keep their mark-up over marginal cost constant, they infer that this implies a 1% reduction in productivity, and decrease their prices by $\left[1/(\eta - \frac{1}{\gamma})\right]\%$. Observe that

¹¹Formally, we take these limits for $\sigma_{t|s}^x$ and $x \in \{M, A\}$ by scaling $\sigma_{x,s,t}$ and σ_t^x by a common factor.

this force implies a downward-sloping supply curve whenever $\eta\gamma < 1$. Intuitively, if $\eta\gamma < 1$, income effects in labor supply are weak and the firm expects a lower real marginal cost when a positive demand shock hits them.

The (Absent) Role of Total Uncertainty. We have so far seen that the nature of uncertainty (idiosyncratic *vs.* aggregate and demand *vs.* productivity) matters. Thus, the *presence* of uncertainty is of central importance to our analysis. However, a distinguishing feature of the theory that we have developed is that the total *level* of uncertainty does not matter. To make this claim formal, fix a scalar $\lambda \geq 0$ and scale all uncertainty in the economy according to:

$$(\sigma_{\vartheta,t}, \sigma_{z,t}, \sigma_{\phi,t}, \sigma_{A,t}, \sigma_{A,s,t}, \sigma_{M,t}, \sigma_{M,s,t}) \mapsto (\lambda\sigma_{\vartheta,t}, \lambda\sigma_{z,t}, \lambda\sigma_{\phi,t}, \lambda\sigma_{A,t}, \lambda\sigma_{A,s,t}, \lambda\sigma_{M,t}, \lambda\sigma_{M,s,t}) \quad (52)$$

In this sense, λ is a measure of the total level of uncertainty faced by firms. Define the correspondence $\mathcal{E}_t^S : \mathbb{R}_+ \rightrightarrows \bar{\mathbb{R}}$, where $\mathcal{E}_t^S(\lambda)$ is the set of equilibrium inverse supply elasticities for the level of uncertainty λ . We observe the following:

Proposition 5 (Invariance to Uncertainty and Discontinuity in the Limit). *For $\lambda > 0$, $\mathcal{E}_t^S(\lambda)$ is a constant function and the equilibrium supply elasticity is invariant to the level of uncertainty. Moreover, $\mathcal{E}_t^S(0) = \{\infty\}$. Therefore, the equilibrium supply elasticity is discontinuous in the zero uncertainty limit:*

$$\lim_{\lambda \rightarrow 0} \mathcal{E}_t^S(\lambda) \neq \mathcal{E}_t^S(0) \quad (53)$$

Proof. See Appendix [A.11](#) □

Thus, the predictions that our theory makes for the slope of the aggregate supply curve are invariant to the level of uncertainty. This distinguishes our theory from models of costly adjustment of prices or quantities, such as menu cost models ([Barro, 1972](#); [Sheshinski and Weiss, 1977](#)). In these theories, the level of uncertainty matters because it both affects the probability that firms hit their adjustment thresholds and firms' optimal adjustment thresholds. For example, in the standard menu cost model with quadratic payoffs and Gaussian shocks, greater risk increases the chance that firms adjust their prices and makes the aggregate supply curve steeper.

Moreover, our model is discontinuous in the zero uncertainty limit. Indeed, $\mathcal{E}_t^S(\lambda)$ is typically neither upper hemi-continuous nor lower hemi-continuous at $\lambda = 0$. Thus, even a vanishingly small level of uncertainty can have significant effects on firm and aggregate behavior.

5 Model Meets Data: Aggregate Supply in the US

In this final section, we study our model’s implications for the slope of aggregate supply in the United States. We find that our model can explain a quantitatively significant portion of the secular flattening of aggregate supply from the 1980s to the Great Moderation due to changing relative uncertainty about inflation versus demand. The model explains even more of this flattening if we allow for an upward trend in market power. Our model remains consistent with a relatively flat aggregate supply curve in the Great Recession, since this period is characterized by a spike in real rather than nominal uncertainty (at the micro and macro levels), and a pronounced increase in the supply curve’s slope in the post-Covid period, due to a resurgence of inflation and cost shocks.

5.1 Mapping the Data to the Model

Our model provides explicit formulae for the firm-level inverse supply elasticity $\omega_{1,t}$ (Theorem 1) and the inverse elasticity of aggregate supply ϵ_t^S (Corollary 4). To obtain estimates of these quantities, we need to know two sets of objects: the structural parameters (η, γ) and firms’ time-varying uncertainty. We summarize our calibration in Table 1 and describe the details below.

Structural Parameters. We set the price elasticity of demand at $\eta = 9$, based on the study of Broda and Weinstein (2006). These authors use comprehensive panel data on US imports to estimate demand curves at the level of disaggregated products. This is, usefully for our purposes, *direct* evidence for the slope of demand curves, as opposed to indirect evidence from matching average product markups under the assumption that firms are full-information price setters. Later, in an extension, we study the effects of a secular downward trend in η (i.e., a secular upward trend in market power) over our studied time period (1960-2022). We set the size of income effects at $\gamma = 0.095$, based on the calibration in Flynn and Sastry (2022) for the cyclicity of US real wages. This is the relevant moment in our model given Proposition 2. We will later perform a full sensitivity analysis of these choices (see Figure 17).

Time-Varying Uncertainty. We next need to estimate firms’ time-varying uncertainty about aggregate prices P_t , demand Ψ_{it} , and real marginal costs \mathcal{M}_{it} . To our knowledge, there are no datasets that directly measure firms’ potentially correlated uncertainty about both microeconomic and macroeconomic objects. Our approach is to proxy for this using time-varying *statistical* uncertainty about macroeconomic aggregates implied by a GARCH model and assumptions, based on the existing literature, about the systematic relation-

ship between macroeconomic and microeconomic uncertainty. This gives us estimates of $(\sigma_{P,t}^2, \sigma_{\Psi,t}^2, \sigma_{\mathcal{M},\Psi,t}, \sigma_{P,\Psi,t}, \sigma_{\mathcal{M},P,t})$. We then choose a time-invariant value of κ^M , the quality of firms’ signal of the money supply, to target a sample-average aggregate supply slope of 0.15.

To estimate our model of macroeconomic uncertainty, we use quarterly-frequency US data on real GDP, the GDP deflator, and capacity-utilization adjusted total factor productivity (TFP) (Basu et al., 2006; Fernald, 2014) from 1960Q1 to 2022 Q4. Thus, our mapping from model to data considers quarterly-frequency decisions.

We map these variables to our general equilibrium model from Section 3 as follows. First, by Proposition 2, the model-implied demand shock is $\Psi_{it} = Y_t \vartheta_{it}$, where Y_t is aggregate real GDP (i.e., “aggregate demand”) and ϑ_{it} is a firm-specific demand shock that is, by construction, orthogonal to aggregate conditions. Thus, we can decompose $\sigma_{\Psi,t}^2 = \sigma_{Y,t}^2 + \sigma_{\vartheta,t}^2$, where the latter two terms are respectively the perceived variances of $\log Y_t$ and $\log \vartheta_{it}$. Second, we may use Proposition 2 to obtain that real marginal costs are $\mathcal{M}_{it} = (\phi_{it} Y^\gamma)/(\zeta_{it} A_t)$. However, as the firm-level factor price shock ϕ_{it} and productivity shock ζ_{it} are idiosyncratic and we only need to measure the covariances of \mathcal{M}_{it} with Ψ_{it} and P_t , we do not need to measure ϕ_{it} or ζ_{it} . Thus, it is sufficient for us to measure the common component of real marginal costs $\mathcal{M}_t = Y_t^\gamma/A_t$. We use capacity-adjusted TFP as our measure of A_t .

Finally, we assume that uncertainty about idiosyncratic demand is directly proportional to uncertainty about aggregate marginal costs, or $\sigma_{\vartheta,t}^2 = R^2 \sigma_{\mathcal{M},t}^2$. We justify this based on the finding of Bloom et al. (2018) that the stochastic volatility of TFPR among manufacturing firms (“micro volatility”) is well modeled as directly proportional to stochastic volatility in aggregate conditions (“macro volatility”). This justification relies on an assumption that TFPR dispersion is primarily driven by demand shocks. This assumption is consistent with the findings of Foster et al. (2008) that cross-firm variation in revenue total factor productivity (TFPR) derives almost exclusively from demand differences rather than marginal cost differences within specific industries. Based on the quantitative findings of Bloom et al. (2018), we take $R = 6.5$ as a baseline. In an extension, to examine robustness to this proportionality assumption, we directly use (annual) data on TFPR dispersion from Bloom et al. (2018) to perform our calculation and find very similar results.

We next estimate time-varying uncertainties regarding inflation, real output, and real marginal costs using a multivariate GARCH model. In particular, letting Z_t denote the vector $(\Delta \log P_t, \Delta \log Y_t, \Delta \log \mathcal{M}_t)$, we model

$$Z_t = AZ_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Sigma_t), \quad \Sigma_t = D_t^{\frac{1}{2}} C D_t^{\frac{1}{2}} \quad (54)$$

where A is a matrix of AR(1) coefficients, D_t is a diagonal matrix of time-varying variances

(and $D_t^{\frac{1}{2}}$ is a diagonal matrix of standard deviations), and C is a static matrix of correlations. We assume that each diagonal element of D_t , denoted as $\sigma_{i,t}^2$, evolves according to:

$$\sigma_{i,t}^2 = s_i + \alpha_i \varepsilon_{i,t-1}^2 + \beta_i \sigma_{i,t-1}^2 \quad (55)$$

with unknown constant s_i and coefficients (α_i, β_i) . Formally, this is a GARCH(1,1) model with constant conditional correlations (Bollerslev, 1990). In our context, the restriction to constant correlations restricts the covariances in Equation 13 to move in proportion to the variances and thus rules out the possibility that the correlation structure among output, prices, and marginal costs varies over time. We estimate all of the parameters via joint maximum likelihood.

From this, we derive maximum-likelihood point estimates of every element of Σ_t , which correspond to the variances in the (Gaussian) conditional forecast of Z_t . Letting $\hat{\sigma}_t$ denote the point estimates of specific elements of that matrix, we directly obtain $\hat{\sigma}_{P,t}^2$ and $\hat{\sigma}_{\mathcal{M},P,t}^2$ from the GARCH model and then we compute:

$$\hat{\sigma}_{\Psi,t}^2 = \hat{\sigma}_{Y,t}^2 + R^2 \hat{\sigma}_{\mathcal{M},t}^2, \quad \hat{\sigma}_{\mathcal{M},\Psi,t} = \hat{\sigma}_{\mathcal{M},Y,t}, \quad \hat{\sigma}_{\Psi,P,t} = \hat{\sigma}_{Y,P,t} \quad (56)$$

We plot our estimates of these objects in Figure 14 in the Appendix. We observe that our estimates of demand uncertainty are an order of magnitude larger than our estimates of other uncertainties. This is natural given our large assumed value of R . But this does not necessarily imply that demand uncertainty is the only influential force shaping the slope of microeconomic or macroeconomic supply, since uncertainties enter our formulas in interaction with the slope of demand. We will return to this point when presenting our results.

Estimates of Model Objects. Armed with these estimates, we calculate our empirical proxies for the firm-level inverse supply elasticity as simple plug-in estimates:

$$\hat{\alpha}_{1,t} = \frac{\eta \hat{\sigma}_{P,t}^2 + \hat{\sigma}_{\mathcal{M},\Psi,t} + \hat{\sigma}_{P,\Psi,t} + \eta \hat{\sigma}_{\mathcal{M},P,t}}{\hat{\sigma}_{\Psi,t}^2 - \eta \hat{\sigma}_{\mathcal{M},\Psi,t} + \eta \hat{\sigma}_{P,\Psi,t} - \eta^2 \hat{\sigma}_{\mathcal{M},P,t}} \quad \text{and} \quad \hat{\omega}_{1,t} = \frac{\hat{\alpha}_{1,t}}{1 + \eta \hat{\alpha}_{1,t}} \quad (57)$$

Our calculation captures uncertainty about outcomes realized in quarter t and is measurable in data from quarter $t - 1$ and prior. It therefore describes incentives of a decisionmaker fixing a choice for quarter t based on their uncertainty at the end of quarter $t - 1$. We similarly compute our estimate of the inverse elasticity of aggregate supply:

$$\hat{\epsilon}_t^S = \gamma \frac{\kappa^M + \frac{\hat{\omega}_{1,t}}{\gamma} (1 - \kappa^M)}{(1 - \hat{\omega}_{1,t} \eta)(1 - \kappa^M)} \quad (58)$$

Table 1: Model Parameters and Estimation Approach

Parameter	Interpretation	Method	Value
η	Price elas. of demand	Match Broda and Weinstein (2006)	9
γ	Income effects	Match Flynn and Sastry (2022)	0.095
κ^M	Prec. of monetary info.	Match average slope of aggregate supply	0.40
$\sigma_{P,t}^2$	Price uncertainty	GARCH Model	Fig. 14
$\sigma_{\Psi,t}^2$	Demand uncertainty	GARCH + match Bloom et al. (2018)	Fig. 14
$\sigma_{\mathcal{M},\Psi,t}$	Cost-demand covariance	GARCH model	Fig. 14
$\sigma_{P,\Psi,t}$	Pice-demand covariance	GARCH model	Fig. 14
$\sigma_{\mathcal{M},P,t}$	Cost-price covariance	GARCH model	Fig. 14

Note: Description of model parameters, how we interpret them, how we estimate them, and their values. The time series for the time-varying uncertainties are presented in [Figure 14](#). The text of [Section 5.1](#) describes our methods in full detail.


Given our earlier estimates, this is pinned down given a single unknown parameter, the precision of private information about the money supply, κ^M . As mentioned earlier, we set κ^M so that the average $\hat{\epsilon}_t^S$ from 1960 to 2022 is 0.15. This yields a value of $\kappa^M = 0.40$.

5.2 Results: Aggregate Supply Over Time

[Figure 4](#) plots our main, quarterly-frequency estimates from [Equation 58](#) from 1960 Q1 to 2022 Q4. Aggregate supply was relatively flat in the 1960s, steepened in the 1970s and 1980s, before flattening again during the Volcker disinflation. Furthermore, aggregate supply was notably flat during the Great Moderation leading up to the financial crisis and then steepened again after the Covid-19 pandemic. [Figure 15](#) plots our secondary, annual-frequency estimates that directly incorporate data on microeconomic uncertainty from [Bloom et al. \(2018\)](#). These estimates imply an even more dramatic flattening from the 1970s to the Great Moderation, although we cannot use them to make predictions for the 1960s or 2020s as the estimates from [Bloom et al. \(2018\)](#) do not cover these periods.

Our estimates are consistent with external estimates of the time-varying slope of aggregate supply. In [Figure 4](#), we indicate estimates of the slope of aggregate supply by [Hazell et al. \(2022\)](#) with a blue dashed line and estimates by [Ball and Mazumder \(2011\)](#) with a red dotted line. The former authors use state-level data on unemployment and inflation and an instrumental variables (IV) strategy based on isolating state-level demand shocks. The latter authors use aggregate data on the output gap and inflation and measure their unconditional relationship; in the presence of confounding supply shocks, this should systematically understate the slope of aggregate supply.

Figure 4: The Slope of Aggregate Supply Over Time



Note: This Figure plots estimates of the inverse elasticity of aggregate supply as measured by Equations 57 and 58. The blue dashed line indicates the estimates of Hazell et al. (2022), based on state-level estimates of consumer prices and unemployment and an identification strategy that isolates local demand shocks (columns 3 and 4, panel B, of Table II). The red dotted line indicates the estimates of Ball and Mazumder (2011), based on aggregate data (column 4 of Table 3).

In Table 2, we quantitatively compare our model’s estimates for the flattening of the AS curve during the Great Moderation with those of the aforementioned references. The model with changing uncertainties but fixed structural parameters can explain 51% of the flattening estimated by Hazell et al. (2022) and 84% of the flattening estimated by Ball and Mazumder (2011). By an equivalent calculation, our model can also explain 54% of the latter authors’ estimated *steepening* of aggregate supply between 1960-1972 and 1973-1984.

Mechanisms: Which Uncertainties Matter More? The time variation in our estimate of ϵ_t^S arises from time-varying uncertainty about several objects. To better understand the role of each component of the calculation, we perform a variant calculation in which we hold each uncertainty term fixed at its sample average, one by one. We plot the results in Figure 5. The combination of time-varying uncertainty about the price-level and time-varying uncertainty about the relationship between prices and (real) marginal costs helps quantitatively explain the overall flattening from the 1970s into the Great Moderation. While uncertainty about demand is large, and features significant high-frequency variation (see Appendix Figure 14), it is not essential for our low-frequency predictions. By con-

Table 2: The Flattening Aggregate Supply Curve: Theory *vs.* Evidence

		HHNS (2022)	BM (2011)
Pre-Period	Period	1978-1990	1973-1984
	Data	0.198	0.371
	Model: Uncertainty Only	0.166	0.223
	Model: + Mkt. Power Trend	0.182	0.262
Post-Period	Period	1991-2018	1985-2007
	Data	0.090	0.136
	Model: Uncertainty Only	0.119	0.104
	Model: + Mkt. Power Trend	0.108	0.104
Flattening	Data	55%	63%
	Model: Uncertainty Only	28%	53%
	Model: + Mkt. Power Trend	41%	60%

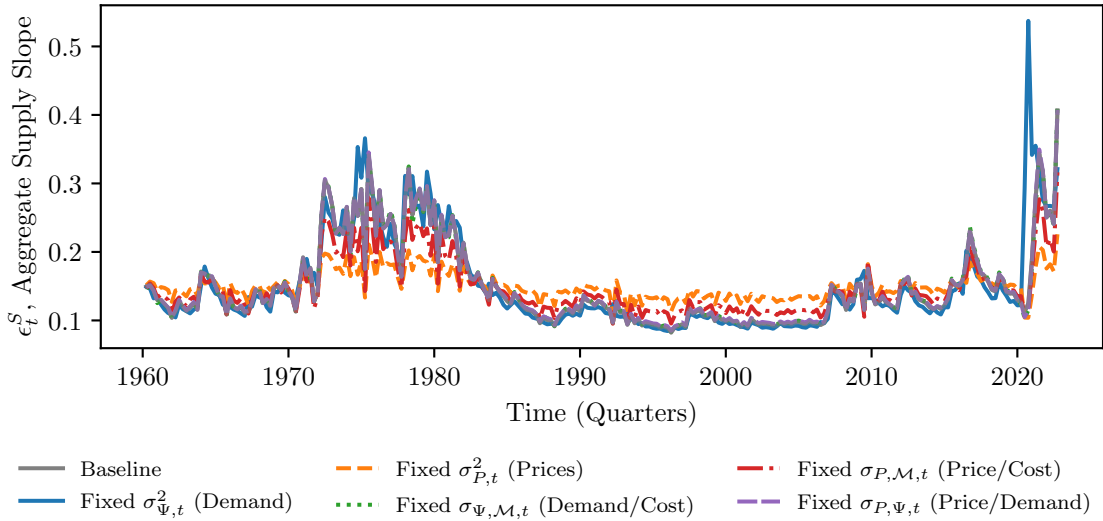
Note: Comparison of model predictions and literature estimates for the long-run flattening of the aggregate supply curve from [Hazell et al. \(2022\)](#) and [Ball and Mazumder \(2011\)](#). The “Uncertainty Only” model is the baseline with fixed elasticity of demand $\eta = 9$ (Figure 4). The “+ Market Power Trend” model is the “Small Change” scenario (Figure 6), which introduces a linear trend from $\eta = 12$ in 1960 to $\eta = 6$ in 2022. “Flattening” is $100 \cdot (1 - \text{SlopePost}/\text{SlopePre})$.

trast, incorporating demand uncertainty is essential to avoid predicting a large spike in the aggregate supply slope in the first several quarters of the Covid-19 lockdown.

Market Power and the Flattening Supply Curve. A recent literature has suggested that market power, as measured by rising mark-ups, has risen throughout time ([De Loecker et al., 2020](#); [Edmond et al., 2023](#); [Demirer, 2020](#)). Combined with our theoretical finding that increased market power flattens aggregate supply under plausible parameter values, this suggests another potentially relevant culprit for the long-run flattening of supply. To study this possibility, we now consider alternative calibrations of the model in which we allow a secular downward trend in the elasticity of demand faced by firms. We consider a “small change” in which η linearly declines from 12 to 6 and a more exaggerated “large change” in which η linearly declines from 15 to 3.

Introducing a decline in market power increases the slope of aggregate supply in the 1970s and decreases the slope in modern times (Figure 6). In Table 2, we observe that adding the “small change” in market power allows us to fit 75% of the flattening measured by [Hazell et al. \(2022\)](#) and 95% of the flattening measured by [Ball and Mazumder \(2011\)](#). The more extreme scenario for market power (the red dotted line in Figure 6) allows our model to explain a greater flattening, but incorrectly predicts a relatively flat aggregate supply curve in the 2020s.

Figure 5: Deconstructing the Slope of Aggregate Supply: Which Uncertainties Matter?




Note: This Figure plots estimates of the inverse elasticity of aggregate supply as measured by Equations 57 and 58, holding fixed one component of uncertainty at a time. The grey solid line corresponds to the baseline estimate from Figure 4.

Supply over a Longer Time Period. In our main analysis, we focus on the period after 1960. In an extension, we consider all data since World War II. This necessitates estimating a different GARCH model for macroeconomic uncertainty, so in principle, it could affect our estimates for the entire sample. We plot the results in Appendix Figure 16. We find very comparable estimates from 1960 onward, and a very large and volatile slope of aggregate supply between 1947 and 1960. The latter finding is consistent with there being large volatility in the money supply and the price level in the wake of the War and the Bretton Woods agreement.

Robustness and Sensitivity Analysis. Finally, in Appendix Figure 17 we report the robustness of these findings to different calibrations of η , R , and γ . Specifically, we recalibrate the model to match the average slope of 0.15 and check how different assumptions affect our model’s predictions for the long-run flattening during the Great Moderation. We predict a larger flattening under larger assumed values of η , which exacerbate the effect of changing inflation uncertainty; smaller R , which allows inflation uncertainty to play a larger role in the calculation; and larger γ , which makes real wages more cyclical, especially in the 1970s.

Figure 6: Rising Market Power and Flattening Aggregate Supply



Note: This Figure plots the inverse elasticity of aggregate supply under different scenarios of declining market power (Panel b). The solid line keeps η constant at 9 and corresponds to our baseline estimates. The blue, dashed line (“small change”) assumes a linear decline in η from 12 to 6 over the time sample. The red, dotted line assumes a linear decline in η from 15 to 3 over the time sample.

6 Conclusion

In this paper, we enrich firms’ supply decisions by allowing them to choose an arbitrary supply function that describes the price charged at each quantity of production. We show how to model supply functions in a macroeconomic setting and characterize how the optimal supply function depends on the elasticity of demand and the nature of uncertainty that firms face. Our framework yields rich implications when embedded in an otherwise standard monetary business cycle model. We find that greater market power and increased uncertainty about the price level relative to demand endogenously steepen the slope of aggregate supply. When mapped to the data, our theory explains the long-run flattening of the aggregate supply curve as an outcome of more hawkish monetary policy and rising market power.

On the basis of our analysis, we argue that supply schedules warrant serious consideration as an alternative model of firm conduct in macroeconomics for three core reasons. First, most existing work assumes that firms set a price in advance and commit to supply at the market-clearing quantity. Our results emphasize that this is not generally an optimal way for a firm to behave and that the macroeconomic conclusions that one draws about the effects of uncertainty, the propagation of monetary and productivity shocks, and the role of market power are highly sensitive to this choice. For example, the price-setting

assumption maximizes the degree of monetary non-neutrality and leaves no role for market power. Second, we have shown that working with supply schedules is analytically tractable under the standard assumptions in the literature, making them a viable alternative. Finally, taking the supply-schedule perspective yields economic predictions that are consistent with broad trends in US aggregate supply over the last 60 years.

Within the context of supply schedules and the macroeconomy, our study is only a first exploration; there remains much to examine, both empirically and theoretically. We highlight two particularly salient implications of our analysis that we leave open to future work. First, our work highlights the importance of firm-level supply elasticities as a critical moment for the business cycle. Empirical work that measures such elasticities and investigates how they systematically vary would be highly valuable for disciplining richer models with supply schedules. Such work would require detailed firm-level data on both prices and quantities. Second, it would be interesting to examine the conduct of optimal monetary policy in a setting with supply schedule choice. We have shown that more volatile monetary policy (perhaps associated with more discretionary monetary policy) can be self-defeating by making the economy endogenously less responsive to monetary stimulus. A complete normative analysis of this issue is an interesting avenue for future research.

References

- Afrouzi, H. and Yang, C. (2021). Dynamic rational inattention and the Phillips curve. Technical report, Columbia University.
- Aiginger, K. (1999). The use of game theoretical models for empirical industrial organization. In Mueller, D., Haid, A., and Weigand, J., editors, *Competition, Efficiency and Welfare—Essays in Honor of Manfred Neumann*, pages 253–277. Kluwer Academic Publishers, Dordrecht.
- Angeletos, G.-M. and Huo, Z. (2021). Myopia and anchoring. *American Economic Review*, 111(4):1166–1200.
- Angeletos, G.-M. and La’O, J. (2010). Noisy business cycles. *NBER Macroeconomics Annual*, 24(1):319–378.
- Angeletos, G.-M. and La’O, J. (2013). Sentiments. *Econometrica*, 81(2):739–779.
- Ball, L. and Mazumder, S. (2011). Inflation dynamics and the great recession. *Brookings Papers on Economic Activity*, 2011(1):337–381.
- Barro, R. J. (1972). A theory of monopolistic price adjustment. *The Review of Economic Studies*, 39(1):17–26.

- Basu, S., Fernald, J. G., and Kimball, M. S. (2006). Are technology improvements contractionary? *American Economic Review*, 96(5):1418–1448.
- Benhabib, J., Wang, P., and Wen, Y. (2015). Sentiments and aggregate demand fluctuations. *Econometrica*, 83(2):549–585.
- Blanchard, O. J. (2016). The Phillips curve: back to the '60s? *American Economic Review*, 106(5):31–34.
- Blanchard, O. J. and Bernanke, B. S. (2023). What caused the US pandemic-era inflation? Working Paper w31417, National Bureau of Economic Research.
- Blanchard, O. J. and Kiyotaki, N. (1987). Monopolistic competition and the effects of aggregate demand. *The American Economic Review*, pages 647–666.
- Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., and Terry, S. J. (2018). Really uncertain business cycles. *Econometrica*, 86(3):1031–1065.
- Bollerslev, T. (1990). Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model. *The Review of Economics and Statistics*, pages 498–505.
- Broda, C. and Weinstein, D. E. (2006). Globalization and the gains from variety. *The Quarterly Journal of Economics*, 121(2):541–585.
- Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398.
- Caplin, A. S. and Spulber, D. F. (1987). Menu costs and the neutrality of money. *The Quarterly Journal of Economics*, 102(4):703–725.
- Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45.
- De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644.
- Demirer, M. (2020). Production function estimation with factor-augmenting technology: An application to markups. Mimeo, MIT.
- Drenik, A. and Perez, D. J. (2020). Price setting under uncertainty about inflation. *Journal of Monetary Economics*, 116:23–38.
- Edmond, C., Midrigan, V., and Xu, D. Y. (2023). How costly are markups? *Journal of Political Economy*, 131(7):000–000.
- Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Accessed from <https://www.johnfernald.net/>.
- Flynn, J. P. and Sastry, K. A. (2022). Attention cycles. Working Paper 3592107, SSRN.
- Foster, L., Haltiwanger, J., and Syverson, C. (2008). Reallocation, firm turnover, and effi-

- ciency: Selection on productivity or profitability? *American Economic Review*, 98(1):394–425.
- Fujiwara, I. and Matsuyama, K. (2022). Competition and the Phillips curve. Discussion Paper DP17521, Centre for Economic Policy Research.
- Garcia, R. and Schaller, H. (2002). Are the effects of monetary policy asymmetric? *Economic Inquiry*, 40(1):102–119.
- Golosov, M. and Lucas, R. E. (2007). Menu costs and Phillips curves. *Journal of Political Economy*, 115(2):171–199.
- Grossman, S. (1976). On the efficiency of competitive stock markets where trades have diverse information. *The Journal of Finance*, 31(2):573–585.
- Grossman, S. J. (1981). Nash equilibrium and the industrial organization of markets with large fixed costs. *Econometrica*, pages 1149–1172.
- Hart, O. (1985). Imperfect competition in general equilibrium: an overview of recent work. In *Frontiers of Economics*. Basil Blackwell, Oxford.
- Hazell, J., Herreno, J., Nakamura, E., and Steinsson, J. (2022). The slope of the Phillips Curve: evidence from US states. *The Quarterly Journal of Economics*, 137(3):1299–1344.
- Hellwig, C. and Veldkamp, L. (2009). Knowing what others know: Coordination motives in information acquisition. *The Review of Economic Studies*, 76(1):223–251.
- Hellwig, C. and Venkateswaran, V. (2009). Setting the right prices for the wrong reasons. *Journal of Monetary Economics*, 56:S57–S77.
- Hellwig, M. F. (1980). On the aggregation of information in competitive markets. *Journal of Economic Theory*, 22(3):477–498.
- Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.
- Klemperer, P. D. and Meyer, M. A. (1986). Price competition vs. quantity competition: the role of uncertainty. *The RAND Journal of Economics*, pages 618–638.
- Klemperer, P. D. and Meyer, M. A. (1989). Supply function equilibria in oligopoly under uncertainty. *Econometrica*, pages 1243–1277.
- Kyle, A. S. (1989). Informed speculation with imperfect competition. *The Review of Economic Studies*, 56(3):317–355.
- Lo, M. C. and Piger, J. (2005). Is the response of output to monetary policy asymmetric? Evidence from a regime-switching coefficients model. *Journal of Money, Credit and Banking*, pages 865–886.
- Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory*, 4(2):103–124.

- Lucas, R. E. (1973). Some international evidence on output-inflation tradeoffs. *The American Economic Review*, pages 326–334.
- Lucas, R. E. (1975). An equilibrium model of the business cycle. *Journal of Political Economy*, 83(6):1113–1144.
- Maćkowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.
- Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips Curve. *The Quarterly Journal of Economics*, 117(4):1295–1328.
- Mongey, S. (2021). Market structure and monetary non-neutrality. Technical Report 29233, National Bureau of Economic Research.
- Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.
- Pavan, A., Sundaresan, S., and Vives, X. (2022). (in)efficiency in information acquisition and aggregation through prices. Discussion Paper DP17025, CEPR.
- Phelps, E. S. (1970). Introduction: The new microeconomics in employment and inflation theory. In *Microeconomic foundations of employment and inflation theory*, volume 1. Norton, New York.
- Poole, W. (1970). Optimal choice of monetary policy instruments in a simple stochastic macro model. *The Quarterly Journal of Economics*, 84(2):197–216.
- Ramey, V. A. (2016). Macroeconomic shocks and their propagation. In *Handbook of Macroeconomics*, volume 2, pages 71–162. Elsevier.
- Reis, R. (2006). Inattentive producers. *The Review of Economic Studies*, 73(3):793–821.
- Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94(4):1055–1084.
- Sheshinski, E. and Weiss, Y. (1977). Inflation and costs of price adjustment. *The Review of Economic Studies*, 44(2):287–303.
- Taylor, J. B. (1980). Aggregate dynamics and staggered contracts. *Journal of Political Economy*, 88(1):1–23.
- Tenreyro, S. and Thwaites, G. (2016). Pushing on a string: US monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, 8(4):43–74.
- Vavra, J. (2014). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *The Quarterly Journal of Economics*, 129(1):215–258.
- Vives, X. (2011). Strategic supply function competition with private information. *Econometrica*, 79(6):1919–1966.

- Vives, X. (2017). Endogenous public information and welfare in market games. *The Review of Economic Studies*, 84(2):935–963.
- Wang, O. and Werning, I. (2022). Dynamic oligopoly and price stickiness. *American Economic Review*, 112(8):2815–2849.
- Weise, C. L. (1999). The asymmetric effects of monetary policy: A nonlinear vector autoregression approach. *Journal of Money, Credit and Banking*, pages 85–108.
- Weitzman, M. L. (1974). Prices vs. quantities. *The Review of Economic Studies*, 41(4):477–491.
- Woodford, M. (2003a). Imperfect common knowledge and the effects of monetary policy. In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Princeton University Press, Princeton.
- Woodford, M. (2003b). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.

Appendices

A Omitted Proofs

A.1 Proof of Proposition 1

To derive the optimal price, we take the first-order condition of Equation 6. This yields:

$$\eta p^{*\eta-1} \mathbb{E} [\Lambda \mathcal{M} P^\eta \Psi] = (\eta - 1) p^{*\eta} \mathbb{E} [\Lambda P^{\eta-1} \Psi] \quad (59)$$

which rearranges to Equation 7. Substituting p^* into Equation 6, we obtain Equation 8:

$$\begin{aligned} V^P &= \mathbb{E} \left[\Lambda \left(\frac{p^*}{P} - \mathcal{M} \right) \Psi \left(\frac{p^*}{P} \right)^{-\eta} \right] \\ &= \mathbb{E} \left[\Lambda \left(\frac{1}{P} \frac{\eta}{\eta-1} \frac{\xi^P}{\zeta^P} - \mathcal{M} \right) \Psi P^\eta \left(\frac{\eta}{\eta-1} \right)^{-\eta} \left(\frac{\xi^P}{\zeta^P} \right)^{-\eta} \right] \\ &= \left(\frac{\eta}{\eta-1} \right)^{-\eta} \left[\frac{\eta}{\eta-1} \left(\frac{\xi^P}{\zeta^P} \right)^{1-\eta} \zeta^P - \left(\frac{\xi^P}{\zeta^P} \right)^{-\eta} \xi^P \right] \\ &= \left(\frac{\eta}{\eta-1} \right)^{-\eta} \left(\frac{\eta}{\eta-1} - 1 \right) \xi^{P^{1-\eta}} \zeta^{P^\eta} = \frac{1}{\eta-1} \left(\frac{\eta}{\eta-1} \right)^{-\eta} \mathbb{E} [\Lambda \mathcal{M} P^\eta \Psi]^{1-\eta} \mathbb{E} [\Lambda P^{\eta-1} \Psi]^\eta \end{aligned} \quad (60)$$

where $\xi^P = \mathbb{E} [\Lambda \mathcal{M} P^\eta \Psi]$ and $\zeta^P = \mathbb{E} [\Lambda P^{\eta-1} \Psi]$.

To derive the optimal quantity, we take the first-order condition of Equation 9. This yields:

$$\mathbb{E} [\Lambda \mathcal{M}] = \frac{\eta-1}{\eta} q^{*\eta-1} \mathbb{E} [\Lambda \Psi^{\frac{1}{\eta}}] \quad (61)$$

which rearranges to Equation 10. Substituting q^* into Equation 9, we obtain Equation 11:

$$\begin{aligned} V^Q &= \mathbb{E} \left[\Lambda \left(\left(\frac{x}{\Psi} \right)^{-\frac{1}{\eta}} - \mathcal{M} \right) q \right] \\ &= \mathbb{E} \left[\Lambda \left(\frac{\eta}{\eta-1} \frac{\xi^Q}{\zeta^Q} \Psi^{\frac{1}{\eta}} - \mathcal{M} \right) \left(\frac{\eta}{\eta-1} \right)^{-\eta} \left(\frac{\xi^Q}{\zeta^Q} \right)^{-\eta} \right] \\ &= \left(\frac{\eta}{\eta-1} \right)^{-\eta} \left[\frac{\eta}{\eta-1} \left(\frac{\xi^Q}{\zeta^Q} \right)^{1-\eta} \zeta^Q - \left(\frac{\xi^Q}{\zeta^Q} \right)^{-\eta} \xi^Q \right] \\ &= \left(\frac{\eta}{\eta-1} \right)^{-\eta} \left[\frac{\eta}{\eta-1} - 1 \right] \xi^{Q^{1-\eta}} \zeta^{Q^\eta} = \frac{1}{\eta-1} \left(\frac{\eta}{\eta-1} \right)^{-\eta} \mathbb{E} [\Lambda \mathcal{M}]^{1-\eta} \mathbb{E} [\Lambda \Psi^{\frac{1}{\eta}}]^\eta \end{aligned} \quad (62)$$

where $\xi^Q = \mathbb{E} [\Lambda \mathcal{M}]$ and $\zeta^Q = \mathbb{E} [\Lambda \Psi^{\frac{1}{\eta}}]$.

To find Δ , we first write $V^P - V^Q$ as:

$$\Delta = \eta (\log \zeta^P - \log \zeta^Q) - (\eta - 1) (\log \xi^P - \log \xi^Q) \quad (63)$$

Thus, it suffices to compute $(\zeta^P, \zeta^Q, \xi^P, \xi^Q)$. Given log-normality of $(\Psi, P, \Lambda, \mathcal{M})$, we may write:

$$\begin{pmatrix} \log \Psi \\ \log P \\ \log \Lambda \\ \log \mathcal{M} \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_\Psi \\ \mu_P \\ \mu_\Lambda \\ \mu_{\mathcal{M}} \end{pmatrix}, \begin{pmatrix} \sigma_\Psi^2 & \sigma_{\Psi,P} & \sigma_{\Psi,\Lambda} & \sigma_{\Psi,\mathcal{M}} \\ \sigma_{\Psi,P} & \sigma_P^2 & \sigma_{P,\Lambda} & \sigma_{P,\mathcal{M}} \\ \sigma_{\Psi,\Lambda} & \sigma_{P,\Lambda} & \sigma_\Lambda^2 & \sigma_{\Lambda,\mathcal{M}} \\ \sigma_{\Psi,\mathcal{M}} & \sigma_{P,\mathcal{M}} & \sigma_{\Lambda,\mathcal{M}} & \sigma_{\mathcal{M}}^2 \end{pmatrix} \right) \quad (64)$$

To compute the second term of Equation 63, we compute:

$$\begin{aligned} \log \xi^P &= \log \mathbb{E} [\Lambda \mathcal{M} P^\eta \Psi] = \log \mathbb{E} [\exp\{\log \Lambda + \log \mathcal{M} + \eta \log P + \log \Psi\}] \\ &= \mu_\Lambda + \mu_{\mathcal{M}} + \eta \mu_P + \mu_\Psi + \frac{1}{2} (\sigma_\Lambda^2 + \sigma_{\mathcal{M}}^2 + \eta^2 \sigma_P^2 + \sigma_\Psi^2) \\ &\quad + \sigma_{\Lambda,\mathcal{M}} + \eta \sigma_{\Lambda,P} + \sigma_{\Lambda,\Psi} + \eta \sigma_{\mathcal{M},P} + \sigma_{\mathcal{M},\Psi} + \eta \sigma_{P,\Psi} \end{aligned} \quad (65)$$

and

$$\begin{aligned} \log \xi^Q &= \log \mathbb{E} [\Lambda \mathcal{M}] = \log \mathbb{E} [\exp\{\log \Lambda + \log \mathcal{M}\}] \\ &= \mu_\Lambda + \mu_{\mathcal{M}} + \frac{1}{2} (\sigma_\Lambda^2 + \sigma_{\mathcal{M}}^2) + \sigma_{\Lambda,\mathcal{M}} \end{aligned} \quad (66)$$

Thus, the second term of Equation 63 is given by:

$$\begin{aligned} (\eta - 1)(\log \xi^P - \log \xi^Q) &= (\eta - 1) \left[\eta \mu_P + \mu_\Psi + \frac{1}{2} (\eta^2 \sigma_P^2 + \sigma_\Psi^2) \right. \\ &\quad \left. + \eta \sigma_{\Lambda,P} + \sigma_{\Lambda,\Psi} + \eta \sigma_{\mathcal{M},P} + \sigma_{\mathcal{M},\Psi} + \eta \sigma_{P,\Psi} \right] \end{aligned} \quad (67)$$

To compute the first term of Equation 63, we compute:

$$\begin{aligned} \log \zeta^P &= \log \mathbb{E} [\Lambda P^{\eta-1} \Psi] = \log \mathbb{E} [\exp\{\log \Lambda + (\eta - 1) \log P + \log \Psi\}] \\ &= \mu_\Lambda + (\eta - 1) \mu_P + \mu_\Psi + \frac{1}{2} (\sigma_\Lambda^2 + (\eta - 1)^2 \sigma_P^2 + \sigma_\Psi^2) \\ &\quad + (\eta - 1) \sigma_{\Lambda,P} + \sigma_{\Lambda,\Psi} + (\eta - 1) \sigma_{P,\Psi} \end{aligned} \quad (68)$$

and

$$\begin{aligned} \log \zeta^Q &= \log \mathbb{E} [\Lambda \Psi^{\frac{1}{\eta}}] = \log \mathbb{E} \left[\exp \left\{ \log \Lambda + \frac{1}{\eta} \log \Psi \right\} \right] \\ &= \mu_\Lambda + \frac{1}{\eta} \mu_\Psi + \frac{1}{2} \left(\sigma_\Lambda^2 + \frac{1}{\eta^2} \sigma_\Psi^2 \right) + \frac{1}{\eta} \sigma_{\Lambda,\Psi} \end{aligned} \quad (69)$$

Thus, the first term of Equation 63 is given by:

$$\begin{aligned} \eta(\log \zeta^P - \log \zeta^Q) = \eta \left[(\eta - 1)\mu_P + \frac{\eta - 1}{\eta}\mu_\Psi + \frac{1}{2} \left((\eta - 1)^2\sigma_P^2 + \left(1 - \frac{1}{\eta^2}\right)\sigma_\Psi^2 \right) \right. \\ \left. + (\eta - 1)\sigma_{\Lambda,P} + \frac{\eta - 1}{\eta}\sigma_{\Lambda,\Psi} + (\eta - 1)\sigma_{P,\Psi} \right] \end{aligned} \quad (70)$$

Taking the difference between the two terms, we obtain Equation 13:

$$\begin{aligned} \Delta &= \frac{1}{2} \left[(\eta(\eta - 1)^2 - \eta^2(\eta - 1))\sigma_P^2 + \left(\eta \left(1 - \frac{1}{\eta^2}\right) - (\eta - 1) \right)\sigma_\Psi^2 \right] \\ &\quad - \eta(\eta - 1)\sigma_{\mathcal{M},P} - (\eta - 1)\sigma_{\mathcal{M},\Psi} \\ &= \frac{1}{2} \left(\frac{\eta - 1}{\eta}\sigma_\Psi^2 - \eta(\eta - 1)\sigma_P^2 - 2(\eta - 1)\sigma_{\mathcal{M},\Psi} - 2\eta(\eta - 1)\sigma_{\mathcal{M},P} \right) \end{aligned} \quad (71)$$

This completes the derivation of Equation 13. The comparative statics in the variance terms are immediate from the expression. The comparative statics for $\text{sgn}(\Delta)$ come from inspection of the expression for $\Delta/(\eta - 1)$, which has the same sign as Δ under the maintained assumption that $\eta > 1$.

A.2 Proof of Theorem 1

Proof. We first derive Equation 15 using variational methods. Consider a variation $\tilde{p}(z) = p(z) + \varepsilon h(z)$. The expected payoff under this variation is:

$$J(\varepsilon; h) = \int_{\mathbb{R}_{++}^4} \Lambda \left(\frac{p(z) + \varepsilon h(z)}{P} - \mathcal{M} \right) z (p(z) + \varepsilon h(z))^{-\eta} dG(\Lambda, P, \mathcal{M}, z) \quad (72)$$

A necessary condition for the optimality of a function p is that $J_\varepsilon(0; h) = 0$ for all F -measurable h . Taking this derivative and setting $\varepsilon = 0$, we obtain:

$$0 = \int_{\mathbb{R}_{++}^4} \left[\Lambda \frac{h(z)}{P} z p(z)^{-\eta} - \eta \Lambda h(z) \left(\frac{p(z)}{P} - \mathcal{M} \right) z p(z)^{-\eta-1} \right] dG(\Lambda, P, \mathcal{M}, z) \quad (73)$$

Consider h functions given by the Dirac delta functions on each z , $h(z) = \delta_z$. This condition becomes:

$$0 = \int_{\mathbb{R}_{++}^3} \left[\Lambda \frac{1}{P} t p(t)^{-\eta} - \eta \Lambda \left(\frac{p(t)}{P} - \mathcal{M} \right) t p(t)^{-\eta-1} \right] g(\Lambda, P, \mathcal{M}, t) d\Lambda dP d\mathcal{M} \quad (74)$$

for all $t \in \mathbb{R}_{++}$. This is equivalent to:

$$\begin{aligned} 0 &= \int_{\mathbb{R}_{++}^3} \left[\Lambda \frac{1}{P} t p(t)^{-\eta} - \eta \Lambda \left(\frac{p(t)}{P} - \mathcal{M} \right) t p(t)^{-\eta-1} \right] g(\Lambda, P, \mathcal{M}|t) d\Lambda dP d\mathcal{M} \\ &= (1 - \eta) \mathbb{E} \left[\Lambda \frac{1}{P} | z = t \right] t p(t)^{-\eta} + \eta \mathbb{E} [\Lambda \mathcal{M} | z = t] t p(t)^{-\eta-1} \end{aligned} \quad (75)$$

Thus, we have that an optimal solution necessarily follows:

$$p(t) = \frac{\eta}{\eta - 1} \frac{\mathbb{E}[\Lambda \mathcal{M} | z = t]}{\mathbb{E}[\Lambda P^{-1} | z = t]} \quad (76)$$

as claimed in Equation 15.

We now evaluate the expectations. Using log-normality,

$$\begin{aligned} \mathbb{E}[\Lambda \mathcal{M} | z = t] &= \exp \left\{ \mu_{\Lambda|z}(t) + \mu_{\mathcal{M}|z}(t) + \frac{1}{2} \sigma_{\Lambda|z}^2 + \frac{1}{2} \sigma_{\mathcal{M}|z}^2 + \sigma_{\Lambda, \mathcal{M}|z} \right\} \\ \mathbb{E}[\Lambda P^{-1} | z = t] &= \exp \left\{ \mu_{\Lambda|z}(t) - \mu_{P|z}(t) + \frac{1}{2} \sigma_{\Lambda|z}^2 + \frac{1}{2} \sigma_{P|z}^2 - \sigma_{\Lambda, P|z} \right\} \end{aligned} \quad (77)$$

where $\mu_{X|z} = \mathbb{E}[\log X | \log z]$ and $\sigma_{X, Y|z} = \text{Cov}[\log X, \log Y | \log z]$. Thus,

$$\frac{\mathbb{E}[\Lambda \mathcal{M} | z = t]}{\mathbb{E}[\Lambda P^{-1} | z = t]} = \exp \left\{ \mu_{\mathcal{M}|z}(t) + \mu_{P|z}(t) + \frac{1}{2} \sigma_{\mathcal{M}|z}^2 - \frac{1}{2} \sigma_{P|z}^2 + \sigma_{\Lambda, \mathcal{M}|z} + \sigma_{\Lambda, P|z} \right\} \quad (78)$$

Using standard formulas for Gaussian conditional expectations,

$$\begin{aligned} \mu_{\mathcal{M}|z}(t) &= \left(1 - \frac{\sigma_{\mathcal{M}, z}}{\sigma_z^2} \right) \mu_{\mathcal{M}} + \frac{\sigma_{\mathcal{M}, z}}{\sigma_z^2} \log t & \mu_{P|z}(t) &= \left(1 - \frac{\sigma_{P, z}}{\sigma_z^2} \right) \mu_P + \frac{\sigma_{P, z}}{\sigma_z^2} \log t \\ \sigma_{\mathcal{M}|z}^2 &= \sigma_{\mathcal{M}}^2 - \frac{\sigma_{\mathcal{M}, z}^2}{\sigma_z^2} & \sigma_{P|z}^2 &= \sigma_P^2 - \frac{\sigma_{P, z}^2}{\sigma_z^2} \\ \sigma_{\Lambda, \mathcal{M}|z} &= \sigma_{\Lambda, \mathcal{M}} - \frac{\sigma_{\Lambda, z} \sigma_{\mathcal{M}, z}}{\sigma_z^2} & \sigma_{\Lambda, P|z} &= \sigma_{\Lambda, P} - \frac{\sigma_{\Lambda, z} \sigma_{P, z}}{\sigma_z^2} \end{aligned} \quad (79)$$

where:

$$\begin{aligned} \sigma_z^2 &= \sigma_{\Psi}^2 + \eta^2 \sigma_P^2 + 2\eta \sigma_{\Psi, P} & \sigma_{P, z} &= \sigma_{P, \Psi} + \eta \sigma_P^2 \\ \sigma_{\mathcal{M}, z} &= \sigma_{\mathcal{M}, \Psi} + \eta \sigma_{\mathcal{M}, P} & \sigma_{\Lambda, z} &= \sigma_{\Lambda, \Psi} + \eta \sigma_{\Lambda, P} \end{aligned} \quad (80)$$

We now combine these expressions with Equation 76 to derive the optimal supply function. We first observe that

$$\log p = \omega_0 + \omega_1 \log t \quad (81)$$

where:

$$\omega_0 = \log \frac{\eta}{\eta - 1} + \left(1 - \frac{\sigma_{\mathcal{M},z}}{\sigma_z^2}\right) \mu_{\mathcal{M}} + \left(1 - \frac{\sigma_{P,z}}{\sigma_z^2}\right) \mu_P + \frac{1}{2} \sigma_{\mathcal{M}|z}^2 - \frac{1}{2} \sigma_{P|z}^2 + \sigma_{\Lambda, \mathcal{M}|z} + \sigma_{\Lambda, P|z} \quad (82)$$

$$\omega_1 = \frac{\sigma_{\mathcal{M},z} + \sigma_{P,z}}{\sigma_z^2} = \frac{\sigma_{\mathcal{M},\Psi} + \eta \sigma_{\mathcal{M},P} + \sigma_{P,\Psi} + \eta \sigma_P^2}{\sigma_{\Psi}^2 + \eta^2 \sigma_P^2 + 2\eta \sigma_{\Psi,P}} \quad (83)$$

Next, using the demand curve, we observe that $z = qp^\eta$. Therefore, $\log t = \log q - \eta \log p$. Substituting this into Equation 81, and re-arranging, we obtain

$$\log p = \alpha_0 + \alpha_1 \log q \quad (84)$$

where:

$$\alpha_0 = \frac{\omega_0}{1 - \eta \omega_1}, \quad \alpha_1 = \frac{\omega_1}{1 - \eta \omega_1} \quad (85)$$

We finally derive the claimed expression for α_1 ,

$$\begin{aligned} \alpha_1 &= \frac{\frac{\sigma_{\mathcal{M},\Psi} + \eta \sigma_{\mathcal{M},P} + \sigma_{P,\Psi} + \eta \sigma_P^2}{\sigma_{\Psi}^2 + \eta^2 \sigma_P^2 + 2\eta \sigma_{\Psi,P}}}{1 - \eta \frac{\sigma_{\mathcal{M},\Psi} + \eta \sigma_{\mathcal{M},P} + \sigma_{P,\Psi} + \eta \sigma_P^2}{\sigma_{\Psi}^2 + \eta^2 \sigma_P^2 + 2\eta \sigma_{\Psi,P}}} \\ &= \frac{\sigma_{\mathcal{M},\Psi} + \eta \sigma_{\mathcal{M},P} + \sigma_{P,\Psi} + \eta \sigma_P^2}{\sigma_{\Psi}^2 + \eta \sigma_{\Psi,P} - \eta \sigma_{\mathcal{M},\Psi} - \eta^2 \sigma_{\mathcal{M},P}} \end{aligned} \quad (86)$$

Completing the proof. □

A.3 Proof of Corollary 1

Proof. If $2\eta \sigma_{\mathcal{M},P} + \sigma_{\mathcal{M},\Psi} \geq \sigma_{P,\Psi}$, then the denominator of Equation 17 is decreasing in η . Moreover, if $\sigma_{\mathcal{M},P} \geq 0$, the numerator is increasing in η . Hence, α_1 is increasing in η whenever $\alpha_1 > 0$. □

A.4 Proof of Proposition 2

Proof. We first derive Equation 32. From Equations 30 and 31, we obtain:

$$\frac{1}{M_t} + \beta \mathbb{E}_t \left[C_{t+1}^{-\gamma} \frac{1}{P_{t+1}} \right] = \beta (1 + i_t) \mathbb{E}_t \left[C_{t+1}^{-\gamma} \frac{1}{P_{t+1}} \right] \quad (87)$$

It follows that:

$$\frac{1}{M_t} = \beta i_t \mathbb{E}_t \left[C_{t+1}^{-\gamma} \frac{1}{P_{t+1}} \right] = \frac{i_t}{1 + i_t} C_t^{-\gamma} \frac{1}{P_t} \quad (88)$$

where the second equality uses Equation 31 once again. This rearranges to Equation 32.

We next derive Equation 35. Substituting equation 32 into Equation 31, we obtain:

$$\frac{1+i_t}{i_t} \frac{1}{M_t} = \beta(1+i_t) \mathbb{E}_t \left[\frac{1+i_{t+1}}{i_{t+1}} \frac{1}{M_{t+1}} \right] \quad (89)$$

Dividing both sides by $(1+i_t)$, multiplying by M_t , and then adding one, we obtain:

$$\frac{1+i_t}{i_t} = 1 + \beta \mathbb{E}_t \left[\frac{1+i_{t+1}}{i_{t+1}} \frac{M_t}{M_{t+1}} \right] = 1 + \beta \mathbb{E}_t \left[\exp\{-\mu - \sigma_t^M \varepsilon_t^M\} \frac{1+i_{t+1}}{i_{t+1}} \right] \quad (90)$$

where the second equality exploits the fact that M_t follows a random walk with drift. If we guess that i_t is deterministic and define $x_t = \frac{1+i_t}{i_t}$, then we obtain that:

$$x_t = 1 + \delta_t x_{t+1} \quad (91)$$

where:

$$\delta_t = \beta \exp \left\{ -\mu + \frac{1}{2} (\sigma_t^M)^2 \right\} \quad (92)$$

We observe that $\delta_t \in [0, \beta]$ for all t due to the assumption that $\frac{1}{2} (\sigma_t^M)^2 \leq \mu_M$. Solving this equation forward, we obtain that:

$$x_t = 1 + \delta_t \sum_{i=1}^{T-1} \prod_{j=1}^i \delta_{t+j} + \delta_t \left(\prod_{j=1}^T \delta_{t+j} \right) x_{t+T+1} \quad (93)$$

Taking the limit $T \rightarrow \infty$, this becomes:

$$x_t = 1 + \delta_t \sum_{i=1}^{\infty} \prod_{j=1}^i \delta_{t+j} + \delta_t \lim_{T \rightarrow \infty} \left(\prod_{j=1}^T \delta_{t+j} \right) x_{t+T+1} \quad (94)$$

where the final term can be bounded using the fact that $\delta_t \in [0, \beta]$:

$$0 \leq \delta_t \lim_{T \rightarrow \infty} \left(\prod_{j=1}^T \delta_{t+j} \right) x_{t+T+1} \leq \lim_{T \rightarrow \infty} \beta^{T+1} x_{t+T+1} \quad (95)$$

The household's transversality condition ensures that this upper bound is zero. Formally, the transversality condition (necessary for the optimality of the household's choices) is that:

$$\lim_{T \rightarrow \infty} \beta^T \frac{C_T^{-\gamma}}{P_T} (M_T + (1+i_T)B_T) = 0 \quad (96)$$

Moreover, as $B_t = 0$ for all $t \in \mathbb{N}$, this reduces to $\lim_{T \rightarrow \infty} \beta^T \frac{C_T^{-\gamma}}{P_T} M_T = 0$. By Equation 88, we have that $\frac{x_t}{M_t} = \frac{C_t^{-\gamma}}{P_t}$. Thus, the transversality condition reduces to $\lim_{T \rightarrow \infty} \beta^T x_T = 0$. Combining this with Equation 95, we have that $\lim_{T \rightarrow \infty} \left(\prod_{j=1}^T \delta_{t+j} \right) x_{t+T+1} = 0$. Equation 35 follows:

$$\frac{1 + i_t}{i_t} = 1 + \beta \exp\left\{-\mu + \frac{1}{2}(\sigma_t^M)^2\right\} \sum_{i=1}^{\infty} \prod_{j=1}^i \beta \exp\left\{-\mu + \frac{1}{2}\sigma_{M,t+j}^2\right\} \quad (97)$$

The formulae in Equation 37 then follow. In particular, $\Psi_{it} = \vartheta_{it} C_t$ follows from comparing Equations 2 and 36. $P_t = \frac{i_t}{1+i_t} C_t^{-\gamma} M_t$ follows from Equation 32. $\Lambda_t = C_t^{-\gamma}$ is the households marginal utility from consumption. Finally, $\mathcal{M}_{it} = \frac{1}{z_{it} A_t} \frac{w_{it}}{P_t} = \frac{\phi_{it} C_t^\gamma}{z_{it} A_t}$ follows from Equation 29. \square

A.5 Proof of Proposition 3

Proof. We suppress dependence on t to ease of notation. Consider a plan:

$$\log p_i = \log \tilde{\alpha}_{0,i} + \alpha_1 \log q_i \quad (98)$$

where $\tilde{\alpha}_{0,i} = e^{\alpha_0, i}$. The demand-supply relationship that the firm faces is:

$$\log p_i = -\frac{1}{\eta} (\log q_i - \log \Psi) + \log P \quad (99)$$

The realized quantity therefore is:

$$\log q_i = \frac{-\eta}{1 + \eta\alpha_1} \log \tilde{\alpha}_{0,i} + \frac{1}{1 + \eta\alpha_1} \log \Psi_i P^\eta \quad (100)$$

and the realized price is:

$$\log p_i = \frac{1}{1 + \eta\alpha_1} \log \tilde{\alpha}_{0,i} + \frac{\alpha_1}{1 + \eta\alpha_1} \log \Psi_i P^\eta \quad (101)$$

It is useful to make the change of variables $\omega_1 = \frac{\alpha_1}{1 + \eta\alpha_1}$, which implies that we may write

$$\log p_i = (1 - \eta\omega_1) \log \tilde{\alpha}_{0,i} + \omega_1 \log \Psi_i P^\eta \quad (102)$$

Our goal is to express dynamics only as a function of ω_1 . We first find the optimal $\alpha_{0,i}$ in terms of ω_1 . The firm therefore solves:

$$\max_{\tilde{\alpha}_{0,i}} \mathbb{E}_i \left[\Lambda \left(\frac{p_i}{P} - \mathcal{M}_i \right) \left(\frac{p_i}{P} \right)^{-\eta} \Psi_i \right] \quad (103)$$

Substituting for the realized price using the demand-supply relationship yields:

$$\max_{\tilde{\alpha}_{0,i}} \mathbb{E} \left[\Lambda \left(\frac{\tilde{\alpha}_{0,i}^{1-\eta\omega_1}}{P} (\Psi_i P^\eta)^{\omega_1} - \mathcal{M}_i \right) \tilde{\alpha}_{0,i}^{\eta^2\omega_1-\eta} (\Psi_i P^\eta)^{1-\eta\omega_1} \right] \quad (104)$$

The optimal $\tilde{\alpha}_{0,i}$ is:

$$\tilde{\alpha}_{0,i}^{1-\eta\omega_1} = \frac{\eta}{\eta-1} \frac{\mathbb{E}_i[\Lambda \mathcal{M}_i (\Psi_i P^\eta)^{1-\eta\omega_1}]}{\mathbb{E}_i[\frac{\Lambda}{P} (\Psi_i P^\eta)^{1-\eta\omega_1+\omega_1}]} \quad (105)$$

Substituting back into the realized price yields:

$$p_i = \frac{\eta}{\eta-1} \frac{\mathbb{E}_i[\Lambda \mathcal{M}_i (\Psi_i P^\eta)^{1-\eta\omega_1}]}{\mathbb{E}_i[\frac{\Lambda}{P} (\Psi_i P^\eta)^{1-\eta\omega_1+\omega_1}]} (\Psi_i P^\eta)^{\omega_1} \quad (106)$$

We may express this only in terms of P by using Proposition 2, where we let $I = \frac{1+i}{i}$ for ease of notation:

$$p_i = \frac{\eta}{\eta-1} \frac{\mathbb{E}_i \left[\phi(z_i A)^{-1} \left(\vartheta_i I^{-\frac{1}{\gamma}} P^{-\frac{1}{\gamma}} M_i^{\frac{1}{\gamma}} P^\eta \right)^{1-\eta\omega_1} \right]}{\mathbb{E}_i \left[I^{1-\frac{1}{\gamma}(1+\omega_1-\eta\omega_1)} M_i^{\frac{1}{\gamma}(1+\omega_1-\eta\omega_1)-1} \vartheta_i^{1+\omega_1-\eta\omega_1} P^{(\eta-\frac{1}{\gamma})(1+\omega_1-\eta\omega_1)} \right]} \quad (107)$$

$$\times \left(\vartheta_i I^{-\frac{1}{\gamma}} M_i^{\frac{1}{\gamma}} P^{\eta-\frac{1}{\gamma}} \right)^{\omega_1}$$

Given the ideal price index formula (Equation 23), P must satisfy the aggregation:

$$P^{1-\eta} = \mathbb{E} [\vartheta_i p_i^{1-\eta}] \quad (108)$$

where the expectation is over the cross-section of firms. We guess and verify that the aggregate price is log-linear in aggregates

$$\log P = \chi_0 + \chi_A \log A + \chi_M \log M \quad (109)$$

Moreover, if the p_i are log-normally distributed (we will verify this below), then:

$$\log P = \mathbb{E}[\log p_i] + \frac{1}{2(1-\eta)} \text{Var}((1-\eta) \log p_i) + \text{constants} \quad (110)$$

We first simplify the numerator of the first term by collecting all the terms involving s_i^A and s_i^M :

$$\begin{aligned} \log \mathbb{E}_i \left[\phi_i(z_i A)^{-1} \left(\vartheta I^{-\frac{1}{\gamma}} P^{-\frac{1}{\gamma}} M^{\frac{1}{\gamma}} P^\eta \right)^{1-\eta\omega_1} \right] &= \left[-\kappa^A + \kappa^A \left(\eta - \frac{1}{\gamma} \right) \chi_A (1 - \eta\omega_1) \right] s_i^A \\ &+ \left[\chi_M \left(\eta - \frac{1}{\gamma} \right) (1 - \eta\omega_1) \kappa^M + \frac{1}{\gamma} (1 - \eta\omega_1) \kappa^M \right] s_i^M + \text{constants} \end{aligned} \quad (111)$$

where the constants are independent of signals. We similarly simplify the denominator of the second term:

$$\begin{aligned} \log \mathbb{E}_i \left[I^{1-\frac{1}{\gamma}(1+\omega_1-\eta\omega_1)} M^{\frac{1}{\gamma}(1+\omega_1-\eta\omega_1)-1} \vartheta^{1+\omega_1-\eta\omega_1} P^{(\eta-\frac{1}{\gamma})(1+\omega_1-\eta\omega_1)} \right] &= \\ \left[\chi_A \left(\eta - \frac{1}{\gamma} \right) (1 + \omega_1 - \eta\omega_1) \kappa^A \right] s_i^A & \\ + \left[\left[\frac{1}{\gamma} (1 + \omega_1 - \eta\omega_1) - 1 \right] (\kappa^M) + \chi_M \left(\eta - \frac{1}{\gamma} \right) (1 + \omega_1 - \eta\omega_1) (\kappa^M) \right] s_i^M & \\ + \text{constants} & \end{aligned} \quad (112)$$

where the constants are again independent of signals. Finally, we can simplify the last term:

$$\log \left(\vartheta_i I^{-\frac{1}{\gamma}} M^{\frac{1}{\gamma}} P^{\eta-\frac{1}{\gamma}} \right)^{\omega_1} = \omega_1 \chi_A \left(\eta - \frac{1}{\gamma} \right) \log A + \omega_1 \left[\chi_M \left(\eta - \frac{1}{\gamma} \right) + \frac{1}{\gamma} \right] \log M + \text{constants} \quad (113)$$

where the constants are independent of the aggregate shocks. Hence, $\log p_i$ is indeed normally distributed and its variance is independent of the realization of aggregate shocks. We can now collect terms to verify our log-linear guess. Substituting the resulting expression for $\log p_i$ and our guess for $\log P$ from Equation 109 into Equation 110, and solving for χ_A by collecting coefficients on $\log A$ yields:

$$\chi_A = -\frac{\kappa^A}{1 - \omega_1 \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa^A)} \quad (114)$$

We may similarly solve for χ_M :

$$\chi_M = \frac{\kappa^M + \frac{\omega_1}{\gamma}(1 - \kappa^M)}{1 - \omega_1 \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa^M)} \quad (115)$$

This proves the dynamics for the price level. The dynamics for consumption then follow from Proposition 2. \square

A.6 Proof of Corollary 4

Proof. Using Equation 42 and market clearing $C_t = Y_t$, we have:

$$\log M_t = \frac{1}{\tilde{\chi}_{M,t}} (\log Y_t - \tilde{\chi}_{A,t} \log A_t - \tilde{\chi}_{0,t}) \quad (116)$$

Substituting for $\log M_t$ in Equation 43 then yields Equation AS. Doing a similar substitution for $\log A_t$ in Equation 42 then yields Equation AD. Note that this can also be derived by taking logarithms of

$$P_t = \frac{i_t}{1 + i_t} \frac{C_t^{-\gamma}}{M_t} \quad (117)$$

in Proposition 2. \square

A.7 Proof of Theorem 2

Proof. We suppress dependence on t for ease of notation. We have χ_M and χ_A as a function of ω_1 from Proposition 3. We also know that:

$$\omega_1 = \frac{\sigma_{M_i,z} + \sigma_{P,z}}{\sigma_z^2} \quad (118)$$

from Equation 83. As $z_i = \vartheta_i \left(\frac{i}{1+i}\right)^{\frac{1}{\gamma}} M^{\frac{1}{\gamma}} P^{\eta-\frac{1}{\gamma}}$ and $\mathcal{M}_i = \phi_i(z_i A)^{-1} \frac{i}{1+i} \frac{M}{P}$, we have that:

$$\begin{aligned}
\sigma_{\mathcal{M}_i, z} &= \text{Cov} \left(-(1 + \chi_A) \log A + (1 - \chi_M) \log M, (\eta - \frac{1}{\gamma}) \chi_A \log A + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \chi_M) \log M \right) \\
&= -(\eta - \frac{1}{\gamma}) \chi_A (1 + \chi_A) \sigma_A^2 + (1 - \chi_M) (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \chi_M) \sigma_M^2 \\
\sigma_{P, z} &= \text{Cov} \left(\chi_A \log A + \chi_M \log M, (\eta - \frac{1}{\gamma}) \chi_A \log A + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \chi_M) \log M \right) \\
&= (\eta - \frac{1}{\gamma}) \chi_A^2 \sigma_A^2 + \chi_M (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \chi_M) \sigma_M^2 \\
\sigma_z^2 &= \sigma_\vartheta^2 + (\eta - \frac{1}{\gamma})^2 \chi_A^2 \sigma_A^2 + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \chi_M)^2 \sigma_M^2
\end{aligned} \tag{119}$$

Thus:

$$\omega_1 = \frac{-(\eta - \frac{1}{\gamma}) \chi_A \sigma_A^2 + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \chi_M) \sigma_M^2}{\sigma_\vartheta^2 + (\eta - \frac{1}{\gamma})^2 \chi_A^2 \sigma_A^2 + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \chi_M)^2 \sigma_M^2} \tag{120}$$

Note that the optimal ω_1 is common across all firms i . We may express this in fully reduced form as:

$$\omega_1 = T(\omega_1) = \frac{(\eta - \frac{1}{\gamma}) \frac{\kappa_A}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_A)} \sigma_A^2 + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \frac{\kappa_M + \frac{\omega_1}{\gamma} (1 - \kappa_M)}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_M)}) \sigma_M^2}{\sigma_\vartheta^2 + (\eta - \frac{1}{\gamma})^2 \left(\frac{\kappa_A}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_A)} \right)^2 \sigma_A^2 + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \frac{\kappa_M + \frac{\omega_1}{\gamma} (1 - \kappa_M)}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_M)})^2 \sigma_M^2} \tag{121}$$

or

$$\omega_1 = T(\omega_1) = \frac{\frac{(\eta - \frac{1}{\gamma}) \kappa_A}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_A)} \sigma_A^2 + \frac{\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \kappa_M}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_M)} \sigma_M^2}{\sigma_\vartheta^2 + \left(\frac{(\eta - \frac{1}{\gamma}) \kappa_A}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_A)} \right)^2 \sigma_A^2 + \left(\frac{\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \kappa_M}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_M)} \right)^2 \sigma_M^2} \tag{122}$$

□

A.8 Proof of Proposition 4

Proof. We first establish equilibrium existence. First, we observe that T_t is a continuous function. The only possible points of discontinuity are: $\omega_{1,t}^M = \frac{1}{(\eta - \frac{1}{\gamma})(1 - \kappa_t^M)}$ and $\omega_{1,t}^A = \frac{1}{(\eta - \frac{1}{\gamma})(1 - \kappa_t^A)}$. However, at these points $\lim_{\omega_{1,t} \rightarrow \omega_{1,t}^M} T_t(\omega_{1,t}) = \lim_{\omega_{1,t} \rightarrow \omega_{1,t}^A} T_t(\omega_{1,t}) = T_t(\omega_{1,t}^M) = T_t(\omega_{1,t}^A) = 0$. Second, we observe that $\lim_{\omega_{1,t} \rightarrow -\infty} T_t(\omega_{1,t}) = \lim_{\omega_{1,t} \rightarrow \infty} T_t(\omega_{1,t}) = 0$. Consider now the function $W_t(\omega_{1,t}) = \omega_{1,t} - T_t(\omega_{1,t})$. This is a continuous function, $\lim_{\omega_{1,t} \rightarrow -\infty} W_t(\omega_{1,t}) = -\infty$, and $\lim_{\omega_{1,t} \rightarrow \infty} W_t(\omega_{1,t}) = \infty$. Thus, by the intermediate value theorem, there exists an $\omega_{1,t}^*$

such that $W_t(\omega_{1,t}^*) = 0$. By Theorem 2, $\omega_{1,t}^*$ defines a log-linear equilibrium.

We now show that there are at most five log-linear equilibria. For $\omega_{1,t} \neq \omega_{1,t}^A, \omega_{1,t}^M$ (neither of which can be a fixed point), we can rewrite Equation 50 as:

$$\begin{aligned}
& \omega_{1,t} \left[\sigma_{\vartheta,t}^2 \left(1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A) \right)^2 \left(1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right)^2 \right. \\
& \quad + (\sigma_{t|s}^A)^2 \left(\eta - \frac{1}{\gamma} \right) \kappa_t^A \left(1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right)^2 \\
& \quad \left. + (\sigma_{t|s}^M)^2 \left(\frac{1}{\gamma} + \left(\eta - \frac{1}{\gamma} \right) \kappa_t^M \right) \left(1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A) \right)^2 \right] \\
& = (\sigma_{t|s}^A)^2 \left(\eta - \frac{1}{\gamma} \right) \kappa_t^A \left(1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right)^2 \left(1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A) \right) \\
& \quad + (\sigma_{t|s}^M)^2 \left(\frac{1}{\gamma} + \left(\eta - \frac{1}{\gamma} \right) \kappa_t^M \right) \left(1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A) \right)^2 \left(1 - \omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right)
\end{aligned} \tag{123}$$

This is a quintic polynomial in $\omega_{1,t}$, which has at most five real roots. Thus, by Theorem 2, there are at most five log-linear equilibria. \square

A.9 Proof of Corollary 7

Proof. We drop time subscripts for ease of notation. Substituting $\eta = \frac{1}{\gamma}$ in Equation 50 yields:

$$\omega_1 = \frac{\frac{1}{\gamma}}{\rho^2 + \left(\frac{1}{\gamma} \right)^2} \tag{124}$$

Substituting this into Equation 44 yields:

$$\epsilon_t^S = \gamma \frac{\kappa_t^M}{(1 - \kappa_t^M)} + \frac{1}{\gamma \rho^2 (1 - \kappa_t^M)} \tag{125}$$

\square

A.10 Proof of Corollary 8

We drop time subscripts for ease of notation. From Equation 50, as $\sigma_{t|s}^M \rightarrow \infty$, $\omega_{1,t}$ must solve:

$$\begin{aligned}\omega_1 &= \frac{1 - \omega_1 \left(\eta - \frac{1}{\gamma}\right) (1 - \kappa^M)}{\frac{1}{\gamma} + \left(\eta - \frac{1}{\gamma}\right) \kappa^M} \\ &= \frac{\gamma}{1 + (\eta\gamma - 1) \kappa^M} + \left(1 - \frac{\eta\gamma}{1 + (\eta\gamma - 1) \kappa^M}\right) \omega_1 \\ &= \frac{1}{\eta}\end{aligned}\tag{126}$$

This proves the first statement. As $\sigma_{t|s} \rightarrow \infty$ and $\eta\gamma \neq 1$, the $\omega_{1,t}$ must solve:

$$\begin{aligned}\omega_1 &= \frac{1 - \omega_1 \left(\eta - \frac{1}{\gamma}\right) (1 - \kappa^A)}{\left(\eta - \frac{1}{\gamma}\right) \kappa^A} \\ &= \frac{\gamma}{(\eta\gamma - 1) \kappa^A} + \left(1 - \frac{1}{\kappa^A}\right) \omega_1 \\ &= \frac{1}{\eta - \frac{1}{\gamma}}\end{aligned}\tag{127}$$

This proves the second statement. The third statement follows directly from Equation 50.

A.11 Proof of Proposition 5

Proof. By Theorem 2, The map describing equilibrium $\omega_{1,t}$ is invariant to λ for $\lambda > 0$. Thus, $\mathcal{E}_t^S(\lambda)$ is constant for $\lambda > 0$. If $\lambda = 0$, there are potentially many equilibria in supply functions. Nevertheless, from the proof of Theorem 1, we have that firms set $p_{it}/P_t = \frac{\eta}{\eta-1} \mathcal{M}_{it} = \frac{\eta}{\eta-1} C_t^\gamma / A_t$ under any optimal supply function. This implies that $\frac{\eta}{\eta-1} C_t^\gamma / A_t = 1$, and so money has no real effects, which implies that $\epsilon_t^S = \infty$. \square

B Supply Function Choice with Multiple Inputs, Decreasing Returns, and Monopsony

In this section, we generalize our baseline model of supply function choice to allow for multiple inputs, decreasing returns, and monopsony. We find that: (i) supply functions remain endogenously log-linear and (ii) decreasing returns and monopsony flatten the optimal supply schedule.

Primitives. Consider the baseline model from Section 2 with two modifications. First, the production function uses multiple inputs with different input shares and possibly features decreasing returns-to-scale:

$$q = \Theta \prod_{i=1}^I x_i^{a_i} \quad (128)$$

where $x_i \in \mathbb{R}_+$, $a_i \geq 0$, and $\sum_{i=1}^I a_i \leq 1$. Moreover, suppose that the producer potentially has monopsony power and faces a cost of acquiring each type of input that is given by:

$$c_i(x_i) = p_{xi} x_i^{b_i} \quad (129)$$

where $p_{xi} \in \mathbb{R}_{++}$ and $b_i \geq 1$. The firm believes that $(\Psi, P, \Lambda, \Theta, p_x)$ is jointly log-normal.

The Firm's Problem. We begin by solving the firm's cost minimization problem:

$$K(q; \Theta, p_x) = \min_x \sum_{i=1}^I p_{xi} x_i^{b_i} \quad \text{s.t.} \quad q = \Theta \prod_{i=1}^I x_i^{a_i} \quad (130)$$

This has first-order condition given by:

$$\lambda = \frac{b_i p_{xi}}{a_i} x_i^{b_i} q^{-1} \quad (131)$$

Which implies that:

$$K(q; \Theta, p_x) = \lambda q \sum_{i=1}^I \frac{a_i}{b_i} \quad (132)$$

Moreover, fixing i , the FOC implies that we may write for all $j \neq i$:

$$x_j = \left(\frac{\frac{b_i p_{xi}}{a_i}}{\frac{b_j p_{xj}}{\alpha_j}} \right)^{\frac{1}{b_j}} x_i^{\frac{b_i}{b_j}} \quad (133)$$

By substituting this into the production function we have that:

$$q = \Theta x_i^{a_i + b_i \sum_{j \neq i} \frac{a_j}{b_j}} \prod_{j \neq i} \left(\frac{b_i p x_i}{\frac{a_i}{b_j p x_j}} \right)^{\frac{1}{b_j}} \quad (134)$$

which implies that:

$$x_i = \left(\frac{q}{\Theta \prod_{j \neq i} \left(\frac{b_i p x_i}{\frac{a_i}{b_j p x_j}} \right)^{\frac{1}{b_j}}} \right)^{\frac{1}{a_i + b_i \sum_{j \neq i} \frac{a_j}{b_j}}} \quad (135)$$

Returning to the FOC, we have that the Lagrange multiplier is given by:

$$\lambda = q^{-1 + \frac{1}{\sum_{i=1}^I \frac{a_i}{b_i}}} \frac{b_i p x_i}{a_i} \left(\Theta \prod_{j \neq i} \left(\frac{b_i p x_i}{\frac{a_i}{b_j p x_j}} \right)^{\frac{1}{b_j}} \right)^{\frac{-1}{\sum_{i=1}^I \frac{a_i}{b_i}}} \quad (136)$$

Which then yields the cost function:

$$K(q; \Theta, p_x) = \mathcal{M} P q^{\frac{1}{\delta}} \quad (137)$$

where:

$$\delta = \sum_{i=1}^I \frac{a_i}{b_i} \quad \text{and} \quad \mathcal{M} = P^{-1} \frac{b_i p x_i}{a_i} \left(\Theta \prod_{j \neq i} \left(\frac{b_i p x_i}{\frac{a_i}{b_j p x_j}} \right)^{\frac{1}{b_j}} \right)^{\frac{-1}{\sum_{i=1}^I \frac{a_i}{b_i}}} \sum_{i=1}^I \frac{a_i}{b_i} \quad (138)$$

and we observe that \mathcal{M} is log-normal given the joint log-normality of (Θ, p_x) .

Turning to the firm's payoff function, we therefore have:

$$\mathbb{E} \left[\Lambda \left(\frac{p}{P} q - \mathcal{M} q^{\frac{1}{\delta}} \right) \right] \quad (139)$$

Thus, the problem with multiple inputs, monopsony, and decreasing returns modifies the firms' original payoff by only introducing the parameter δ . Helpfully, observe that $\delta = 1$ when: (i) there are constant returns to scale $\sum_{i=1}^I a_i = 1$ and (ii) there is no monopsony $b_i = 1$ for all i .

Given this, we can write the firm's objective as:

$$J(\hat{p}) = \int_{\mathbb{R}_{++}^4} \Lambda \left(\frac{\hat{p}(z)^{1-\eta}}{P} z - \mathcal{M} z^{\frac{1}{\delta}} \hat{p}(z)^{-\frac{\eta}{\delta}} \right) dG(\Lambda, P, \mathcal{M}, z) \quad (140)$$

And, as before, we study the problem:

$$\sup_{\hat{p}: \mathbb{R}_+ \rightarrow \mathbb{R}_{++}} J(\hat{p}) \quad (141)$$

By doing this, we obtain a modified formula for the optimal supply function:

Proposition 6 (Optimal Supply Schedule With Multiple Inputs, Decreasing Returns, and Monopsony). *Any optimal supply schedule is almost everywhere given by:*

$$f(p, q) = \log p - \frac{\omega_0 - \log \delta}{1 - \eta\omega_1} - \frac{\eta \left(\omega_1 + \frac{1-\delta}{\delta} \right)}{1 - \eta\omega_1} \log q \quad (142)$$

where ω_0 and ω_1 are the same as those derived in Theorem 1. Thus, the optimal inverse supply elasticity is given by:

$$\hat{\alpha}_1 = \frac{\eta\sigma_P^2 + \sigma_{\mathcal{M},\Psi} + \sigma_{P,\Psi} + \eta\sigma_{\mathcal{M},P}}{\sigma_{\Psi}^2 - \eta\sigma_{\mathcal{M},\Psi} + \eta\sigma_{P,\Psi} - \eta^2\sigma_{\mathcal{M},P}} \left(1 + \frac{1-\delta}{\delta} \frac{\sigma_{\Psi}^2 + \eta^2\sigma_P^2 + 2\eta\sigma_{\Psi,P}}{\sigma_{\mathcal{M},\Psi} + \eta\sigma_{\mathcal{M},P} + \sigma_{P,\Psi} + \eta\sigma_P^2} \right) \quad (143)$$

Proof. Applying the same variational arguments as in the Proof of Theorem 1, we obtain that $\hat{p}(t)$ must solve:

$$(\eta - 1)\mathbb{E}[\Lambda P^{-1}|z = t]t\hat{p}(t)^{-\eta} = \frac{\eta}{\delta}\mathbb{E}[\Lambda \mathcal{M}|z = t]t^{\frac{1}{\delta}}\hat{p}(z)^{-\frac{\eta}{\delta}-1} \quad (144)$$

Which yields:

$$\hat{p}(t) = \left(\delta^{-1} \frac{\eta}{\eta - 1} \frac{\mathbb{E}[\Lambda \mathcal{M}|z = t]}{\mathbb{E}[\Lambda P^{-1}|z = t]} \right)^{\frac{1}{1+\eta\left(\frac{1-\delta}{\delta}\right)}} t^{\frac{\frac{1-\delta}{\delta}}{1+\eta\left(\frac{1-\delta}{\delta}\right)}} \quad (145)$$

Thus, we have that:

$$\log p = \frac{1}{1 + \eta\left(\frac{1-\delta}{\delta}\right)} (\omega_0 - \log \delta) + \frac{1}{1 + \eta\left(\frac{1-\delta}{\delta}\right)} \left(\omega_1 + \frac{1-\delta}{\delta} \right) \log z \quad (146)$$

where ω_0 and ω_1 are as in Theorem 1. Rewriting as a supply schedule, we obtain:

$$\log p = \frac{\frac{1}{1+\eta\left(\frac{1-\delta}{\delta}\right)} (\omega_0 - \log \delta)}{1 - \frac{\eta}{1+\eta\left(\frac{1-\delta}{\delta}\right)} \left(\omega_1 + \frac{1-\delta}{\delta} \right)} + \frac{\frac{\eta}{1+\eta\left(\frac{1-\delta}{\delta}\right)} \left(\omega_1 + \frac{1-\delta}{\delta} \right)}{1 - \frac{\eta}{1+\eta\left(\frac{1-\delta}{\delta}\right)} \left(\omega_1 + \frac{1-\delta}{\delta} \right)} \log q \quad (147)$$

Which reduces to the claimed formula. \square

Thus, when the supply curve is initially upward-sloping ($\omega_1 \in [0, \eta^{-1}]$), the introduction of decreasing returns and/or monopsony unambiguously reduces the supply elasticity and makes firms closer to quantity-setting.

C Prices *vs.* Quantities in General Equilibrium

C.1 Model and Equilibrium

All of the model primitives are as in Section 3. However, we now restrict firms to follow either price-setting or quantity-setting at all times. We define equilibrium in two steps. We first fix firms’ “choice of choices” at each date t to define a rational expectations *temporary equilibrium*:

Definition 2 (Temporary Equilibrium). *A temporary equilibrium is a partition of \mathbb{N} into two sets \mathcal{T}^P and \mathcal{T}^Q and a collection of variables*

$$\left\{ \{p_{it}, q_{it}, C_{it}, N_{it}, L_{it}, w_{it}, \phi_{it}, \vartheta_{it}, z_{it}, \Pi_{it}\}_{i \in [0,1]}, C_t, P_t, M_t, A_t, B_t, N_t, \Lambda_t, \sigma_t^\phi, \sigma_t^\vartheta, \sigma_t^z, \sigma_t^A, \sigma_t^M \right\}_{t \in \mathbb{N}}$$

such that:

1. In periods $t \in \mathcal{T}^P$, all firms choose their prices p_{it} to maximize expected real profits under the household’s real stochastic discount factor.
2. In periods $t \in \mathcal{T}^Q$, all firms choose their quantities q_{it} to maximize expected real profits under the household’s real stochastic discount factor.
3. In all periods, the household chooses consumption C_{it} , labor supply N_{it} , money holdings M_t , and bond holdings B_t to maximize their expected utility subject to their lifetime budget constraint, while Λ_t is the household’s marginal utility of consumption.
4. In all periods, money supply M_t and productivity A_t and evolve exogenously via Equations 25 and 27.
5. In all periods, firms’ and consumers’ expectations are consistent with the equilibrium law of motion.
6. In all periods, the markets for the intermediate goods, final good, labor varieties, bonds, and money balances all clear.

In a temporary equilibrium, firms set either prices or quantities, but the choice between the two is not necessarily optimal. We define an *equilibrium* as a temporary equilibrium in which the choice between price and quantity-setting is optimal at all times:

Definition 3 (Equilibrium). *An equilibrium is a temporary equilibrium in which:*

1. If $t \in \mathcal{T}^P$, all firms find price-setting optimal. That is, expected real profits under the household’s real stochastic discount factor are weakly higher under price-setting than quantity-setting.

2. If $t \in \mathcal{T}^Q$, all firms find quantity-setting optimal. That is, expected real profits under the household's real stochastic discount factor are weakly higher under price-setting than quantity-setting.

We now study the equilibrium properties of the model.

C.2 Prices *vs.* Quantities in Equilibrium: Incentives and Strategic Interactions

We first describe dynamics in temporary equilibria (Definition 2) in which all firms set prices or quantities *by assumption*. To this end, we use Proposition 3 to express the dynamics of consumption and prices under pure price-setting and pure quantity-setting. Concretely, we assume that the dynamics for consumption are log-linear in the aggregate shocks:

$$\log C_t = \tilde{\chi}_{0,t}^X + \chi_{A,t}^X \log A_t + \chi_{M,t}^X \log M_t \quad (148)$$

where we allow the coefficients to be *regime-specific* $X \in \{Q, P\}$. We then verify that dynamics take a log-linear form.

Corollary 9. *Output and the price level under pure price-setting follow in the unique log-linear equilibrium of the economy follow:*

$$\log C_t = \chi_{0,t}^P + \frac{1}{\gamma} \kappa_t^A \log A_t + \frac{1}{\gamma} (1 - \kappa_t^M) \log M_t \quad (149)$$

$$\log P_t = \tilde{\chi}_{0,t}^P - \kappa_t^A \log A_t \quad (150)$$

Output and the price level under pure quantity-setting in the unique log-linear equilibrium of the economy follow:

$$\log C_t = \chi_{0,t}^Q + \frac{1}{\gamma} \frac{\kappa_t^A}{1 - \left(1 - \frac{1}{\eta\gamma}\right) (1 - \kappa_t^A)} \log A_t \quad (151)$$

$$\log P_t = \tilde{\chi}_{0,t}^Q - \frac{\kappa_t^A}{1 - \left(1 - \frac{1}{\eta\gamma}\right) (1 - \kappa_t^A)} \log A_t + \log M_t \quad (152)$$

where $\chi_{0,t}^P$, $\tilde{\chi}_{0,t}^P$, $\chi_{0,t}^Q$, and $\tilde{\chi}_{0,t}^Q$ are constants independent of A_t and M_t .

Proof. Setting $\omega_{1,t} = 0$ in Proposition 3 yields dynamics under pure price-setting. Setting $\omega_{1,t} = 1/\eta$ in Proposition 3 yields dynamics under pure quantity-setting. \square

Observe that money is neutral under pure quantity-setting, but not under pure price-setting, as discussed in Corollary 6. Moreover, productivity shocks have a larger effect on output in a quantity-setting regime relative to a price-setting regime if and only if $\eta\gamma > 1$. Intuitively, it is the elasticity of substitution that mediates how much firms will adjust their output in response to a *perceived* increase in productivity under a quantity-setting regime. In contrast, it is real money balances (that is independent of η) that mediates the output response in a pure price-setting regime.

We now ask when would firms prefer to set prices or quantities in equilibrium. To study this, we first derive an expression for Δ in terms of uncertainty about equilibrium objects. We combine Proposition 1 with (i) Proposition 2 and (ii) the observation that consumption is log-linear in both temporary equilibria to derive that:

$$\Delta_t = \frac{1}{2}(\eta - 1) \left(\frac{1}{\eta} \sigma_{\vartheta,t}^2 + \frac{1}{\eta} (1 - \eta\gamma)^2 \sigma_{C,t}^2 - \eta(\sigma_t^M)^2 + 2(1 - \eta\gamma)\sigma_{C,A,t} \right) \quad (153)$$

where $\sigma_{C,t}^2$ is the firm's posterior variance for output, $\sigma_{C,A,t}$ is the firm's posterior covariance for output and productivity, $\sigma_{\vartheta,t}^2$ is the variance of idiosyncratic demand shocks, and $(\sigma_t^M)^2$ is the variance of money supply innovations. Higher uncertainty about idiosyncratic demand shocks and lower uncertainty about the money supply provide *exogenous* incentives for price-setting. Higher uncertainty about consumption provides an *endogenous* incentive that unambiguously favors price-setting. This is the net effect of two forces that favor price-setting – increasing demand uncertainty and decreasing the covariance of prices and marginal costs – with two forces that favor quantity-setting – increasing price uncertainty and increasing the covariance between demand and marginal costs. Higher covariance between consumption and productivity favors price-setting if $\eta\gamma < 1$ and quantity-setting otherwise. In the former case, the dominant effect of this covariance is to lower the covariance of marginal costs and demand (favoring price-setting); in the latter case, the dominant effect is to raise the covariance of marginal costs and the price level (favoring quantity-setting). Finally, the variance of idiosyncratic productivity shocks and the variance of idiosyncratic labor supply (factor price) shocks drop out, because they do not induce *covariance* between marginal costs and either demand or the price level.

We now combine the previous observation with the equilibrium dynamics (Corollary 9) to fully describe Δ_t in terms of primitives in each regime:

Lemma 1 (Prices *vs.* Quantities in Equilibrium). *If all firms set quantities, then the com-*

parative advantage of price-setting is:

$$\begin{aligned} \Delta_t^Q &= \frac{1}{2}(\eta - 1) \left(\frac{1}{\eta} \sigma_{\vartheta,t}^2 - \eta \kappa_t^M \sigma_{M,s}^2 \right. \\ &\quad \left. + \left(\frac{1}{\eta} (1 - \eta\gamma) \frac{\eta \kappa_t^A}{1 - \kappa_t^A (1 - \eta\gamma)} + 2 \right) (1 - \eta\gamma) \frac{\eta (\kappa_t^A)^2}{1 - \kappa_t^A (1 - \eta\gamma)} \sigma_{A,s}^2 \right) \end{aligned} \quad (154)$$

Moreover, all firms can set quantities in equilibrium at time t if and only if $\Delta_t^Q \leq 0$.

If all firms set prices, then the comparative advantage of price-setting is:

$$\begin{aligned} \Delta_t^P &= \frac{1}{2}(\eta - 1) \left(\frac{1}{\eta} \sigma_{\vartheta,t}^2 + \left(-\eta + \frac{1}{\eta} (1 - \eta\gamma)^2 \left(\frac{1 - \kappa_t^M}{\gamma} \right)^2 \right) \kappa_t^M \sigma_{M,s}^2 \right. \\ &\quad \left. + \left(\frac{1}{\eta} (1 - \eta\gamma) \frac{\kappa_t^A}{\gamma} + 2 \right) (1 - \eta\gamma) \frac{(\kappa_t^A)^2}{\gamma} \sigma_{A,s}^2 \right) \end{aligned} \quad (155)$$

Proof. From Proposition 1, we have that:

$$\Delta_t = \frac{1}{2}(\eta - 1) \left(\frac{1}{\eta} \sigma_{\Psi,t}^2 - \eta \sigma_{P,t}^2 - 2\sigma_{\Psi,\mathcal{M},t} - 2\eta \sigma_{P,\mathcal{M},t} \right) \quad (156)$$

where, now, all the variances are time-dependent. Applying Proposition 2 to obtain expressions for (Ψ, P, \mathcal{M}) in equilibrium, and exploiting the log-linearity of each expression, we have that:

$$\begin{aligned} \sigma_{\Psi,t}^2 &= \sigma_{\vartheta,t}^2 + \sigma_{C,t}^2 \\ \sigma_{P,t}^2 &= \gamma^2 \sigma_{C,t}^2 + (\sigma_t^M)^2 - 2\gamma \sigma_{C,M,t} \\ \sigma_{\Psi,\mathcal{M},t} &= \gamma \sigma_{C,t}^2 - \sigma_{C,A,t} \\ \sigma_{P,\mathcal{M},t} &= \gamma \sigma_{C,A,t} - \gamma^2 \sigma_{C,t}^2 + \gamma \sigma_{C,M,t} \end{aligned} \quad (157)$$

where $\sigma_{X,t}^2$ denotes the firm's posterior variance of variable X at time t and $\sigma_{X,Y,t}$ denotes the firm's posterior covariance of variables X and Y at time t . Substituting these formulae, we obtain Equation 153:

$$\Delta_t = \frac{1}{2}(\eta - 1) \left(\frac{1}{\eta} \sigma_{\vartheta,t}^2 + \frac{1}{\eta} (1 - \eta\gamma)^2 \sigma_{C,t}^2 - \eta (\sigma_t^M)^2 + 2(1 - \eta\gamma) \sigma_{C,A,t} \right) \quad (158)$$

Moreover, applying Proposition 3, we have that these variances for the firm in each of the

price-setting and quantity-setting regimes are given in each regime $X \in \{Q, P\}$ by

$$\begin{aligned}(\sigma_{C,t}^X)^2 &= (\chi_{A,t}^X)^2 \sigma_{A|s,t}^2 + (\chi_{M,t}^X)^2 \sigma_{M|s,t}^2 \\(\sigma_{C,A,t}^X)^2 &= \chi_{A,t}^X \sigma_{A|s,t}^2\end{aligned}\tag{159}$$

Substituting Equation 159 into Equation 153, we obtain the following expression for Δ_t indexed by the regime $X \in \{Q, P\}$:

$$\begin{aligned}\Delta_t^X &= \frac{1}{2}(\eta - 1) \left(\frac{1}{\eta} \sigma_{\vartheta,t}^2 + \left(-\eta + \frac{1}{\eta} (1 - \eta\gamma)^2 (\chi_{M,t}^X)^2 \right) \sigma_{M|s,t}^2 \right. \\ &\quad \left. + \left(\frac{1}{\eta} (1 - \eta\gamma) \chi_{A,t}^X + 2 \right) (1 - \eta\gamma) \chi_{A,t}^X \sigma_{A|s,t}^2 \right)\end{aligned}\tag{160}$$

We now derive the two desired expressions for Δ_t , splitting the calculation into the quantity-setting and price-setting cases.

Quantity-Setting. Substituting $\chi_{A,t}^Q$ and $\chi_{M,t}^Q$ (quantity-setting) from Corollary 9 and exploiting the fact that the conditional variances are given by $\sigma_{A|s,t}^2 = \kappa_t^A \sigma_{A,s}^2$ and $\sigma_{M|s,t}^2 = \kappa_t^M \sigma_{M,s}^2$, we obtain:

$$\begin{aligned}\Delta_t^Q &= \frac{1}{2}(\eta - 1) \left(\frac{1}{\eta} \sigma_{\vartheta,t}^2 - \eta \kappa_t^M \sigma_{M,s}^2 \right. \\ &\quad \left. + \left(\frac{1}{\eta} (1 - \eta\gamma) \frac{\eta \kappa_t^A}{1 - \kappa_t^A (1 - \eta\gamma)} + 2 \right) (1 - \eta\gamma) \frac{\eta \kappa_t^A}{1 - \kappa_t^A (1 - \eta\gamma)} \kappa_t^A \sigma_{A,s}^2 \right)\end{aligned}\tag{161}$$

as desired.

Price-Setting. Mirroring the steps above using the coefficients from Corollary 9, we obtain

$$\begin{aligned}\Delta_t^P &= \frac{1}{2}(\eta - 1) \left(\frac{1}{\eta} \sigma_{\vartheta,t}^2 + \left(-\eta + \frac{1}{\eta} (1 - \eta\gamma)^2 \left(\frac{1 - \kappa_t^M}{\gamma} \right)^2 \right) \kappa_t^M \sigma_{M,s}^2 \right. \\ &\quad \left. + \left(\frac{1}{\eta} (1 - \eta\gamma) \frac{\kappa_t^A}{\gamma} + 2 \right) (1 - \eta\gamma) \frac{\kappa_t^A}{\gamma} \kappa_t^A \sigma_{A,s}^2 \right)\end{aligned}\tag{162}$$

yielding the claimed expressions.

Next, consider the comparative statics for Δ_t^P . First,

$$\frac{\partial \Delta_t^P}{\partial \kappa_t^M} = \sigma_{M,s}^2 \left(-\eta + \frac{1}{\eta} (1 - \eta\gamma)^2 \left(\frac{1 - \kappa_t^M}{\gamma} \right)^2 + \frac{2}{\eta\gamma^2} (1 - \eta\gamma)^2 (1 - \kappa_t^M) \kappa_t^M \right)\tag{163}$$

The condition $\frac{\partial \Delta_t^P}{\partial \kappa_t^M} > 0$ corresponds to

$$\left(\frac{\eta\gamma}{1-\eta\gamma}\right)^2 < (1-\kappa_t^M)^2 + 2(1-\kappa_t^M)\kappa_t^M = 1 - (\kappa_t^M)^2 \quad (164)$$

Re-arranging gives, as desired, $\kappa_t^M < \sqrt{1 - \left(\frac{\eta\gamma}{1-\eta\gamma}\right)^2}$. Next, Δ_t^P is strictly increasing in κ_t^A if and only if:

$$\left(\frac{1}{\eta}(1-\eta\gamma)\frac{\kappa_t^A}{\gamma} + 2\right)(1-\eta\gamma)\frac{(\kappa_t^A)^2}{\gamma} \quad (165)$$

is strictly increasing in κ_t^A . As argued above, this condition holds if $\eta\gamma < 1$ and does not if $\eta\gamma > 1$. □

Importantly, since $\Delta_t^Q \neq \Delta_t^P$ in general, others' choice of whether to set prices or quantities affects any given firm's incentives to set prices or quantities. Does the fact that others set prices (quantities) increase or decrease my own desire to set prices (quantities)? Strikingly, we find that these decisions are always *strategic complements*. That is, when all other firms set prices, a given firm has stronger incentives to set prices:

Proposition 7 (Complementarity in Choices of Choices). *The decision to set a price or a quantity is one of strategic complements, i.e., $\Delta_t^P \geq \Delta_t^Q$, with strict inequality whenever $\eta\gamma \neq 1$.*

Proof. Define $\Delta\Delta_t = \Delta_t^P - \Delta_t^Q$ and observe that:

$$\begin{aligned} \Delta\Delta_t = \frac{1}{2}(\eta-1) & \left[\frac{1}{\eta}(1-\eta\gamma)^2 \left(\frac{1-\kappa_t^M}{\gamma}\right)^2 \kappa_t^M \sigma_{M,s}^2 \right. \\ & \left. + \left(\frac{1}{\eta}(1-\eta\gamma)^2 (\chi_{A,t}^{P^2} - \chi_{A,t}^{Q^2}) + 2(1-\eta\gamma)(\chi_{A,t}^P - \chi_{A,t}^Q)\right) \kappa_t^A \sigma_{A,s}^2 \right] \end{aligned} \quad (166)$$

First, when $\eta\gamma = 1$, we have that $\Delta\Delta_t = 0$. Second, suppose that $\eta\gamma < 1$. We observe that the first term in brackets is strictly positive. Turning to the second term, as $\eta\gamma < 1$, we have that $\Delta\Delta_t > 0$ if and only if $\chi_{A,t}^P > \chi_{A,t}^Q$. This inequality is equivalent to:

$$\frac{\kappa_t^A}{\gamma} > \frac{\eta\kappa_t^A}{1-\kappa_t^A(1-\eta\gamma)} \quad (167)$$

As $\eta\gamma < 1$ and $\kappa_t^A \in (0, 1)$, we have that the denominator on the right-hand side is positive.

Thus, we can re-express this required inequality as:

$$1 - \eta\gamma > \kappa_t^A(1 - \eta\gamma) \quad (168)$$

which is true as $\eta\gamma < 1$ and $\kappa_t^A \in (0, 1)$. Thus, $\Delta\Delta_t > 0$ when $\eta\gamma < 1$. Third, suppose that $\eta\gamma > 1$. Once again, the first term in brackets is strictly positive. Thus, it suffices to show that:

$$\frac{1}{\eta}(1 - \eta\gamma)^2 (\chi_{A,t}^{P^2} - \chi_{A,t}^{Q^2}) + 2(1 - \eta\gamma)(\chi_{A,t}^P - \chi_{A,t}^Q) > 0 \quad (169)$$

See that we can factor the left-hand side of this expression as:

$$(1 - \eta\gamma)(\chi_{A,t}^P - \chi_{A,t}^Q) \left(\frac{1}{\eta}(1 - \eta\gamma)(\chi_{A,t}^P + \chi_{A,t}^Q) + 2 \right) \quad (170)$$

By the reverse of the logic in part two, we have that $\chi_{A,t}^P < \chi_{A,t}^Q$. Thus, the expression in question is strictly positive if and only if:

$$2 > \frac{1}{\eta}(\eta\gamma - 1)(\chi_{A,t}^P + \chi_{A,t}^Q) \quad (171)$$

We now observe that $\chi_{A,t}^P + \chi_{A,t}^Q < 2\chi_{A,t}^Q$. Moreover, $\chi_{A,t}^Q$ is an increasing function of κ_t^A and is therefore bounded above by $\frac{\eta}{1+\eta\gamma-1} = \frac{1}{\gamma}$. Thus, we have that:

$$\frac{1}{\eta}(\eta\gamma - 1)(\chi_{A,t}^P + \chi_{A,t}^Q) < \frac{2}{\eta\gamma}(\eta\gamma - 1) = 2 - \frac{2}{\eta\gamma} < 2 \quad (172)$$

This establishes that $\Delta\Delta_t > 0$ if $\eta\gamma > 1$. Taken together, we have shown that $\Delta\Delta_t \geq 0$ and $\Delta\Delta_t > 0$ if and only if $\eta\gamma \neq 1$, establishing the claim. \square

To give the intuition for this result, we first consider the case when $\eta\gamma < 1$. In this case, it can easily be seen that consumption responds more to productivity shocks under price-setting. Moreover, regardless of the value of $\eta\gamma$, consumption responds more to monetary shocks under price-setting. Therefore, others being price-setters increases both the variance of consumption and the covariance of consumption with productivity. Both of these forces favor price-setting, as shown in Equation 153. In summary, others setting prices induces aggregate volatility which makes it more attractive for any given firm to also set a price. In the case of $\eta\gamma \geq 1$, consumption is more responsive to monetary shocks but less responsive to productivity shocks under price-setting versus quantity-setting. In the proof, we show how these effects net out in Equation 153 in the direction of making price-setting more attractive when other firms set prices.

We now use this result to show the existence of equilibria in which all firms *optimally* choose to set prices or quantities:

Corollary 10 (Existence of Pure Equilibria). *There exists an equilibrium in which, at each date t , either all firms set prices or all firms set quantities.*

To prove this result, we consider two cases at each date t . First, suppose that firms prefer to set prices if others set quantities, or $\Delta_t^Q \geq 0$. In this case, they even more strongly prefer to set prices if others set prices, or $\Delta_t^P \geq \Delta_t^Q \geq 0$. Therefore, choosing to set prices is consistent with equilibrium. Conversely, suppose that firms prefer to set quantities when others set quantities, $\Delta_t^Q < 0$. In this case, choosing to set quantities is consistent with equilibrium. As these cases are exhaustive, a pure equilibrium exists. Note that this logic heavily relies on our finding that the decision to set prices was one of complements; if it were one of substitutes, then pure equilibria could fail to exist.¹²

C.3 Time-Varying Uncertainty and Regime Switches

As shown in Lemma 1, the comparative advantage of price-setting *vs.* quantity-setting changes over time because firms' uncertainty about microeconomic and macroeconomic variables changes over time. This observation, combined with Corollary 10, implies the existence of equilibria in which time-varying volatility induces time-varying uncertainty and *regime changes* between price- and quantity-setting. These regime changes, in turn, affect the propagation of aggregate shocks as summarized in Corollary 9. Thus, "uncertainty shocks" that affect exogenous volatility have further effects on the volatility of endogenous outcomes (income, prices) due to the endogenous "choice of choices."

To better understand these forces, we now study the comparative statics of (Δ_t^Q, Δ_t^P) in the parameters for time-varying volatility. We start by studying uncertainty about the aggregate productivity state A_t . Higher aggregate productivity uncertainty pushes toward either price- or quantity-setting depending on the parameter condition $\eta\gamma \gtrless 1$:

Corollary 11. *If $\eta\gamma > 1$, then both Δ_t^Q and Δ_t^P are decreasing in κ_t^A and in σ_t^A . If $\eta\gamma < 1$, then both Δ_t^Q and Δ_t^P are increasing in κ_t^A and in σ_t^A . If $\eta\gamma = 1$, then Δ_t^Q and Δ_t^P are equal and invariant to κ_t^A and σ_t^A .*

Proof. We consider the three cases $\eta\gamma = 1$, $\eta\gamma < 1$, and $\eta\gamma > 1$ separately.

1. $\eta\gamma = 1$. By Lemma 1, we have that $\Delta_t^Q = \Delta^Q(0)$ and $\Delta_t^P = \Delta^Q(0)$, which are both independent of κ_t^A .

¹²A mixed equilibrium would always exist.

2. $\eta\gamma < 1$. By Lemma 1, we have that Δ_t^Q is strictly increasing in κ_t^A if and only if

$$\left(\frac{1}{\eta}(1 - \eta\gamma) \frac{\eta\kappa_t^A}{1 - \kappa_t^A(1 - \eta\gamma)} + 2 \right) (1 - \eta\gamma) \frac{\eta(\kappa_t^A)^2}{1 - \kappa_t^A(1 - \eta\gamma)} \quad (173)$$

is strictly increasing in κ_t^A . As $\eta\gamma < 1$ and $\frac{\eta\kappa_t^A}{1 - \kappa_t^A(1 - \eta\gamma)}$ is strictly increasing in κ_t^A and strictly positive, we have that the term in parentheses is strictly positive and strictly increasing. The term outside parentheses is strictly increasing and strictly positive for the same reasons. Moreover, Δ_t^P is strictly increasing in κ_t^A if and only if:

$$\left(\frac{1}{\eta}(1 - \eta\gamma) \frac{\kappa_t^A}{\gamma} + 2 \right) (1 - \eta\gamma) \frac{(\kappa_t^A)^2}{\gamma} \quad (174)$$

is strictly increasing in κ_t^A . As $\eta\gamma < 1$, this is immediate.

3. $\eta\gamma > 1$. By Lemma 1, we have that Δ_t^Q is strictly decreasing in κ_t^A if and only if Expression 173 is strictly decreasing in κ_t^A . Define $\omega = 1 - \eta\gamma$ and observe that we need to show that:

$$\left(\frac{\omega\kappa_t^A}{1 - \omega\kappa_t^A} + 2 \right) \frac{\omega(\kappa_t^A)^2}{1 - \omega\kappa_t^A} \quad (175)$$

is a strictly decreasing function of κ_t^A . Taking the derivative of this expression and rearranging, we require that:

$$\omega\kappa_t^A \left(\omega^2 (\kappa_t^A)^2 - 3\omega\kappa_t^A + 4 \right) < 0 \quad (176)$$

As $\omega < 0$, we require that $\omega^2 (\kappa_t^A)^2 - 3\omega\kappa_t^A + 4 > 0$. This is positive if the quadratic on the left-hand side has no real roots. As $9\omega^2 - 16\omega^2 < 0$, the quadratic indeed has no real roots and so Δ_t^Q is strictly decreasing in κ_t^A .

Δ_t^P is strictly decreasing in κ_t^A if and only:

$$\left(\frac{\omega}{1 - \omega} \kappa_t^A + 2 \right) \frac{\omega}{1 - \omega} (\kappa_t^A)^2 \quad (177)$$


is strictly decreasing in κ_t^A . Taking the derivative of this expression and rearranging, we require that:

$$\kappa_t^A < \frac{4\omega - 1}{3\omega} \quad (178)$$

which is always satisfied as $\omega < 0$.

As κ_t^A is increasing in σ_t^A , this establishes the result. \square

Figure 7: Equilibrium with Changing Productivity Uncertainty




Note: This figure illustrates firms' equilibrium incentives for price-setting as uncertainty about productivity changes. In each panel, we plot Δ^Q (dashed line) and Δ^P (dotted line) as a function of κ^A , fixing all other parameter values. In Example A, we use parameters such that $\eta\gamma < 1$. In Example B, we use parameters such that $\eta\gamma > 1$. We shade the region with only a quantity-setting equilibrium blue, the region with only a price-setting equilibrium orange, and the region with both equilibria red.

When $\eta\gamma < 1$, the dominant effects of productivity uncertainty are to increase aggregate demand uncertainty and to lower the covariance between demand and marginal costs. When $\eta\gamma > 1$, the dominant effect is to increase the covariance between marginal costs and the price level. Finally, in the special case in which $\eta\gamma = 1$, these forces net out to zero.

We illustrate this result and its implications for equilibrium regime-switching in a numerical example. In Figure 7, we plot Δ^Q and Δ^P as a function of κ^A for two different calibrations, corresponding to $\eta\gamma < 1$ and $\eta\gamma > 1$. We shade regions of the parameter space in which only one equilibrium exists (blue for quantity-setting and orange for price-setting) and in which both equilibria exist (red). In the economies corresponding to each parameter case, as κ_t^A moves exogenously (because of underlying movements in σ_t^A), the equilibrium transitions between quantity-setting and price-setting. For example, in the left panel with $\eta\gamma < 1$, periods of high productivity uncertainty (high κ_t^A) correspond to price-setting and periods of low productivity uncertainty (high κ_t^A) correspond to quantity-setting. If κ_t^A lies in the middle, red-shaded region in any period t , there exists an equilibrium in which firms set prices in that period as well as one in which they set quantities in that period. In this way, even if κ_t^A were constant over time but lying in this multiple-equilibrium region, there could be self-fulfilling macroeconomic volatility that arises from endogenous regime shifts.

Figure 8: Equilibrium with Changing Money-Supply Uncertainty



Note: This figure illustrates firms’ equilibrium incentives for price-setting as uncertainty about the money supply changes. We plot Δ^Q (dashed line) and Δ^P (dotted line) as a function of κ^M , fixing all other parameter values. We use parameters such that $\eta\gamma > \frac{1}{2}$, so both functions are monotone decreasing (see Corollary 13). We shade the region with only a quantity-setting equilibrium blue, the region with only a price-setting equilibrium orange, and the region with both equilibria red.

We next study the role of idiosyncratic uncertainty. We find that idiosyncratic demand uncertainty unambiguously favors quantity-setting, while idiosyncratic uncertainty about productivity and factor prices (via labor supply) do not matter:

Corollary 12. *Both Δ_t^Q and Δ_t^P are increasing in $\sigma_{\vartheta,t}^2$ and neither depends on $\sigma_{z,t}^2$ or $\sigma_{\phi,t}^2$.*

This result is immediate from inspection of the formulas in Lemma 1 and is not intermediated by equilibrium forces. Economically, it implies that “uncertainty shocks” that increase idiosyncratic variation in firms’ demand unambiguously push the economy toward price-setting. In light of empirical evidence that (i) idiosyncratic volatility in firms’ revenue TFP rises dramatically in recessions (*e.g.*, Bloom et al., 2018) and (ii) a majority of revenue TFP variation arises from demand rather than productivity shocks (Foster et al., 2008), Corollary 12 suggests a powerful force for regime switches that line up with the business cycle. By contrast, uncertainty about idiosyncratic productivity and factor prices does *not* behave symmetrically to uncertainty about demand. This follows from our original observation that the uncertainty about marginal costs matters only through its covariance with demand and the price level, and not through its variance (Proposition 1).

We finally study uncertainty about the money supply M_t . As with uncertainty about productivity, understanding its effect requires disciplining opposing equilibrium forces:

Corollary 13. Δ_t^Q is always decreasing in κ_t^M and σ_t^M . If $\eta\gamma \geq \frac{1}{2}$, then Δ_t^P is strictly decreasing in κ_t^M and σ_t^M . If $\eta\gamma < \frac{1}{2}$, then there exists a $\bar{\kappa}^M \in [0, 1/3]$ such that Δ_t^P is increasing for $\kappa_t^M < \bar{\kappa}^M$ and decreasing for $\kappa_t^M > \bar{\kappa}^M$.

Proof. The fact that Δ_t^Q is decreasing in κ_t^M is immediate from Lemma 1. Moreover, from Lemma 1, Δ_t^P is decreasing in κ_t^M if and only if

$$\left(-\eta + \frac{1}{\eta}(1 - \eta\gamma)^2 \left(\frac{1 - \kappa_t^M}{\gamma} \right)^2 \right) \kappa_t^M \quad (179)$$

is decreasing in κ_t^M . Taking the derivative of this expression, this is equivalent to

$$\frac{(1 - \eta\gamma)^2}{(\eta\gamma)^2} \left(\frac{1 - 2\eta\gamma}{(1 - \eta\gamma)^2} - 4\kappa_t^M + 3(\kappa_t^M)^2 \right) < 0 \quad (180)$$

This is a strictly convex quadratic. Hence, if we show that this expression is weakly negative evaluated at $\kappa_t^M = 0$ and $\kappa_t^M = 1$, it will be strictly negative for all $\kappa_t^M \in (0, 1)$. A sufficient condition for this expression to be weakly negative at $\kappa_t^M = 0$ is that

$$\frac{1 - 2\eta\gamma}{(1 - \eta\gamma)^2} \leq 0 \quad (181)$$

which occurs if and only if $\eta\gamma \geq 1/2$. It is easily verified that $\eta\gamma \geq 1/2$ also makes the expression strictly negative at $\kappa_t^M = 1$. This proves that Δ_t^P is strictly decreasing for all $\kappa_t^M \in (0, 1)$ whenever $\eta\gamma \geq 1/2$.

We next study the case in which $\eta\gamma < \frac{1}{2}$. We re-arrange condition 180 to

$$\kappa_t^M(4 - 3\kappa_t^M) > \frac{1 - 2\eta\gamma}{(1 - \eta\gamma)^2} \quad (182)$$

We first observe that this condition always holds at $\kappa_t^M = 1$, as the left-hand-side is 1 and the right-hand-side, given $\eta\gamma < 1/2$, is bounded above by 1. Therefore, the critical value $\bar{\kappa}^M$ is the smaller root of the quadratic equation $\kappa_t^M(4 - 3\kappa_t^M) - \frac{1 - 2\eta\gamma}{(1 - \eta\gamma)^2} = 0$. By direct calculation,

this is

$$\bar{\kappa}^M = \frac{1}{3} \left(2 - \sqrt{4 - 3 \frac{1 - 2\eta\gamma}{(1 - \eta\gamma)^2}} \right) \quad (183)$$

$$= \frac{1}{3} \left(2 - \sqrt{1 - \left(\frac{\eta\gamma}{1 - \eta\gamma} \right)^2} \right) \quad (184)$$

$$= \frac{2}{3} - \sqrt{\frac{1}{9} + \frac{1}{3} \left(\frac{\eta\gamma}{1 - \eta\gamma} \right)^2} \quad (185)$$

where, in the second equality, we use the fact that $\frac{1-2\eta\gamma}{(1-\eta\gamma)^2} = 1 - \left(\frac{\eta\gamma}{1-\eta\gamma}\right)^2$. We finally note that $\bar{\kappa}^M$ is monotone increasing in $\eta\gamma$, is minimized at 0 when $\eta\gamma = 1/2$, and is maximized at $\frac{1}{3}$ when $\eta\gamma = 0$. \square


Under quantity-setting, because monetary shocks are neutral for output, increasing the volatility of money-supply shocks (lowering κ_t^M) serves only to increase the volatility of the price level and further favor quantity-setting. Under price-setting, because monetary shocks are not neutral for output, there is a countervailing effect from increasing the volatility of aggregate demand. Therefore, when $\eta\gamma$ is sufficiently low, the effect of money-supply uncertainty is ambiguous.

We illustrate this result numerically in Figure 8. We focus on a calibration in which $\eta\gamma > \frac{1}{2}$, so both Δ_t^Q and Δ_t^P are monotone decreasing in κ_t^M . In periods of low money-supply uncertainty, firms have stronger incentives to set prices; in periods of high money-supply uncertainty, firms have stronger incentives to set quantities. Moreover, in the quantity-setting regimes, the aggregate price level responds more to money-supply innovations (Corollary 9) which further sharpens the incentives for quantity-setting. This positive feedback loop underlies our comparative statics result. We will explore the further implications of this logic for systematic monetary policy in the next section.

D Prices *vs.* Quantities in the Data

If a researcher wishes to focus on price-setting or quantity-setting model, they should wish to do so in the macroeconomic circumstances in which each is actually preferable to firms. In this appendix, we develop an empirical approach to determining which is a more preferable choice and test its predictive power for the pass-through of monetary policy.

Figure 9: The Relative Benefit of Price-Setting in US Data



Note: This figure plots our empirical estimate of $\hat{\Delta}_t$ (the comparative advantage of price-setting relative to quantity-setting) and its components, as defined in Proposition 1 (Equation 13). The black line plots $\hat{\Delta}_t$, in units of expected percent profit improvement (100 times log points). The blue (dashed), orange (dotted), green (dashed), and red (dash-dotted) lines plot each of the four components of $\hat{\Delta}_t$, corresponding to uncertainty about different variables. The grey shading denotes periods in which $\hat{\Delta}_t < 0$ and thus, according to Proposition 1, quantity-setting is optimal for firms. As described in Section 5, the calculation uses estimates of time-varying volatilities from a CCC GARCH(1,1) model and a calibrated demand elasticity of $\eta = 9$. The demand component exceeds the scale of the figure in Q2 and Q3 of 2020.

D.1 When Might Price-Setting Be a Better Approximation Than Quantity-Setting?

Using our GARCH model from the main text and Proposition 1, we can compute an empirical estimate for when firms should prefer price-setting:

$$\hat{\Delta}_t = \frac{1}{2}(\eta - 1) \left(\frac{1}{\eta} \hat{\sigma}_{\Psi,t}^2 - \eta \hat{\sigma}_{P,t}^2 - 2\hat{\sigma}_{\Psi, \mathcal{M}, t} - 2\eta \hat{\sigma}_{P, \mathcal{M}, t} \right) \quad (186)$$

Our calculation captures uncertainty about outcomes realized in quarter t , and is measurable in data from quarter $t - 1$ and earlier. It therefore describes incentives of a decisionmaker fixing a choice for quarter t based on their uncertainty at the beginning of the quarter, before data are realized.

We plot our calculation of $\hat{\Delta}_t$ in Figure 9. We show our overall calculation in black and each component in color. We shade periods which favor quantity-setting, or for which

$\hat{\Delta}_t < 0$.

Strikingly, both quantity- and price-setting are optimal at different points in the sample. Thus, viewed through the lens of our model and its mapping to the data, firms may be either price- or quantity-setters depending on the macroeconomic context. Moreover, through the same lens, this evidence rules out the conventional assumption that firms always choose prices or always choose quantities.

Price-setting is optimal in most of the sample, or 219 of 251 quarters. This notably comprises the 1960s and the Great Moderation, in which both inflation and demand variance were relatively tame, and the Great Recession and the onset of the Covid-19 Lockdown Recession (Q2 2020), when demand variance abruptly spiked.


Quantity-setting is optimal intermittently between 1972 Q2 and 1981 Q2, for a total of 25 of the possible 37 quarters in this period, and continuously between 2021 Q2 and the end of the sample. These all correspond to periods of particularly high contributions of the terms corresponding to inflation variance and inflation-marginal-cost covariance. Through the lens of the model, firms would prefer to set quantities in these periods to hedge against the increase in uncertainty about joint movements in inflation and marginal costs. Our calculation weighs this consideration against demand risk, which favors price-setting and may also be elevated in recessions. For example, in 1975 Q2 and 2021 Q1, demand uncertainty is sufficiently high to outweigh elevated inflation and inflation-marginal-cost uncertainty, and our calculation favors price-setting on net ($\hat{\Delta}_t > 0$).

We finally note that the advantage of one method over another is always relatively small in payoff terms. In our sample, this advantage peaks at 0.48% (0.0048 log points) in Q3 of 2020. In all periods excluding Q2 and Q3 of 2020, the difference peaks at 0.16%. This is a striking juxtaposition with the model prediction that a change in firm behavior between price- and quantity-setting can have large effects on equilibrium outcomes.

Robustness to Parameter Values and Measurement Strategies. Two parameters that were central to our calculation, but difficult to pin down in the data, were the price elasticity of demand ($\eta = 9$) and the ratio of micro to macro volatility ($R = 6.5$). In Figure 10, we plot the implied time series for Δ under specific alternative assumptions for each parameter. In Figure 11, we vary both parameters continuously over a larger grid and plot “heat maps” for the average value of $\hat{\Delta}_t$ and the percentage of the sample with $\hat{\Delta}_t > 0$.

Decreasing the elasticity of demand favors price-setting, while increasing the elasticity of demand favors quantity-setting (left panel). The primary reason, quantitatively, is that highly inelastic demand curves amplify the effects of demand shocks on prices for fixed quantities, and hence increase potential losses from quantity-setting. In the data, this further pushes toward price-setting, especially in time periods with especially high demand volatility.

Figure 10: The Relative Benefit of Price-Setting Under Alternative Parameters




Note: Both panels plot our empirical estimate of $\hat{\Delta}_t$ defined in Proposition 1 (Equation 13) under alternative assumptions for the elasticity of substitution η (left) and the micro-to-macro volatility ratio R (right). In both plots, our baseline estimate corresponds to the solid black line.

Increasing demand risk favors price-setting by construction (right panel). In particular, increasing the extent of microeconomic volatility by 50% favors price-setting in all periods (orange dotted line), while decreasing this parameter by 50% implies quantity-setting in a majority of periods (blue dashed line). As noted by Bloom et al. (2018), calibrating this parameter on the basis of observed variances in *measured* firm-level fundamentals requires modeling choices. In particular, one must take a stand on what fraction of measured volatility corresponds to measurement error and what fraction of volatility from an econometrician’s perspective is unknown to firm managers, who likely have superior information.


As an alternative strategy to measure the contribution of idiosyncratic volatility, we can use the direct measurements of Bloom et al. (2018) based on annual data from manufacturing establishments from 1972 to 2010, along with assumptions about measurement error and observability of shocks. To accommodate this variant calculation, we re-estimate the VAR(1) CCC GARCH(1,1) model on annual data for the same macro time series. We then use the Bloom et al. (2018) estimates of the cross-sectional standard deviation of manufacturing TFP along with those authors’ quantitative assumption that 45.4% of this measured volatility (standard deviation) corresponds to measurement error. We make the intentionally extreme assumption that all of this remaining variance is unforecastable by firms. Figure 12 shows our results. This calculation echoes the conclusion that the 1970s were favorable

Figure 11: The Relative Benefit of Price-Setting in an Alternative, Annual Calculation



Note: This figure summarizes the relative advantage of price-setting for alternative values of the elasticity of demand (horizontal axis) and the ratio of microeconomic to macroeconomic volatility (vertical axis). The left panel plots the average advantage of price-setting over the sample, in units of 100 times log points (percent). The right panel plots the fraction of the sample in which price-setting is optimal, or in which $\hat{\Delta}_t > 0$. In both panels, our baseline calibration is indicated with a solid dot.

Figure 12: The Relative Benefit of Price-Setting in an Alternative, Annual Calculation



Note: This figure plots our empirical estimate of $\hat{\Delta}_t$ (the comparative advantage of price-setting relative to quantity-setting) and its components, as defined in Proposition 1 (Equation 13), under a variant method with annual-frequency data and direct measurement of micro volatility from Bloom et al. (2018). Note that the time period (1972-2010) and time-frequency (annual) differs from that in Figures 9 and 10 (quarterly, 1960 Q1 to 2022 Q4). The black line plots $\hat{\Delta}_t$, in units of expected percent profit improvement (100 times log points). The blue (dashed), orange (dotted), green (dashed), and red (dash-dotted) lines plot each of the four components of $\hat{\Delta}_t$, corresponding to uncertainty about different variables. The grey shading denotes periods in which $\hat{\Delta}_t < 0$ and thus, according to Proposition 1, quantity-setting is optimal for firms. As described in Section 5, the calculation uses estimates of time-varying volatilities from an annual-frequency CCC GARCH(1,1) model and a calibrated elasticity of demand $\eta = 9$

to quantity-setting due to the relatively high inflation volatility and relatively low demand volatility.

Comparison to External Evidence. An alternative way to gauge the plausibility of firms’ entertaining both price- and quantity-setting plans is via direct survey evidence. As observed by Reis (2006), Aiginger (1999) collected data on this topic. In a survey of managers of Austrian manufacturing firms, he asked: “What is your main strategic variable: do you decide to produce a specific quantity, thereafter permitting demand to decide upon price conditions, or do you set the price, with competitors and the market determining the quantity sold?” Among managers, 32% said that they use the quantity plan and 68% said that they use the price plan. We interpret this as additional evidence that neither price nor quantity plans are obviously favored in practice.

D.2 Testing the Model: Asymmetric Effects of Monetary Policy in Price and Quantity Regimes

Our model predicts that expansionary monetary shocks have muted effects on real output and exaggerated effects on prices in a quantity-setting regime compared to a price-setting regime (Corollary 9). The model also predicts that incentives for price-setting are shaped by the volatility of macroeconomic and microeconomic aggregates in a specific way (Proposition 1 and Lemma 1). Crucially, both predictions rely purely on the premise of “choice of choices” and not on specific parameter restrictions.¹³ Thus, we can use them to derive an empirical test.

In this final section, we provide suggestive evidence consistent with these predictions. In particular, using local projection regressions, we find that output responds more negatively and prices respond less negatively to contractionary Romer and Romer (2004) monetary policy shocks in price-setting regimes relative to quantity-setting regimes, measured using the method of Section 5.

Hypotheses and Strategy. We test the main model prediction that monetary shocks are: (i) neutral for output under quantity-setting regimes, (ii) contractionary for output under price-setting regimes, and (iii) more inflationary under quantity-setting regimes. Formally, the theory implies that the following relationships hold:

$$\log Y_t = \chi_M^P \mathbb{I}[\Delta_t > 0] \log M_t + \varepsilon_t^Y \quad (187)$$

$$\log P_t = \log M_t - \tilde{\chi}_M^P \mathbb{I}[\Delta_t > 0] \log M_t + \varepsilon_t^P \quad (188)$$

where $\chi_M^P > 0$, $\tilde{\chi}_M^P \in (0, 1)$, and $\Delta_t, \log M_t \perp \varepsilon_t^Y, \varepsilon_t^P$. Thus, given a measure of Δ_t and exogenous monetary shocks M_t , we can estimate these equations consistently via ordinary least squares. Here the model-implied definition of an exogenous monetary shock is one that is not a response by the central bank to either endogenous or exogenous economic circumstances.

Measurement and Empirical Specification. We measure monetary policy shocks using the methodology of Romer and Romer (2004). These authors residualize changes in the Federal Funds Rate on the Federal Reserve’s macroeconomic projections reported in the Greenbook. Specifically, we use the updated series reported in Ramey (2016) which spans March 1969 to December 2007. We aggregate these shocks to a quarterly-frequency variable,

¹³By contrast, the sign of the differential response to productivity shocks in each regime depends on the parametric condition $\eta\gamma \gtrless 1$.

MonShock $_t$ by summing. The key quantity-setting regimes that overlap with the studied sample of Romer and Romer (2004) shocks are primarily in the 1970s.

To proxy for whether the economy is in a price-setting or quantity-setting regime, we translate $\hat{\Delta}_t$ into a binary variable PriceSet $_t = 1_{\Delta_{t+1} > 0}$. In the model, this object determines whether decisionmakers who observe data before and during time t would set prices as their decision variable for period $t + 1$. This timing convention is appropriate since we will focus on how macroeconomic aggregates at time $t + 1$ and onward respond to shocks at time t .

To estimate an empirical analog of Equations 187 and 188, we proxy for real output with real GDP and the price level with the GDP deflator. We estimate the state-dependent response of outcomes $Z_t \in \{\text{RealGDP}_t, \text{GDPDeflator}_t\}$ to the variable MonShock $_t$ by running the following local projection regressions for each horizon $h \in \{1, \dots, 12\}$:

$$Z_{t+h} = \beta_h \cdot \text{MonShock}_t + \gamma_h \cdot \text{PriceSet}_t + \phi_h \cdot (\text{MonShock}_t \times \text{PriceSet}_t) + \tau' X_t + \varepsilon_{t,h} \quad (189)$$

As control variables, we include the contemporaneous and lagged values of real GDP, GDP deflator, and utilization-adjusted TFP, and interactions of all of these variables with PriceSet $_t$.¹⁴ Including the interaction variables is consistent with our model’s implications that the joint dynamics of macroeconomic variables may change between the two regimes.¹⁵ In all reported results, we report frequentist confidence intervals based on Newey and West (1987) standard errors with a six-quarter bandwidth.


The coefficients $\{\beta_h\}_{h=1}^H$ measure the response of output (or prices) to monetary shocks in the quantity regime. We predict that $\beta_h < 0$ for both outcomes. The coefficients $\{\phi_h\}_{h=1}^H$ measure the differential response of output (or prices) to monetary shocks in the price regime, compared to the quantity regime. We predict that $\phi_h < 0$ when real GDP is the outcome and $\phi_h > 0$ when the GDP deflator is the outcome.

Results: State-Dependent Effects of Monetary Policy. We show our results graphically in Figure 13. We first consider our results for output (top row). We find on average a zero response to monetary shocks in quantity-setting regimes ($\beta_h = 0$; first column). This average zero response belies weak evidence of a negative response at shorter horizons ($h < 6$) and a positive response at longer horizons ($h \in \{7, 8\}$). By contrast, we find a consistently

¹⁴As observed by Ramey (2016), including contemporaneous values amounts to assuming a zero contemporaneous response of macroeconomic quantities to shocks on impact, as is typical in the structural VAR literature (e.g., Christiano et al., 2005). This is also consistent with our conditioning on Δ_t , which is measurable in time- t macroeconomic aggregates. Results are very similar when we do not control for contemporaneous values, suggesting that this timing assumption is close to correct in the data.

¹⁵Tenreyro and Thwaites (2016) make a similar observation about the necessity of these controls in a local-projections estimation of whether monetary policy shocks, also measured as in Romer and Romer (2004), have different effects in recessions.

Figure 13: IRFs to Monetary Shocks in Price-Setting and Quantity-Setting Regimes



Note: These plots display our estimates of the state-dependent response to monetary policy shocks from Equation 189. The outcome variable is real GDP in the top row and GDP deflator in the bottom row. The columns respectively show our estimates of β_h , the response under quantity-setting; $\beta_h + \phi_h$, the response under price-setting; and ϕ_h , the difference between the price-setting and quantity-setting responses. In each plot, the solid line gives the point estimates, the dark-shaded region gives 68% confidence intervals, and the light-shaded region gives 95% confidence intervals, where the latter two are based on Newey and West (1987) standard errors with a six-quarter bandwidth.

negative response under price-setting for all horizons $h > 4$. This is statistically significant at the 68% level for $h \geq 6$ and at the 95% level for $h \in \{10, 11\}$. The difference between these responses is also negative ($\phi < 0$) at these longer horizons, and statistically significant at the 95% level for $h \geq 7$. These results, taken together, are consistent with our theory: in price-setting regimes, contractionary monetary policy has considerably more power to shape real outcomes.

We next consider our results for prices (bottom row). We find a small negative response in quantity-setting regimes (column 1) and a significant positive response in price-setting regimes (column 2). The second prediction violates the theory in the direction of the familiar “price puzzle” (see, e.g., Ramey, 2016). But our prediction for the difference of

coefficients is consistent with the theory ($\phi_h > 0$, column 3): under quantity-setting regimes, contractionary policy is more able to control the price level.

Robustness. In Table 3, we probe the robustness of these findings on three margins. When reporting these results, we focus on the interactive coefficients ϕ at the horizon $h = 12$. First, we vary the timing of our measurement of PriceSet, since the mapping from theory to data imperfectly captured the realistic delays in the effects of monetary policy. When we replace PriceSet with a four-quarter backward-looking average (model (2) of Table 3) or a four-quarter forward-looking average (model (3) of Table 3), we continue to find $\phi^{RGDP} < 0$ and $\phi^{PGDP} > 0$. Next, we parametrize the model using the continuous measure of Δ rather than the binary measure of PriceSet. This guards against the possibility that our binary transformation masked non-monotone effects. We again find $\phi^{RGDP} < 0$ and $\phi^{PGDP} > 0$.

Comparison to the Literature. Existing work draws a mixed conclusion on whether monetary policy is more or less powerful in “downturns,” broadly defined. Weise (1999) finds weaker price effects and stronger output effects when output is initially low; Garcia and Schaller (2002) and Lo and Piger (2005) find stronger responses of output in recessions; and Tenreyro and Thwaites (2016) find *weaker* responses of output and prices in recessions. Our analysis differs both because (i) it conditions on a different variable, the model’s prediction for whether firms set prices or quantities, which itself depends on uncertainties rather than means; and (ii) it tests for differences in both the response of output to monetary shocks and the response of prices to monetary shocks, as predicted by the theory.

Table 3: Asymmetric Effects of Monetary Policy, Robustness

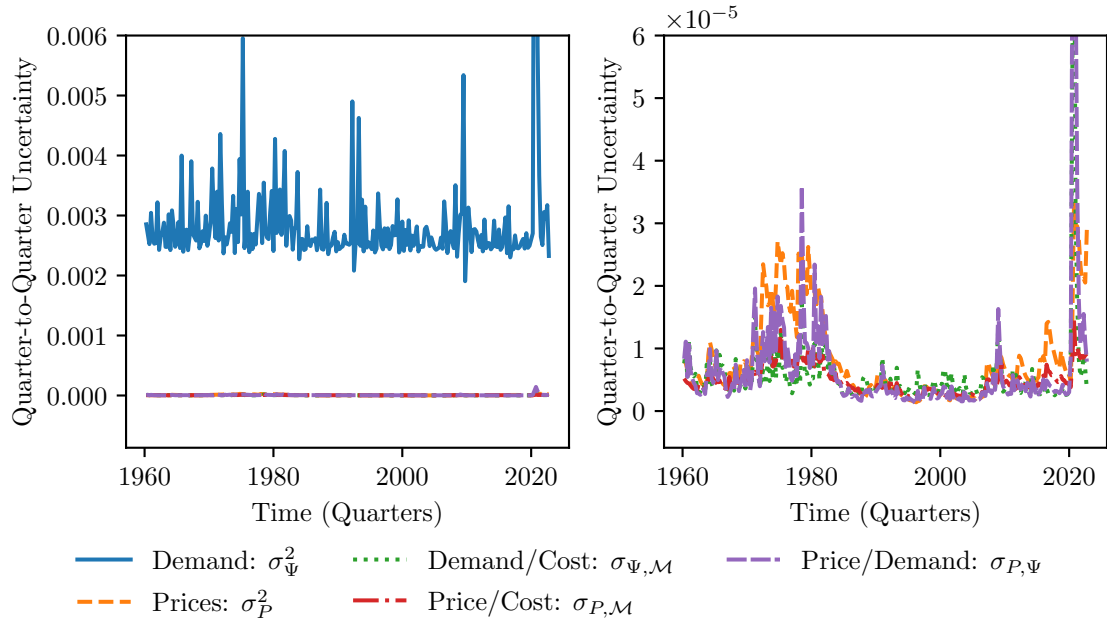
Panel (a): Outcome is log RealGDP _{t+12}				
	(1)	(2)	(3)	(4)
	Baseline	Lag Avg.	Lead Avg.	Continuous
MonShock _t × PriceSet _t	-0.0172 (0.0077)	-0.0071 (0.0101)	-0.0253 (0.0162)	
MonShock _t × $\hat{\Delta}_{t+1}$				-28.50 (13.51)
MonShock _t	✓	✓	✓	✓
PriceSet _t	✓	✓	✓	
$\hat{\Delta}_{t+1}$				✓
Macro Controls	✓	✓	✓	✓
Macro Controls × PriceSet _t	✓	✓	✓	
Macro Controls × $\hat{\Delta}_{t+1}$				✓
<i>N</i>	156	156	156	156

Panel (b): Outcome is log GDPDeflator _{t+12}				
	(1)	(2)	(3)	(4)
	Baseline	Lag Avg.	Lead Avg.	Continuous
MonShock _t × PriceSet _t	0.0120 (0.0053)	0.0042 (0.0059)	0.0184 (0.0083)	
MonShock _t × $\hat{\Delta}_{t+1}$				1.018 (10.38)
MonShock _t	✓	✓	✓	✓
PriceSet _t	✓	✓	✓	
$\hat{\Delta}_{t+1}$				✓
Macro Controls	✓	✓	✓	✓
Macro Controls × PriceSet _t	✓	✓	✓	
Macro Controls × $\hat{\Delta}_{t+1}$				✓
<i>N</i>	156	156	156	156

Note: This Table shows results from estimating Equation 189 at the 12-quarter horizon with different constructions of the interactive variable, focusing only on estimates of the interaction coefficient. In Panel (a) the outcome variable is Real GDP and in Panel (b) the outcome variable is GDP Deflator. Model (1) is our baseline. Model (2) uses a four-quarter lagged average of PriceSet, or sets $\text{PriceSet}_t = (\sum_{j=0}^3 \text{PriceSetBaseline}_{t-j})/4$. Model (3) uses a four-quarter lead average of PriceSet, or sets $\text{PriceSet}_t = (\sum_{j=0}^3 \text{PriceSetBaseline}_{t+j})/4$. Model (4) uses the continuous variable $\hat{\Delta}_{t+1}$. In all cases, we control for interactions of the macroeconomic variables (contemporaneous and lagged values of log Real GDP, log GDP Deflator, and log TFP) interacted with the variant construction of PriceSet or Δ . Standard errors in parentheses are based on the method of Newey and West (1987) with a six-quarter bandwidth.


E Additional Tables and Figures

Figure 14: Estimates of Time-Varying Uncertainty




Note: Both panels plot our quarterly time-series estimates of uncertainty, estimated as described in Section 5.1. All lines except for “Demand” (solid blue) are one-quarter-ahead volatility predictions from a constant conditional correlations (CCC) GARCH model. The “Demand” estimates combine the GARCH model’s predictions for real GDP uncertainty and TFP uncertainty with an assumption about the relationship between microeconomic (demand) volatility and macroeconomic (productivity) uncertainty, as described in the main text. The left plot shows all series on a common scale, and the right plot zooms in on the series other than demand. Both plots feature spikes that are off the scale of the graph during the Covid-19 lockdown.

Figure 15: The Slope of Aggregate Supply (Annual-Frequency Calculation)



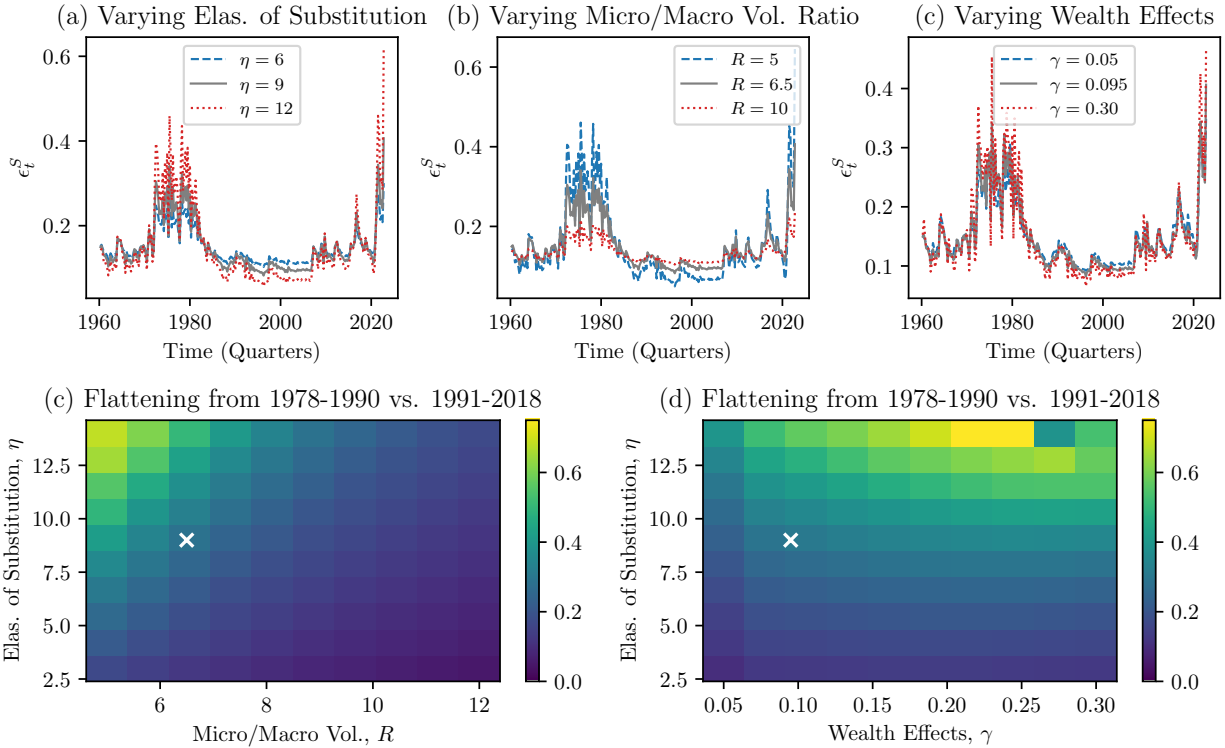
Note: This Figure plots estimates of the inverse elasticity of aggregate supply as measured by Equations 57 and 58. These estimates correspond to our secondary calculation using an annual-frequency GARCH model and a direct measure of microeconomic (demand) uncertainty from Bloom et al. (2018). The blue dashed line indicates the estimates of Hazell et al. (2022), based on state-level estimates of consumer prices and unemployment and an identification strategy that isolates local demand shocks (columns 3 and 4, panel B, of Table II). The red dotted line indicates the estimates of Ball and Mazumder (2011), based on aggregate data (column 4 of Table 3).

Figure 16: The Slope of Aggregate Supply Since World War II



Note: This Figure plots estimates of the inverse elasticity of aggregate supply as measured by Equations 57 and 58. In this calculation, in contrast to the calculation of Figure 4, we extend the calculation back to 1947:Q2. The elasticity of aggregate supply is off the scale of the graph from 1951:Q1 to 1952:Q2. Note that the estimates from 1960 onward numerically differ from the ones in Figure 2 because we re-estimate the GARCH model for macroeconomic uncertainty over the larger sample.

Figure 17: Sensitivity of Aggregate Supply Estimates to Parameters



Note: This Figure plots the sensitivity of the inverse elasticity of aggregate supply to the elasticity of substitution η (Panel a), to varying ratios of microeconomic to macroeconomic uncertainty R (Panel b), and to varying calibrations of wealth effects (Panel c). Panels (c) and (d) each plot a heat map for the difference in the average slope for aggregate supply from 1970-1980 to 1991-2018, varying different pairs of parameters. In each of these panels, the white “x” denotes our baseline calibration.