

# A Theory of Supply Function Choice and Aggregate Supply

Joel P. Flynn\*  
Yale

George Nikolakoudis†  
Princeton

Karthik A. Sastry‡  
Princeton

April 1, 2024

## Abstract

Modern theories of aggregate supply are built on the foundation that firms set prices and commit to producing whatever the market demands. We remove this strategic restriction and allow firms to choose *supply functions*, mappings that describe the prices charged at each quantity of production. Theoretically, we characterize firms' optimal supply function choices in general equilibrium and study the resulting implications for aggregate supply. Aggregate supply flattens under lower inflation uncertainty, higher idiosyncratic demand uncertainty, and less elastic demand. Quantitatively, our theory explains the flattening of aggregate supply during the Great Moderation and steepening during the 1970s and 2020s.

---

\*Yale University Department of Economics, 30 Hillhouse Avenue, New Haven, CT, 06511. Email: [joel.flynn@yale.edu](mailto:joel.flynn@yale.edu).

†Princeton University Department of Economics, Julis Romo Rabinowitz Building, Princeton, NJ 08544. Email: [nikolakoudis@princeton.edu](mailto:nikolakoudis@princeton.edu).

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

We are grateful to Hassan Afrouzi, Narek Alexanian, George-Marios Angeletos, Gideon Bornstein (discussant), V.V. Chari, Basil Halperin, Joe Hazell, Nobuhiro Kiyotaki, Ricardo Lagos, Jennifer La'O, Ernest Liu, N. Gregory Mankiw (discussant), Guido Menzio, Alessandro Pavan, Michael Peters, Ricardo Reis, Richard Rogerson, Kunal Sangani, Aleh Tsyvinski, Michael Woodford, Venky Venkateswaran, Gianluca Violante, Pierre Yared and seminar participants at The University of California Berkeley, Columbia University, the University of Pennsylvania, the NBER Summer Institute, the NBER DSGE Conference, the NBER Monetary Economics Conference, Carnegie Mellon University (Tepper), The Minneapolis Fed, the LAEF Junior Macroeconomics Workshop, Yale University, Princeton University, UCLA, 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 specific class: setting a price and committing to produce whatever the market demands. 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; Hellwig and Venkateswaran, 2009). The price-setting assumption is also part of the foundation of the ubiquitous New Keynesian framework (Woodford, 2003b) and its modern incarnations (*e.g.*, Wolf, 2023; Dávila and Schaab, 2023).

However, as has long been recognized (see *e.g.*, Grossman, 1981b), price-setting is *not* typically an optimal way for a firm to behave and is, at some level, *ad hoc*. Why should firms not be able to raise their prices when goods are flying off the shelves? Of course, in practice, they can and do: firms use temporary sales and surge pricing to navigate changing demand conditions. Such “dynamic pricing” policies, in which prices explicitly respond to demand, are commonly used in many sectors (see Den Boer, 2015, for a review), including electricity (Joskow and Wolfram, 2012), gasoline (Borenstein and Shepard, 1996), air travel (Williams, 2022), retail (*e.g.*, grocery stores with sales), ride-sharing (*e.g.*, Uber), e-commerce (*e.g.*, Amazon), food (*e.g.*, Wendy’s), and entertainment (*e.g.*, Disney World or Ticketmaster).

In this paper, we remove external restrictions on the nature of firms’ pricing strategies and instead allow firms to choose any *supply function*: a mapping that describes the price charged at each quantity of production.<sup>1</sup> Supply function choice is a standard approach in microeconomic theory to model firms’ ability to adjust decisions to realized demand without imposing *ad hoc* strategic restrictions, while remaining consistent with a foundation of information, contracting, or organizational frictions (*e.g.*, Grossman, 1981b; Hart, 1985; Klemperer and Meyer, 1989; Vives, 2011, 2017; Pavan et al., 2022; Rostek and Yoon, 2023). However, supply function choice has not yet been studied in general equilibrium, macroeconomic models. Our goal is to understand how this enriched model of pricing and production at the microeconomic level affects our understanding of the macroeconomy.

We find that introducing supply functions in an otherwise standard monetary business cycle model yields an aggregate supply curve with an *endogenous* slope. That is, the relative response of the price level and real output to an aggregate demand shock depends on the interaction between uncertainty and market structure, precisely because these forces affect

---

<sup>1</sup>This is different from *nonlinear pricing* (as recently studied in a macroeconomic context by Bornstein and Peter, 2022), 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.

firms' optimal supply strategies. Specifically, we show that aggregate supply flattens, or aggregate demand shocks have bigger real and smaller nominal effects, under lower inflation uncertainty, higher idiosyncratic demand uncertainty, and less elastic demand. Quantitatively, our model generates variation in the slope of aggregate supply that is consistent with empirical evidence in the US. Thus, we find that supply functions provide a realistic, tractable, and quantifiable foundation for a state-dependent aggregate supply curve.

**Supply Function Choice of a Single Firm.** We begin our analysis in partial equilibrium. We study a firm that 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 beliefs, the firm chooses a *supply function*  $f : \mathbb{R}_{++}^2 \rightarrow \mathbb{R}$ . This function defines the firm's supply curve as the locus of prices ( $p$ ) and quantities ( $q$ ) that solves  $f(p, q) = 0$ . Because the market clears, the firm produces and prices where the market demand curve intersects its supply curve. Internalizing this, the firm chooses its optimal, non-parametric supply function to maximize its expected real profits under the stochastic discount factor. This interpretation is in line with the ECON 101 notion of a supply curve: a systematic relationship between the price that firms charge and the quantity that they produce. We show formally that this model is isomorphic to one in which firms can condition their prices on realized demand. This interpretation links our model theoretically to the notion of rational expectations equilibrium (Lucas, 1972; Grossman, 1981a) and practically to the aforementioned examples of dynamic pricing. The model moreover nests two strategic restrictions imposed by previous studies. Price-setting is nested by functions of the form  $f(p) = 0$ , or perfectly elastic supply. Quantity-setting is nested by functions of the form  $f(q) = 0$ , or perfectly inelastic supply (*e.g.*, as in Jaimovich and Rebelo, 2009; Angeletos and La'O, 2010, 2013; Benhabib et al., 2015). By relaxing strategic restrictions on firms, we allow firms to choose potentially preferable strategies.

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 changes in market demand is described by its optimally chosen inverse supply elasticity,  $\alpha_1$ : 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 price elasticity of demand (which measures its market power in our model) 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.

Three comparative statics are particularly important for our macroeconomic analysis. First, higher uncertainty about firm-level demand pushes toward a lower  $\alpha_1$ , or firms behav-

ing more like price-setters. The limit case of price-setting perfectly insulates firms against demand shocks, as the optimal response of a firm to changing demand conditions is to set its relative price equal to a constant markup on its real marginal cost. Second, higher uncertainty about competitors' prices pushes toward a higher  $\alpha_1$ , or firms' behaving more like quantity-setters. The limit case of 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, a lower elasticity of demand pushes toward a lower  $\alpha_1$ , or firms behaving more like price-setters. More market power, thus defined, 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 next embed our framework in a monetary business-cycle model with incomplete information (Woodford, 2003a; Hellwig and Venkateswaran, 2009).

We first characterize aggregate outcomes given fixed firm-level supply functions. This allows us to isolate the importance of supply functions for shock transmission, before studying equilibrium choice of supply functions. We show that, in the unique log-linear equilibrium, the price level and real output follow an *aggregate supply and aggregate demand* representation. There is a well-defined “slope of aggregate supply,” which also corresponds to the relative responses of the price level and real output to an aggregate demand (money supply) shock. This slope depends critically on the slope of firms' supply functions.

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 and only if firms are price-setters. These results are disquieting—the specification of price-setting or quantity-setting has substantial implications for basic macroeconomic properties of the model, such as whether money has real effects. A key benefit of the supply function approach is that the analyst does not inadvertently impose restrictions on firms' supply function choices, but allows these choices to be made optimally. Indeed, between those extremes, the slope of aggregate supply is monotone increasing in the slope of firm-level supply.

Finally, a lower elasticity of demand flattens the aggregate supply curve. This effect is present as long as firms are *not* pure price-setters. Intuitively, a higher elasticity of demand increases how much a given change in the *aggregate* price level moves any given firm's demand curve.

We next characterize how the slope of aggregate supply is endogenously determined, via the fixed point relating macroeconomic uncertainty to 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.

For the sharpest closed-form illustration of these ideas, we first derive the slope of aggregate supply under a parameter restriction that balances strategic complementarity (from aggregate demand externalities) with substitutability (from wage pressure). This slope 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 volatile money supply, features a flatter aggregate supply curve and therefore an endogenously smaller effect of aggregate demand shocks on the price level. In that sense, this economy has “more stable prices” for two reinforcing reasons. First, there are fewer aggregate demand shocks. Second, firms respond to more stable prices by flattening their supply curves, which endogenously reduces the responsiveness of prices to aggregate demand shocks. This observation is 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 flattening the aggregate supply curve. Moreover, this implies the following trade-off for policymakers: maintaining a high degree of monetary discretion, which induces greater monetary uncertainty, endogenously makes monetary policy less effective at influencing the real economy.

An economy with higher idiosyncratic demand variation also features a flatter aggregate supply curve. Combining this with the empirical observation that firms’ idiosyncratic uncertainty rises substantially in recessions (Bloom et al., 2018), 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.

Away from the special case of balanced strategic interaction, the model makes richer predictions in which the elasticity of demand and the volatility of productivity shocks also affect the slope of aggregate supply. Our quantitative exercise will suggest that both forces play an important role in determining the slope of aggregate supply in the US.

**Quantitative Analysis.** We finally study the model’s implications for the slope of aggregate supply in the United States. To do so, we estimate time-varying uncertainty for macroeconomic aggregates using a Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model on aggregate time series for output, the price level, and productivity. We combine these estimates with the model to generate a historical time series for the slope of aggregate supply in the US. This allows us to study whether the model’s predictions for the state-varying slope of aggregate supply are quantitatively reasonable.

We find that the model helps explain three empirically documented phenomena related to the changing slope of aggregate supply in the US that existing models have struggled to rationalize. First, the model explains a quantitatively significant portion of the steeping of

aggregate supply from the 1960s to the 1970s and the flattening of aggregate supply from the 1970s to the Great Moderation (as estimated by, *e.g.*, [Ball and Mazumder, 2011](#)). The changes are primarily driven by changes in inflation uncertainty. When inflation uncertainty is low, as during the Great Moderation, firms choose flatter supply functions (*i.e.*, closer to price-setting). This flattening at the micro level translates to a flatter aggregate supply curve. Second, our model rationalizes why aggregate supply remained flat during the Great Recession, a period characterized by a spike in real, rather than nominal, uncertainty. Thus, our model is consistent with both the missing disinflation ([Coibion and Gorodnichenko, 2015](#)) and missing inflation ([Bobeica and Jarociński, 2019](#)) puzzles: inflation did not fall during the Great Recession and rise thereafter by as much as standard macroeconomic models imply. Finally, the model explains the steepening of aggregate supply in the post-Covid period (as has been estimated by [Cerrato and Gitti, 2022](#)) as a consequence of a surge in inflation uncertainty. This implies both that post-pandemic aggregate demand shocks may have had large effects on prices and that contractionary monetary policy might be able to rein in inflation with relatively low costs to output and usher in a “soft landing”.

**Related Literature.** Our methodological contribution is to derive aggregate supply in a business-cycle model from a foundation of supply function competition. Supply function competition has been extensively studied in microeconomic theory, industrial organization, and finance ([Grossman, 1981b](#); [Hart, 1985](#); [Klemperer and Meyer, 1989](#); [Kyle, 1989](#); [Vives, 2017](#); [Pavan et al., 2022](#); [Rostek and Yoon, 2023](#)). We contribute to this theoretical literature by analytically characterizing equilibrium supply functions with several new features: non-quadratic preferences; imperfect substitutability; multiple, correlated sources of uncertainty; and, most importantly, general equilibrium interactions in both input and product markets. Moreover, by using estimated macroeconomic volatility to place discipline on the set of possible information structures that are consistent with the observed data, our sufficient statistics approach allows us to make sharp quantitative statements about the macroeconomic implications of supply functions. This approach allows us to overcome the [Bergemann et al. \(2021\)](#) critique that supply functions have few predictions that hold over *all* information structures.

The closest analysis in the macroeconomics literature on optimal supply decisions is performed by [Reis \(2006\)](#), who compares the polar extremes of price-setting and quantity-setting for a rationally inattentive firm in partial equilibrium. Our analysis goes beyond Reis’ by studying completely flexible supply schedule choice and removing all *ad hoc* strategic restrictions on firms’ choices. Moreover, we characterize general equilibrium, as opposed to [Reis’ \(2006\)](#) partial equilibrium analysis.<sup>2</sup> This allows us to study the equilibrium relationship

---

<sup>2</sup>In an earlier version of this project ([Flynn et al., 2024](#)), we studied the restricted problem of prices *vs.* quantities choice in general equilibrium.

between supply function choice, macroeconomic dynamics, and the slope of aggregate supply. In so doing, we also allow for multiple, correlated shocks to firms. The main upshot from this perspective is our new theory of aggregate supply and its implications for macro dynamics.

Our finding that uncertainty shapes the slope of aggregate supply is shared with the classic “islands model” analysis of Lucas (1972, 1973, 1975). A shared methodological premise is that economic agents act on what they learn from endogenous objects, as in the broader literature on rational expectations equilibrium (Grossman, 1981a; Grossman and Stiglitz, 1980). Our ultimate results for the slope of the aggregate supply curve differ substantially for two reasons. First, we study producers with market power (monopolistic competition), consistent with modern macroeconomic theory and evidence, instead of price-taking producers in competitive markets. Second, 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). Rather, firms use the demand for their product as a noisy signal to infer their optimal price.

Our work is also distinguished from a literature that has pursued other avenues to reconcile Lucas’ insights with non-competitive markets. Unlike Woodford (2003a), which restricts firms to price-setting, we allow firms to choose flexible schedules. This restores the spirit of Lucas’ insight that firms can learn from market conditions in rational expectations equilibrium. Our analysis also suggests that existing conclusions about the link between information frictions and monetary non-neutrality are sensitive to strategic restrictions on firms: for example, if firms were restricted to set quantities in our model, money would be neutral despite information frictions. Hellwig and Venkateswaran (2009) share our premise of allowing firms to learn from demand conditions, but do not study the static fixed point that supply functions generate between firms’ decisions and market information. This two-way feedback is at the core of our mechanism and our predictions for how information and market power shape the slope of aggregate supply.<sup>3</sup>

**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 quantifies the model’s predictions for the slope of aggregate supply. Section 6 concludes.

---

<sup>3</sup>Lucas and Woodford (1993) and Eden (1994) study markets with *ex ante* capacity investment and sequential transactions as a way to model learning from demand conditions. These authors also do not study the static fixed-point between uncertainty and market information.

## 2 Supply Function Choice in Partial Equilibrium

In this section, we introduce our model of supply function choice for a single firm making decisions under uncertainty. We show that supply function choice is formally equivalent to allowing firms to learn from their demand and update their pricing strategies accordingly. Our main result in this section shows that optimal supply functions are log-linear and characterizes their slope in terms of firms' uncertainty and the elasticity of demand.

### 2.1 The Firm's Problem

**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  $\eta$  is high, the quantity demanded is more sensitive to the price. 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  $\mu$  and variance  $\Sigma$ .<sup>4</sup> 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  $(\Lambda, P, \mathcal{M}, p, q)$ .

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  $f$ , allowing for possible non-monotonicity and discontinuities.

---

<sup>4</sup>Of course,  $\mathcal{M}$  is log-normal so long as  $P$ ,  $\Theta$ , and  $p_x$  are log-normal.


After choosing a supply function  $f$ , and following the realization of  $\Psi$  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 firm's price is given by some solution  $\hat{p}$  to the equation  $f(\hat{p}, z\hat{p}^{-\eta}) = 0$  with the realized quantity being  $\hat{q} = z\hat{p}^{-\eta}$ . We assume that the firm chooses the profit-maximizing selection from the set of solutions if there are many and does not produce if there is no solution. Given a supply function  $f$ , we let  $H(f)$  be the induced joint distribution over  $(\Lambda, P, \mathcal{M}, p, q)$  given the firm's prior beliefs. The firm's problem of choosing an optimal supply function is therefore equivalently stated as either of the following maximization problems:

$$\sup_{f: \mathbb{R}_{++}^2 \rightarrow \mathbb{R}} \mathbb{E}_{H(f)} \left[ \Lambda \left( \frac{p}{P} - \mathcal{M} \right) q \right] \iff \sup_{\hat{p}(z)} \mathbb{E} \left[ \Lambda \left( \frac{\hat{p}(z)}{P} - \mathcal{M} \right) z \hat{p}(z)^{-\eta} \middle| z \right] \text{ for all } z \in \mathbb{R}_+ \quad (4)$$

While mathematically equivalent, these two formulations of the problem provide two different economic intuitions for how the firm behaves. Under the first formulation, the interpretation is that the firm chooses its supply curve *ex ante*, knowing that it will price and produce where its supply curve meets the demand curve. Under the second formulation, the interpretation is that the firm prices in the *interim*: it is *as if* the firm sees the state of its demand, updates its beliefs, and then sets its optimal price.

In Figure 1, we visually illustrate how different supply functions translate into different outcomes for the firm. In the first row, we show each supply function (solid black line) and two demand curves, corresponding to a high demand realization  $z_1$  (red dashed line) and a low demand realization  $z_0$  (blue dotted line). The dots indicate the intersections of supply and demand, or realized quantity-price pairs in these states. In the second row, we illustrate the induced joint distribution of prices and quantities. Panel (a) shows a “price-setting” supply function,  $f(p, q) = 1 - p$ , which allows only quantities to vary with realized demand. This might describe a firm that responds to low demand by producing less and responds to high demand by producing more, allowing the market to clear at a fixed price. The “quantity-setting” policy (panel (b)),  $f(p, q) = 1 - q$ , does the opposite: it might describe a retailer that aggressively decreases the price of low-demand goods and increases the price of high-demand goods to fix the quantity sold. The third supply function (panel (c)),  $f(p, q) = 1 - \frac{p}{q}$ , allows both prices and quantities to increase with higher demand. This might describe a retailer with less extreme dynamic pricing: high-demand states have higher prices and volumes, and low-demand states have lower prices and volumes. To evaluate a supply function, the firm evaluates its payoffs given the induced joint distribution of prices and quantities with competitors' prices, real marginal costs, and the stochastic discount factor.

**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  $z$ .

**Interpreting Supply Functions.** The *ex ante* interpretation of a supply function is familiar from ECON 101: it is the price that the firm plans to charge given any level of production. At a basic level, setting a supply function allows the firm to change its price and level of production as a function of the demand that it faces. This allows our model to formalize the idea that a producer may want to raise their price and increase production when demand is high and goods are flying off the shelves and lower their price and reduce production when demand is weak.

The *interim* interpretation of a supply function is that the firm learns from its demand and chooses its optimal behavior incorporating all of this possible information. This is exactly the idea that underpins “learning from prices” in rational expectations equilibrium (REE), as studied (among others) by Lucas (1972) and Grossman (1981a). To assume that the firm does not learn from this information and update its behavior is to assume that it either disregards valuable information—which is in principle reasonable, but at odds with the modern paradigm of studying REE—or that firms for some exogenous reason *cannot* incorporate this information into their decisions. Even in models with positive but finite adjustment costs, this second assumption is violated. Thus, the only tenable arguments against allowing for supply functions are that firms do not learn in the manner required by

REE or that adjustments are infinitely costly.

The interim perspective makes clear that supply functions have an appealing property: they do not require any commitment on the part of the firm. Indeed, as the interim perspective makes clear, the defining property of an optimal supply curve is that the firm will *ex post, i.e.*, upon observing demand, actually wish to charge the price and supply the quantity specified by the optimal supply curve. By contrast, the standard price-setting assumption does require a commitment on the part of the firm to ignore the information that it learns from the state of demand and not change prices in response to demand conditions.

As a practical matter, there are many examples in which firms implement supply schedules *outright* precisely because they learn from demand and seek to adjust prices to maximize profits by using this information (see *e.g.*, Farias and Van Roy, 2010; Dutta and Mitra, 2017). This is common for many goods and services that feature “dynamic pricing,” which include but are not limited to: electricity (Joskow and Wolfram, 2012), air travel (Williams, 2022), gasoline (Borenstein and Shepard, 1996), e-commerce (*e.g.*, Amazon), ride-sharing (*e.g.*, Uber and Lyft), food (*e.g.*, Wendy’s), entertainment and sports events (*e.g.*, Ticketmaster), and theme parks (*e.g.*, Disney World).

Moreover, demand-responsive pricing is no new innovation. Posted prices have not been the norm throughout human history and only became widespread after the invention of the price tag in the mid-19th century (Phillips, 2012). The practical importance of supply functions has long been recognized by industrial organization economists. Writing in 1989, Klemperer and Meyer provide two particularly concrete examples of firms that *de facto* implement supply schedules. In the first, they describe how service providers (specifically, management consultants) vary the prices that they charge in response to the quantity of services provided.<sup>5</sup> In the second, they describe how airlines use computer software to put seats on discount depending on how many are currently sold. In this case, the firm explicitly uses technology to implement a supply schedule. Of course, even if firms do not implement supply schedules *explicitly*, they may still implicitly use demand to inform their optimal pricing strategy and therefore act *as if* they are choosing a supply function.

In summary, as REE and non-infinite adjustment costs imply that firms must choose supply functions, we argue that there is a strong theoretical basis for their study. Moreover,

---

<sup>5</sup>They write: “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—some firms hold the real price very close to the quoted one by choosing very rigid rules about accurately reporting the hours worked to the client, while others allow individual managers far more discretion.”

as the pricing decisions of many firms are transparently and explicitly designed as supply functions, there is a strong practical basis for studying supply functions.

## 2.2 The Optimal Supply Function

We now study the globally optimal supply function, the solution to Problem 4. The following result characterizes the firm's optimal policy in closed form and allows us to illustrate comparative statics in the extent of uncertainty and the price elasticity of demand.

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

$$f(p, q) = \log p - \alpha_0 - \alpha_1 \log q \quad (5)$$

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

$$\alpha_1 = \frac{\eta\sigma_P^2 + \sigma_{\mathcal{M},\Psi} + \sigma_{P,\Psi} + \eta\sigma_{\mathcal{M},P}}{\sigma_{\Psi}^2 - \eta\sigma_{\mathcal{M},\Psi} + \eta\sigma_{P,\Psi} - \eta^2\sigma_{\mathcal{M},P}} \quad (6)$$

*Proof.* See Appendix A.1. □

To provide intuition for this result, it is helpful to first sketch its proof. As observed above (Equation 4), the problem of choosing an optimal supply function *ex ante* can be recast as a problem of choosing price-quantity pairs  $(p(z), q(z))$  that are indexed by the realization of the nominal demand state  $z = \Psi P^\eta$  and are such that the market clears:  $(p(z), q(z)) = (p(z), zp(z)^{-\eta})$ . Intuitively, when setting a supply schedule, the firm anticipates that it will produce where the demand curve hits the supply function. Thus, as the demand curve is indexed by  $z$ , it is as if the firm chooses a  $z$ -contingent price-quantity plan. Under price-setting, for instance, the price is fixed at  $\bar{p} \in \mathbb{R}_{++}$  and the quantity adjusts to clear the market,  $(p(z), q(z)) = (\bar{p}, z\bar{p}^{-\eta})$ . Similarly, under quantity-setting, the quantity is fixed at  $\bar{q} \in \mathbb{R}_{++}$  and the price adjusts to clear the market,  $(p(z), q(z)) = (z^{1/\eta}\bar{q}, \bar{q})$ . In a general problem of supply function choice, the only difference is that the contingency of prices and quantities on realized demand, which was also present under both price-setting and quantity-setting, is chosen optimally to maximize payoffs.

A necessary condition for optimality is that, for any given realization  $z = t$ , there is no local benefit to changing the price  $p(t)$ . Taking a first-order condition of the firm's maximization problem implies that the following equations must hold for almost all  $t \in \mathbb{R}_{++}$ :

$$p(t) = \frac{\eta}{\eta - 1} \frac{\mathbb{E}[\Lambda \mathcal{M} \mid z = t]}{\mathbb{E}[\Lambda P^{-1} \mid z = t]} \quad \text{and} \quad q(t) = tp(t)^{-\eta} \quad (7)$$

This resembles the standard optimality condition for monopolistic price-setting (“markup over marginal cost”), 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. We are then able to solve for the optimal supply function in closed form, despite the infinite-dimensionality of Problem 4, because Equation 7 reduces to a log-linear relation between  $p$  and  $q$  given lognormal beliefs.

It remains to explain why the optimal inverse supply elasticity takes the form given in Equation 6. This specific form arises because  $\alpha_1$  is the relative rate at which the firm wants log prices and log quantities to increase with the nominal demand state  $\log z$ :

$$\alpha_1 = \frac{d \log p}{d \log z} \bigg/ \frac{d \log q}{d \log z} = \frac{\text{Cov}[\log z, \log p^{**}]}{\text{Cov}[\log z, \log q^{**}]} \quad (8)$$

where  $p^{**}$  and  $q^{**}$  are the optimal *ex post* prices and quantities that the firm would set with full information:

$$p^{**} = \frac{\eta}{\eta - 1} \mathcal{M}P \quad \text{and} \quad q^{**} = \left( \frac{\eta}{\eta - 1} \right)^{-\eta} \frac{z}{(\mathcal{M}P)^\eta} \quad (9)$$

An econometric metaphor illustrates why this is the optimal way to set  $\alpha_1$ . By Equation 8, the firm’s optimal policy is equivalent to running the following two-stage least squares (2SLS) regression: the firm estimates how its optimal price should change with its optimal quantity, using 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, this corresponds to a large coefficient in the “reduced form” regression of  $p^{**}$  on  $z$ . The supply function is flat ( $|\alpha_1|$  is small) if nominal demand predicts large movements in the *ex post* optimal quantity. In the 2SLS metaphor, this corresponds to a large coefficient in the “first stage” regression of  $q^{**}$  on  $z$ .

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 ( $\sigma_p^2$  increases) steepens the optimal supply schedule. Intuitively, not knowing the prices of your competitors makes more aggressive dynamic pricing (*i.e.*, a strategy closer to quantity-setting) attractive because this allows one’s *relative* price to adjust *ex post*. On the other hand, greater demand uncertainty ( $\sigma_\Psi^2$  increases) flattens the optimal supply schedule. Intuitively, demand uncertainty favors a strategy closer to a fixed price 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.

**The Elasticity of Demand and Supply Functions.** The elasticity of demand plays two roles in determining the optimal (inverse) elasticity of supply. The first relates to *payoffs*: when  $\eta$  is high, *ex post* optimal quantities are more sensitive to changes in nominal marginal costs (holding fixed nominal demand). Intuitively, 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  $\eta$  is high, nominal demand contains relatively more information about the price level  $P$  and less about real demand  $\Psi$ . 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 market power in the macroeconomy.

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

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

$$\alpha_1 \geq 0, \sigma_{M,P} \geq 0, 2\eta\sigma_{M,P} + \sigma_{M,\Psi} \geq \sigma_{P,\Psi} \quad (10)$$

*Proof.* See Appendix A.2. □

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 demand is highly sensitive to relative prices, 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 quantitative 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 a lower elasticity of demand flattens firms’ optimal supply function.

**Pure Price- and Quantity-Setting Obtain in Extreme Limits.** The previous result and discussion make clear that pure price-setting (in which the price is unresponsive to demand while the quantity produced is) and quantity-setting (in which the quantity produced is unresponsive to demand while the price charged is) are “edge cases” in the larger space of supply functions. Moreover, they are almost never optimal. We observe below that they are 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.

Thus, focusing on price- and quantity-setting is justified when *and only when* 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.

**Extensions: Multiple Inputs, Decreasing Returns to Scale, Monopsony, and Beyond Isoelastic Demand.** We have made many simplifying assumptions for expositional simplicity. In Appendix B, we generalize this analysis in two directions. First, we 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 change the analysis solely by introducing a single composite parameter that aggregates the decreasing returns and monopsony forces across inputs. We show in Proposition 4 that the optimal supply function remains optimally log-linear and uncertainty enters in a similar way. Moreover, as is perhaps intuitive, decreasing returns to scale and monopsony power (that may arise because of adjustment costs in production, for instance) both reduce the optimal supply elasticity of the firm and push the firm toward quantity-setting. Second, we allow for demand that is not iso-elastic and separates the firm’s own-price elasticity of demand from the firm’s cross-price elasticity of demand. We solve for the optimal supply curve in this case in Proposition 5. We show that uncertainty enters in a similar way but the optimal supply curve ceases to be log-linear as the optimal markup is endogenous to the scale of production. Under all such extensions, the qualitative insights remain the same.

### 3 Supply Functions in a 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](#)). These microfoundations will allow for a closed-form analysis and highlight the core economics of supply functions without any approximations.<sup>6</sup> In this context, we will be interested in understanding three things: (i) how the microeconomic inverse supply elasticity maps into the elasticity of aggregate supply, (ii) how equilibrium macroeconomic dynamics endogenously influence the optimal microeconomic supply elasticity, and (iii) how these two channels interact to determine equilibrium macroeconomic dynamics.

#### 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 Venkateswaran, 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,  $N_{it}$ :

$$\mathbb{E}_0 \left[ \sum_{t=0}^{\infty} \beta^t \left( \frac{C_t^{1-\gamma}}{1-\gamma} + \ln \frac{M_t}{P_t} - \int_{[0,1]} \phi_{it} N_{it} \, di \right) \right] \quad (11)$$

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  $\eta > 1$ :

$$C_t = \left( \int_{[0,1]} \vartheta_{it}^{\frac{1}{\eta}} c_{it}^{\frac{\eta-1}{\eta}} \, di \right)^{\frac{\eta}{\eta-1}} \quad (12)$$

where  $\vartheta_{it}$  is an IID preference shock that is also lognormal with time-dependent variance, or

---

<sup>6</sup>In Section 4.5, we describe how to perform the same analysis in a large family of linearized macroeconomic models following [McKay and Wolf \(2023\)](#).

$\log \vartheta_{it} \sim N(\mu_\vartheta, \sigma_{\vartheta,t}^2)$ . We also define the corresponding ideal price index:

$$P_t = \left( \int_{[0,1]} \vartheta_{it} p_{it}^{1-\eta} di \right)^{\frac{1}{1-\eta}} \quad (13)$$

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  $\Pi_{it}$ . Thus, the household faces the following budget constraint at each time  $t$ :

$$M_t + B_t + \int_{[0,1]} p_{it} C_{it} di = M_{t-1} + (1 + i_{t-1})B_{t-1} + \int_{[0,1]} w_{it} N_{it} di + \int_{[0,1]} \Pi_{it} di \quad (14)$$

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  $\mu_M$  and time-dependent volatility  $\sigma_t^M$ :

$$\log M_t = \log M_{t-1} + \mu_M + \sigma_t^M \varepsilon_t^M \quad (15)$$

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  $t \in \mathbb{N}$ .

### 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:

$$q_{it} = \zeta_{it} A_t L_{it} \quad (16)$$

where  $L_{it}$  is the amount of labor employed,  $\zeta_{it}$  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  $\sigma_t^A$ :

$$\log A_t = \rho \log A_{t-1} + \sigma_t^A \varepsilon_t^A \quad (17)$$

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  $\Lambda_t$ . 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  $s_{it}^M$ :

$$\begin{aligned} s_{it}^A &= \log A_t + \sigma_{A,s,t} \varepsilon_{it}^{s,A} \\ s_{it}^M &= \log M_t + \sigma_{M,s,t} \varepsilon_{it}^{s,M} \end{aligned} \tag{18}$$

where the signal noise is IID and follows  $\varepsilon_{it}^{s,A}, \varepsilon_{it}^{s,M} \sim N(0, 1)$ . Firms are uncertain about the idiosyncratic productivity shock  $z_{it}$ , demand shock  $\vartheta_{it}$ , and labor supply shock  $\phi_{it}$ .<sup>7</sup>

### 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  $\{f_{it} : \mathbb{R}_{++}^2 \rightarrow \mathbb{R}\}_{i \in [0,1], t \in \mathbb{N}}$  such that, in all periods:

1. All firms choose their supply function  $f_{it}$  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  $\Lambda_t$  is the household's marginal utility of consumption.
3. Money supply  $M_t$  and productivity  $A_t$  evolve exogenously via Equations 15 and 17.
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.

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.

---

<sup>7</sup>It is not important that firms are fully uninformed about these quantities. The model's predictions would be identical if firms also received noisy signals about their idiosyncratic shocks.

## 4 Supply Function Choice and Aggregate Supply

We now study the model's equilibrium predictions, focusing on the equilibrium determination of the aggregate supply curve. We proceed in three steps. First, we solve for all equilibrium conditions except for the firm's supply-function decision. Second, we show that, fixing any log-linear supply schedule, the economy admits a unique log-linear equilibrium that has a simple Aggregate Supply and Aggregate Demand representation. Importantly, the slope of aggregate supply depends on the slope of firm-level supply, in conjunction with other parameters. Third, we combine this with our solution for optimal supply schedules from Theorem 1 and fully characterize equilibrium in terms of a single, scalar fixed-point equation for the firm-level supply elasticity. We study how strategic interactions, the elasticity of demand, and the combination of microeconomic demand uncertainty alongside aggregate productivity and monetary uncertainty affect the equilibrium aggregate supply elasticity. Finally, we show how supply function choice can be tractably incorporated in a larger class of dynamic general equilibrium models.

### 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. We do so by deriving the household's Euler equations for bonds, money, and labor supply. We summarize the results of this below:

**Proposition 1** (Firm-Level Shocks in General Equilibrium). *In any temporary equilibrium, demand shocks, aggregate price shocks, stochastic discount factor shocks, and marginal cost shocks follow:*

$$\Psi_{it} = \vartheta_{it} C_t, \quad P_t = \frac{i_t}{1 + i_t} C_t^{-\gamma} M_t, \quad \Lambda_t = C_t^{-\gamma}, \quad \mathcal{M}_{it} = \frac{\phi_{it} C_t^\gamma}{z_{it} A_t} \quad (19)$$

where  $i_t$  is a deterministic function of only exogenous parameters that we provide in the Appendix.

*Proof.* See Appendix A.3. □

Each of these expressions is intuitive given the general equilibrium structure of the model. First, the firm's demand shock is the product between its idiosyncratic demand shock and aggregate demand, which is familiar from Blanchard and Kiyotaki (1987). Second, the demand for real money balances is decreasing in the interest rate as this determines the opportunity cost of holding money (which itself depends on the future path of monetary volatility, the

drift of the money supply, and the household’s discount factor). Moreover, this demand is increasing in the household’s level of consumption because of an income effect, which is governed by the curvature of consumption utility  $\gamma$ . 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  $\gamma$  is large. Third, the SDF is simply the marginal utility of consumption. Finally, the real marginal cost of firms is increasing in the level of consumption because of the same income effect, and decreasing in the productivity of the firm.

The uncertainty the firm faces in light of Proposition 1 concerns endogenous objects. This introduces strategic uncertainty (*i.e.*, payoff-relevant uncertainty about other firms’ choices).<sup>8</sup> Moreover, firms’ uncertainty is correlated across variables due to macroeconomic linkages in the product, money, and labor markets.

An important technical implication of Proposition 1 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 our 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 start by assuming that firms’ exogenously set log-linear supply functions:

$$\log p_{it} = \alpha_{0t,i}^*(\alpha_{1,t}) + \alpha_{1,t} \log q_{it} \tag{20}$$

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.<sup>9</sup> This optimal intercept depends on the slope  $\alpha_{1,t}$ , the firm’s beliefs, and realized demand, but not (independently) on the realized quantity. This has two purposes. First, this assumption allows us to explore what happens in temporary equilibrium when firms use a given supply function. This is useful for understanding what strategic restrictions on firms’ pricing strategies (*e.g.*, exogenously imposing price-setting) imply for macroeconomic dynamics. Second, this assumption is our “guess” about what firms’ supply function will be in equilibrium, which we will later “verify” as correct. This allows us to

---

<sup>8</sup>One interesting implication of Proposition 1 is that nominal wages,  $w_{it} = \frac{i_t}{1+i_t} \phi_{it} M_t$ , provide information *only* about exogenous objects. A stronger implication is that a model in which firms draw inferences from both output-market prices and input-market prices has identical predictions to our studied model.

<sup>9</sup>We will later verify that all firms use a common slope in equilibrium. In light of Theorem 1, this is because all firms are exposed to uncertainty in the same way.

understand the ultimate macroeconomic implications of optimal supply function choice.

Conditional on these supply functions, we guess and verify that there exists an equilibrium in which aggregate consumption and the price level are log-linear in aggregate shocks:

$$\begin{aligned}\log P_t &= \chi_{0,t}(\alpha_{1,t}) + \chi_{A,t}(\alpha_{1,t}) \log A_t + \chi_{M,t}(\alpha_{1,t}) \log M_t \\ \log C_t &= \tilde{\chi}_{0,t}(\alpha_{1,t}) + \tilde{\chi}_{A,t}(\alpha_{1,t}) \log A_t + \tilde{\chi}_{M,t}(\alpha_{1,t}) \log M_t\end{aligned}\tag{21}$$

To this end, we define the posterior weight on firms' signals of productivity and the aggregate money supply as, respectively,  $\kappa_t^A = \left(1 + (\sigma_{A,s,t}/\sigma_t^A)^2\right)^{-1}$  and  $\kappa_t^M = \left(1 + (\sigma_{M,s,t}/\sigma_t^M)^2\right)^{-1}$ . Moreover, define the slope of supply functions in terms of  $\log z_{it} = \eta \log P_t + \log \Psi_{it}$  as:<sup>10</sup>

$$\omega_{1,t} = \frac{\alpha_{1,t}}{1 + \eta\alpha_{1,t}}\tag{22}$$

We now characterize equilibrium macroeconomic dynamics with fixed supply functions. We show that macroeconomic dynamics in log-linear general equilibrium are equivalent to those that would be generated by an Aggregate Demand and Aggregate Supply (AD/AS) model, in which productivity shocks shift the AS curve and money shocks shift the AD curve. Critically, the slope of aggregate supply depends on the slope of firms' supply schedules.

**Theorem 2** (Equilibrium and AD/AS Representation). *There is a unique log-linear temporary equilibrium. The behavior of aggregate prices and output in this temporary equilibrium is equivalent to those generated by the following ‘‘Aggregate Demand/Aggregate Supply’’ model:*

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

$$\log P_t = \log \bar{P}_t + \epsilon_t^S \log Y_t + \delta_t \log A_t\tag{AS}$$

where the inverse supply and demand elasticities are given by:

$$\epsilon_t^S = \gamma \frac{\kappa_t^M + \frac{\omega_{1,t}}{\gamma}(1 - \kappa_t^M)}{(1 - \omega_{1,t}\eta)(1 - \kappa_t^M)} \quad \text{and} \quad \epsilon_t^D = \gamma\tag{23}$$

and the interest rate  $i_t$ , the intercept for the price level  $\log \bar{P}_t$ , and the partial equilibrium effect of productivity shocks  $\delta_t$  do not depend on  $(\log P_t, \log Y_t, \log M_t, \log A_t)$ .<sup>11</sup>

*Proof.* See Appendix A.4. □

<sup>10</sup>So everything remains well defined, we will impose that  $\omega_{1,t} \neq \left(\eta - \frac{1}{\gamma}\right)(1 - \kappa_t^x)$  for  $x \in \{A, M\}$ . Our analysis verifies that these values of  $\omega_{1,t}$  cannot occur in log-linear equilibrium (see the proof of Theorem 3).

<sup>11</sup>See Appendix A.4 for explicit formulae for these terms.

In this representation, the aggregate demand curve 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 has an elasticity of  $1/\gamma$  to changes in real money balances. Thus, the “inverse elasticity of aggregate demand” in our model is  $\gamma$ . 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.

We illustrate this representation in Figure 2. An “aggregate demand shock,” an increase of the money supply by  $\log M_1 - \log M_0 = \Delta \log M > 0$ , shifts up the AD curve. This has an effect of  $\frac{\Delta \log M}{\epsilon^D + \epsilon^S}$  on real output and  $\epsilon^S \frac{\Delta \log M}{\epsilon^D + \epsilon^S}$  on the price level. In particular, the price effect is larger and the quantity effect is smaller if  $\epsilon^S$  is large. This calculation also makes clear that  $\epsilon^S$  measures the relative effect of an aggregate demand shock on the price level versus real output.


**The Propagation of Demand Shocks.** To obtain more intuition for the propagation of shocks via firms’ supply schedules, we expand the response of the price level to a money shock into a partial equilibrium effect and a series of higher-order general equilibrium effects:<sup>12</sup>

$$\frac{\Delta \log P}{\Delta \log M} = \frac{\epsilon_t^S}{\epsilon_t^D + \epsilon_t^S} = \underbrace{\left( \kappa_t^M + \frac{\omega_{1,t}}{\gamma} (1 - \kappa_t^M) \right)}_{\text{Partial Equilibrium}} \times \underbrace{\sum_{j=0}^{\infty} \left( \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right)^j}_{\text{General Equilibrium}} \quad (24)$$

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  $\omega_{1,t}$ , 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

<sup>12</sup>This expression is derived by multiplying the numerator and denominator of  $\epsilon_t^S / (\epsilon_t^D + \epsilon_t^S)$  by  $(1 - \omega_{1,t}\eta)(1 - \kappa_t^M)$  and expanding it into a geometric summation. The summation only converges when  $|\omega_{1,t}(\eta - 1/\gamma)(1 - \kappa_t^M)| < 1$ . Our fixed point arguments establish that the claimed formulae hold more generally whenever  $|\omega_{1,t}(\eta - 1/\gamma)(1 - \kappa_t^M)| \neq 1$ . The proof of Theorem 3 shows the final case of  $|\omega_{1,t}(\eta - 1/\gamma)(1 - \kappa_t^M)| = 1$  cannot happen in equilibrium.

**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 Theorem 2).

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 markup 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 24.

A novel implication of our model is that the extent of general-equilibrium strategic complementarity hinges critically on the slope of the supply function. Starkly, general-equilibrium interactions would be entirely absent (*i.e.*, pricing decisions would be neither complements nor substitutes) 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%. Through this lens, predictions for complementarity in benchmark price-setting (Woodford, 2003a) and quantity-setting (Angeletos and La’O, 2010) models are joint predictions of the economic environment and an exogenous restriction on firms’ strategy space.

**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 transmission of supply shocks. In the AD/AS representation, a positive shock to  $\log A_t$  corresponds to an outward shift of the AS curve, which raises real output and lowers the price level. While the relative effect on the price level and on real output is fixed at  $\epsilon^D = \gamma$ , the level of these responses varies with the slope of firms’ supply functions,  $\omega_{1,t}$ .

To understand the reason for this, we can, just as above, decompose the effect into partial and general equilibrium components:

$$\frac{\Delta \log P}{\Delta \log A} = \underbrace{\kappa_t^A}_{\text{PE}} \times \underbrace{\sum_{j=0}^{\infty} \left( \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A) \right)^j}_{\text{GE}} \quad (25)$$

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  $1 - \kappa_t^M$ . In sum, 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 Slope of Aggregate Supply in Temporary Equilibrium

Having demonstrated that the response of the economy to demand and supply shocks can be understood in terms of the slope of the aggregate supply curve, we now formally investigate how various microeconomic forces affect it. The following result, the proof of which follows immediately from differentiation of Equation 23, describes how the slope of aggregate supply depends on four key parameters, holding fixed the others:

**Corollary 3** (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. *Precision of private information about money steepens 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. *A higher elasticity of demand steepens the AS curve:  $\partial \epsilon_t^S / \partial \eta \geq 0$ .*

To understand the first statement, observe that a steeper microeconomic supply function makes prices more responsive to realized quantities *ex post*. At the aggregate level, this implies that the price level is also more responsive to changes in output. Second, more precise private information about the money supply steepens the AS curve because firms respond to the perceived increase in the money supply by increasing *average* prices (as modulated through the intercept  $\alpha_{0,t,i}^*$ ). 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  $\gamma$  (cf. Proposition 1). Consequently, a higher  $\gamma$  steepens the AS curve.


Finally, a lower elasticity of demand 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  $\eta$ . 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. This prediction is opposite to the prediction that Woodford (2003b) obtains: in a New Keynesian model with decreasing returns to scale, the slope of the Phillips curve is lower when demand is more elastic.<sup>13</sup>

**Aggregate Supply Under Price-Setting and Quantity Setting.** We can illustrate some of these effects even more sharply by describing the slope of aggregate supply under the extreme, but common, assumptions of pure price-setting and quantity-setting. We find that the aggregate supply curve is vertical under quantity-setting and maximally flat under price-setting:

---

<sup>13</sup>Moreover, the interaction between market power and the slope of aggregate supply arises for completely different reasons. In the New Keynesian model, the logic is that: when demand is very elastic, higher prices translate to much lower quantities and, under decreasing returns, much lower marginal costs. This dampens the desired price change in response to a nominal cost shock.

**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 Theorem 2) under price-setting (panel a) and quantity-setting (panel b).

**Corollary 4** (Aggregate Supply Under Price- and Quantity-Setting). *If firms engage in price-setting ( $\omega_{1,t} = 0$ ), then:*

$$\epsilon_t^S = \gamma \frac{\kappa_t^M}{1 - \kappa_t^M} \quad (26)$$

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

$$\epsilon_t^S = \infty \quad (27)$$

We illustrate these two “extreme” predictions for aggregate supply and demand in Figure 3. Since  $\epsilon_t^S$  is increasing in  $\omega_{1,t}$ , the price-setting case provides a *lower bound* on the inverse elasticity of the aggregate supply curve and therefore maximizes the real effects of demand shocks. Moreover, as mentioned above, the slope is invariant to the elasticity of demand only in this case. The case of price-setting recovers the aggregate supply elasticity of Lucas (1972) with the same insight that more precise information about the money supply leads to a steeper aggregate supply curve.

In sharp contrast, the AS curve is vertical under quantity-setting and money has no 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. As a simple example, setting  $\log q_{it} = s_{it}^M$  is feasible for firms and this would imply that money has real effects:  $C_t \propto M_t$ . The second part of Corollary 4 follows from the fact that if firms set quantities, then there is no equilibrium in which firms’ quantities depend on the monetary signal.

These results emphasize that the kinds of strategies firms use have large macroeconomic consequences. It may be unappealing that the choice of the economic analyst about what kinds of strategies firms use has such large macroeconomic implications. A key benefit of the supply functions approach is that it allows the analyst to avoid imposing such restrictions and the potentially unintended consequences for macroeconomic predictions that follow.

#### 4.4 The Equilibrium Slope of Aggregate Supply

We now endogenize the firm-level inverse supply elasticity as a best response to equilibrium macroeconomic dynamics. We have verified that if firms use log-linear supply functions, then aggregate dynamics are endogenously log-linear (by Theorem 2). Moreover, we have verified that if aggregate dynamics are log-linear, then firms' uncertainty is endogenously log-normal (by Proposition 1). Thus, we have shown that firms' supply curves are endogenously log-linear in a log-linear equilibrium (by Theorem 1). By combining 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 the slopes of firms' supply functions:

**Theorem 3** (Equilibrium Supply Elasticity Characterization). *All (and only all) solutions  $\omega_{1,t} \in \mathbb{R}$  of the following equation correspond to transformed inverse supply elasticities in log-linear equilibrium:*

$$\omega_{1,t} = T_t(\omega_{1,t}) \equiv \frac{\frac{(\eta - \frac{1}{\gamma})\kappa_t^A}{1 - \omega_{1,t}(\eta - \frac{1}{\gamma})(1 - \kappa_t^A)}(\sigma_{t|s}^A)^2 + \frac{\frac{1}{\gamma} + (\eta - \frac{1}{\gamma})\kappa_t^M}{1 - \omega_{1,t}(\eta - \frac{1}{\gamma})(1 - \kappa_t^M)}(\sigma_{t|s}^M)^2}{\sigma_{\vartheta,t}^2 + \left(\frac{(\eta - \frac{1}{\gamma})\kappa_t^A}{1 - \omega_{1,t}(\eta - \frac{1}{\gamma})(1 - \kappa_t^A)}\right)^2 (\sigma_{t|s}^A)^2 + \left(\frac{\frac{1}{\gamma} + (\eta - \frac{1}{\gamma})\kappa_t^M}{1 - \omega_{1,t}(\eta - \frac{1}{\gamma})(1 - \kappa_t^M)}\right)^2 (\sigma_{t|s}^M)^2} \quad (28)$$

where  $(\sigma_{t|s}^A)^2 = (1 - \kappa_t^A)(\sigma_t^A)^2$  and  $(\sigma_{t|s}^M)^2 = (1 - \kappa_t^M)(\sigma_t^M)^2$ .

*Proof.* See Appendix A.5. □

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.

This result makes clear that our model has different implications than those that study monetary non-neutrality with endogenous information acquisition. In our model, firms learn via the endogenous signals produced by the market mechanism. This differs from the premise of rational inattention models, wherein firms have unrestricted access to information but can only process it at a cost. This difference in information structures drives significant differences in results. In the rational inattention model of Maćkowiak and Wiederholt (2009), for example, *any* increase in idiosyncratic uncertainty lowers the responsiveness of prices to aggregate shocks. In our model, idiosyncratic productivity and cost uncertainty are not directly relevant for the slope of aggregate supply, whereas idiosyncratic demand uncertainty is directly relevant. Moreover, in our framework, the slope of aggregate supply depends on whether aggregate uncertainty is driven by real or nominal shocks. Thus, the information that arises endogenously through the market mechanism (firms' observation of their demand) is fundamentally different from the information that firms are restricted to obtain under existing models of costly information acquisition with price-setting firms.

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

*Proof.* See Appendix A.6 □

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 slope of aggregate supply under the parametric condition  $\eta\gamma = 1$ . Recall from our discussion in Section 4.2 that  $\eta$  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 28 considerably.

**Corollary 5** (Idiosyncratic *vs.* Aggregate Demand Uncertainty). *When  $\eta\gamma = 1$ , the unique inverse elasticity of aggregate supply is*

$$\epsilon_t^S = \gamma \frac{\kappa_t^M}{1 - \kappa_t^M} \left( 1 + \frac{1}{\gamma^2 \rho_t^2 \kappa_t^M} \right) \quad (29)$$

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

*Proof.* See Appendix A.7 □

First, observe that uncertainty about *aggregate* productivity does not enter the slope of aggregate supply when  $\eta\gamma = 1$ . 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 7) is not useful for conducting inference about productivity and so  $\kappa_t^A$  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 nominal marginal costs when the money supply is uncertain.

Second, as  $\rho_t \rightarrow \infty$ , the inverse elasticity of aggregate supply approaches  $\frac{1}{\gamma} \frac{\kappa_t^M}{1 - \kappa_t^M}$ . This is the AS curve slope under price-setting (Equation 26). 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 approaches infinity. Consequently, aggregate supply is perfectly inelastic and money has no real effects on output. This is the AS curve that arises from quantity-setting ( $\omega_{1,t} = \frac{1}{\eta}$ ). 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 slope of aggregate supply. Moreover, this feature only becomes relevant once firms are allowed to optimally choose their supply functions. As Corollary 4 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.

Thus, supply function choice implies, as a positive matter, a thorny trade-off for monetary policymakers. If the central bank wishes to maintain discretion to use monetary policy to

affect the real economy, this will increase uncertainty about the money supply. In turn, this will steepen the equilibrium aggregate supply curve and make money less effective in guiding real economic outcomes. Therefore, as in the classic “islands” model of Lucas (1972), discretionary monetary policy may be, at least partially, self-defeating.

**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.<sup>14</sup> These results hold for any values of  $\eta > 1$  and  $\gamma > 0$ , in contrast to the analysis with balanced strategic interactions above.

**Corollary 6** (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.8 □

The intuition for this result mirrors that of Corollary 5. As idiosyncratic uncertainty about demand becomes dominant, firms find it optimal to set prices to keep their markup over real marginal costs constant. As prior 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 markup 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 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 after a positive demand shock.

**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

---

<sup>14</sup>Formally, we take these limits for  $\sigma_{t|s}^x$  and  $x \in \{M, A\}$  by scaling  $\sigma_x^{s,t}$  and  $\sigma_t^x$  by a common factor.

matter. To make this claim formal, fix a scalar  $\lambda \geq 0$  and scale all uncertainty in the economy according to:

$$(\sigma_{\vartheta,t}, \sigma_{z,t}, \sigma_{\phi,t}, \sigma_{A,t}, \sigma_{A,s,t}, \sigma_{M,t}, \sigma_{M,s,t}) \mapsto (\lambda\sigma_{\vartheta,t}, \lambda\sigma_{z,t}, \lambda\sigma_{\phi,t}, \lambda\sigma_{A,t}, \lambda\sigma_{A,s,t}, \lambda\sigma_{M,t}, \lambda\sigma_{M,s,t}) \quad (30)$$

In this sense,  $\lambda$  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  $\lambda$ . We observe the following:

**Proposition 3** (Invariance to Uncertainty and Discontinuity in the Limit). *For  $\lambda > 0$ ,  $\mathcal{E}_t^S(\lambda)$  is constant 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:*

$$\lim_{\lambda \rightarrow 0} \mathcal{E}_t^S(\lambda) \neq \mathcal{E}_t^S(0) \quad (31)$$

*Proof.* See Appendix A.9 □

There are two important implications of this result. First, the total level of uncertainty does not matter for the slope of the aggregate supply curve. This constitutes a significant difference between our model and models with menu costs. Concretely, in menu cost models, any increase in uncertainty regarding the optimal reset price raises firms' private benefits of price flexibility without affecting the private costs, which are assumed to be fixed. Thus, increases in uncertainty lead to more variable prices at the micro level and more monetary neutrality at the macro level. By contrast, in our model, the level of uncertainty does not matter—only the relative magnitudes of uncertainty matter. One important implication of this difference is that idiosyncratic productivity uncertainty has *no effect* on the slope of aggregate supply in our model, while it would steepen aggregate supply in menu cost models that share our primitive economic assumptions on preferences and technology (such as Golosov and Lucas, 2007; Alvarez et al., 2016). Another important implication is that while idiosyncratic demand variation *flattens* aggregate supply in our model, it would have no effect in these menu cost models.

Second, the slope of the aggregate supply curve 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. This again represents a substantial difference to menu cost models, in which a small level of uncertainty has small effects on aggregate behavior and not the discontinuity that our model generates.<sup>15</sup> Importantly, this means that even

---

<sup>15</sup>Similarly, models with information acquisition and nominal rigidities (Afrouzi et al., 2024) are also different from our model in that they do not feature this discontinuity in the limit.

in environments with low levels of uncertainty, the economic mechanisms that underlie our analysis are unchanged.

## 4.5 A General Framework for Macroeconomic Analysis

Our preceding analysis tractably illustrated the effect of supply functions in a fully non-linear fashion. To do so, we made a number of simplifying assumptions on utility and the nature of firms' production functions. However, we emphasize that our analysis can readily be extended to general linearized macroeconomic environments of the kind that are commonly studied in both state-of-the-art theoretical and quantitative work (see *e.g.*, Wolf, 2023; McKay and Wolf, 2023). We now describe a general class of models in which the study of supply functions is tractable. We note that this is not meant as being exhaustive of the set of models in which supply functions are tractable or reasonable macroeconomic models.

Consider a model which generates a demand function for products given by  $q_{i,t} = d(p_{it}, z_{it}^D)$ , where the random variable  $z_{it}^D$  can depend on other, potentially endogenous variables of the model as well as exogenous stochastic processes. Assume further that the value  $V$  a firm derives from setting a price  $p_{it}$  and selling a quantity  $q_{it}$  is given by  $V(p_{it}, q_{it}, \mathbf{z}_{it}^V)$ , where  $\mathbf{z}_{it}^V$  is an  $n_V$ -sized vector of (potentially endogenous) variables that affect firm's profits at time  $t$ . Given the discussion in Section 2, the firm's problem is to choose a price that is contingent on demand  $q_{it}$ , which is further equivalent to choosing a price contingent on the demand state  $z_{it}^D$ . That is, for each state realization  $z_{it}^D$ , the firm chooses a price  $p_{it}$  that maximizes the conditional expected value

$$p_{it}(z_{it}^D) = \arg \max \mathbb{E}_{it}[V(p_{it}, d(p_{it}, z_{it}^D), \mathbf{z}_{it}^V) | z_{it}^D] \quad (32)$$

where we have substituted firms' demand into the value function. We now consider a log-linear approximation around a deterministic steady state of this model, using hats to denote log deviations. The approximated policy function must evidently satisfy

$$\hat{p}_{it} = \tilde{\omega}'_{1,it} \mathbb{E}_{it} [\hat{\mathbf{z}}_{it}^V | \hat{z}_{it}^D] \quad (33)$$

for some  $n_V$ -sized vector  $\tilde{\omega}_{1,it}$ . Under the assumption that the shocks  $\hat{\mathbf{z}}_{it}^V$  and  $\hat{z}_{it}^D$  are normally distributed, optimal prices can further be written as

$$\hat{p}_{it} = \omega_{1,it} \hat{z}_{it}^D \quad (34)$$

for some scalar  $\omega_{1,it}$ . The coefficients  $\omega_{1,it}$ , the slopes of firms' supply functions in their

demand states, then determine the motion for the economy’s log-linearized ideal price index when averaged across firms, *i.e.*,  $\hat{P}_t = \int_0^1 \hat{p}_{it} di$ . This concludes the “firm block” of the model.

Following McKay and Wolf (2023) or Wolf (2023), we assume that the aggregate dynamics of our economy can be summarized as

$$\mathcal{H}_x \mathbf{x} + \mathcal{H}_\varepsilon \varepsilon = \mathbf{0} \tag{35}$$

where  $x_t$  is an  $n_x$ -dimensional vector of endogenous variables (such as the ideal price index  $\hat{P}_t$ ),  $\varepsilon_t$  is an  $n_\varepsilon$ -dimensional vector of Gaussian structural shocks, and  $\mathcal{H}_z$  and  $\mathcal{H}_\varepsilon$  are conforming matrices.<sup>16</sup> Equation 35, for example, contains the relevant first-order conditions and market-clearing conditions that determine the dynamics of an economy. Of course, the matrices  $\mathcal{H}_x$  and  $\mathcal{H}_\varepsilon$  are dependent on firms’ supply functions through the scalars  $\omega_{1,it}$ .

Given  $\omega_{1,it}$ , we can solve for the equilibrium dynamics of the system summarized by 35. Our additional “rational expectations” restriction then imposes that the value of  $\omega_{1,it}$  is consistent with the equilibrium law of motion for prices given by Equation 33.<sup>17</sup> As argued by McKay and Wolf, many of the parametric structural models commonly used for counterfactual analysis fit into the general framework of Equation 35. Our supply function approach simply asserts that the coefficients  $\mathcal{H}_x$  and  $\mathcal{H}_\varepsilon$  are consistent with the information underlying firm decision-making. We thus argue that supply functions can be embedded and studied within a general class of commonly used macroeconomic models. Indeed, in ongoing work, Nikolakoudis (2024) incorporates supply function choice in a production network economy using the method described above.

## 5 Quantitative Analysis: 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 employ a sufficient-statistics approach, which allows us to bypass any issues of equilibrium selection. We find that our model generates a historical time series for the slope of aggregate supply in the US that is consistent with external empirical evidence. Concretely, 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 also predicts a

---

<sup>16</sup>Following McKay and Wolf (2023), we use boldface notation to stack the time paths for all variables (*e.g.*  $\mathbf{x} = (\mathbf{x}'_1, \dots, \mathbf{x}'_{n_x})'$ ). The matrices  $\mathcal{H}_x$  and  $\mathcal{H}_\varepsilon$  are conformable matrices that map bounded sequences to the space of bounded sequences.

<sup>17</sup>The Gaussian nature of shocks implies that the initial assumption of normality is verified.

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 inflationary shocks.

## 5.1 Estimating 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  $\epsilon_t^S$  (Theorem 2). To obtain estimates of these quantities, we need to know two sets of objects: the structural parameters  $(\eta, \gamma)$  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 consider an alternative calibration with a secular downward trend in  $\eta$  (*i.e.*, a secular upward trend in market power) over our studied time period (1960-2022). For our analysis,  $\gamma$  affects only how real wages respond to output in equilibrium. Thus, the cyclicity of US real wages is the relevant moment in our model given Proposition 1 to which  $\gamma$  should be calibrated. Because of this, we set  $\gamma = 0.095$ , based on the calibration in Flynn and Sastry (2022) for the cyclicity of US real wages. We will later perform a full sensitivity analysis of these choices (see Figure 12).

**Time-Varying Uncertainty.** We next need to estimate firms’ time-varying uncertainty about aggregate prices  $P_t$ , demand  $\Psi_{it}$ , 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 relationship 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  $\kappa^M$ , 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 1960 Q1 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 1, the model-implied demand shock is  $\Psi_{it} = Y_t \vartheta_{it}$ , where  $Y_t$  is aggregate real GDP (*i.e.*, “aggregate demand”) and  $\vartheta_{it}$  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 1 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  $\phi_{it}$  and productivity shock  $\zeta_{it}$  are idiosyncratic and we only need to measure the covariances of  $\mathcal{M}_{it}$  with  $\Psi_{it}$  and  $P_t$ , we do not need to measure  $\phi_{it}$  or  $\zeta_{it}$ . 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

$$Z_t = AZ_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \Sigma_t), \quad \Sigma_t = D_t^{\frac{1}{2}} C D_t^{\frac{1}{2}} \quad (36)$$

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:

$$\sigma_{i,t}^2 = s_i + \alpha_i \varepsilon_{i,t-1}^2 + \beta_i \sigma_{i,t-1}^2 \quad (37)$$

with unknown constant  $s_i$  and coefficients  $(\alpha_i, \beta_i)$ . Formally, this is a GARCH(1,1) model

with constant conditional correlations (Bollerslev, 1990). We estimate all of the parameters via joint maximum likelihood.

Our goal is to capture the broad trends in *relative* uncertainty about macroeconomic variables. There are, of course, many possible statistical models for macroeconomic uncertainty. We use the GARCH approach as it is a standard approach in the literature.<sup>18</sup> In the GARCH model we estimate, uncertainty is high when there is a large prediction error in one of the equations (if  $\alpha_i > 0$ ) and if uncertainty was high previously (if  $\beta_i > 0$ ). The restriction to constant correlations restricts all covariances 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. In return, this significantly reduces the number of estimated parameters and improves the convergence of the maximum likelihood algorithm.

Using the GARCH estimates, we derive maximum-likelihood point estimates of every element of  $\Sigma_t$ , 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:

$$\hat{\sigma}_{\Psi,t}^2 = \hat{\sigma}_{Y,t}^2 + R^2 \hat{\sigma}_{\mathcal{M},t}^2, \quad \hat{\sigma}_{\mathcal{M},\Psi,t} = \hat{\sigma}_{\mathcal{M},Y,t}, \quad \hat{\sigma}_{\Psi,P,t} = \hat{\sigma}_{Y,P,t} \quad (38)$$

We plot our estimates of these objects in Figure 7 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 elasticity of demand  $\eta$ . We will return to this point when presenting our results.

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

$$\hat{\alpha}_{1,t} = \frac{\eta \hat{\sigma}_{P,t}^2 + \hat{\sigma}_{\mathcal{M},\Psi,t} + \hat{\sigma}_{P,\Psi,t} + \eta \hat{\sigma}_{\mathcal{M},P,t}}{\hat{\sigma}_{\Psi,t}^2 - \eta \hat{\sigma}_{\mathcal{M},\Psi,t} + \eta \hat{\sigma}_{P,\Psi,t} - \eta^2 \hat{\sigma}_{\mathcal{M},P,t}} \quad \text{and} \quad \hat{\omega}_{1,t} = \frac{\hat{\alpha}_{1,t}}{1 + \eta \hat{\alpha}_{1,t}} \quad (39)$$

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  $t - 1$ .

This sufficient statistics approach has four potentially appealing features compared to full

---

<sup>18</sup>Another option would have been to employ a latent state model that allows volatility to be directly affected by shocks. As Jurado et al. (2015) find very similar uncertainty series in latent-state and GARCH models in the US data both “qualitatively and quantitatively,” we employ a GARCH approach for simplicity.

**Table 1:** Model Parameters and Estimation Approach

Parameter	Interpretation	Method	Value
$\eta$	Price elas. of demand	Match <a href="#">Broda and Weinstein (2006)</a>	9
$\gamma$	Income effects	Match <a href="#">Flynn and Sastry (2022)</a>	0.095
$\kappa^M$	Prec. of monetary info.	Match average slope of aggregate supply	0.40
$\sigma_{P,t}^2$	Price uncertainty	GARCH model	<a href="#">Fig. 7</a>
$\sigma_{\Psi,t}^2$	Demand uncertainty	GARCH + match <a href="#">Bloom et al. (2018)</a>	<a href="#">Fig. 7</a>
$\sigma_{\mathcal{M},\Psi,t}$	Cost-demand covariance	GARCH model	<a href="#">Fig. 7</a>
$\sigma_{P,\Psi,t}$	Pice-demand covariance	GARCH model	<a href="#">Fig. 7</a>
$\sigma_{\mathcal{M},P,t}$	Cost-price covariance	GARCH model	<a href="#">Fig. 7</a>

*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 7](#). The text of [Section 5.1](#) describes our methods in full detail.

structural estimation of our model. First, we do not need to estimate stochastic processes for the underlying structural shocks. Second, we bypass any issues of equilibrium selection. Third, these estimates of  $\alpha_1$  and  $\omega$  are valid in a larger class of models that use our “supply block” but close the model with different “demand blocks,” as described in [Section 4.5](#). Finally, by disciplining beliefs with sufficient statistics, we overcome the [Bergemann et al. \(2021\)](#) critique that supply functions have few robust predictions across *all* potential beliefs that firms might have. Because of this, our approach could be useful in other areas of economics that use supply functions, such as industrial organization and finance (see [Rostek and Yoon, 2023](#), for a review).

We finally compute our estimate of the inverse elasticity of aggregate supply:


$$\hat{\epsilon}_t^S = \gamma \frac{\kappa^M + \frac{\hat{\omega}_{1,t}}{\gamma}(1 - \kappa^M)}{(1 - \hat{\omega}_{1,t}\eta)(1 - \kappa^M)} \quad (40)$$

Given our earlier estimates, this is pinned down given a single unknown parameter, the precision of private information about the money supply,  $\kappa^M$ . As mentioned earlier, we set  $\kappa^M$  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 estimates of the slope of Aggregate Supply from 1960 to 2022. Quantitatively, the model predicts that 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

**Figure 4:** The Slope of Aggregate Supply Over Time



*Note:* This Figure plots estimates of the inverse elasticity (“slope”) of aggregate supply as measured by Equations 39 and 40. The units can be interpreted as the percentage change in the price level associated with a demand shock that changes real GDP by 1%. 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).

crisis and then steepened again after the Covid-19 pandemic.

Our estimates are consistent with external estimates for medium-frequency changes in the slope of aggregate supply in the United States. 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.<sup>19</sup> In Table 2, we quantitatively compare our model’s estimates with those of the aforementioned references. Our model 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. These changes arise in our model as a consequence of changing

<sup>19</sup>In the presence of confounding supply shocks, this understates the slope of aggregate supply.

**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	0.175	0.254
Post-Period	Period	1991-2018	1985-2007
	Data	0.090	0.136
	Model: Uncertainty Only	0.119	0.104
	Model: + Mkt. Power Trend	0.103	0.092
Flattening	Data	55%	63%
	Model: Uncertainty Only	28%	53%
	Model: + Mkt. Power Trend	41%	64%


*Note:* This table compares the 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  $100 \cdot (1 - \text{SlopePost}/\text{SlopePre})$ .

uncertainties, holding fixed all other structural parameters.

Our finding that aggregate supply is relatively flat during the Great Recession is consistent with the “missing disinflation” of the crash and “missing inflation” of the recovery, phenomena which standard macroeconomic models have struggled to explain ([Coibion and Gorodnichenko, 2015](#); [Bobeica and Jarociński, 2019](#)). An additional implication of our analysis is that a spike in microeconomic demand uncertainty, as documented for example by [Bloom et al. \(2018\)](#), may have contributed toward flattening the aggregate supply curve and hence lessening the response of prices to aggregate demand shocks.

Finally, our finding that the slope of aggregate supply spiked following the Covid Crisis is consistent with the conclusions of an emerging empirical literature on that topic. Using an empirical strategy that isolates local demand shocks in MSA-level data, [Cerrato and Gitti \(2022\)](#) estimate that supply steepened by a factor of 3.4 between the “pre-Covid” period of January 1990 to February 2020 and the “post-Covid” period of March 2021 to the present. The comparable number generated by our model is 2.5, or about 3/4 of this estimate. Our estimating steepening is consistent with both large effects of aggregate demand shocks on inflation and with a relatively “soft landing” for monetary policy—that is, disinflation at relatively low output cost. Note that these are *conditional moment* predictions about the response to inflation to demand shocks, which map to the empirical estimates of [Cerrato and Gitti \(2022\)](#) and the discussion of a monetary “soft landing.” This does not directly

**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 39 and 40, holding fixed one component of uncertainty at a time. The grey solid line corresponds to the baseline estimate from Figure 4.

attribute *realized* inflation to demand shocks versus supply shocks—our sufficient statistics approach does not require taking a stance on this issue. Nonetheless, one implication of a steep aggregate supply curve is that smaller demand shocks would be required to rationalize the same changes in inflation post-Covid compared to other episodes.

To probe the robustness of these findings, Figure 8 in the Appendix 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. Finally, in the Appendix we also plot our estimates of the slope of firm-level supply (Figure 9), which as mentioned above would be valid even under different closures of the model’s “demand block” as described in Section 4.5. As the estimated slope of firm-level supply has the same basic patterns as the estimated slope of aggregate supply, our main quantitative lessons would go through under any closure that preserves the intuitive link between steeper microeconomic supply and steeper macroeconomic supply.

**Mechanisms: Which Uncertainties Matter More?** The time variation in our estimate of  $\epsilon_t^S$  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 7), it is not essential for our low-frequency predictions. By contrast, 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 markups, 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 consider alternative calibrations of the slope of aggregate supply in which we allow a secular downward trend in the elasticity of demand. We consider a “small change” in which  $\eta$  linearly declines from 12 to 6 and a more exaggerated “large change” in which  $\eta$  linearly declines from 15 to 3. These exercises are *not* counterfactuals, which would require fully estimating the model and accounting for the effects of market power on macroeconomic uncertainty. Instead, they are alternative calibrations that would be more appropriate than our baseline if the elasticity of demand has truly fallen over time, which is a prominent hypothesis.

Introducing a decline in market power increases the slope of aggregate supply in the 1970s and decreases the slope in modern times (Figure 6). Calibrating to the “small change” in market power allows us to fit three-quarters of the flattening measured by Hazell et al. (2022) and all of the flattening measured by Ball and Mazumder (2011) (Table 2). The more extreme scenario for market power allows our model to explain a greater flattening, but potentially incorrectly predicts a relatively flat aggregate supply curve in the 2020s.

**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 10. 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.

**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  $\eta$  constant at 9 and corresponds to our baseline estimates. The blue, dashed line (“small change”) assumes a linear decline in  $\eta$  from 12 to 6 over the time sample. The red, dotted line assumes a linear decline in  $\eta$  from 15 to 3 over the time sample.

**Robustness and Sensitivity Analysis.** In Appendix Figure 12 we report the robustness of these findings to different calibrations of  $\eta \in [6, 12]$ ,  $R \in [5, 10]$ , and  $\gamma \in [0.05, 1.00]$ . Specifically, we re-calibrate 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  $\eta$ , which exacerbate the effect of changing inflation uncertainty; smaller  $R$ , which allows inflation uncertainty to play a larger role in the calculation; and larger  $\gamma$ , which makes real wages more cyclical, especially in the 1970s. We note in particular that the effects of higher  $\gamma$  (more cyclical wages) are very muted: we conclude that equilibrium effects through wages are relatively unimportant for our main findings about the slope of aggregate supply.

In Appendix Figure 11, we report results from a pseudo-out-of-sample variant of our main calculation, in which we use data up to quarter  $t - 1$  to forecast the conditional variance of variables at quarter  $t$ . We find very similar broad patterns, although estimates in the early part of the sample are, unsurprisingly, noisier.

## 6 Conclusion

In this paper, we enrich firms' supply decisions by allowing them to choose supply functions that describe 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 a higher elasticity of demand and increased uncertainty about the price level relative to demand endogenously steepen aggregate supply. When mapped to the data, our model generates variation in the slope of aggregate supply that is consistent with empirical evidence in the US.

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 and can be done in a large class of linearized macroeconomic models of the kind studied by, for example, [Wolf \(2023\)](#) and [McKay and Wolf \(2023\)](#). 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., Flynn, J. P., and Yang, C. (2024). What can measured beliefs tell us about monetary non-neutrality? Working paper, Columbia University.
- Alvarez, F., Le Bihan, H., and Lippi, F. (2016). The real effects of monetary shocks in sticky price models: a sufficient statistic approach. *American Economic Review*, 106(10):2817–2851.
- 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.
- Bergemann, D., Heumann, T., and Morris, S. (2021). Information, market power, and price volatility. *The RAND Journal of Economics*, 52(1):125–150.
- Blanchard, O. J. and Kiyotaki, N. (1987). Monopolistic competition and the effects of aggregate demand. *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.
- Bobeica, E. and Jarociński, M. (2019). Missing disinflation and missing inflation: A VAR perspective. *International Journal of Central Banking*, 15(1):199–232.
- 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.
- Borenstein, S. and Shepard, A. (1996). Dynamic pricing in retail gasoline markets. *The RAND Journal of Economics*, 27(3):429–451.
- Bornstein, G. and Peter, A. (2022). Nonlinear pricing and misallocation. Working paper, NYU.
- 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.
- Cerrato, A. and Gitti, G. (2022). Inflation since Covid: Demand or supply. Working Paper 4193594, SSRN.
- Coibion, O. and Gorodnichenko, Y. (2015). Is the Phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232.
- Dávila, E. and Schaab, A. (2023). Optimal monetary policy with heterogeneous agents: Discretion, commitment, and timeless policy. Working Paper 30961, National Bureau of Economic Research.
- 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. Working paper, MIT.
- Den Boer, A. V. (2015). Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in Operations Research and Management Science*, 20(1):1–18.
- Drenik, A. and Perez, D. J. (2020). Price setting under uncertainty about inflation. *Journal of Monetary Economics*, 116:23–38.
- Dutta, G. and Mitra, K. (2017). A literature review on dynamic pricing of electricity. *Journal of the Operational Research Society*, 68(10):1131–1145.
- Eden, B. (1994). The adjustment of prices to monetary shocks when trade is uncertain and sequential. *Journal of Political Economy*, 102(3):493–509.
- Edmond, C., Midrigan, V., and Xu, D. Y. (2023). How costly are markups? *Journal of Political Economy*, 131(7):000–000.
- Farias, V. F. and Van Roy, B. (2010). Dynamic pricing with a prior on market response. *Operations Research*, 58(1):16–29.
- Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Accessed from <https://www.johnfernald.net/>, Federal Reserve Bank of San Francisco.
- Flynn, J. P., Nikolakoudis, G., and Sastry, K. A. (2024). Prices vs. quantities: A macroeconomic analysis. Manuscript.
- 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 efficiency: 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.
- Golosov, M. and Lucas, R. E. (2007). Menu costs and Phillips curves. *Journal of Political Economy*, 115(2):171–199.
- Grossman, S. J. (1981a). An introduction to the theory of rational expectations under asymmetric information. *The Review of Economic Studies*, 48(4):541–559.
- Grossman, S. J. (1981b). Nash equilibrium and the industrial organization of markets with large fixed costs. *Econometrica*, pages 1149–1172.
- Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408.
- Hart, O. (1985). Imperfect competition in general equilibrium: an overview of recent work. In *Frontiers of Economics*. Basil Blackwell, Oxford.
- Hazell, J., Herreño, 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 Venkateswaran, V. (2009). Setting the right prices for the wrong reasons. *Journal of Monetary Economics*, 56:S57–S77.
- Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.
- Joskow, P. L. and Wolfram, C. D. (2012). Dynamic pricing of electricity. *American Economic Review*, 102(3):381–385.
- Jurado, K., Ludvigson, S. C., and Ng, S. (2015). Measuring uncertainty. *American Economic Review*, 105(3):1177–1216.
- Kimball, M. S. (1995). The quantitative analytics of the basic Neomonetarist model. *Journal of Money, Credit and Banking*, 27(4):1241–1277.
- 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.
- 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. *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.
- Lucas, R. E. and Woodford, M. (1993). Real effects of monetary shocks in an economy with sequential purchases. Working Paper w4250, National Bureau of Economic Research.

- 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.
- McKay, A. and Wolf, C. K. (2023). What can time-series regressions tell us about policy counterfactuals? *Econometrica*, 91(5):1695–1725.
- Nikolakoudis, G. (2024). Incomplete information and supply chain risk. In Progress.
- 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.
- Phillips, R. (2012). Why Are Prices Set the Way They Are? In *The Oxford Handbook of Pricing Management*. Oxford University Press.
- Reis, R. (2006). Inattentive producers. *The Review of Economic Studies*, 73(3):793–821.
- Rostek, M. J. and Yoon, J. H. (2023). Equilibrium theory of financial markets: Recent developments. Working Paper 3710206, SSRN.
- Taylor, J. B. (1980). Aggregate dynamics and staggered contracts. *Journal of Political Economy*, 88(1):1–23.
- 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.
- Williams, K. R. (2022). The welfare effects of dynamic pricing: Evidence from airline markets. *Econometrica*, 90(2):831–858.
- Wolf, C. K. (2023). The missing intercept: A demand equivalence approach. *American Economic Review*, 113(8):2232–2269.
- 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 Theorem 1

*Proof.* Fix a supply function  $f$ . The realized price of the firm in state  $z$  solves  $f(\hat{p}(z), z\hat{p}(z)^{-\eta}) = 0$ . 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:

$$J(\hat{p}) = \int_{\mathbb{R}_{++}^4} \Lambda \left( \frac{\hat{p}(z)}{P} - \mathcal{M} \right) z\hat{p}(z)^{-\eta} dG(\Lambda, P, \mathcal{M}, z) \quad (41)$$

where  $G$  is the cumulative distribution function representing the firm's beliefs. We therefore study the problem:

$$\sup_{\hat{p}: \mathbb{R}_+ \rightarrow \mathbb{R}_{++}} J(\hat{p}) \quad (42)$$

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.

We first derive Equation 7 using variational methods. Consider a variation  $\tilde{p}(z) = p(z) + \varepsilon h(z)$ . The expected payoff under this variation is:

$$J(\varepsilon; h) = \int_{\mathbb{R}_{++}^4} \Lambda \left( \frac{p(z) + \varepsilon h(z)}{P} - \mathcal{M} \right) z(p(z) + \varepsilon h(z))^{-\eta} dG(\Lambda, P, \mathcal{M}, z) \quad (43)$$

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:

$$0 = \int_{\mathbb{R}_{++}^4} \left[ \Lambda \frac{h(z)}{P} z p(z)^{-\eta} - \eta \Lambda h(z) \left( \frac{p(z)}{P} - \mathcal{M} \right) z p(z)^{-\eta-1} \right] dG(\Lambda, P, \mathcal{M}, z) \quad (44)$$

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

$$0 = \int_{\mathbb{R}_{++}^3} \left[ \Lambda \frac{1}{P} t p(t)^{-\eta} - \eta \Lambda \left( \frac{p(t)}{P} - \mathcal{M} \right) t p(t)^{-\eta-1} \right] g(\Lambda, P, \mathcal{M}, t) d\Lambda dP d\mathcal{M} \quad (45)$$

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

$$\begin{aligned} 0 &= \int_{\mathbb{R}_{++}^3} \left[ \Lambda \frac{1}{P} t p(t)^{-\eta} - \eta \Lambda \left( \frac{p(t)}{P} - \mathcal{M} \right) t p(t)^{-\eta-1} \right] g(\Lambda, P, \mathcal{M}|t) d\Lambda dP d\mathcal{M} \\ &= (1 - \eta) \mathbb{E} \left[ \Lambda \frac{1}{P} | z = t \right] t p(t)^{-\eta} + \eta \mathbb{E} [\Lambda \mathcal{M} | z = t] t p(t)^{-\eta-1} \end{aligned} \quad (46)$$

Thus, we have that an optimal solution necessarily follows:

$$p(t) = \frac{\eta}{\eta - 1} \frac{\mathbb{E}[\Lambda \mathcal{M} | z = t]}{\mathbb{E}[\Lambda P^{-1} | z = t]} \quad (47)$$

as claimed in Equation 7.

We now evaluate the expectations. Using log-normality,

$$\begin{aligned} \mathbb{E}[\Lambda \mathcal{M} | z = t] &= \exp \left\{ \mu_{\Lambda|z}(t) + \mu_{\mathcal{M}|z}(t) + \frac{1}{2} \sigma_{\Lambda|z}^2 + \frac{1}{2} \sigma_{\mathcal{M}|z}^2 + \sigma_{\Lambda, \mathcal{M}|z} \right\} \\ \mathbb{E}[\Lambda P^{-1} | z = t] &= \exp \left\{ \mu_{\Lambda|z}(t) - \mu_{P|z}(t) + \frac{1}{2} \sigma_{\Lambda|z}^2 + \frac{1}{2} \sigma_{P|z}^2 - \sigma_{\Lambda, P|z} \right\} \end{aligned} \quad (48)$$

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

$$\frac{\mathbb{E}[\Lambda \mathcal{M} | z = t]}{\mathbb{E}[\Lambda P^{-1} | z = t]} = \exp \left\{ \mu_{\mathcal{M}|z}(t) + \mu_{P|z}(t) + \frac{1}{2} \sigma_{\mathcal{M}|z}^2 - \frac{1}{2} \sigma_{P|z}^2 + \sigma_{\Lambda, \mathcal{M}|z} + \sigma_{\Lambda, P|z} \right\} \quad (49)$$

Using standard formulas for Gaussian conditional expectations,

$$\begin{aligned} \mu_{\mathcal{M}|z}(t) &= \mu_{\mathcal{M}} + \frac{\sigma_{\mathcal{M},z}}{\sigma_z^2} (\log t - \mu_z) & \mu_{P|z}(t) &= \mu_P + \frac{\sigma_{P,z}}{\sigma_z^2} (\log t - \mu_z) \\ \sigma_{\mathcal{M}|z}^2 &= \sigma_{\mathcal{M}}^2 - \frac{\sigma_{\mathcal{M},z}^2}{\sigma_z^2} & \sigma_{P|z}^2 &= \sigma_P^2 - \frac{\sigma_{P,z}^2}{\sigma_z^2} \\ \sigma_{\Lambda, \mathcal{M}|z} &= \sigma_{\Lambda, \mathcal{M}} - \frac{\sigma_{\Lambda,z} \sigma_{\mathcal{M},z}}{\sigma_z^2} & \sigma_{\Lambda, P|z} &= \sigma_{\Lambda, P} - \frac{\sigma_{\Lambda,z} \sigma_{P,z}}{\sigma_z^2} \end{aligned} \quad (50)$$

where:

$$\begin{aligned} \sigma_z^2 &= \sigma_{\Psi}^2 + \eta^2 \sigma_P^2 + 2\eta \sigma_{\Psi, P} & \sigma_{P,z} &= \sigma_{P, \Psi} + \eta \sigma_P^2 \\ \sigma_{\mathcal{M},z} &= \sigma_{\mathcal{M}, \Psi} + \eta \sigma_{\mathcal{M}, P} & \sigma_{\Lambda,z} &= \sigma_{\Lambda, \Psi} + \eta \sigma_{\Lambda, P} \end{aligned} \quad (51)$$

We now combine these expressions with Equation 47 to derive the optimal supply function. We first observe that

$$\log p = \omega_0 + \omega_1 \log t \quad (52)$$

where:

$$\omega_0 = \log \frac{\eta}{\eta - 1} + \mu_{\mathcal{M}} + \mu_P - \omega_1 \mu_z + \frac{1}{2} \sigma_{\mathcal{M}|z}^2 - \frac{1}{2} \sigma_{P|z}^2 + \sigma_{\Lambda, \mathcal{M}|z} + \sigma_{\Lambda, P|z} \quad (53)$$

$$\omega_1 = \frac{\sigma_{\mathcal{M}, z} + \sigma_{P, z}}{\sigma_z^2} = \frac{\sigma_{\mathcal{M}, \Psi} + \eta \sigma_{\mathcal{M}, P} + \sigma_{P, \Psi} + \eta \sigma_P^2}{\sigma_{\Psi}^2 + \eta^2 \sigma_P^2 + 2\eta \sigma_{\Psi, P}} \quad (54)$$

Next, using the demand curve, we observe that  $z = qp^n$ . Therefore,  $\log t = \log q - \eta \log p$ . Substituting this into Equation 52, and re-arranging, we obtain

$$\log p = \alpha_0 + \alpha_1 \log q \quad (55)$$

where:

$$\alpha_0 = \frac{\omega_0}{1 - \eta \omega_1}, \quad \alpha_1 = \frac{\omega_1}{1 - \eta \omega_1} \quad (56)$$

We finally derive the claimed expression for  $\alpha_1$ ,

$$\begin{aligned} \alpha_1 &= \frac{\frac{\sigma_{\mathcal{M}, \Psi} + \eta \sigma_{\mathcal{M}, P} + \sigma_{P, \Psi} + \eta \sigma_P^2}{\sigma_{\Psi}^2 + \eta^2 \sigma_P^2 + 2\eta \sigma_{\Psi, P}}}{1 - \eta \frac{\sigma_{\mathcal{M}, \Psi} + \eta \sigma_{\mathcal{M}, P} + \sigma_{P, \Psi} + \eta \sigma_P^2}{\sigma_{\Psi}^2 + \eta^2 \sigma_P^2 + 2\eta \sigma_{\Psi, P}}} \\ &= \frac{\sigma_{\mathcal{M}, \Psi} + \eta \sigma_{\mathcal{M}, P} + \sigma_{P, \Psi} + \eta \sigma_P^2}{\sigma_{\Psi}^2 + \eta \sigma_{\Psi, P} - \eta \sigma_{\mathcal{M}, \Psi} - \eta^2 \sigma_{\mathcal{M}, P}} \end{aligned} \quad (57)$$

Completing the proof. □

## A.2 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 6 is decreasing in  $\eta$ . Moreover, if  $\sigma_{\mathcal{M}, P} \geq 0$ , the numerator is increasing in  $\eta$ . Hence,  $\alpha_1$  is increasing in  $\eta$  whenever  $\alpha_1 > 0$ . □

## A.3 Proof of Proposition 1

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

$$\frac{p_{it}}{P_t} = \left( \frac{c_{it}}{\vartheta_{it} C_t} \right)^{-\frac{1}{\eta}} \quad (58)$$

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  $\phi_{it}$ . Thus, variety-specific wages are given by

$$w_{it} = \phi_{it} P_t C_t^\gamma \quad (59)$$

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:

$$C_t^{-\gamma} \frac{1}{P_t} = \frac{1}{M_t} + \beta \mathbb{E}_t \left[ C_{t+1}^{-\gamma} \frac{1}{P_{t+1}} \right] \quad (60)$$

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:

$$C_t^{-\gamma} \frac{1}{P_t} = \beta(1 + i_t) \mathbb{E}_t \left[ C_{t+1}^{-\gamma} \frac{1}{P_{t+1}} \right] \quad (61)$$

From Equations 60 and 61, we obtain:

$$\frac{1}{M_t} + \beta \mathbb{E}_t \left[ C_{t+1}^{-\gamma} \frac{1}{P_{t+1}} \right] = \beta(1 + i_t) \mathbb{E}_t \left[ C_{t+1}^{-\gamma} \frac{1}{P_{t+1}} \right] \quad (62)$$

It follows that:

$$\frac{1}{M_t} = \beta i_t \mathbb{E}_t \left[ C_{t+1}^{-\gamma} \frac{1}{P_{t+1}} \right] = \frac{i_t}{1 + i_t} C_t^{-\gamma} \frac{1}{P_t} \quad (63)$$

where the second equality uses Equation 61 once again. This rearranges to:

$$C_t = \left( \frac{i_t}{1 + i_t} \right)^{\frac{1}{\gamma}} \left( \frac{M_t}{P_t} \right)^{\frac{1}{\gamma}} \quad (64)$$

We next derive the interest rate. Substituting equation 64 into Equation 61, we obtain:

$$\frac{1 + i_t}{i_t} \frac{1}{M_t} = \beta(1 + i_t) \mathbb{E}_t \left[ \frac{1 + i_{t+1}}{i_{t+1}} \frac{1}{M_{t+1}} \right] \quad (65)$$

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

$$\frac{1 + i_t}{i_t} = 1 + \beta \mathbb{E}_t \left[ \frac{1 + i_{t+1}}{i_{t+1}} \frac{M_t}{M_{t+1}} \right] = 1 + \beta \mathbb{E}_t \left[ \exp\{-\mu_M - \sigma_{t+1}^M \varepsilon_{t+1}^M\} \frac{1 + i_{t+1}}{i_{t+1}} \right] \quad (66)$$

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:

$$x_t = 1 + \delta_t x_{t+1} \quad (67)$$

where:

$$\delta_t = \beta \exp \left\{ -\mu_M + \frac{1}{2}(\sigma_{t+1}^M)^2 \right\} \quad (68)$$

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 for  $T \geq 2$ :

$$x_t = 1 + \delta_t \left( 1 + \sum_{i=1}^{T-1} \prod_{j=1}^i \delta_{t+j} \right) + \delta_t \left( \prod_{j=1}^T \delta_{t+j} \right) x_{t+T+1} \quad (69)$$

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

$$x_t = 1 + \delta_t \left( 1 + \sum_{i=1}^{\infty} \prod_{j=1}^i \delta_{t+j} \right) + \delta_t \lim_{T \rightarrow \infty} \left( \prod_{j=1}^T \delta_{t+j} \right) x_{t+T+1} \quad (70)$$

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

$$0 \leq \delta_t \lim_{T \rightarrow \infty} \left( \prod_{j=1}^T \delta_{t+j} \right) x_{t+T+1} \leq \lim_{T \rightarrow \infty} \beta^{T+1} x_{t+T+1} \quad (71)$$

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:

$$\lim_{T \rightarrow \infty} \beta^T \frac{C_T^{-\gamma}}{P_T} (M_T + (1 + i_T) B_T) = 0 \quad (72)$$

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 63, 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 71, we have that  $\lim_{T \rightarrow \infty} \left( \prod_{j=1}^T \delta_{t+j} \right) x_{t+T+1} = 0$ . An explicit formula for the interest rate follows:

$$\frac{1 + i_t}{i_t} = 1 + \beta \exp \left\{ -\mu_M + \frac{1}{2}(\sigma_{t+1}^M)^2 \right\} \left( 1 + \sum_{i=1}^{\infty} \prod_{j=1}^i \beta \exp \left\{ -\mu_M + \frac{1}{2}(\sigma_{t+j+1}^M)^2 \right\} \right) \quad (73)$$

The formulae in Equation 19 then follow. In particular,  $\Psi_{it} = \vartheta_{it} C_t$  follows from comparing Equations 2 and 58.  $P_t = \frac{i_t}{1+i_t} C_t^{-\gamma} M_t$  follows from Equation 64.  $\Lambda_t = C_t^{-\gamma}$  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 59.  $\square$

## A.4 Proof of Theorem 2

*Proof.* We begin by characterizing log-linear equilibria, which is achieved by the following Lemma:

**Lemma 1** (Macroeconomic Dynamics with Supply Functions). *If all firms use log-linear supply functions of the form in Equation 20, output in the unique log-linear temporary equilibrium follows:*

$$\log C_t = \tilde{\chi}_{0,t} + \frac{1}{\gamma} \frac{\kappa_t^A}{1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A)} \log A_t + \frac{1}{\gamma} \frac{(1 - \kappa_t^M)(1 - \eta\omega_{1,t})}{1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M)} \log M_t \quad (74)$$

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

$$\log P_t = \chi_{0,t} - \frac{\kappa_t^A}{1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A)} \log A_t + \frac{\kappa_t^M + \frac{\omega_{1,t}}{\gamma} (1 - \kappa_t^M)}{1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M)} \log M_t \quad (75)$$

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.* We suppress dependence on  $t$  for ease of notation. Consider a plan:

$$\log p_i = \log \tilde{\alpha}_{0,i} + \alpha_1 \log q_i \quad (76)$$

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

$$\log p_i = -\frac{1}{\eta} (\log q_i - \log \Psi) + \log P \quad (77)$$

The realized quantity therefore is:

$$\log q_i = \frac{-\eta}{1 + \eta\alpha_1} \log \tilde{\alpha}_{0,i} + \frac{1}{1 + \eta\alpha_1} \log \Psi_i P^\eta \quad (78)$$

and the realized price is:

$$\log p_i = \frac{1}{1 + \eta\alpha_1} \log \tilde{\alpha}_{0,i} + \frac{\alpha_1}{1 + \eta\alpha_1} \log \Psi_i P^\eta \quad (79)$$

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

$$\log p_i = (1 - \eta\omega_1) \log \tilde{\alpha}_{0,i} + \omega_1 \log \Psi_i P^\eta \quad (80)$$

Our goal is to express dynamics only as a function of  $\omega_1$ . We first find the optimal  $\alpha_{0,i}$  in terms of  $\omega_1$ . The firm therefore solves:

$$\max_{\tilde{\alpha}_{0,i}} \mathbb{E}_i \left[ \Lambda \left( \frac{p_i}{P} - \mathcal{M}_i \right) \left( \frac{p_i}{P} \right)^{-\eta} \Psi_i \right] \quad (81)$$

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

$$\max_{\tilde{\alpha}_{0,i}} \mathbb{E} \left[ \Lambda \left( \frac{\tilde{\alpha}_{0,i}^{1-\eta\omega_1}}{P} (\Psi_i P^\eta)^{\omega_1} - \mathcal{M}_i \right) \tilde{\alpha}_{0,i}^{\eta^2\omega_1-\eta} (\Psi_i P^\eta)^{1-\eta\omega_1} \right] \quad (82)$$

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

$$\tilde{\alpha}_{0,i}^{1-\eta\omega_1} = \frac{\eta}{\eta-1} \frac{\mathbb{E}_i[\Lambda \mathcal{M}_i (\Psi_i P^\eta)^{1-\eta\omega_1}]}{\mathbb{E}_i[\frac{\Lambda}{P} (\Psi_i P^\eta)^{1-\eta\omega_1+\omega_1}]} \quad (83)$$

Substituting back into the realized price yields:

$$p_i = \frac{\eta}{\eta-1} \frac{\mathbb{E}_i[\Lambda \mathcal{M}_i (\Psi_i P^\eta)^{1-\eta\omega_1}]}{\mathbb{E}_i[\frac{\Lambda}{P} (\Psi_i P^\eta)^{1-\eta\omega_1+\omega_1}]} (\Psi_i P^\eta)^{\omega_1} \quad (84)$$

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

$$p_i = \frac{\eta}{\eta-1} \frac{\mathbb{E}_i \left[ \phi(z_i A)^{-1} \left( \vartheta_i I^{-\frac{1}{\gamma}} P^{-\frac{1}{\gamma}} M^{\frac{1}{\gamma}} P^\eta \right)^{1-\eta\omega_1} \right]}{\mathbb{E}_i \left[ I^{1-\frac{1}{\gamma}(1+\omega_1-\eta\omega_1)} M^{\frac{1}{\gamma}(1+\omega_1-\eta\omega_1)-1} \vartheta^{1+\omega_1-\eta\omega_1} P^{(\eta-\frac{1}{\gamma})(1+\omega_1-\eta\omega_1)} \right]} \times \left( \vartheta_i I^{-\frac{1}{\gamma}} M^{\frac{1}{\gamma}} P^{\eta-\frac{1}{\gamma}} \right)^{\omega_1} \quad (85)$$

Given the ideal price index formula (Equation 13),  $P$  must satisfy the aggregation:

$$P^{1-\eta} = \mathbb{E} [\vartheta_i p_i^{1-\eta}] \quad (86)$$

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

$$\log P = \chi_0 + \chi_A \log A + \chi_M \log M \quad (87)$$

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

$$\log P = \mathbb{E}[\log p_i] + \frac{1}{2(1-\eta)} \text{Var}((1-\eta) \log p_i) + \text{constants} \quad (88)$$

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

$$\begin{aligned} \log \mathbb{E}_i \left[ \phi_i(z_i A)^{-1} \left( \vartheta I^{-\frac{1}{\gamma}} P^{-\frac{1}{\gamma}} M^{\frac{1}{\gamma}} P^\eta \right)^{1-\eta\omega_1} \right] &= \left[ -\kappa^A + \kappa^A \left( \eta - \frac{1}{\gamma} \right) \chi_A (1 - \eta\omega_1) \right] s_i^A \\ &+ \left[ \chi_M \left( \eta - \frac{1}{\gamma} \right) (1 - \eta\omega_1) \kappa^M + \frac{1}{\gamma} (1 - \eta\omega_1) \kappa^M \right] s_i^M + \text{constants} \end{aligned} \quad (89)$$

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

$$\begin{aligned} \log \mathbb{E}_i \left[ I^{1-\frac{1}{\gamma}(1+\omega_1-\eta\omega_1)} M^{\frac{1}{\gamma}(1+\omega_1-\eta\omega_1)-1} \vartheta^{1+\omega_1-\eta\omega_1} P^{(\eta-\frac{1}{\gamma})(1+\omega_1-\eta\omega_1)} \right] &= \\ \left[ \chi_A \left( \eta - \frac{1}{\gamma} \right) (1 + \omega_1 - \eta\omega_1) \kappa^A \right] s_i^A & \\ + \left[ \left[ \frac{1}{\gamma} (1 + \omega_1 - \eta\omega_1) - 1 \right] (\kappa^M) + \chi_M \left( \eta - \frac{1}{\gamma} \right) (1 + \omega_1 - \eta\omega_1) (\kappa^M) \right] s_i^M & \\ + \text{constants} & \end{aligned} \quad (90)$$

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

$$\log \left( \vartheta_i I^{-\frac{1}{\gamma}} M^{\frac{1}{\gamma}} P^{\eta-\frac{1}{\gamma}} \right)^{\omega_1} = \omega_1 \chi_A \left( \eta - \frac{1}{\gamma} \right) \log A + \omega_1 \left[ \chi_M \left( \eta - \frac{1}{\gamma} \right) + \frac{1}{\gamma} \right] \log M + \text{constants} \quad (91)$$

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 87 into Equation 88, and solving for  $\chi_A$  by collecting coefficients on  $\log A$  yields:

$$\chi_A = - \frac{\kappa^A}{1 - \omega_1 \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa^A)} \quad (92)$$

We may similarly solve for  $\chi_M$ :

$$\chi_M = \frac{\kappa^M + \frac{\omega_1}{\gamma}(1 - \kappa^M)}{1 - \omega_1 \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa^M)} \quad (93)$$

This proves the dynamics for the price level. The dynamics for consumption then follow from Proposition 1.  $\square$

With this characterization in hand, by Equation 74 and market clearing  $C_t = Y_t$ , we have:

$$\log M_t = \frac{1}{\tilde{\chi}_{M,t}} (\log Y_t - \tilde{\chi}_{A,t} \log A_t - \tilde{\chi}_{0,t}) \quad (94)$$

Substituting for  $\log M_t$  in Equation 75 and defining  $\log \bar{P}_t = \chi_{0,t} - \epsilon_t^S \tilde{\chi}_{0,t}$  and  $\delta_t = \chi_{A,t} - \epsilon_t^S \tilde{\chi}_{A,t}$  then yields Equation AS:

$$\log P_t = \log \bar{P}_t + \epsilon_t^S \log Y_t + \delta_t \log A_t \quad (95)$$

Doing a similar substitution for  $\log A_t$  in Equation 74 then yields Equation AD:

$$\log P_t = \log \left( \frac{i_t}{1 + i_t} \right) - \epsilon_t^D \log Y_t + \log M_t \quad (96)$$

Completing the proof.  $\square$

## A.5 Proof of Theorem 3

*Proof.* We suppress dependence on  $t$  for ease of notation. We have  $\chi_M$  and  $\chi_A$  as a function of  $\omega_1$  from Lemma 1. We also know that:

$$\omega_1 = \frac{\sigma_{M_i,z} + \sigma_{P,z}}{\sigma_z^2} \quad (97)$$

from Equation 54. 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:

$$\begin{aligned}
\sigma_{\mathcal{M}_i, z} &= \text{Cov} \left( -(1 + \chi_A) \log A + (1 - \chi_M) \log M, \left( \eta - \frac{1}{\gamma} \right) \chi_A \log A + \left( \frac{1}{\gamma} + \left( \eta - \frac{1}{\gamma} \right) \chi_M \right) \log M \right) \\
&= - \left( \eta - \frac{1}{\gamma} \right) \chi_A (1 + \chi_A) \sigma_A^2 + (1 - \chi_M) \left( \frac{1}{\gamma} + \left( \eta - \frac{1}{\gamma} \right) \chi_M \right) \sigma_M^2 \\
\sigma_{P, z} &= \text{Cov} \left( \chi_A \log A + \chi_M \log M, \left( \eta - \frac{1}{\gamma} \right) \chi_A \log A + \left( \frac{1}{\gamma} + \left( \eta - \frac{1}{\gamma} \right) \chi_M \right) \log M \right) \\
&= \left( \eta - \frac{1}{\gamma} \right) \chi_A^2 \sigma_A^2 + \chi_M \left( \frac{1}{\gamma} + \left( \eta - \frac{1}{\gamma} \right) \chi_M \right) \sigma_M^2 \\
\sigma_z^2 &= \sigma_\vartheta^2 + \left( \eta - \frac{1}{\gamma} \right)^2 \chi_A^2 \sigma_A^2 + \left( \frac{1}{\gamma} + \left( \eta - \frac{1}{\gamma} \right) \chi_M \right)^2 \sigma_M^2
\end{aligned} \tag{98}$$

Thus:

$$\omega_1 = \frac{-(\eta - \frac{1}{\gamma}) \chi_A \sigma_A^2 + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \chi_M) \sigma_M^2}{\sigma_\vartheta^2 + (\eta - \frac{1}{\gamma})^2 \chi_A^2 \sigma_A^2 + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \chi_M)^2 \sigma_M^2} \tag{99}$$

Note that the optimal  $\omega_1$  is common across all firms  $i$ . We may express this in fully reduced form as:

$$\omega_1 = T(\omega_1) = \frac{(\eta - \frac{1}{\gamma}) \frac{\kappa_A}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_A)} \sigma_A^2 + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \frac{\kappa_M + \frac{\omega_1}{\gamma} (1 - \kappa_M)}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_M)}) \sigma_M^2}{\sigma_\vartheta^2 + (\eta - \frac{1}{\gamma})^2 \left( \frac{\kappa_A}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_A)} \right)^2 \sigma_A^2 + (\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \frac{\kappa_M + \frac{\omega_1}{\gamma} (1 - \kappa_M)}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_M)})^2 \sigma_M^2} \tag{100}$$

or

$$\omega_1 = T(\omega_1) = \frac{\frac{(\eta - \frac{1}{\gamma}) \kappa_A}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_A)} \sigma_A^2 + \frac{\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \kappa_M}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_M)} \sigma_M^2}{\sigma_\vartheta^2 + \left( \frac{(\eta - \frac{1}{\gamma}) \kappa_A}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_A)} \right)^2 \sigma_A^2 + \left( \frac{\frac{1}{\gamma} + (\eta - \frac{1}{\gamma}) \kappa_M}{1 - \omega_1 (\eta - \frac{1}{\gamma}) (1 - \kappa_M)} \right)^2 \sigma_M^2} \tag{101}$$

□

## A.6 Proof of Proposition 2

*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 3,  $\omega_{1,t}^*$  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 28 as:

$$\begin{aligned}
& \omega_{1,t} \left[ \sigma_{\vartheta,t}^2 \left( 1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A) \right)^2 \left( 1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right)^2 \right. \\
& \quad + (\sigma_{t|s}^A)^2 \left( \eta - \frac{1}{\gamma} \right) \kappa_t^A \left( 1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right)^2 \\
& \quad \left. + (\sigma_{t|s}^M)^2 \left( \frac{1}{\gamma} + \left( \eta - \frac{1}{\gamma} \right) \kappa_t^M \right) \left( 1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A) \right)^2 \right] \\
& = (\sigma_{t|s}^A)^2 \left( \eta - \frac{1}{\gamma} \right) \kappa_t^A \left( 1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right)^2 \left( 1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A) \right) \\
& \quad + (\sigma_{t|s}^M)^2 \left( \frac{1}{\gamma} + \left( \eta - \frac{1}{\gamma} \right) \kappa_t^M \right) \left( 1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^A) \right)^2 \left( 1 - \omega_{1,t} \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa_t^M) \right)
\end{aligned} \tag{102}$$

This is a quintic polynomial in  $\omega_{1,t}$ , which has at most five real roots. Thus, by Theorem 3, there are at most five log-linear equilibria.  $\square$

## A.7 Proof of Corollary 5

*Proof.* We drop time subscripts for ease of notation. Substituting  $\eta = \frac{1}{\gamma}$  in Equation 28 yields:

$$\omega_1 = \frac{\frac{1}{\gamma}}{\rho^2 + \left( \frac{1}{\gamma} \right)^2} \tag{103}$$

Substituting this into Equation 23 yields:

$$\epsilon_t^S = \gamma \frac{\kappa_t^M}{(1 - \kappa_t^M)} + \frac{1}{\gamma \rho^2 (1 - \kappa_t^M)} \tag{104}$$

$\square$

## A.8 Proof of Corollary 6

We drop time subscripts for ease of notation. The first statement follows directly from Equation 28. Furthermore, using Equation 28, as  $\sigma_{t|s}^M \rightarrow \infty$ ,  $\omega_1$  must solve:

$$\begin{aligned}\omega_1 &= \frac{1 - \omega_1 \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa^M)}{\frac{1}{\gamma} + \left( \eta - \frac{1}{\gamma} \right) \kappa^M} \\ &= \frac{\gamma}{1 + (\eta\gamma - 1) \kappa^M} + \left( 1 - \frac{\eta\gamma}{1 + (\eta\gamma - 1) \kappa^M} \right) \omega_1 \\ &= \frac{1}{\eta}\end{aligned}\tag{105}$$

This proves the second statement. As  $\sigma_{t|s}^A \rightarrow \infty$  and  $\eta\gamma \neq 1$ ,  $\omega_1$  must solve:

$$\begin{aligned}\omega_1 &= \frac{1 - \omega_1 \left( \eta - \frac{1}{\gamma} \right) (1 - \kappa^A)}{\left( \eta - \frac{1}{\gamma} \right) \kappa^A} \\ &= \frac{\gamma}{(\eta\gamma - 1) \kappa^A} + \left( 1 - \frac{1}{\kappa^A} \right) \omega_1 \\ &= \frac{1}{\eta - \frac{1}{\gamma}}\end{aligned}\tag{106}$$

This proves the third statement.

## A.9 Proof of Proposition 3

*Proof.* By Theorem 3, The map describing equilibrium  $\omega_{1,t}$  is invariant to  $\lambda$  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, Monopsony, and Beyond Isoelastic Demand

In this appendix, we generalize the firm’s partial-equilibrium supply schedule problem in two ways. First, we enrich both its technology and inputs by allowing for many inputs, decreasing returns to scale, and monopsony power. Second, we enrich the demand it faces by decoupling the own-price elasticity and the cross-price elasticity and allowing for non-isoelastic demand curves that feature endogenous markups (allowing for Marshall’s Second and Third laws of demand). In both cases, we show that our core insights generalize. In the interests of brevity, we leave embedding these generalizations in general equilibrium to future research, though it is clear to see how one could do this by leveraging our main analysis.<sup>20</sup>

### B.1 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:

$$q = \Theta \prod_{i=1}^I x_i^{a_i} \quad (107)$$

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 an upward-sloping factor price curve such that the price of acquiring any input  $i$  when the firm demands  $x_i$  units is given by  $\tilde{p}_i(x_i) = p_{xi} x_i^{b_i-1}$ , where  $p_{xi} \in \mathbb{R}_{++}$  and  $b_i \geq 1$ . The case of no monopsony, or price-taking in the input market, occurs when  $b_i = 1$ . Thus, the cost of acquiring each type of input is given by:

$$c_i(x_i) = p_{xi} x_i^{b_i} \quad (108)$$

The firm believes that  $(\Psi, P, \Lambda, \Theta, p_x)$  is jointly log-normal.

---

<sup>20</sup>The only complication with endogenous markups would be the endogenous non-log-linearity of the optimal supply curve. This would have to be dealt with via either approximation arguments or numerical methods, or both.

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

$$K(q; \Theta, p_x) = \min_x \sum_{i=1}^I p_{x_i} x_i^{b_i} \quad \text{s.t.} \quad q = \Theta \prod_{i=1}^I x_i^{a_i} \quad (109)$$

This has first-order condition given by:

$$\lambda = \frac{b_i p_{x_i}}{a_i} x_i^{b_i} q^{-1} \quad (110)$$

Which implies that:

$$K(q; \Theta, p_x) = \lambda q \sum_{i=1}^I \frac{a_i}{b_i} \quad (111)$$

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

$$x_j = \left( \frac{\frac{b_i p_{x_i}}{a_i}}{\frac{b_j p_{x_j}}{\alpha_j}} \right)^{\frac{1}{b_j}} x_i^{\frac{b_i}{b_j}} \quad (112)$$

By substituting this into the production function we have that:

$$q = \Theta x_i^{a_i + b_i \sum_{j \neq i} \frac{a_j}{b_j}} \prod_{j \neq i} \left( \frac{\frac{b_i p_{x_i}}{a_i}}{\frac{b_j p_{x_j}}{\alpha_j}} \right)^{\frac{\alpha_j}{b_j}} \quad (113)$$

which implies that:

$$x_i = \left( \frac{q}{\Theta \prod_{j \neq i} \left( \frac{\frac{b_i p_{x_i}}{a_i}}{\frac{b_j p_{x_j}}{\alpha_j}} \right)^{\frac{\alpha_j}{b_j}}} \right)^{\frac{1}{a_i + b_i \sum_{j \neq i} \frac{a_j}{b_j}}} \quad (114)$$

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

$$\lambda = q^{-1 + \frac{1}{\sum_{i=1}^I \frac{a_i}{b_i}}} \frac{b_i p_{x_i}}{a_i} \left( \Theta \prod_{j \neq i} \left( \frac{\frac{b_i p_{x_i}}{a_i}}{\frac{b_j p_{x_j}}{\alpha_j}} \right)^{\frac{\alpha_j}{b_j}} \right)^{\frac{-1}{\sum_{i=1}^I \frac{a_i}{b_i}}} \quad (115)$$

Which then yields the cost function:

$$K(q; \Theta, p_x) = \mathcal{M} P q^{\frac{1}{\delta}} \quad (116)$$

where:

$$\delta = \sum_{i=1}^I \frac{a_i}{b_i} \quad \text{and} \quad \mathcal{M} = P^{-1} \left( \Theta \prod_{i=1}^I \left( \frac{b_i p x_i}{\alpha_i} \right)^{\frac{\alpha_i}{b_i}} \right)^{\frac{1}{\sum_{i=1}^I \frac{\alpha_i}{b_i}}} \sum_{i=1}^I \frac{a_i}{b_i} \quad (117)$$

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

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

$$\mathbb{E} \left[ \Lambda \left( \frac{p}{P} q - \mathcal{M} q^{\frac{1}{\delta}} \right) \right] \quad (118)$$

Thus, the problem with multiple inputs, monopsony, and decreasing returns modifies the firms' original payoff by only introducing the parameter  $\delta$ . 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:

$$J(\hat{p}) = \int_{\mathbb{R}_{++}^4} \Lambda \left( \frac{\hat{p}(z)^{1-\eta}}{P} z - \mathcal{M} z^{\frac{1}{\delta}} \hat{p}(z)^{-\frac{\eta}{\delta}} \right) dG(\Lambda, P, \mathcal{M}, z) \quad (119)$$

And, as before, we study the problem:

$$\sup_{\hat{p}: \mathbb{R}_+ \rightarrow \mathbb{R}_{++}} J(\hat{p}) \quad (120)$$

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

**Proposition 4** (Optimal Supply Schedule With Multiple Inputs, Decreasing Returns, and Monopsony). *Any optimal supply schedule is almost everywhere given by:*

$$f(p, q) = \log p - \frac{\omega_0 - \log \delta}{1 - \eta \omega_1} - \frac{\eta \left( \omega_1 + \frac{1-\delta}{\delta} \right)}{1 - \eta \omega_1} \log q \quad (121)$$

where  $\omega_0$  and  $\omega_1$  are the same as those derived in Theorem 1. Thus, the optimal inverse supply elasticity is given by:

$$\hat{\alpha}_1 = \frac{\eta \sigma_P^2 + \sigma_{\mathcal{M}, \Psi} + \sigma_{P, \Psi} + \eta \sigma_{\mathcal{M}, P}}{\sigma_{\Psi}^2 - \eta \sigma_{\mathcal{M}, \Psi} + \eta \sigma_{P, \Psi} - \eta^2 \sigma_{\mathcal{M}, P}} \left( 1 + \frac{1-\delta}{\delta} \frac{\sigma_{\Psi}^2 + \eta^2 \sigma_P^2 + 2\eta \sigma_{\Psi, P}}{\sigma_{\mathcal{M}, \Psi} + \eta \sigma_{\mathcal{M}, P} + \sigma_{P, \Psi} + \eta \sigma_P^2} \right) \quad (122)$$

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

$$(\eta - 1) \mathbb{E}[\Lambda P^{-1} | z = t] t \hat{p}(t)^{-\eta} = \frac{\eta}{\delta} \mathbb{E}[\Lambda \mathcal{M} | z = t] t^{\frac{1}{\delta}} \hat{p}(z)^{-\frac{\eta}{\delta} - 1} \quad (123)$$

Which yields:

$$\hat{p}(t) = \left( \delta^{-1} \frac{\eta}{\eta - 1} \frac{\mathbb{E}[\Lambda \mathcal{M} | z = t]}{\mathbb{E}[\Lambda P^{-1} | z = t]} \right)^{\frac{1}{1+\eta(\frac{1-\delta}{\delta})}} t^{\frac{\frac{1-\delta}{\delta}}{1+\eta(\frac{1-\delta}{\delta})}} \quad (124)$$

Thus, we have that:

$$\log p = \frac{1}{1 + \eta \left( \frac{1-\delta}{\delta} \right)} (\omega_0 - \log \delta) + \frac{1}{1 + \eta \left( \frac{1-\delta}{\delta} \right)} \left( \omega_1 + \frac{1-\delta}{\delta} \right) \log z \quad (125)$$

where  $\omega_0$  and  $\omega_1$  are as in Theorem 1. Rewriting as a supply schedule, we obtain:

$$\log p = \frac{\frac{1}{1+\eta(\frac{1-\delta}{\delta})} (\omega_0 - \log \delta)}{1 - \frac{\eta}{1+\eta(\frac{1-\delta}{\delta})} (\omega_1 + \frac{1-\delta}{\delta})} + \frac{\frac{1}{1+\eta(\frac{1-\delta}{\delta})} (\omega_1 + \frac{1-\delta}{\delta})}{1 - \frac{\eta}{1+\eta(\frac{1-\delta}{\delta})} (\omega_1 + \frac{1-\delta}{\delta})} \log q \quad (126)$$

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.

## B.2 Beyond Isoelastic Demand

Isoelastic demand imposes both that the firm's own price elasticity of demand and its cross-price elasticity of demand are constant. In this appendix, we show how to derive optimal supply functions in closed form when the firm's own price elasticity of demand varies. This allows the demand curve to satisfy Marshall's second law of demand that the price elasticity of demand is increasing in the price as well as Marshall's third law of demand that the rate of increase of the price elasticity goes down with the price. We show that uncertainty about demand, prices, and marginal costs continue to operate in a very similar fashion. However, due to endogeneity of the optimal markup, the optimal supply schedule now ceases to be log-linear.

To capture these features, suppose that demand is *multiplicatively separable*:  $d(p, \Psi, P) = z(\Psi, P)\phi(p)$  for some function  $\phi$  such that  $p\phi''(p)/\phi'(p) < -2$ . This latter condition is satisfied by isoelastic demand exactly under the familiar condition that  $\eta > 1$  and ensures the existence of a unique optimal price. We further assume that  $z(\Psi, P) = \nu_0 \Psi^{\nu_1} P^{\nu_2}$  for  $\nu_0, \nu_1, \nu_2 \in \mathbb{R} \setminus \{0\}$ . This makes firms' uncertainty about the location of their demand curve log-normal. This assumption does rule out non-separable demand, such as the demand system proposed by [Kimball \(1995\)](#). However, it is important to note that this demand system is motivated by evidence on the firm's own price elasticity, which is governed by

$\phi$ , and not the cross-price elasticity, which is governed by  $\nu_2$ . Thus, our proposed demand system is equally able to capture facts about the firms' own price elasticity as the one proposed in [Kimball \(1995\)](#) (or the richer structures proposed by [Fujiwara and Matsuyama, 2022](#); [Wang and Werning, 2022](#)).

Under this demand system, we can derive a modified formula for the optimal supply curve which is now no longer log-linear, but continues to be governed by similar forces:

**Proposition 5.** *If demand is multiplicatively separable, then any optimal supply function is almost everywhere given by:*

$$f(p, q) = \log q + \hat{\alpha}_0 - \log \left( \phi(p) \left\{ p \left[ 1 + \frac{\phi(p)}{p\phi'(p)} \right] \right\}^{\frac{1}{\hat{\omega}_1}} \right) \quad (127)$$

where:

$$\hat{\omega}_1 = \frac{\nu_1(\sigma_{\mathcal{M},\Psi} + \sigma_{P,\Psi}) + \nu_2(\sigma_P^2 + \sigma_{\mathcal{M},P})}{\nu_1^2\sigma_\Psi^2 + \nu_2^2\sigma_P^2 + 2\nu_1\nu_2\sigma_{\Psi,P}} \quad (128)$$

*Proof.* Applying the same variational arguments as in [Theorem 1](#), we obtain that:

$$\hat{p}(z) + \frac{\phi(\hat{p}(z))}{\phi'(\hat{p}(z))} = \frac{\mathbb{E}[\Lambda\mathcal{M}|z]}{\mathbb{E}[\Lambda P^{-1}|z]} \quad (129)$$

where the condition  $p\phi''(p)/\phi'(p) < -2$  yields strict concavity of the objective and makes  $\hat{p}(z)$  the unique maximizer. Taking logarithms of both sides and evaluating the conditional expectations as per [Theorem 1](#), we obtain that:


$$\log \left( \hat{p}(z) \left[ 1 + \frac{\phi(\hat{p}(z))}{\hat{p}(z)\phi'(\hat{p}(z))} \right] \right) = \hat{\omega}_0 + \hat{\omega}_1 \log z \quad (130)$$

where  $\hat{\omega}_1 = \frac{\sigma_{\mathcal{M},z} + \sigma_{P,z}}{\sigma_z^2}$ , which yields [Equation 128](#). Using  $\log z = \log q - \log \phi(p)$  and rearranging yields [Equation 127](#).  $\square$

Demand uncertainty and price uncertainty enter the same way as before, via  $\hat{\omega}_1$ , and the intuition is the same. However, there are now two distinct notions of market power and they therefore operate in a more subtle way. First, consider the role of the cross-price elasticity of demand  $\nu_2$ . When  $\nu_2$  is higher, the firm's price is *ex post* more responsive to changes in others' prices. Second, consider the role of the own-price elasticity of demand  $\left(\frac{p\phi'(p)}{\phi(p)}\right)^{-1}$ . This induces non-linearity of the optimal supply schedule to the extent that it is not constant. This is because the firm's optimal markup changes as it moves along its demand curve.


## C Additional Tables and Figures

**Figure 7:** 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 8:** The Slope of Aggregate Supply (Annual-Frequency Calculation)




*Note:* This Figure plots estimates of the inverse elasticity of aggregate supply as measured by Equations 39 and 40. 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 9:** The Slope of Microeconomic Supply




*Note:* This Figure plots estimates of the slope of microeconomic supply as measured by Equation 39. These estimates are an input to our calculation of the slope of aggregate supply (Figure 4).

**Figure 10:** 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 39 and 40. 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 11:** Aggregate Supply with Pseudo-out-of-sample Uncertainty



*Note:* This Figure plots estimates of the inverse elasticity of aggregate supply as measured by Equations 39 and 40. In this calculation, in contrast to the calculation of Figure 4, we estimate uncertainty about variables in quarter  $t$  using a maximum-likelihood estimate of the model with data only up to quarter  $t-1$ . To cover the beginning of the sample, we use data starting in 1947:Q2.

**Figure 12:** 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  $\eta$  (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.