American Economic Review 2026, 116(2): 710–748
https://doi.org/10.1257/aer.20240443

A Theory of Supply Function Choice and Aggregate Supply†

By Joel P. Flynn, George Nikolakoudis, and Karthik A. Sastry*

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 can rationalize the flattening of aggregate supply during the Great Moderation and steepening during the 1970s and 2020s. (*JEL* D21, D43, E13, E23, E31, E32)

At the heart of modern macroeconomic models are monopolistic firms that make decisions under uncertainty. It is common to restrict these firms' decisions to a 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 (Caplin and Spulber 1987; Golosov and Lucas, 2007), and limited information (Mankiw and Reis 2002; Woodford 2003a). Price-setting firms are also at the core of the ubiquitous New Keynesian framework (Woodford 2003b).

However, as has long been recognized (e.g., Grossman 1981), price setting is *not* typically an optimal way for a firm to behave and is, at some level, an ad hoc modeling assumption. Why should firms not be able to raise their prices when goods are flying off the shelves? In practice, they can and do: firms use policies like temporary sales and surge pricing to navigate changing demand conditions (Den Boer 2015).

*Flynn: Yale University (email: joel.flynn@yale.edu). Nikolakoudis: University of Chicago (email: nikolakoudis@uchicago.edu). Sastry: Princeton University (email: ksastry@princeton.edu). Yuriy Gorodnichenko was the coeditor for this article. 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 NYU, CREI, 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.

†Go to https://doi.org/10.1257/aer.20240443 to visit the article page for additional materials and author disclosure statement(s).

In this paper, we remove external restrictions on firms' pricing strategies and instead allow firms to choose any *supply function*: a mapping that describes the price charged at each quantity of production. 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 (e.g., Grossman 1981; Klemperer and Meyer 1989). However, supply function choice has not previously 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.

Introducing supply functions in an otherwise standard business-cycle model yields an aggregate supply curve with an endogenous slope. That is, the relative response of inflation and output to an aggregate demand shock depends on the interaction between uncertainty and market structure, precisely because these forces affect firms' optimal supply functions. 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, the model's predictions for the slope of aggregate supply are consistent with time-series and cross-sectional evidence. Thus, supply functions provide a realistic and tractable 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 lognormal uncertainty about its competitors' prices, demand, productivity, input prices, and the stochastic discount factor. Given its beliefs, the firm chooses a *supply function* f: $\mathbb{R}^2_{++} \to \mathbb{R}$ that 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, nonparametric supply function to maximize its expected real profits under the stochastic discount factor. The ex ante interpretation of the firm's choice is in line with the ECON 101 notion of a supply curve: a systematic plan relating the price that a firm charges and the quantity that it produces. We show that the model has a complementary interpretation in terms of the firm's information set for price and quantity choices: Allowing firms to choose supply functions is isomorphic to allowing firms to condition prices and quantities on realized demand. This interpretation links our model theoretically to the notion of rational expectations equilibrium (Lucas 1972) and practically to the aforementioned examples of firms adjusting prices in response to current demand conditions. We finally note that the model nests the polar cases of price and quantity setting studied in previous literature. By relaxing strategic restrictions, we allow firms to choose strategies that are potentially preferable to these extremes.

We solve 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 α_1, the percentage by which the firm increases prices in response to a one percent increase in production. In turn, this elasticity depends on the firm's price elasticity of demand (i.e., its market power) and its relative uncertainty about demand, competitors' prices, and real marginal costs. These relationships arise because uncertainty

and market power shape firms' relative desires to hedge against different types of shocks.

Three comparative statics are particularly important for our macroeconomic analysis. First, higher uncertainty about firm-level demand pushes toward a lower α_1, or firms behaving 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 α_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 α_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.—We next embed supply-function choice in a business-cycle model with incomplete information (Woodford 2003a; Hellwig and Venkateswaran 2009). The model features a representative household that demands differentiated consumption goods, demands money balances, and supplies labor; macroeconomic shocks to the money supply and productivity; microeconomic shocks to firm-specific wages, productivity, and demand; time-varying volatility for these shocks; and intermediate goods firms that choose supply functions in the face of endogenous uncertainty. Because of this uncertainty, shocks to the money supply can affect real aggregate output, as in Lucas (1972).

We first characterize aggregate outcomes given fixed firm-level supply functions. The unique log-linear equilibrium has an aggregate supply and aggregate demand representation. That is, real GDP and the price level can be determined as the intersection of two curves: an *aggregate supply* curve that is affected by productivity shocks but not monetary shocks and an *aggregate demand* curve that is affected by monetary shocks but not technology shocks. The "slope of aggregate supply" determines the relative within-period responses of the price level and real output to an aggregate demand shock.

The slope of aggregate supply depends critically on the slope of firms' microeconomic 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 nonneutral as possible—if and only if firms are price setters. These results are disquieting in light of the standard approach of *assuming* that firms set either prices or quantities. 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. Between those extremes, the slope of aggregate supply is increasing in the slope of firm-level supply.

We next characterize how the slope of aggregate supply is endogenously determined via a 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 a closed-form illustration, we first study a special case that balances strategic

complementarity (from aggregate demand externalities) with substitutability (from wage pressure). In this case, the slope of aggregate supply decreases in a sufficient statistic for firms' uncertainty: their relative uncertainty about idiosyncratic demand shocks versus aggregate demand shocks. 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.

We finally observe that microeconomic and macroeconomic supply depend on relative rather than absolute uncertainties. Strikingly, even a vanishing amount of uncertainty can be consistent with *any* level of monetary nonneutrality, depending on the exact composition of uncertainty across different microeconomic and macroeconomic factors.

Our model predicts that the transmission of aggregate demand shocks is *state-dependent* because of firms' endogenous supply-function choices. This has several implications. First, more volatility in aggregate demand steepens aggregate supply. In our model, firms desire prices to respond more to demand in environments with more nominal uncertainty, and therefore endogenously lead demand shocks to have larger nominal effects and smaller real effects. Second, more volatility in *idiosyncratic* demand conditions has the opposite effect, flattening aggregate supply. Thus, countercyclical "risk shocks" at the microeconomic level (Bloom et al. 2018) increase the real effect of aggregate demand shocks. Third, a lower elasticity of demand (i.e., more market power) flattens aggregate supply, due to forces missing *only* in the knife-edge case of exogenously assumed price setting.

Quantification.—To illustrate the plausibility and macroeconomic relevance of these predictions, we quantify their implications in two simple calculations. We first study the model's predictions across time in the United States. We calculate a time-varying slope of aggregate supply using a standard calibration for macroeconomic parameters plus estimates of time-varying volatility from a simple statistical model. Our quantification implies that the slope of aggregate supply is relatively flat in normal times but spikes dramatically during the inflationary episodes of the 1970s and 2020s, consistent with empirical findings (e.g., Ball and Mazumder 2011; Cerrato and Gitti 2022). The key mechanism is that a spike in inflation uncertainty triggers firms to choose different supply functions, more aggressively varying prices with realized demand. In a second exercise, we study our model's implications for inflation-output trade-offs across countries in the spirit of Lucas (1973). Our model predicts vastly heterogeneous slopes of aggregate supply that correlate positively with simple empirical proxies, but *not* with realized inflation or inflation volatility. Thus, our model's predictions based on relative uncertainties account for variation in the slope of aggregate supply beyond what is explained by models that tie the slope solely to the level of inflation and absolute inflation uncertainty.

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 (e.g., Grossman 1981; Klemperer and

Meyer 1989; Vives 2017). We contribute to this literature by analytically characterizing equilibrium supply functions with several new features: nonquadratic preferences; imperfect substitutability; multiple, correlated sources of uncertainty; and general equilibrium interactions in both input and product markets.

The closest analysis in the macroeconomics literature is performed by Reis (2006), who compares the 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, removing all ad hoc strategic restrictions on firms' choices, and studying general equilibrium.[1] The latter feature allows us to study the equilibrium relationship between supply function choice, macroeconomic dynamics, and the slope of aggregate supply.

Our finding that uncertainty shapes the slope of aggregate supply is shared with the classic "islands model" analysis of Lucas (1972). A shared methodological premise is that economic agents act on what they learn from endogenous objects. Our results for the slope of the aggregate supply curve differ substantially for two reasons. First, we study producers with market power, 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 migration or physically separated markets. 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 noncompetitive 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 nonneutrality 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.[2]

Outline.—Section I solves for the firm's optimal supply function in partial equilibrium. Section II introduces a monetary business cycle model with supply functions. Section III characterizes equilibrium with supply function choice and shows how supply function choices affect aggregate supply. Section IV quantifies the model's predictions. Section V concludes.

[1] In Flynn, Nikolakoudis and Sastry (2024), we study the problem of prices *versus* quantities choice in general equilibrium and draw out implications for monetary policy.

[2] 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.

I. 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.

A. *The Firm's Problem*

Environment.—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 curve given by

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

where $p \in \mathbb{R}_+$ is the market price, $\Psi \in \mathbb{R}_{++}$ is a demand shifter, $P \in \mathbb{R}_{++}$ is the aggregate price level, and $\eta > 1$ is the price elasticity of demand. We interpret the elasticity of demand as an (inverse) measure of market power: when η is high, 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 $\left(\Psi, \mathcal{M}, P, \Lambda\right)$ follows a lognormal distribution with mean μ and variance Σ. 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 $\left(\Lambda, P, \mathcal{M}, \Psi, p, q\right)$.

Supply-Function Choice.—The firm implements price-quantity pairs described by the implicit equation $f\left(p,q\right) = 0$ where $f: \mathbb{R}^2_{++} \to \mathbb{R}$. We will refer to f as the supply function. Price setting is nested as a case in which $f\left(p,q\right) \equiv f^P\left(p\right)$. Quantity setting is nested as a case in which $f\left(p,q\right) \equiv f^Q\left(q\right)$. More generally, we allow plans to be given by any nonparametric function f, allowing for possible nonmonotonicity and discontinuities.

After choosing a supply function f, and following the realization of Ψ and P, the firm produces at a point where f intersects the demand curve. That is, the market clears. To formalize this, 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 $(\Lambda, P, \mathcal{M}, \Psi, p, q)$ 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:

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

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. This is like the ECON 101 notion of a supply curve as the firm's plan linking production and prices. 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. That is, supply functions allow the firm to condition its price and quantity on the strength of demand.

In Figure 1, we illustrate how different supply functions translate into price-quantity outcomes. The first row illustrates market clearing by showing different supply functions intersected with two demand curves, corresponding to "high" and "low" realizations. The second row illustrates the induced joint distribution of prices and quantities. Panel A shows a "price-setting" supply function, $f(p,q) = 1 - p$. In this case, a firm responds to low demand by producing less and responds to high demand by producing more in order to maintain a fixed price. The "quantity-setting" policy (panel B), $f(p,q) = 1 - q$, does the opposite: in this case, the firm aggressively decreases the price of low-demand goods and increases the price of high-demand goods to fix the quantity sold. The supply function in panel C, $f(p,q) = 1 - \frac{p}{q}$, allows both prices and quantities to increase with demand. This describes a firm with less extreme dynamic pricing: high-demand states have higher prices and volumes, and low-demand states have lower prices and volumes. In our model, the firm picks the optimal supply function given its uncertainty about economic conditions.

Interpreting Supply Functions.—We argue that there are strong theoretical and empirical grounds to study supply function choice as a benchmark model of firm choice under uncertainty.

First, the supply function model is consistent with firms' using all valuable information revealed by market clearing. This premise is central to the modern paradigm of rational expectations equilibrium (REE, Lucas 1972). It is of course possible

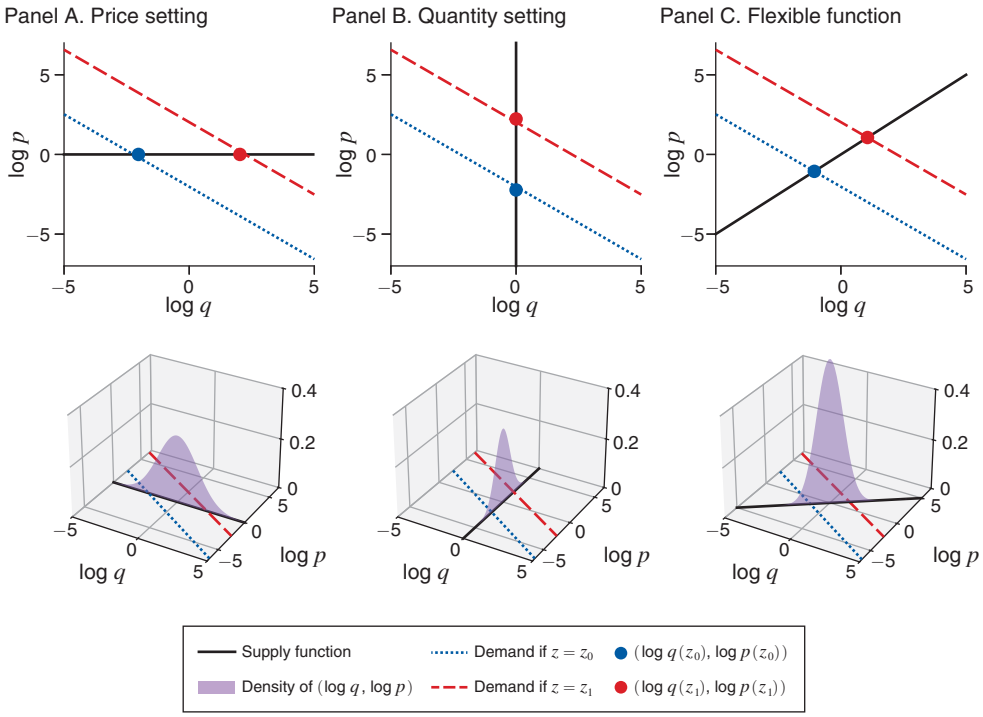

FIGURE 1. AN ILLUSTRATION OF SUPPLY-FUNCTION CHOICE

Notes: 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 lognormal uncertainty about z.

that firms are hindered in incorporating this information by ex post costs to adjust quantities and prices, for example due to technological constraints or nominal rigidities. But, as we further clarify in Section IC, any frictions that induce *finite* adjustment costs are conceptually straightforward to incorporate into the supply-function model—whereas only *infinite* adjustment costs on specific margins can justify the standard models of price setting and quantity setting. From this perspective, 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.

Second, supply function choice does not require that firms commit to plans that are revealed to be suboptimal by demand conditions. For example, in the typical model of price setting, a firm could lose money (to an arbitrary extent) to honor its fixed-price commitment in a state of the world in which goods are "flying off the shelves." Our supply function model can be thought of as formalizing the scenario in which a firm can reconsider its plans after every realization of demand.

Third, the supply function model decouples "price inertia" from "price stickiness," but is not incompatible with the latter. As in other models of aggregate supply based on incomplete information (e.g., Lucas 1972; Mankiw and Reis 2002; Woodford 2003a), our model predicts that the prices charged by an individual firm change in every period. We adopt this intentionally extreme view to highlight how

the novel mechanisms of our model can translate "flexible" prices into sluggish responses of the aggregate price level to shocks (see Section III). Nonetheless, it is conceptually straightforward to extend the model to allow for microeconomic price rigidity (see Section IC). In this case, the supply functions model still has appealing theoretical properties relative to the benchmarks of price and quantity setting—no firm regrets its price after observing market clearing—as well as materially different predictions, upon which we later elaborate.

Fourth, as a practical matter, there are many pricing strategies that could be well described as supply functions that incorporate information from contemporaneous demand conditions into prices. This is common for many goods and services that feature "dynamic pricing," for example in markets for electricity, gasoline, e-commerce, ride-sharing, and entertainment (e.g., Den Boer 2015). Klemperer and Meyer (1989) provide two concrete examples of firms that *de facto* implement supply schedules: management consultants who do not post prices, but instead vary them as a function of the quantity of services provided, and airlines that use computer software to put seats on discount depending on how many are currently sold. Crucially, in these examples, firms' choices do not seem completely constrained by technological necessity: they could in principle vary prices more or less aggressively depending on what is more profitable. Finally, varying prices with demand is not a new innovation: negotiated rather than posted prices were the norm throughout human history until the invention of the price tag in the mid-19th century (Phillips 2012).

B. *The Optimal Supply Function*

We now study the optimal supply function. 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 function is almost everywhere given by*

$$f(p,q) \;=\; \log p \;-\; \alpha_0 \;-\; \alpha_1 \log q, \tag{5}$$

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

$$\alpha_1 \;=\; \frac{\eta\sigma_P^2 + \sigma_{\mathcal{M},\Psi} + \sigma_{P,\Psi} + \eta\sigma_{\mathcal{M},P}}{\sigma_\Psi^2 - \eta\sigma_{\mathcal{M},\Psi} + \eta\sigma_{P,\Psi} - \eta^2\sigma_{\mathcal{M},P}}. \tag{6}$$

PROOF:

See Supplemental Appendix A.1.

Understanding the Result.—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 $\big(p(z), q(z)\big)$ that are indexed by the realization of the nominal demand state

$z = \Psi P^{\eta}$ and are such that the market clears: $\big(p(z), q(z)\big) = \big(p(z), zp(z)^{-\eta}\big)$. Intuitively, when setting a supply schedule, the firm anticipates that it will produce where the demand curve intersects 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. Importantly, it does so *without additional constraints*. We emphasize that this differs from the standard case of price and quantity setting due to the *lack of constraints* and not the contingency of choices on demand. Under price setting, firms implement z-contingent price-quantity pairs of the constrained form $\big(p(z), q(z)\big) = \big(\bar{p}, z\bar{p}^{-\eta}\big)$, where $\bar{p}$ is a fixed price and $z\bar{p}^{-\eta}$ is the quantity that clears the market; and similarly, under quantity setting, firms implement z-contingent price-quantity pairs of the constrained form $\big(p(z), q(z)\big) = \big(z^{1/\eta}\bar{q}, \bar{q}\big)$.

When choosing flexible supply schedules, firms can therefore freely incorporate information from the nominal demand state z into their optimal choices. To see this, we note that 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 at each $z = t$, we find that the firm equates the marginal revenue and cost effects of raising the price:

$$\underbrace{\mathbb{E}\big[(\eta - 1)z\Lambda P^{-1} p(z)^{-\eta} \big| z = t\big]}_{\text{Expected revenue effect}} = \underbrace{\mathbb{E}\big[\eta z \Lambda \mathcal{M} p(z)^{-\eta-1} \big| z = t\big]}_{\text{Expected cost effect .}}. \tag{7}$$

This principle of equating expected benefits and costs would apply in a broader class of models with different production technologies and demand curves (see Section IC), although the specific expressions for these terms in equation (7) rely on our exact assumptions. Rearranging terms, we observe that, for almost all $t \in \mathbb{R}_{++}$, the optimal price must satisfy

$$p(t) = \frac{\eta}{\eta - 1} \frac{\mathbb{E}\big[\Lambda \mathcal{M} \big| z = t\big]}{\mathbb{E}\big[\Lambda P^{-1} \big| z = t\big]}, \text{ and } q(t) = tp(t)^{-\eta}. \tag{8}$$

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 outcomes under optimal price setting due to the firm's ability to make inferences about the stochastic discount factor, real marginal costs, and the price level. Finally, we exploit the joint lognormality of the variables in the firm's problem to solve equation (8) for an exact log-linear relationship between prices and quantities. Without lognormality, equation (8) could be solved using numerical techniques.

It remains to explain why the optimal inverse supply elasticity takes the form given in equation (6). This specific form arises because α_1 is the relative rate at which the firm wants log prices and log quantities to increase with the nominal demand state $\log z$:

$$\alpha_1 = \frac{d \log p}{d \log z} \Big/ \frac{d \log q}{d \log z} = \frac{\mathrm{Cov}\big[\log z, \log p^{**}\big]}{\mathrm{Cov}\big[\log z, \log q^{**}\big]}. \tag{9}$$

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

$$p^{**} = \frac{\eta}{\eta - 1}\mathcal{M}P, \quad \text{and} \quad q^{**} = \left(\frac{\eta}{\eta - 1}\right)^{-\eta} \frac{z}{(\mathcal{M}P)^{\eta}}. \tag{10}$$

An econometric metaphor illustrates why this is the optimal way to set α_1. By equation (9), 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.

The Effects of Uncertainty.—A critical determinant of the firm's optimal responsiveness of prices to quantities is their relative uncertainty about the price level, real marginal costs, and demand. To build intuition for this, we first focus on the case in which the firm's supply schedule is upward-sloping. This occurs if $0 \leq \text{Cov}\left[\log z, \log(\mathcal{M}P)\right] \leq \frac{1}{\eta}\text{Var}\left[\log z\right]$: High demand predicts that nominal costs are higher, but not too much higher. In this case, greater price-level uncertainty (σ_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 (σ_Ψ^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 overproducing 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 firm may prefer a *downward* sloping supply function in order to hedge against high costs in high-demand states.

The Effect of Market Power.—The elasticity of demand plays two roles in determining the optimal (inverse) elasticity of supply. The first relates to *payoffs*: When η is high, ex post optimal quantities 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 η is high, nominal demand contains relatively more information about the price level P and less about real demand Ψ.

In general, the interaction of these two forces can make the optimal supply function steepen or flatten when η increases. Below, we describe an intuitive 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* $\partial\alpha_1/\partial\eta > 0$, *is that each of the following four inequalities holds:*

$$\alpha_1 \geq 0,\ \eta\sigma_P^2 + \sigma_{\mathcal{M},\Psi} + \sigma_{P,\Psi} + \eta\sigma_{\mathcal{M},P} \geq 0, \tag{11}$$

$$\sigma_{\mathcal{M},P} \geq 0,\ 2\eta\sigma_{\mathcal{M},P} + \sigma_{\mathcal{M},\Psi} \geq \sigma_{P,\Psi}.$$

PROOF:

See Supplemental 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. Later, in our quantitative analysis (Section IV), we find that the condition of Corollary 1 always holds in US data since 1960 as long as $\eta > 2.5$. 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 makes clear that pure price setting and quantity setting are two isolated points 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.*

(*i*) *As* $\sigma_P^2 \to \infty, |\alpha_1| \to \infty$ *and the optimal plan converges to quantity setting.*

(*ii*) *As* $\sigma_\Psi^2 \to \infty, \alpha_1 \to 0$ *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.

C. *Supply Functions in Other Environments*

We have made many specific assumptions on technology and demand for exposition simplicity. As we will show in the remainder of the paper, these same assumptions will also allow us to tractably study a general-equilibrium environment with an endogenous feedback loop between supply-function choice and endogenous uncertainty.

Nonetheless, the basic economic logic of supply-function choice extends to a much broader class of models. The unifying theme is the observation that choosing supply functions is tantamount to incorporating information regarding the desired price and quantity from market clearing. This basic observation did not rely on our specific description of technology and demand. As such, firms' optimal supply functions in other environments inherit the following fundamental logic of Theorem 1: firms want prices to increase steeply in quantities when high demand predicts a high desired price but not a high desired quantity. Moreover, taking the perspective of supply functions versus the traditional view of price or quantity setting introduces novel channels through which firms' uncertainty shapes their choices.

In Supplemental Appendix B, we illustrate this point by characterizing optimal supply functions in four extensions that explore different assumptions about real and nominal rigidities. We briefly summarize all four extensions below.

Multiple Inputs, Decreasing Returns-to-Scale, and Monopsony.—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). 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. 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 more rigid quantities.

Endogenous Markups.—Second, we allow for demand that is not iso-elastic in a class that 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. Intuitively, this allows the model to capture the possibility that "goods flying off the shelves" is informative about the desired markup.

Additional Choice Variables.—Third, we allow for the firm to choose additional variables beyond prices and quantities that may affect the joint distribution of all variables relevant to the firm, such as marginal costs and demand. One example is a firm that can invest in improving the quality of its product. We show in Proposition 6 that the firm's supply function remains optimally log-linear but with a slope that depends on the uncertainty that is induced by their choice of additional variables. Using this, we characterize the value of any choice of additional variables and show

in an example how such a firm would optimally choose the quality of its product at a cost.

Sticky Prices.—Fourth, we enrich the firm's problem with price stickiness as in Calvo (1983) to capture the empirically relevant possibility that the firm's prices may be fixed for multiple periods. We show in Proposition 7 that the firm's supply function remains optimally log-linear but with a slope that depends on how the firm learns from its demand today about the full future sequence of its nominal marginal costs. This demonstrates the simplicity with which the supply function approach could be integrated into macroeconomic models with sticky prices, like the textbook New Keynesian model (Woodford 2003a).

II. Supply Functions in a Macroeconomic Model

We now embed supply-function choice in a monetary macroeconomic model. We otherwise use intentionally standard microfoundations (e.g., Woodford 2003b; Hellwig and Venkateswaran 2009). These microfoundations will allow for a closed-form analysis and highlight the core economics of supply functions without any approximations. 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.

A. *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}:

$$\mathbb{E}_0\left[\sum_{t=0}^{\infty}\beta^t\left(\frac{C_t^{1-\gamma}}{1-\gamma}+\log\frac{M_t}{P_t}-\int_{[0,1]}\phi_{it}N_{it}di\right)\right], \tag{12}$$

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 $\log \phi_{it} \sim N\left(\mu_\phi, \sigma^2_{\phi,t}\right)$. The consumption aggregate C_t is a constant-elasticity-of-substitution aggregate of the individual consumption varieties with elasticity of substitution $\eta > 1$:

$$C_t = \left(\int_{[0,1]}\vartheta_{it}^{\frac{1}{\eta}}c_{it}^{\frac{\eta-1}{\eta}}di\right)^{\frac{\eta}{\eta-1}}, \tag{13}$$

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

$$P_t = \left(\int_{[0,1]} \vartheta_{it} p_{it}^{1-\eta} di\right)^{\frac{1}{1-\eta}}. \tag{14}$$

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

$$\begin{aligned} M_t + B_t + \int_{[0,1]} p_{it} C_{it} di = M_{t-1} + \left(1 + i_{t-1}\right) B_{t-1} + \int_{[0,1]} w_{it} N_{it} di \\ + \int_{[0,1]} \Pi_{it} di + T_t, \end{aligned} \tag{15}$$

where p_{it} is the price of variety i, w_{it} is a variety-specific nominal wage, and T_t are transfers of money from the government.

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

$$\log M_t = \log M_{t-1} + \mu_M + \sigma_t^M \varepsilon_t^M, \tag{16}$$

where the monetary innovation is an IID random variable that follows $\varepsilon_t^M \sim N\left(0,1\right)$. So that interest rates remain strictly positive, we assume that $\frac{1}{2}\left(\sigma_t^M\right)^2 \leq \mu_M$ for all $t \in \mathbb{N}$.

B. *Firms*

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

$$q_{it} = \zeta_{it} A_t L_{it}, \tag{17}$$

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

$$\log A_t = \rho \log A_{t-1} + \sigma_t^A \varepsilon_t^A, \tag{18}$$

where the productivity innovations are IID and follow $\varepsilon_t^A \sim N\left(0,1\right)$. For our main analysis, we assume that innovations to aggregate productivity and the money supply are independent. In Supplemental Appendix C, we generalize our results to allow for arbitrary correlation between these shocks. When the firm sells output at price p_{it} and hires labor at wage w_{it}, its nominal profits are given by $\Pi_{it} = p_{it} q_{it} - w_{it} L_{it}$. Since firms are owned by the representative household, their objective is to

maximize expectations of real profits, discounted by a real stochastic discount factor Λ_t. Thus, the firm's payoff is $\frac{\Lambda_t}{P_t}\Pi_{it}$.

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

$$s_{it}^A = \log A_t + \sigma_{A,s,t}\varepsilon_{it}^{s,A} \tag{19}$$

$$s_{it}^M = \log M_t + \sigma_{M,s,t}\varepsilon_{it}^{s,M},$$

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 ζ_{it}, demand shock ϑ_{it}, and labor supply shock ϕ_{it}.[3]

C. *Markets and Equilibrium*

In each period, conditional on the aforementioned information set, firms choose a supply function. As in Section I, 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*

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

and a sequence of supply functions $\left\{f_{it}\colon \mathbb{R}_{++}^2 \to \mathbb{R}\right\}_{i\in[0,1],t\in\mathbb{N}}$ *such that, in all periods:*

(*i*) *All firms choose their supply function* f_{it} *to maximize expected real profits under the household's stochastic discount factor.*

(*ii*) *The household chooses consumption* C_{it}*, labor supply* N_{it}*, money holdings* M_t*, and bond holdings* B_t *to maximize their expected utility subject to their lifetime budget constraint, while* Λ_t *is the household's marginal utility of consumption.*

(*iii*) *Money supply* M_t *and productivity* A_t *evolve exogenously via equations* (*16*) *and* (*18*).

[3] 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.

(*iv*) *Firms' and consumers' expectations are consistent with the equilibrium law of motion*.

(*v*) *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.

III. 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. 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 (GE) models.

A. *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 I: 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.*

$$\Psi_{it} = \vartheta_{it} C_t,\ P_t = \frac{i_t}{1+i_t} C_t^{-\gamma} M_t,\ \Lambda_t = C_t^{-\gamma},\ \mathcal{M}_{it} = \frac{\phi_{it} C_t^{\gamma}}{\zeta_{it} A_t}, \tag{20}$$

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

PROOF:

See Supplemental 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. 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 γ. Intuitively, when consumption utility has greater curvature, income effects in money demand are larger and money demand is more responsive to changes in consumption. Thus, consumption responds less to real money balances when γ is large. Third, the SDF is 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 their productivity.

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).[4] 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 lognormal, then so too is $\left(\Psi_{it}, P_t, \Lambda_t, \mathcal{M}_{it}\right)$. This follows from the fact that all four expressions are log-linear and all other fundamentals $(A_t, M_t, \vartheta_{it}, \phi_{it}, \zeta_{it})$ are lognormal by assumption. Therefore, if we can find that C_t is lognormal in equilibrium, our Theorem 1 can be directly applied to determine the optimal supply function in general equilibrium in our fully nonlinear setting. We will call an equilibrium in which $\log C_t$ is linear in $\left(\log A_t, \log M_t\right)$ a *log-linear equilibrium*.

B. *From Supply to Aggregate Supply with Fixed Functions*

We start by assuming that firms' exogenously set log-linear supply functions:

$$\log p_{it} = \alpha^*_{0t,i}\left(\alpha_{1,t}\right) + \alpha_{1,t} \log q_{it}, \tag{21}$$

where $\alpha_{1,t} \in \mathbb{R}$ is a fixed parameter and $\alpha^*_{0t,i}\left(\alpha_{1,t}\right)$ is the profit-maximizing "intercept" conditional on this slope.[5] 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 understand the ultimate macroeconomic implications of optimal supply function choice.

[4] 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.

[5] 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.

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:

$$\log P_t = \chi_{0,t}(\alpha_{1,t}) + \chi_{A,t}(\alpha_{1,t})\log A_t + \chi_{M,t}(\alpha_{1,t})\log M_t \tag{22}$$

$$\log C_t = \tilde{\chi}_{0,t}(\alpha_{1,t}) + \tilde{\chi}_{A,t}(\alpha_{1,t})\log A_t + \tilde{\chi}_{M,t}(\alpha_{1,t})\log M_t.$$

To this end, we define the posterior weight on firms' signals of productivity and the aggregate money supply as, respectively, $\kappa_t^A = \left(1 + \left(\sigma_{A,s,t}/\sigma_t^A\right)^2\right)^{-1}$ and $\kappa_t^M = \left(1 + \left(\sigma_{M,s,t}/\sigma_t^M\right)^2\right)^{-1}$. Moreover, define the slope of supply functions in terms of $\log z_{it} = \eta \log P_t + \log \Psi_{it}$ as[6]

$$\omega_{1,t} = \frac{\alpha_{1,t}}{1 + \eta\alpha_{1,t}}. \tag{23}$$

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 that 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

$$\epsilon_t^S = \gamma \frac{\kappa_t^M + \frac{\omega_{1,t}}{\gamma}\left(1 - \kappa_t^M\right)}{\left(1 - \omega_{1,t}\eta\right)\left(1 - \kappa_t^M\right)} \quad \textit{and} \quad \epsilon_t^D = \gamma, \tag{24}$$

and the interest rate i_t, the intercept for the price level $\log \bar{P}_t$, and the partial equilibrium effect of productivity shocks δ_t do not depend on $\left(\log P_t, \log Y_t, \log M_t, \log A_t\right)$.[7]

PROOF:

See Supplemental Appendix A.4.

[6] So everything remains well defined, we will impose that $\omega_{1,t} \neq \left(\eta - 1/\gamma\right)\left(1 - \kappa_t^x\right)$ 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).

[7] See Supplemental Appendix A.4 for explicit formulae for these terms.

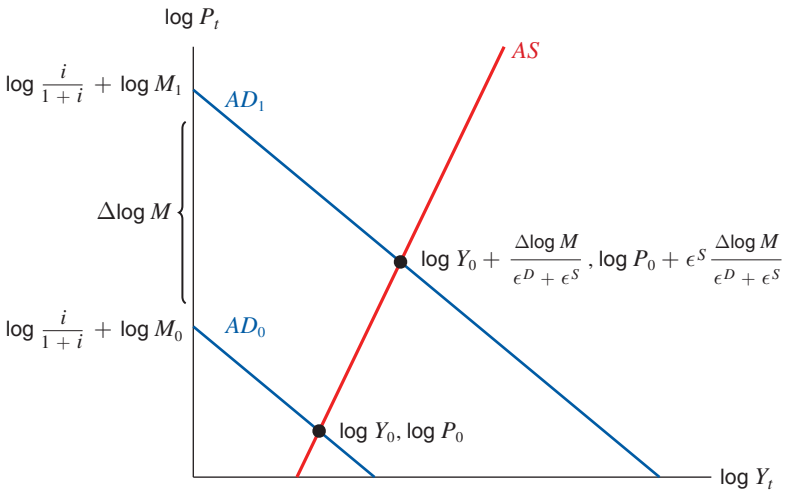

FIGURE 2. AN AGGREGATE SUPPLY AND DEMAND REPRESENTATION

Note: An aggregate supply and demand illustration of dynamics after a shock of size $\Delta \log M$ to the money supply (see Theorem 2).

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. The slope (or inverse elasticity) of aggregate demand in our model is γ. The AS 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. The price effect is larger and the quantity effect is smaller if ϵ^S is large. This calculation also makes clear that ϵ^S is 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 (PE) and a series of higher-order general equilibrium (GE) effects:[8]

$$(25)\quad \frac{\Delta \log P}{\Delta \log M} = \frac{\epsilon_t^S}{\epsilon_t^D + \epsilon_t^S} = \underbrace{\left(\kappa_t^M + \frac{\omega_{1,t}}{\gamma}\left(1 - \kappa_t^M\right)\right)}_{\text{Partial equilibrium}} \times \underbrace{\sum_{j=0}^{\infty}\left(\omega_{1,t}\left(\eta - \frac{1}{\gamma}\right)\left(1 - \kappa_t^M\right)\right)^j}_{\text{General equilibrium}}.$$

[8] 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.

It is first helpful to understand the case of pure price setting ($\omega_{1,t} = 0$). In this case, an increase in M by 1 percent raises real money balances by 1 percent, which increases consumption demand by $1/\gamma\%$. From the households' labor supply condition, this increases real marginal costs by $\gamma \times 1/\gamma = 1\%$. As the firm wishes to set its relative price equal to a constant markup over real marginal costs, the direct partial equilibrium effect of the shock is for the firms to increase prices by 1 percent. Next, in general equilibrium (i.e., accounting for the effect of the shock on the aggregate price level), the 1 percent increase in the price level induces a given firm to increase prices by 1 percent. However, higher prices also reduce real consumption demand, which lowers marginal costs by $\gamma \times 1/\gamma\%$. Thus, the general equilibrium effects net out, and an increase in M increases prices one-for-one. Finally, when firms are imperfectly informed of the money supply, they perceive that real marginal costs increase by only $\kappa_t^M\%$ on average. Thus, the price level increases by $\kappa_t^M\%$, which is obtained in equation (25) by setting $\omega_{1,t} = 0$.

Consider now the case with general supply functions. A 1 percent increase in the money supply induces the firm to experience a $1/\gamma\%$ demand shock. As the firm has an inverse supply elasticity of $\omega_{1,t}$, this leads to an additional effect in which the firm increases prices by $\omega_{1,t}/\gamma\%$. In order to keep its markup over real marginal costs fixed, firms reduce prices by $\kappa_t^M \times \omega_{1,t}/\gamma\%$ on average (via their intercept), which is their perceived increase in their demand following an increase in the money supply. Thus, the PE effect involves an addition term relative to the case of pure price setting, given by $\omega_{1,t}\left(1 - \kappa_t^M\right)/\gamma$.

With supply functions, there are also general equilibrium effects: a 1 percent in the aggregate price level leads all firms to experience a demand shock of η percent (as the prices of their competitors have increased). But it also reduces real consumption demand by $1/\gamma\%$, since real money balances fall. Together, these effects lead firms to increase their prices by $\omega_{1,t} \times \left(\eta - 1/\gamma\right)\%$. To keep the markup over their nominal marginal cost constant, firms reduce prices by $\kappa_t^M \times \omega_{1,t} \times \left(\eta - 1/\gamma\right)\%$ on average, which is their perceived increase in their demand. In total, out of a monetarily induced 1 percent increase in the aggregate price level, the average increase in firms' prices is therefore $\omega_{1,t} \times \left(\eta - 1/\gamma\right) \times \left(1 - \kappa_t^M\right)\%$. Iterating this to all subsequent price increases in general equilibrium GE yields equation (25).

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 percent. 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 $\epsilon^D = \gamma$, the level of these responses varies with the slope of 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:

$$\frac{\Delta \log P}{\Delta \log A} = \underbrace{-\kappa_t^A}_{\text{PE}} \times \underbrace{\sum_{j=0}^{\infty} \left(\omega_{1,t} \left(\eta - \frac{1}{\gamma} \right) \left(1 - \kappa_t^A \right) \right)^j}_{\text{GE}}. \tag{26}$$

The PE effect is immediate: firms perceive a κ_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 to those under monetary shocks, 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 slopes of microeconomic supply curves in the economy.

C. *The Slope of Aggregate Supply in Temporary Equilibrium*

We now study how various microeconomic forces affect the slope of aggregate supply.

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:* ∂

(*i*) *Steeper microeconomic supply steepens the AS curve:* $\partial \epsilon_t^S / \partial \omega_{1,t} \geq 0$.

(*ii*) *Precision of private information about money steepens the AS curve:* $\partial \epsilon_t^S / \partial \kappa_t^M \geq 0$.

(*iii*) *Income effects steepen the AS curve:* $\partial \epsilon_t^S / \partial \gamma \geq 0$.

(*iv*) *A higher elasticity of demand steepens the AS curve:* $\partial \epsilon_t^S / \partial \eta \geq 0$.

PROOF:

Follows immediately from differentiation of equation (24).

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 $\alpha_{0t,i}^*$). 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 γ (see Proposition 1). Consequently, a higher γ steepens the AS curve.

Finally, a lower elasticity of demand flattens the AS curve. Crucially, this effect is nonzero if and only if $\omega_{1,t} \neq 0$, i.e., firms do not undertake pure *price setting*. This flattening operates through the general equilibrium transmission mechanisms of the model. When other firms raise their prices in response to a money supply shock, firm-level demand increases because the firm's *relative* price is now lower. The magnitude of this demand change is exactly parameterized by the elasticity of substitution η. If the responsiveness of prices to quantities at the firm level is nonzero, 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.[9]

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 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:

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

$$\epsilon_t^S = \gamma \frac{\kappa_t^M}{1 - \kappa_t^M}. \tag{27}$$

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

$$\epsilon_t^S = \infty. \tag{28}$$

We illustrate these two "extreme" predictions for aggregate supply and demand in Figure 3. Since ϵ_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 (among all upward sloping supply functions) 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.

[9] 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.

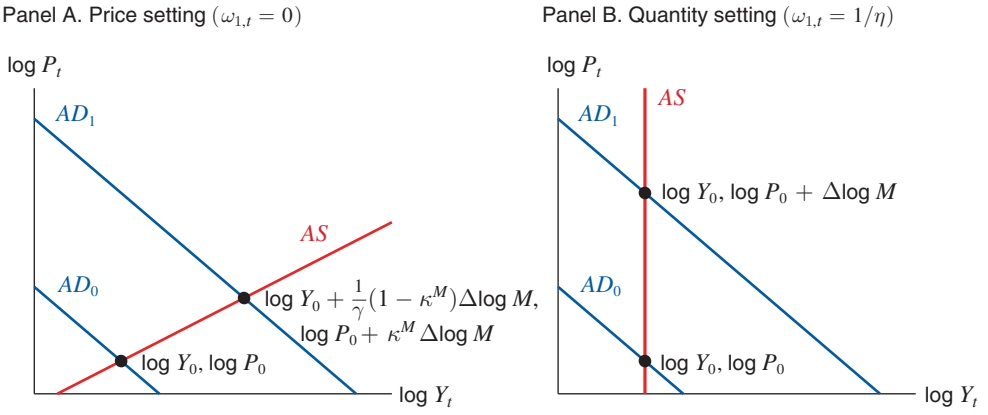

FIGURE 3. AGGREGATE SUPPLY UNDER PRICE SETTING AND QUANTITY SETTING

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).

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.

D. *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 lognormal (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 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:*

$$(29)\quad \omega_{1,t} = T_t(\omega_{1,t})$$

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

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

PROOF:

See Supplemental 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 does 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 nonneutrality 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 Supplemental Appendix A.6.

To understand why there are possibly multiple equilibria, let us compare two situations: one where firms set flat supply functions and one where firms set steep supply functions. If firms set steep supply functions, then prices are highly sensitive to demand, and so the aggregate supply curve is steep (recall Theorem 2). Conversely, if firms set flat supply functions, then the aggregate supply curve is flat. When the economy is hit by larger monetary shocks than productivity shocks, an economy with a steeper aggregate supply curve will generate greater variability in the aggregate price level. Moreover, greater variability in the aggregate price level leads to large losses from price-setting behavior and thus leads to firms wanting to set steeper supply functions (recall Theorem 1). Because of this, there can exist a feedback mechanism whereby steeper supply functions can generate greater equilibrium price volatility, which in turn reinforces the benefits of setting a steeper supply function. If this feedback mechanism is sufficiently strong, then there may exist multiple equilibria. Supplemental Appendix D.4 provides a quantitative illustration of the potential multiplicity of equilibria.

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 IIIB that η parameterizes the strength of *strategic complementarities*: the additional increase in demand that 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 that 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, $\eta\gamma = 1$ corresponds to the case in which these forces exactly balance. This allows us to simplify the fixed point in equation (29) considerably.

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

$$\epsilon_t^S = \gamma \frac{\kappa_t^M}{1 - \kappa_t^M}\left(1 + \frac{1}{\gamma^2 \rho_t^2 \kappa_t^M}\right), \tag{30}$$

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

PROOF:

See Supplemental 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 (8)) is not useful for conducting inference about productivity and so κ_t^A does not enter the fixed point. The same is not true for uncertainty about the money supply, as it induces direct variation in 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 $\rho_t \to \infty$, the inverse elasticity of aggregate supply approaches $\gamma\frac{\kappa_t^M}{1-\kappa_t^M}$. This is the AS curve slope under price setting ($\omega_{1,t} = 0$). Intuitively, *idiosyncratic* demand conditions do not affect a given firm's marginal cost. Hence, as idiosyncratic demand becomes relatively more volatile, the firm optimally sets a constant price to keep its markup over marginal cost constant. Had the firm chosen $\omega_{1,t} \neq 0$, the firm would induce unprofitable variation in its price by responding to idiosyncratic demand conditions.

Third, as $\rho_t \to 0$, 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 *versus* 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 the discretion to surprise private agents via its policy actions, 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, maintaining monetary policy discretion 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.[10] These results hold for any values of $\eta > 1$ and $\gamma > 0$, in contrast to the analysis above under balanced strategic interactions.

[10] Formally, we take these limits for $\sigma^x_{\eta|s}$ and $x \in \{M, A\}$ by scaling $\sigma_{x,s,t}$ and σ^x_t by a common factor.

COROLLARY 6 (Dominant-Shock Limits): *The following statements are true:*

(*i*) *As* $\sigma_{\vartheta,t} \to \infty$, *in any equilibrium* $\omega_{1,t} \to 0$ (*price setting*)

(*ii*) *As* $\sigma^M_{t|s} \to \infty$, *in any equilibrium* $\omega_{1,t} \to \frac{1}{\eta}$ (*quantity setting*)

(*iii*) *As* $\sigma^A_{t|s} \to \infty$ *and* $\eta\gamma \neq 1$, *in any equilibrium* $\omega_{1,t} \to \frac{1}{\eta - \frac{1}{\gamma}}$.

PROOF:

See Supplemental 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 nominal 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 percent decrease in productivity would imply that firms raise their prices by 1 percent. This translates to an $\left(\eta - \frac{1}{\gamma}\right)\%$ increase in demand for a given firm. Since firms believe that all fluctuations in demand are driven by productivity shocks, they set their optimal supply function slope to $\omega_{1,t} = \left[1/\left(\eta - \frac{1}{\gamma}\right)\right]$. This ensures that a 1 percent increase in productivity will reduce prices by 1 percent, thus keeping their markup over nominal marginal costs constant. 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 versus aggregate and demand versus productivity) matters. Thus, the *presence* of uncertainty is of central importance to our analysis. However, a distinguishing feature of the theory that we have developed is that the total *level* of uncertainty does not matter. To make this claim formal, fix a scalar $\lambda \geq 0$ and scale all uncertainty in the economy according to

$$(31)\quad \left(\sigma_{\vartheta,t}, \sigma_{z,t}, \sigma_{\phi,t}, \sigma_{A,t}, \sigma_{A,s,t}, \sigma_{M,t}, \sigma_{M,s,t}\right) \mapsto \left(\lambda\sigma_{\vartheta,t}, \lambda\sigma_{z,t}, \lambda\sigma_{\phi,t}, \lambda\sigma_{A,t}, \lambda\sigma_{A,s,t}, \lambda\sigma_{M,t}, \lambda\sigma_{M,s,t}\right).$$

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

PROPOSITION 3 (Invariance to Uncertainty and Discontinuity in the Limit): *For* $\lambda > 0$, $\mathcal{E}^S_t(\lambda)$ *is constant and the equilibrium supply elasticity is invariant to the level of uncertainty. Moreover*, $\mathcal{E}^S_t(0) = \{\infty\}$. *Therefore*, *the equilibrium supply elasticity is discontinuous in the zero uncertainty limit:*

$$(32)\qquad \lim_{\lambda \to 0} \mathcal{E}^S_t(\lambda) \neq \mathcal{E}^S_t(0).$$

PROOF:

See Supplemental 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). 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.[11] Importantly, this means that even in environments with low levels of uncertainty, the economic mechanisms that underlie our analysis are unchanged.

E. *A General Framework for Macroeconomic Analysis*

Our preceding analysis tractably illustrated the effect of supply functions in a fully nonlinear fashion. To do so, we made a number of simplifying assumptions regarding preferences and production functions. However, 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 (e.g., 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 the set of reasonable macroeconomic models.

Consider a model which generates a demand function for products given by $q_{i,t} = d\left(p_{it}, z_{it}^D\right)$, 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, in this dynamic setting, that the forward-looking value function V that the firm derives from setting a price p_{it} and selling a quantity q_{it} is given by $V\left(p_{it}, q_{it}, \mathbf{z}_{it}^V\right)$, where $\mathbf{z}_{it}^V$ is an n_V-sized vector of (potentially endogenous) variables

[11] Similarly, models with information acquisition and nominal rigidities (Afrouzi, Flynn, and Yang 2024) are also different from our model in that they do not feature this discontinuity in the limit.

that affect firm's value at time t. Candidate modifications to our framework that could be incorporated in this fashion include decreasing returns-to-scale, monopsony, endogenous markups, price stickiness, investments, and endogenous quality choice (see Supplemental Appendix B).

As in our model from Section I, the firm's optimal supply function problem is to choose a price that is contingent on demand q_{it}. As we have shown, this 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

$$p_{it}\left(z_{it}^D\right) = \arg\max_{p_{it}} \mathbb{E}_{it}\left[V\left(p_{it}, d\left(p_{it}, z_{it}^D\right), \mathbf{z}_{it}^V\right) | z_{it}^D\right], \tag{33}$$

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

$$\hat{p}_{it} = \tilde{\omega}_{1,it}' \mathbb{E}_{it}\left[\hat{\mathbf{z}}_{it}^V | \hat{z}_{it}^D\right] \tag{34}$$

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

$$\hat{p}_{it} = \omega_{1,it} \hat{z}_{it}^D \tag{35}$$

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, that is, $\hat{P}_t = \int_0^1 \hat{p}_{it} di$. This concludes the "firm block" of the model.

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

$$\mathcal{H}_x \mathbf{x}_t + \mathcal{H}_\varepsilon \boldsymbol{\varepsilon}_t = \mathbf{0}, \tag{36}$$

where x_t is an n_x-dimensional vector of endogenous variables (such as the ideal price index $\hat{P}_t$), ε_t is an n_ε-dimensional vector of Gaussian structural shocks, and $\mathcal{H}_z$ and $\mathcal{H}_\varepsilon$ are conforming matrices.[12] Equation (36), for example, contains the relevant first-order conditions and market-clearing conditions that determine the dynamics of an economy. Of course, the matrices $\mathcal{H}_x \mathbf{x}_t$ and $\mathcal{H}_\varepsilon \boldsymbol{\varepsilon}_t$ 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 equation (36). 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 (34). As argued by McKay and Wolf (2023), many of the

[12] Following McKay and Wolf (2023), we use boldface notation to stack the time paths for all variables (e.g., $\mathbf{x} = \left(\mathbf{x}_1', \ldots, \mathbf{x}_{n_x}'\right)'$). The matrices $\mathcal{H}_x$xt and $\mathcal{H}_{\varepsilon\varepsilon t}$ are conformable matrices that map bounded sequences to the space of bounded sequences.

parametric structural models commonly used for counterfactual analysis fit into the general framework of equation (36). 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. Validating the utility of this approach, Nikolakoudis (2025) incorporates supply functions along the input margin in a production network economy using these methods.

IV. Quantification

In this final phase, we quantify our model's implications for inflation-output trade-offs. In particular, we compute the model's predictions for the slope of aggregate supply as a function of state-dependent uncertainty. In the United States, aggregate supply dramatically steepens during the 1970s and post-COVID crisis, consistent with empirical evidence (Ball and Mazumder 2011; Hazell et al. 2022; Cerrato and Gitti 2022). Across countries, relative uncertainty can help account for the vast and well-documented differences in the relationship between inflation and real outcomes (Lucas 1973; Ball, Mankiw, and Romer 1988).

A. *Methods*

Our goal is to calculate the model-implied elasticity of aggregate supply and compare this to empirical evidence. To do this, we use Theorems 1 and 2 to write the slope of aggregate supply in terms of three parameters as well as *sufficient statistics* summarizing uncertainty about endogenous variables. That is,

$$\hat{\epsilon}_t^S = \gamma \frac{\kappa^M + \frac{\hat{\alpha}_{1,t}}{\gamma(1+\eta\hat{\alpha}_{1,t})}\left(1-\kappa^M\right)}{\left(1-\frac{\eta\hat{\alpha}_{1,t}}{1+\eta\hat{\alpha}_{1,t}}\right)\left(1-\kappa^M\right)}, \tag{37}$$

where

$$\hat{\alpha}_{1,t} = \frac{\eta\hat{\sigma}^2_{P,t} + \hat{\sigma}_{\mathcal{M},\Psi,t} + \hat{\sigma}_{P,\Psi,t} + \eta\hat{\sigma}_{\mathcal{M},P,t}}{\hat{\sigma}^2_{\Psi,t} - \eta\hat{\sigma}_{\mathcal{M},\Psi,t} + \eta\hat{\sigma}_{P,\Psi,t} - \eta^2\hat{\sigma}_{\mathcal{M},P,t}},$$

where the $\hat{\sigma}_{\cdot,t}$ denote period-by-period uncertainty about demand shifters Ψ, the price level P, and real marginal costs $\mathcal{M}$.

Our approach is to calibrate $\left(\gamma, \eta, \kappa^M\right)$ and directly measure the uncertainties. Compared to the alternative of structurally estimating all of the model's deep parameters and the stochastic processes for all exogenous shocks, this approach has two main benefits. First, there is a relatively simple mapping from data to theory. In particular, our estimation of firms' supply function slopes is independent of the general equilibrium "block" of the economy and relies only on firms' beliefs about payoff-relevant variables. Second, we bypass the theoretically possible issue of multiple equilibria: our calculations are valid in the model regardless of whether measured uncertainty is due to variation in fundamentals or due to equilibrium selection. A limitation is that the method precludes us from studying counterfactuals in which,

for example, firms' uncertainty endogenously responds to policy changes. We leave this kind of counterfactual analysis to future work. Nevertheless, in Supplemental Appendix D.4, we outline a method for counterfactual analysis and provide further intuition for when the fixed point of Theorem 3 may feature multiple equilibria.

Structural Parameters.—The period length controls the decision horizon of the firm. We set this to one quarter to be consistent with findings in US micro-data that the duration of the median posted price is four months (Bils and Klenow 2004). The elasticity of substitution between goods, η, controls the elasticity of demand faced by firms and, therefore, the relative importance of nominal and real shocks for the firm (see Corollary 1). We calibrate $\eta = 8$ to match the own-price elasticities of demand estimated by Hottman, Redding, and Weinstein (2016) using retail scanner data in the United States. The parameter γ controls the elasticity of real marginal costs to real output. We set $\gamma = 0.11$ to match the estimates of Gagliardone et al. (2025), who use micro-data from Belgian firms to estimate the elasticity of real marginal costs to output. The precision of firms' signals about the money supply, summarized by the signal-to-noise ratio κ^M, shifts the average slope of aggregate supply. We calibrate this to match an average aggregate supply slope of 0.11 as estimated by Hazell et al. (2022), yielding $\kappa^M = 0.29$. This allows us to isolate all time variation in the slope of aggregate supply as a consequence of time-varying uncertainty.

Uncertainty.—The crucial remaining ingredient is firms' uncertainty about demand, aggregate prices, and real marginal costs. To our knowledge, there are no datasets that directly measure firms' multidimensional uncertainty about these objects over a long period of time. Therefore, we proxy for firms' subjective uncertainty using a simple statistical model. Specifically, we use quarterly data on real GDP, inflation (GDP deflator), and capacity-utilization-adjusted TFP from 1960:I to 2024:IV.[13] Using the model's structure, we construct the aggregate component of real marginal costs as $\mathcal{M}_t = Y_t^{\gamma}/A_t$, where Y_t is real GDP, $\gamma = 0.11$ controls wealth effects in labor supply, and A_t is TFP. We model real GDP growth, inflation, and real marginal cost growth via a GARCH model (see Supplemental Appendix D.1 for details). This gives us a quarterly measure of one-step-ahead uncertainty regarding macroeconomic variables. This corresponds to the uncertainty an economic agent would have about current-quarter macroeconomic variables if they observed past economic history and interpreted it using the statistical model. Our methods are also compatible with other measures of uncertainty, such as regime-specific covariances from a vector autoregression (VAR) model.

To map our estimated uncertainties to equation (37), we take two additional steps. First, we observe that real marginal cost uncertainty enters firms' supply-function choice only via the covariances of real marginal costs with demand and the price level (see Theorem 1). These covariances, in turn, depend only on uncertainty about the *aggregate* component of real marginal costs. This allows us to directly apply the GARCH estimates. Second, we observe that the firm-level demand shock is $\Psi_{it} = Y_t \vartheta_{it}$, where ϑ_{it} is an idiosyncratic demand shock. Since we lack direct estimates

[13] The original data are from US Bureau of Economic Analysis (2025a, 2025b) and Federal Reserve Bank of San Francisco (2025).

of firms' uncertainty about idiosyncratic demand shocks, we assume that idiosyncratic demand uncertainty is directly proportional to aggregate TFP uncertainty. We justify this based on the finding of Bloom et al. (2018) that the time-varying volatility of revenue-based TFP (TFPR) among manufacturing firms can be accurately modeled as directly proportional to time-varying volatility in aggregate conditions.[14] That is, we set $\hat{\sigma}^2_{\Psi,t} = \hat{\sigma}^2_{Y,t} + R^2\hat{\sigma}^2_{\mathcal{M},t}$, where the $\hat{\sigma}_{\cdot,t}$ are the estimated GARCH variances and $R = 6.5$ from the quantitative estimates of Bloom et al. (2018).[15]

B. *The Slope of Aggregate Supply in the United States*

Our estimates imply that the slope of aggregate supply has varied significantly in the United States since 1960. Panel A of Figure 4 reports an annual time-series of our estimated slope, and Panel B reports average values over macroeconomic regimes. Aggregate supply significantly steepens in two episodes: the 1970s, coincident with the Oil Crisis, and the 2020s, coincident with the post-COVID inflation. Aggregate supply is flatter and roughly constant in the 1960s, the Great Moderation, the Great Recession, and the recovery from the Great Recession.

Comparison to Empirical Estimates.—The slope of aggregate supply ϵ^S_t is defined as the relative response of the price level and real GDP to a monetary expansion in the same quarter. In other words, ϵ^S_t measures how inflationary are aggregate demand shocks in the short-run or the inverse "sacrifice ratio" for monetary policymakers. Therefore, the object of comparison in the data is the relative elasticity of inflation and real outcomes to an identified aggregate demand shock or, equivalently, the elasticity of inflation to real outcomes when the latter is instrumented with an identified aggregate demand shock.

A typical observation in research on US business cycles is that aggregate demand shocks in the data induce puzzlingly little short-run inflation: that is, in our language, short-run aggregate supply is relatively flat. For example, Ramey (2016) summarizes a large literature documenting relatively small responses of inflation to externally identified monetary policy shocks and Hazell et al. (2022) use panel data from US states to show that there is limited pass-through from unemployment to price inflation. Two particularly important observations in these studies of US economic history are the lack of deflation in response to the crash of aggregate demand in the Great Recession and the lack of inflation after the large subsequent monetary expansion (see also Coibion and Gorodnichenko 2015; Bobeica and Jarociński 2019). Our calculations are consistent with a low average slope of aggregate supply as well as a low slope during the Great Recession and recovery.

Nonetheless, studies that have estimated a *time-varying* slope of aggregate supply in the United States have found larger values in two periods, the 1970s and 2020s. Our model generates these dynamics as an endogenous consequence of abrupt

[14] Foster, Haltiwanger and Syverson (2008) show that cross-firm variation in revenue total factor productivity (TFPR) derives almost exclusively from demand differences rather than marginal cost differences within specific industries, justifying our assumption that most "micro-volatility" is demand volatility.

[15] Figure A1 in the Supplemental Appendix plots each component of estimated uncertainty over time.

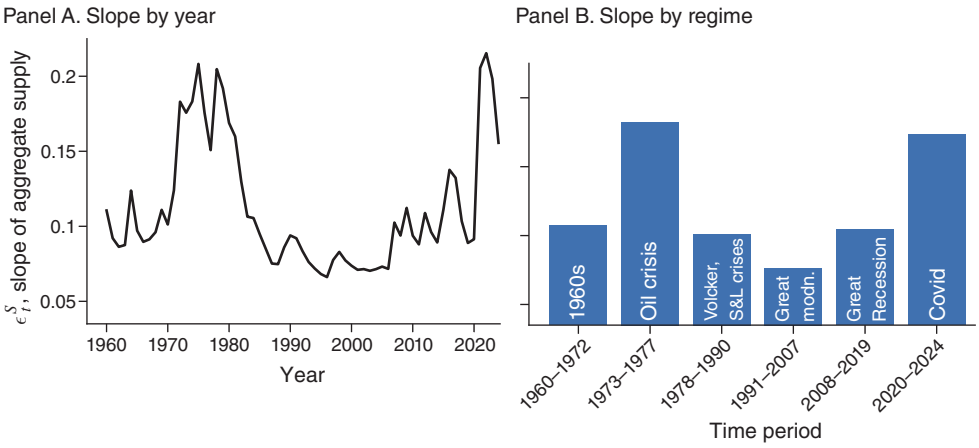

FIGURE 4. THE SLOPE OF AGGREGATE SUPPLY IN THE UNITED STATES OVER TIME

Notes: This figure shows the model implied slope of aggregate supply in the United States. The slope ϵ_t^S is defined by equation (37). As described in Section IVA, we set $\eta = 8$ and $\gamma = 0.11$ based on the literature, calibrate $\kappa^M = 0.29$ to match an average slope of 0.11 (Hazell et al. 2022), and estimate time-varying uncertainties from a GARCH model of quarterly frequency GDP growth, inflation, and real marginal cost growth. Panel A displays the estimated slope averaged to the annual level. Panel B shows average values of the estimated slope in different macroeconomic regimes, defined on the horizontal axis and labeled inside the bars.

changes in the composition of uncertainty. We make this comparison quantitatively in Table 1, which compares our predicted dynamics in the slope of aggregate supply to external estimates. The first panel is based on the study of Ball and Mazumder (2011), who estimate a time-varying relationship between inflation and the output gap.[16] Our model qualitatively matches a steepening of aggregate supply during the oil crisis and Volcker disinflation (1973–1984), while aggregate supply is relatively flat before and after. The second panel is based on the estimates of Hazell et al. (2022) who, as mentioned above, use state-level data in the United States to isolate demand-driven variation in real conditions. Our model can account for about one-half of the authors' estimated flattening of the Phillips curve from 1978–1990 to 1991–2018.

The third panel shows the consistency of our results with the behavior of aggregate supply before and after the COVID-19 crisis based on the estimates of Cerrato and Gitti (2022), who use data from US metropolitan statistical areas (MSAs) to isolate demand shocks. Quantitatively, our model accounts for almost all of the steepening of aggregate supply between the pre-COVID and post-COVID periods. Both the model predictions and empirical evidence imply that aggregate demand shocks were relatively inflationary during this period—consistent with both a large inflationary effect (and small real effect) of fiscal expansion and a large disinflationary effect (and small real effect) of monetary tightening.

[16] These authors construct the output gap using the Congressional Budget Office's estimate of potential real GDP. This corresponds to variation in real GDP induced by aggregate demand shocks if aggregate supply shocks move output together with potential output.

Table 1—Comparing the Model's Prediction to External Estimates

	Change in slope of aggregate supply relative to base period, compared to estimates from						
	A. Ball and Mazumder (2011)			B. Hazell et al. (2022)		C. Cerrato and Gitti (2022)	
	1960–1972	1973–1984	1985–2007	1978–1990	1991–2018	1991–2019	2021–2023
Data	–	+175%	−32%	–	−51%	–	+145%
Model	–	+53%	−27%	–	−27%	–	+138%

Notes: This table compares the model's estimates for the slope of aggregate supply with external empirical estimates. The values are percent changes in the slope of aggregate supply, relative to the base period (i.e., the first period in each panel). The three panels correspond to comparisons with different studies: in Ball and Mazumder (2011), column 4 of Table 3; in Hazell et al. (2022), panel B, columns 3 and 4 of Table II; in Cerrato and Gitti (2022), column 2 of Table 2.

These results suggest that our model's qualitative predictions for variation in the slope of aggregate supply are consistent with US empirical evidence. Of course, our baseline model is lacking several features required to make the model fully consistent with price stickiness at the microeconomic level and realistically sluggish inflation dynamics at the macroeconomic level. Theoretically, we have shown that such features can be integrated into the supply-function approach, while maintaining the essential properties of the latter (Section IC).[17] We leave a more complete quantitative exploration in a richer model to further research.

Rising Markups and Flattening Supply.—While our analysis thus far has focused on time variation in only uncertainty, in principle time variation in other structural parameters could affect the slope of aggregate supply in our model. In Supplemental Appendix D.2, we show that a version of our quantification that feeds in a secular decline of the elasticity of substitution, or a secular increase in average markups, predicts an even more pronounced flattening of the aggregate supply curve over time (Supplemental Figure A3). We leave a fuller empirical investigation of the model's joint predictions for aggregate supply and markups to future work.

C. *The Slope of Aggregate Supply around the World*

A cursory glance at Figure 4 suggests that aggregate supply is relatively steep in times during "inflationary crises" in the United States and relatively flat otherwise. Our quarterly time-series for ϵ_t^S has a correlation of 0.93 with one-quarter-ahead *uncertainty* regarding inflation and a correlation of 0.62 with the quarterly *level* of inflation (Supplemental Appendix Figure A2). The deviation between the level of inflation and inflation uncertainty primarily occurs near "turning points" of the inflation time series: for example, uncertainty is high despite relatively tame inflation in 2021:I, on the heels of the COVID-19 lock-down, while uncertainty is only moderate despite very high inflation in 1981:II, near the peak of the Volcker tightening cycle. Nonetheless, the high correlation between these variables presents a challenge for differentiating the *qualitative* predictions of our model from alternatives in

[17] The extension to sticky prices (Supplemental Appendix B.4) would assist on both fronts. Another possibility, which is conceptually straightforward to integrate into the model, is to assume that firms do not observe the history of macroeconomic aggregates and therefore form beliefs using a latent-state (Kalman filtering) model. As shown by Woodford (2003a) in a model of price setting, the combination of incomplete information with strategic interaction can generate realistically slow responses of inflation to shocks.

which price rigidity depends on the baseline rate of inflation or a one-dimensional summary of inflation volatility (e.g., Ball, Mankiw, and Romer 1988). These include models of endogenous price reset probabilities and models of menu costs.

To help distinguish between the roles of the level of inflation and uncertainty about inflation, we conduct a cross-country analysis in the spirit of Lucas (1973) and Ball, Mankiw, and Romer (1988). We compile annual data on real GDP growth, GDP deflator inflation, and TFP growth from 1960–2019 for countries in the OECD.[18] In each country, we model these variables as a one-lag VAR in first differences. Using the same structural parameters and mapping to the model described above, we construct country-level measures of uncertainty regarding demand, inflation, and real marginal costs and a model-implied slope of aggregate supply.[19]

Our first finding is that our cross-country estimates of the slope of aggregate supply are *not* positively predicted by average inflation (panel A of Figure 5). This is in sharp contrast to our within-country estimates for the United States. Internationally, the slope of aggregate supply is predicted to be *negative* for the countries with the highest rate of inflation—a consequence of the fact that the same countries have very highly correlated uncertainty about the price level and real marginal costs.

We next show that our model's predictions for cross-country relationships between inflation and output growth line up with the data. Panel B of Figure 5 shows a positive relationship between our predicted slope and the year-on-year regression coefficient of inflation on real output growth: for each country c, $\hat{\beta}_c^S = \frac{\text{Cov}[\Delta \log Y_{ct}, \Delta \log P_{ct}]}{\text{Var}[\Delta \log Y_{ct}]}$. Thus, in the raw data, output growth is more inflationary in countries for which we predict a steep slope of aggregate supply. Of course, the classic conceptual issue with interpreting inflation-output "correlations" as the slope of aggregate supply is that macroeconomic outcomes are determined by both demand and supply shocks. In Supplemental Appendix D.3, we construct a model-derived instrument which isolates exogenous variation in the money supply. Our IV estimate of the slope of aggregate supply is also positively correlated with our predicted slope.

While these findings on a small set of countries are merely suggestive, they imply that our model and its emphasis on *relative* uncertainties might help account for the enormous heterogeneity in inflation-output dynamics across countries. Moreover, the model has predictive power over and above models in which the slope of aggregate supply depends positively on only the mean or volatility of inflation (e.g., Ball, Mankiw, and Romer 1988).[20]

[18] We construct TFP as a Solow residual with elasticities $1/3$ and $2/3$ on the real capital stock and total labor hours, respectively. We drop three outliers from our calculation, Greece, Iceland, and Sweden, for which we calculate a slope of aggregate supply and/or inflation-output relationship more than 3 standard deviations away from the median All raw data are from the most recent edition of the *Penn World Tables* (Groningen Growth and Development Centre 2025); see also Feenstra, Inklaar, and Timmer (2015a, b).

[19] While our calculation embodies numerous assumptions, we argue that the main ingredients capture intuitive differences across countries. For example, uncertainty about inflation is the lowest among OECD countries in the United States, Germany, and Canada, and the highest in Chile, Israel, and Mexico.

[20] In Supplemental Appendix Table A2, we more formally demonstrate this in a cross-country regression analysis.

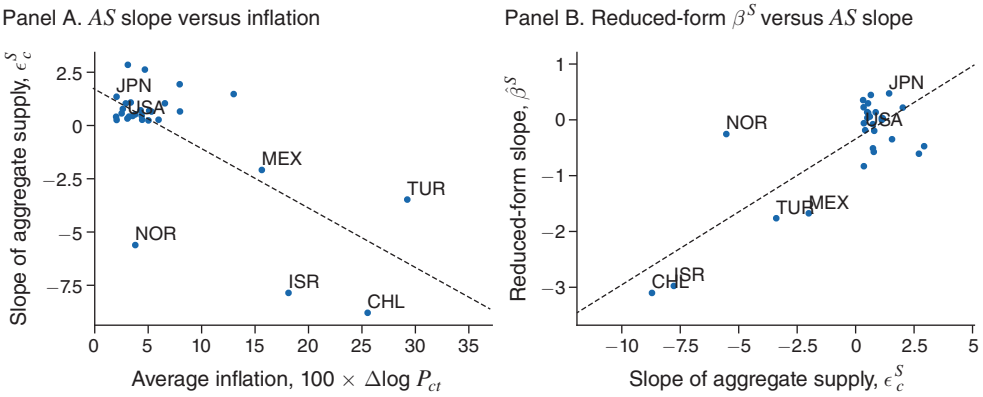

FIGURE 5. INTERNATIONAL EVIDENCE ON THE SLOPE OF AGGREGATE SUPPLY

Notes: This figure summarizes our international estimates of the slope of aggregate supply. In each panel, an observation is an OECD country and the country-level slopes of aggregate supply are calculated using the methods described in Section IVC. Panel A plots the relationship of the model-implied slopes versus mean levels of inflation (GDP deflator) from 1960–2019. Panel B plots the relationship between the model-implied slope of aggregate supply and the "reduced-form" slope estimate, $\hat{\beta}_c^S = \mathrm{Cov}[\Delta \log Y_{ct}, \Delta \log P_{ct}] / \mathrm{Var}[\Delta \log Y_{ct}]$. Dashed lines are from linear regressions and the labels identify selected countries by three-letter codes.

V. 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 within the United States and across countries.

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 nonneutrality 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, 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 and cross-sectional patterns in aggregate supply around the world.

Within the context of supply schedules and the macroeconomy, our study is only a first exploration; there remains much to examine. Recent work has expanded upon our analysis of the macroeconomic implications of supply function equilibrium to consider learning from input markets. Of particular note, Hellwig and Venkateswaran (2025) study a dispersed information model in which firms nonetheless have enough information to implement their full-information pricing policy. Their analysis also shows how our model of supply function competition in general equilibrium can be fruitfully integrated into other dynamic macroeconomic models. Moreover, Nikolakoudis (2025) considers a framework in which firms learn from input prices in a production network. Further exploration of how different market structures affect firms' supply decisions appears a promising avenue for future research.

REFERENCES

Afrouzi, Hassan, Joel P. Flynn, and Choongryul Yang. 2024. "What Can Measured Beliefs Tell Us About Monetary Non-Neutrality?" NBER Working Paper 32541.

Angeletos, George-Marios and Jennifer La'O. 2010. "Noisy Business Cycles." *NBER Macroeconomics Annual* 24 (1): 319–378.

Ball, Laurence, N. Gregory Mankiw, and David Romer. 1988. "The New Keynesian Economics and the Output-Inflation Trade-Off." *Brookings Papers on Economic Activity* 18 (1): 1–82.

Ball, Laurence, and Sandeep Mazumder. 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity* 41 (1): 337–81.

Bils, Mark, and Peter J. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112 (5): 947–85.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2018. "Really Uncertain Business Cycles." *Econometrica* 86 (3): 1031–65.

Bobeica, Elena, and Marek Jarociński. 2019. "Missing Disinflation and Missing Inflation: A VAR Perspective." *International Journal of Central Banking* 15 (1): 199–232.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

Caplin, Andrew S. and Daniel F. Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics* 102 (4): 703–25.

Cerrato, Andrea, and Giulia Gitti. 2022. "Inflation since Covid: Demand or Supply." SSRN Working Paper 4193594.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7 (1): 197–232.

Den Boer, Arnoud V. 2015. "Dynamic Pricing and Learning: Historical Origins, Current Research, and New Directions." *Surveys in Operations Research and Management Science* 20 (1): 1–18.

Eden, Benjamin. 1994. "The Adjustment of Prices to Monetary Shocks when Trade is Uncertain and Sequential." *Journal of Political Economy* 102 (3): 493–509.

Federal Reserve Bank of San Francisco. 2025. *Total Factor Productivity*. https://www.frbsf.org/research-and-insights/data-and-indicators/total-factor-productivity-tfp/ (accessed February 20, 2025).

Feenstra, Robert C., Robert Inklaar and Marcel P. Timmer. 2015a. "The Next Generation of the Penn World Table" *American Economic Review* 105 (10): 3150-3182.

Feenstra, Robert C., Inklaar, Robert, and Timmer, Marcel P. 2015b. *Replication data for: "The Next Generation of the Penn World Table."* Nashville, TN: American Economic Association [publisher], 2015. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2019-10-12. https://doi.org/10.3886/E112937V1.

Flynn, Joel P., George Nikolakoudis, and Karthik A. Sastry. 2024. "Prices vs. Quantities: A Macroeconomic Analysis." Unpublished.

Flynn, Joel P., George Nikolakoudis, and Karthik A. Sastry. 2026. *Data and Code for: "A Theory of Supply Function Choice and Aggregate Supply."* American Economic Association; distributed by Inter-university Consortium for Political and Social Research. https://doi.org/10.3886/E237501V1.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review* 98 (1): 394–425.
Gagliardone, Luca, Mark Gertler, Simone Lenzu, and Joris Tielens. 2025. "Anatomy of the Phillips Curve: Micro Evidence and Macro Implications." *American Economic Review* 115 (11): 3941–74
Golosov, Mikhail, and Robert E. Lucas. 2007. "Menu Costs and Phillips Curves." *Journal of Political Economy* 115 (2): 171–99.
Groningen Growth and Development Centre. 2025. *Penn World Table version 11.0.* DataverseNL, V2 https://doi.org/10.34894/FABVLR.
Grossman, Sanford J. 1981. "Nash Equilibrium and the Industrial Organization of Markets with Large Fixed Costs." *Econometrica* 49 (5): 1149–72.
Hazell, Jonathon, Juan Herreño, Emi Nakamura, and Jón Steinsson. 2022. "The Slope of the Phillips Curve: Evidence from U.S. States." *Quarterly Journal of Economics* 137 (3): 1299–1344.
Hellwig, Christian, and Venky Venkateswaran. 2009. "Setting the Right Prices for the Wrong Reasons." *Journal of Monetary Economics* 56: S57–77.
Hellwig, Christian, and Venky Venkateswaran. 2025. "Dispersed Information, Nominal Rigidities and Monetary Business Cycles: A Hayekian Perspective." *Journal of Monetary Economics* 155:103843.
Hottman, Colin J., Stephen J. Redding, and David E. Weinstein. 2016. "Quantifying the Sources of Firm Heterogeneity." *Quarterly Journal of Economics* 131 (3): 1291–1364.
Klemperer, Paul D., and Margaret A. Meyer. 1989. "Supply Function Equilibria in Oligopoly under Uncertainty." *Econometrica* 57 (6): 1243–77.
Lucas, Robert E. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4 (2): 103–24.
Lucas, Robert E. 1973. "Some International Evidence on Output-Inflation Tradeoffs." *American Economic Review* 63 (3): 326–34.
Lucas, Robert E., and Michael Woodford. 1993. "Real Effects of Monetary Shocks in an Economy with Sequential Purchases." NBER Working Paper 4250.
Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices Under Rational Inattention." *American Economic Review* 99 (3): 769–803.
Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–1328.
McKay, Alisdair, and Christian K. Wolf. 2023. "What Can Time-Series Regressions Tell Us about Policy Counterfactuals?" *Econometrica* 91 (5): 1695–1725.
Nikolakoudis, George. 2025. "Incomplete Information in Production Networks." Unpublished.

Phillips, Robert. 2012. "Why Are Prices Set the Way They Are?" In *The Oxford Handbook of Pricing Management*, edited by Özalp Özer and Robert Phillips, 13–44. Oxford University Press.
Ramey, V. A. 2016. "Macroeconomic Shocks and their Propagation." In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 71–162. Elsevier.
Reis, Ricardo. 2006. "Inattentive Producers." *Review of Economic Studies* 73 (3): 793–821.
Taylor, John B. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy* 88 (1): 1–23.
US Bureau of Economic Analysis. 2025a. *Gross Domestic Product: Implicit Price Deflator* [*GDPDEF*]. Retrieved from FRED, Federal Reserve Bank of St. Louis. https://red.stlouisfed.org/series/GDPDEF (accessed July 7, 2025).
US Bureau of Economic Analysis. 2025b. *Real Gross Domestic Product* [*GDPC1*]. Retrieved from FRED, Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/GDPC1 (accessed July 7, 2025).
Vives, Xavier. 2017. "Endogenous Public Information and Welfare in Market Games." *Review of Economic Studies* 84 (2): 935–63.
Woodford, Michael. 2003a. "Imperfect Common Knowledge and the Effects of Monetary Policy" In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, edited by Philippe Aghion, Roman Frydman, Joseph E. Stiglitz, and Michael Woodford, 25–58. Princeton University Press, Princeton.
Woodford, Michael. 2003b. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.