

Bates College

SCARAB

All Faculty Scholarship

Departments and Programs

7-1-2019

LEARNING, HEDGING, AND THE NATURAL RATE HYPOTHESIS

Thomas E. Cone

The College at Brockport, State University of New York

Paul Shea

Bates College, pshea@bates.edu

Follow this and additional works at: https://scarab.bates.edu/faculty_publications

Recommended Citation

Cone, T. and Shea, P. 2019. "Learning, Hedging, and the Natural Rate Hypothesis." *Macroeconomic Dynamics*. 23(5): 2009-2034. <https://doi.org/10.1017/S1365100517000578>

This Article is brought to you for free and open access by the Departments and Programs at SCARAB. It has been accepted for inclusion in All Faculty Scholarship by an authorized administrator of SCARAB. For more information, please contact batesscarab@bates.edu.

Learning, Hedging, and the Natural Rate Hypothesis*

Thomas E. Cone[†]

Paul Shea[‡]

SUNY Brockport

Bates College

May 19, 2017

Abstract

We assume that firms are more risk averse than households and that they manage their risk through a financial sector, which consists of learning and hedging. Firms that learn (by observing demand shocks) face less uncertainty and produce more than firms that hedge (by selling future production at a fixed price). If a policy or parameter change stabilizes the economy, then there is less learning and usually less production. Welfare, however, is usually maximized when the financial sector, which requires inputs but does not directly provide utility or affect production, is smallest. Monetary policy can improve welfare by either taxing learning or subsidizing hedging. If firms are risk averse over nominal profits instead of real profits, then interest rate policy can also improve welfare by stabilizing prices and thus minimizing the size of the financial sector.

JEL Classification: E32, E38, E52, E58.

Keywords: risk management, bounded rationality, heterogeneous beliefs, optimal monetary policy.

*We thank two anonymous referees, George Evans, as well as seminar participants at the Stanford Institute for Theoretical Economics, SUNY Brockport, the University of Kentucky, and Bates College for helpful comments. All errors are the authors' responsibility.

[†]E-mail: tcone@brockport.edu

[‡]Email: pshea@bates.edu

1 Introduction

Contemporary business cycle models typically assume that firms do not actively manage risk. This paper develops a framework where demand-shocked firms exhibit additional risk aversion, and where they use learning (acquiring information) and hedging (sharing risk) as alternative risk management strategies. Our approach has implications for monetary policy which differ from those of standard frameworks such as the New Keynesian framework. Monetary policy has unusually large effects on mean consumption and output that dominate the more common stabilization motive.

The monetary authority would ideally like to attain high levels of average consumption, but it would also like to minimize the scope of financial activities (learning and hedging) that use resources but do not directly yield utility. When firms are risk averse over real profits, then monetary policy is limited to taxing or subsidizing financial instruments. The optimal policy is to heavily tax learning. Doing so sacrifices a small amount of mean production and consumption in exchange for minimizing the size of the financial sector. If firms are instead risk averse over nominal profits, then an interest rate rule that stabilizes prices is optimal even though it entails a large loss of production and consumption in order to minimize financial activities.

Firms are likely risk averse both because they incorporate the concavity of households' utility functions and for other reasons as identified by the related literature. Econometric evidence from Hall and Liebman (1998) and Guay (1999) suggests that compensation schemes induce firm managers to exhibit excess risk aversion. Tufano (1996) studies the gold-mining industry and finds that the degree of firm risk aversion increases with the amount of stock options held by the firm's manager. Other explanations include either the effect of risk on external finance or heightened costs of financial distress.¹

The evidence is also clear that firms actively manage risk. Bodnar, Hayt, and Marston (1998) find that 50% of corporations, and 83% of large ones, hedge risk through the use of derivatives. Geczy, Minton, and Schrand (1997) find that large corporations also hedge against exchange rate risk through the use of currency derivatives. May (1995) examines the risk reduction behavior of American corporations and finds that those with CEOs who have significant wealth vested in the firm engage in more risk reduction.

Our framework introduces a financial sector that allows firms to actively manage their risk through learning and hedging. Both types of financial activities provide disutility by requiring labor supply, and neither directly contributes to production or consumption. Learning firms pay a cost to contemporaneously observe the economy's stochastic demand shocks. Hedging firms cannot observe these shocks, but instead manage risk by selling some of their production at a pre-determined price. We

¹See, for example, Amihud and Lev (1981) and Choudhary and Levine (2010).

introduce this learning-hedging choice into a micro-founded model of monopolistic competition.

When a parameter or policy change affects volatility, we identify three effects. The first two affect firms' production which equals households' consumption. First, as profits become more stable, hedging firms respond by increasing their production. We denote this as the *supply side effect*. Second, more stable profits incentivize fewer firms to learn. But learning firms produce more than hedging firms on average, and production thus declines. We denote this as the *learning effect*. It works in the opposite direction of the supply side effect. The third effect refers to the change in financial activities that results from less learning and changes in hedging. We denote this as the *financialization effect*. Throughout the paper, the size of the financialization effect closely tracks changes in learning as opposed to hedging.

Whenever a policy or parameter change stabilizes profits, we find that firms invest less in learning and the financial sector shrinks. These changes, however, cause output to exhibit a less efficient response to demand shocks. Stabilization, however, is usually of secondary importance. Because financial activities require labor which yields disutility, optimality usually occurs when the financial sector is minimized. We identify several cases where this entails less mean production and consumption.

For most of the paper, we assume that firms are risk averse over real profits. Because the model includes no nominal distortion, the role of monetary policy is limited to taxing or subsidizing financial instruments. Here, the policy maker can improve welfare by disincentivizing learning. One approach is to subsidize hedging. This produces both a small increase in mean production and a smaller financial sector. It is better, however, to tax learning (at a rate of at least 25% in the baseline case). Here, optimal policy accepts lower levels of both production and output² by minimizing volatility in exchange for less risk management activity.

We also consider a case where firms are risk averse over nominal profits. Here, the optimal policy is an interest rate rule that achieves a highly stable price level. Although this is similar to a basic New Keynesian model, the motivation is very different.³ The monetary authority accepts a large reduction in production in order to reduce learning. The financialization effect is again critical.

Our paper contributes to a literature that suggests that monetary policy may have supply side effects by stabilizing the business climate.⁴ For example, in the theory of Huizinga (1993), increased

²Because output includes the financial sector, we use the term "production" (Y_t) to refer only to consumption goods produced by firms. Production and household consumption are equal.

³See Benigno and Paciello (2014) for helpful discussion of this issue in the New Keynesian framework. In a basic New Keynesian model, price stabilization minimizes inefficient price dispersion and is thus optimal. This motivation dominates any effect on mean output which is often assumed to be at its efficient level through the use of a subsidy on employment.

⁴For a non-academic example, see Edward Prescott, "Five Macroeconomic Myths." *Wall Street Journal*, 12/11/06.

volatility makes the profitability of current investment decisions less predictable, which lowers investment and therefore output growth. Other papers have demonstrated that monetary policy may have supply side effects through very different mechanisms. The most common approach is to assume that firms must carry sufficient working capital to pay their wage bill.⁵ In these papers, however, the supply side effects of monetary policy on output are limited to its second moment. Our paper, however, identifies a novel motivation: monetary policy should focus on minimizing the average level of financial activities.

The paper is organized as follows. Section 2 develops our basic model. Section 3 introduces our calibrations and studies monetary policy. Section 4 discusses the effects of several alternate modeling approaches. Section 5 concludes.

2 Model

Households

Our modeling of households is standard. A unit continuum of identical households maximize utility through consumption of a composite good (C_t) and the supply of labor (N_t). Households take the nominal wage (W_t) and price index (P_t) as given. We assume the following functional form of the representative household's utility function.⁶

$$\text{Max}_{C_t, N_t} E_t \left[\sum_{i=0}^{\infty} \beta^i \left(\frac{C_{t+i}^{1-\sigma}}{(1-\sigma)U_{t+i}} - \gamma N_{t+i} \right) \right] \quad (1)$$

s.t.

$$b_t + P_t C_t = R_{t-1} b_{t-1} + N_t W_t + \Pi_t. \quad (2)$$

U_t is an exogenous *iid* preference shock that is uniformly distributed between $1 - \frac{\eta}{2}$ and $1 + \frac{\eta}{2}$ where $\eta \in (0, 2)$. Firm profits, Π_t , are returned to the representative household, which takes them as given. Households may purchase riskless one-period bonds (b_t) from each other at the interest rate ($R_t - 1$), but in equilibrium, $b_t = 0$. We assume that the monetary authority is able to set R_t at the level of its choosing.

⁵See, for example, Christiano, Eichenbaum, and Evans (1997), and Cooley and Nam (1998).

⁶In addition, the representative's household's objective function may also include real money balances. The resulting money demand equation then allows the monetary authority to determine the equilibrium interest rate. Because money plays no additional role in the model, however, we instead simply assume that the monetary authority may directly determine the nominal interest rate.

Optimization yields a standard Euler Equation and a first-order condition for labor supply.

$$\frac{1}{C_t^\sigma U_t} = E_t \left[\frac{\beta R_t}{\pi_{t+1} C_{t+1}^\sigma U_{t+1}} \right] \quad (3)$$

$$\frac{W_t}{U_t C_t^\sigma P_t} = \gamma \quad (4)$$

where $\pi_{t+1} \equiv P_{t+1}/P_t$.

Firms

Consumption goods are produced by a unit continuum of *ex-ante* identical firms. Each firm i produces a differentiated good; their goods are then costlessly transformed into a final composite good using the standard indices:⁷

$$Y_t = \left(\int_0^1 Y_{i,t}^{\frac{\epsilon-1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}} \quad (5)$$

$$P_t = \left(\int_0^1 P_{i,t}^{1-\epsilon} di \right)^{\frac{1}{1-\epsilon}} \quad (6)$$

with $\epsilon > 1$. The inverse demand function for firm i 's production is:

$$P_{i,t} = P_t \left(\frac{Y_t}{Y_{i,t}} \right)^{1/\epsilon} . \quad (7)$$

Firms choose their risk management strategies, and possibly their production, prior to observing a demand shock that affects the price of their production. There are three methods of risk management. First, firms may eliminate risk by acquiring better information, performing market research, or employing consultants who are able to reduce uncertainty, etc. We refer to this type of risk management as *learning*. Firms that learn employ a fixed amount of labor to observe the demand shock before choosing their production. Second, firms can *hedge*. Hedging firms do not observe the shock, but instead contract to sell a portion of their production at its expected price. We refer to both learning and hedging as “financial” activities. Third, firms may reduce their risk by decreasing their production.

We denote a learning firm's information set at the start of date t by $I_{l,t}$ and a hedging firm's by $I_{h,t}$. If a firm learns, its information set includes U_t . Because the model exhibits no serial correlation, past variables are irrelevant at the start of date t , so we suppress them from the information sets.

The following process determines production. First, all firms choose whether or not to learn. The act of learning requires that a firm employ a fixed amount of labor, κ , beyond what is needed for

⁷One interpretation of the composite good is that a competitive retail sector costlessly transforms intermediate goods into final goods using (5). See Woodford (2003).

production. Learning does not affect the amount of output available for consumption, and its direct effect on the representative household's utility is thus limited to increased labor supply. For notational convenience, we refer to all firms that choose not to learn as hedging firms, even though they may choose to hedge none of their output.

Hedging firms then make their choices without knowing the realized value of the preference shock, U_t . We assume that firms choose their production and that prices then endogenously adjust to clear the market. Section 4 briefly discusses a version of the model where firms instead choose their prices and where production endogenously adjusts. The results are similar.

In addition to choosing their production ($Y_{h,t}$), hedgers also choose how much of their production to hedge (Z_t), taking the per-unit price of hedging, h_t , as given (we model the hedging supply sector later in this section). Thus, $I_{h,t} = \{\text{the model, } Y_{h,t}, Z_t, h_t\}$; as noted later, stationarity makes h_t/P_t a known constant in equilibrium. By hedging, a firm assures itself of receiving the rational expectation of the price of a hedging firm's production, $E_{t-1}[P_{h,t}]$, per unit of production that it hedges.⁸ (Because goods are not perfect substitutes, and hedging and learning firms generically produce different quantities, the price of hedging firms' and learning firms' production are generally not equal.)

The demand shock is then realized and learning firms choose their production knowing U_t . Throughout this section, we assume that learning eliminates all uncertainty, although Section 4 discusses an alternate case where learning provides only a noisy signal of U_t .⁹ Since U_t is the only shock, observing it lets learners infer the values of all endogenous variables, thus $I_{l,t} = \{\text{the model, } Y_{l,t}, U_t, Y_t, P_t, W_t, R_t, I_{h,t}\}$.

Because our topic involves risk management, we need a tractable way of incorporating risk aversion into the model. We assume all firms have mean-variance preferences, an assumption frequently made in the economics literature (*e.g.* Parke and Waters (2007), and Branch and Evans (2011)) and the standard assumption in much of the finance literature. In addition to simply being a tractable way to model risk management, the inclusion of mean-variance preferences in our model may thus be interpreted as the additional risk aversion exhibited by firms compared to households. This interpretation is defensible if firm managers are compensated in a manner that creates additional risk aversion. Persuasive empirical evidence from Hall and Liebman (1998), and Guay (1999) suggests that stock options have caused many managers to exhibit excess risk aversion. Other explanations for this type

⁸This approach is identical to hedging firms contracting in advance to sell their hedged output at a price of $E_{t-1}[P_{h,t}] - h_t$. Notably, the price of hedged output equals the price of the average hedging firm's production, not the specific firm that contracts to hedge its output. This distinction prevents hedging firms from attempting to profit by producing additional output in order to drive their firm specific price below that of the average hedging firm.

⁹This ordering of the production process creates an informational structure where firms cannot extract information about the demand shock based on the behavior of their competitors. Section 4 discusses alternate, more complex, version of the model where each firm's demand depends both on a fully observable aggregate shock and a firm-specific unobservable shock.

of risk aversion may include uncertainty over training, and search costs.¹⁰

Hedging Firms

The representative hedging firm thus solves the following problem:

$$\text{Max}_{Y_{i,t}, Z_t} E_{t-1} \left[\frac{\Pi_t^h}{P_t} - \phi \text{Var}(\Pi_t^h / P_t | I_{h,t}) \right] \quad (8)$$

$$\Pi_t^h = Y_{i,t}^{\frac{\epsilon-1}{\epsilon}} P_t Y_t^{1/\epsilon} - c Y_{i,t} W_t - Z_t (h_t - E_{t-1}[P_{i,t}] + P_{i,t}) \quad (9)$$

$$Y_{i,t} = \frac{N_{i,t}}{c} \quad (10)$$

where Equation (7) has been used in the second expression, c is the constant amount of labor needed to produce one unit, and ϕ represents the level of risk aversion. In equilibrium, $P_{i,t} = P_{h,t}$ for a hedging firm i , where $P_{h,t}$ is the common price of hedging firms' production.

Here, firms are averse to the variance of real profits (Π_t^h / P_t). As a result, monetary policy cannot affect the real economy through interest rate policy. We show, however, in Section 3.1 that non-conventional monetary policy has important effects through the taxation or subsidizing of financial instruments. Section 3.2 also shows that if firms instead care about the variance of nominal profits, then they are then averse to inflation risk and interest rate policy has important effects.

The representative hedger's revenue may be broken down into the certain hedged component $Z_t E_{t-1}[P_{h,t}]$, and the uncertain non-hedged component, $(Y_{i,t} - Z_t) P_{h,t}$. The firm's wage bill $c Y_{i,t} W_t$ is uncertain, while the firm's hedging costs are known. Inserting Equations (4) and (7) into Equation (9), and using $Y_t = C_t$, demonstrates that the random component of the hedger's real profits equals $Y_{i,t}^{-1/\epsilon} (Y_{i,t} - Z_t) Y_t^{1/\epsilon} - c \gamma Y_{i,t} U_t Y_t$. Decomposing the variance allows us to rewrite the hedger's optimization problem.

$$\text{Max}_{Y_{i,t}, Z_t} E_{t-1} \left[\frac{\Pi_t^h}{P_t} - \phi Y_{i,t}^{-2/\epsilon} (Y_{i,t} - Z_t)^2 \Omega_1 - \phi c^2 \gamma^2 Y_{i,t}^2 \Omega_2 + 2 \phi c \gamma Y_{i,t}^{\frac{\epsilon-1}{\epsilon}} (Y_{i,t} - Z_t) \Omega_3 \right] \quad (11)$$

where $\Omega_1 = \text{Var}(Y_t^{1/\epsilon} | I_{h,t})$, $\Omega_2 = \text{Var}(Y_t^\sigma U_t | I_{h,t})$, and $\Omega_3 = \text{Covar}(Y_t^{1/\epsilon}, Y_t^\sigma U_t | I_{h,t})$.

Ω_1 , Ω_2 , and Ω_3 are constants conditional on a hedging firm's information set. Differentiating with respect to $Y_{i,t}$ and Z_t yields two first-order conditions:

¹⁰See Choudhary and Levine (2009).

$$\begin{aligned} & \frac{\epsilon-1}{\epsilon} Y_{h,t}^{-\frac{1}{\epsilon}} E_{t-1}[Y_t^{\frac{1}{\epsilon}}] - c\gamma E_{t-1}[Y_t U_t] - \phi \left(\frac{2(\epsilon-1)}{\epsilon} Y_{h,t}^{\frac{\epsilon-2}{\epsilon}} - \frac{2(\epsilon-2)}{\epsilon} Y_{h,t}^{-\frac{2}{\epsilon}} Z_t - \frac{2}{\epsilon} Y_{h,t}^{-\frac{2-\epsilon}{\epsilon}} Z_t^2 \right) \Omega_1 \dots \\ & - 2\phi c^2 \gamma^2 Y_{h,t} \Omega_2 + 2\phi c \gamma \left(\frac{2\epsilon-1}{\epsilon} Y_{h,t}^{\frac{\epsilon-1}{\epsilon}} - \frac{\epsilon-1}{\epsilon} Y_{h,t}^{-1/\epsilon} Z_t \right) \Omega_3 = 0 \end{aligned} \quad (12)$$

$$\frac{h_t}{2P_t} - \phi \Omega_1 \left(Y_{h,t}^{\frac{\epsilon-2}{\epsilon}} - Y_{h,t}^{-\frac{2}{\epsilon}} Z_t \right) + \phi c \gamma Y_{h,t}^{\frac{\epsilon-1}{\epsilon}} \Omega_3 = 0. \quad (13)$$

In our simulations, the main source of risk is the wage bill, represented by Ω_2 .¹¹ Equation (12) shows that increasing production therefore increases hedging firms' exposure to this risk. Parameter or policy changes that stabilize wages (and thus reduce Ω_2) therefore provide an incentive to both become a hedging firm, and for hedging firms to produce more.

Learning Firms

The representative learning firm's problem is simpler because the informational structure implies that learning firms can solve for W_t , Y_t , and P_t while choosing their production. *Ex ante*, all firms have the same mean-variance objective function. Because learning causes the conditional variance to equal zero, we suppress this term while presenting the optimization problem:

$$Max_{Y_{i,t}} Y_{i,t}^{\frac{\epsilon-1}{\epsilon}} P_t Y_t^{1/\epsilon} - W_t \kappa - c W_t Y_{i,t} \quad (14)$$

where κ is the amount of labor needed to learn. Note that (14) is the problem a learning firm solves after it has observed U_t . *A priori*, there are two possible dates at which the variance might be deleterious in firms' objective function. One possibility is that the relevant variance is the variance before firms choose their output, and the other is the variance after they choose their output and while they're waiting for variables such as P and W , etc., to be realized. In the absence of any reason to prefer one or the other, we assume the second variance is the one that influences firm decisions.¹² At the start of the period, when firms are deciding whether to learn, they know they will not suffer from the uncertainty if they choose to learn. Below, this leads to the disappearance of the variance term from the right side of (16), which equates *ex ante* hedging and learning objectives.

Problem (14) yields the following first-order condition:

$$Y_{i,t} = \left(\frac{\epsilon-1}{\epsilon c \gamma} \right)^\epsilon Y_t^{1-\epsilon} U_t^{-\epsilon} \quad (15)$$

¹¹We later calibrate $\epsilon = 7.67$. $Y^{\frac{1}{\epsilon}}$ is therefore always close to 1 and Ω_1 , which includes $Var(Y_t^{\frac{1}{\epsilon}})$, is thus small.

¹²There is also this argument in favor of our approach: Before firms have made decisions on output they retain their flexibility, so it is plausible that uncertainty is not deleterious at that time. After they have chosen output, they have lost their flexibility and thus feel the full weight of the uncertainty.

where $Y_{l,t}$ is a learning firm's output.

To endogenize q_t , the fraction of firms that learn, we equate the expected mean-variance profits of learning and hedging. Any interior value of q_t satisfies:

$$E_{t-1} \left[\frac{\Pi_t^h}{P_t} - \phi \text{Var}(\Pi_t^h / P_t | I_{h,t}) \right] = E_{t-1} \left[\frac{\Pi_t^l}{P_t} \right] \quad (16)$$

where the conditional variance of profits for learning firms is zero and is thus omitted.

Hedging Suppliers and Equilibrium

We assume the existence of a separate set of firms that are barred from producing. Instead, these hedging suppliers agree to purchase the output of hedging producers prior to observing U_t . Hedging suppliers pay a price $E_{t-1}[P_{h,t}]$ per unit of output, and also receive the hedging price, h_t , per unit purchased. Having purchased this output, hedgers then sell it on the open market and return their profits to the representative household.¹³

We assume that hedging suppliers have the same mean-variance objective function as firms that produce. Hedging is a method of sharing risk within the economy. Because hedging firms produce while hedging suppliers do not, the former initially assume more risk than the latter. Like learning, the act of hedging requires the employment of labor. We assume that each hedging supplier incurs a fixed cost, ιW_t , and must pay variable labor costs that are proportional to the amount of hedged production that it purchases, $v Z_{i,t} W_t$.¹⁴ Like producers that hedge, hedging suppliers face risk from both the wage rate and the equilibrium price of hedgers' production. The optimization problem for a hedging supplier is thus:

$$\begin{aligned} & \text{Max}_{Z_{i,t}} E_{t-1} \left[\frac{h_t Z_{i,t} - \iota W_t - v Z_{i,t} W_t}{P_t} \dots \right. \\ & \left. - \phi P_t^{-2} (Z_{i,t}^2 \text{Var}(P_{h,t} | I_{h,t}) + \phi (v Z_{i,t} + \iota)^2 \text{Var}(W_t | I_{h,t}) - 2\phi Z_{i,t} (v Z_{i,t} + \iota) \text{Cov}(P_{h,t}, W_t | I_{h,t})) \right]. \quad (17) \end{aligned}$$

Manipulation of (17) and imposing a zero-profits condition, which determines the number of hedging suppliers, yields the following equations. The details of this derivation are shown in Appendix 1.

$$E_{t-1} \left[\frac{h}{P_t} \right] = E_{t-1} \left[\gamma v Y_t U_t + 2\phi \gamma \iota (\gamma v \Omega_2 - Y_{h,t}^{\frac{-1}{\epsilon}} \Omega_3) \dots \right]$$

¹³It makes no difference, however, if hedging suppliers and hedging producers engage only in paper transactions or if hedging suppliers directly return Z_t to the representative household.

¹⁴Without a small fixed cost, the number of hedging suppliers will approach infinity, and the amount that each supplies and the price of hedging approach zero.

$$+2\sqrt{\phi\iota\gamma(Y_{h,t}^{\frac{-2}{\epsilon}}\Omega_1 + \gamma^2v^2\Omega_2 - 2\gamma vY_{h,t}^{\frac{-1}{\epsilon}}\Omega_3)(Y_tU_t + \iota\gamma\phi\Omega_2)} \quad (18)$$

$$Z_{i,t} = E_{t-1} \left[\sqrt{\frac{\iota\gamma(Y_tU_t + \iota\gamma\phi\Omega_2)}{\phi(Y_{h,t}^{\frac{-2}{\epsilon}}\Omega_1 + \gamma^2v^2\Omega_2 - 2\gamma vY_{h,t}^{\frac{-1}{\epsilon}}\Omega_3)}} \right]. \quad (19)$$

Because hedging firms and hedging suppliers are equally risk averse, h/P and Z are largely insensitive to changes in risk. Furthermore, the size of the financial sector closely tracks q_l and does not depend much on the behavior of hedging firms. As a result, the model's results are similar to simply treating h/P as exogenous.

Because hedging and learning require additional labor rather than using production, $Y_{i,t} = C_{i,t}$ and $Y_t = C_t$ in equilibrium. An equilibrium is any sequence of $q_l, Y_{l,t}, Y_{h,t}, P_{l,t}, P_{h,t}, Z_t, Y_t, Z_{i,t}, h_t,$ and P_t such that Equations (3), (5)-(7), (12)-(13), (15)-(16), and (18)-(19) are satisfied. Because the interest rate rule only affects P_t , we assume that $R_t = \beta^{-1}$ for now.

The lack of serial correlation greatly improves the model's tractability. To solve the model, we rely on a first-order log-linearization around a zero inflation steady state where $\pi_t = \frac{P_t}{P_{t-1}} = 1$. We linearize equations that include variables that are time dependent in equilibrium.

Log-linearizing the Euler Equation, (3):

$$\sigma\tilde{Y}_t = E_t[\sigma\tilde{Y}_{t+1} + \tilde{\pi}_{t+1}] - \tilde{R}_t - \tilde{u}_t. \quad (20)$$

Log-linearizing a learning firm's first order condition, (15):

$$\tilde{Y}_{l,t} = (1 - \epsilon)\tilde{Y}_t - \epsilon\tilde{u}_t. \quad (21)$$

Log-linearizing the output index, (5):

$$\bar{Y}^{\frac{\epsilon-1}{\epsilon}}\tilde{Y}_t = q_l\bar{Y}_l^{\frac{\epsilon-1}{\epsilon}}\tilde{Y}_{l,t} + (1 - q_l)\bar{Y}_h^{\frac{\epsilon-1}{\epsilon}}\tilde{Y}_{h,t} \quad (22)$$

where \bar{X} is the steady state value of X_t and $\tilde{X}_t = \frac{X_t - \bar{X}}{\bar{X}}$.

The lack of serial correlation allows us to treat $Y_{h,t}, Z_t, Z_{i,t}, h, \Omega_1, \Omega_2, \Omega_3,$ and q_l as constants. We therefore do not need to log-linearize the first order conditions for hedging firms or hedging suppliers.

Notably, total output includes not just production but financial activities. Defining the financial sector to include both risk management activities (learning and hedging), total output equals:

$$GDP_t = Y_t + P_t^{-1}[q_l\kappa W_t + (1 - q_l)h_tZ_t]. \quad (23)$$

The financial sector is included in output but it does not produce production that is consumed. It does, however, require additional labor. Typically, the financial sector will be too large, resulting in reduced welfare through increased labor disutility.

Combining Equations (15), (21), and (22), and using the method of undetermined coefficients yields a log-linearized expression for output.¹⁵

$$\tilde{Y}_t = \frac{\epsilon q_l}{(1 - \epsilon)q_l - \left(\frac{\epsilon-1}{\epsilon c \gamma}\right)^{1-\epsilon} \bar{Y}^{\epsilon-1}} \tilde{u}_t = \tau \tilde{u}_t \quad (24)$$

If $q_l = 0$, then no firms observe the shock and $\tau = 0$ so that production is constant. As q_l approaches one, \bar{Y} approaches $\left(\frac{\epsilon-1}{\epsilon c \gamma}\right)$, and τ approaches -1 so that demand shocks are fully passed on to production. Restating Equation (20) and taking the informed expectation:

$$\tilde{Y}_t = -\frac{\tilde{P}_t + \tilde{u}_t}{\sigma} \quad (25)$$

$$\tilde{P}_t = -(\tau\sigma + 1)\tilde{u}_t \quad (26)$$

Ω_1 , Ω_2 , and Ω_3 may then be approximated as complicated functions of ϵ , η , τ , and \bar{Y} through integration by parts using (24). Solving for equilibrium thus consists of jointly calculating the model's steady state, τ , Ω_1 , Ω_2 , and Ω_3 .

Because different parameterizations or policies affect the steady state, we cannot analyze welfare by comparing the second moment of output and inflation. Instead, we integrate the representative household's utility function, (1), over the distribution of \tilde{u}_t in order to report expected utility.

3 Results

We begin by considering different exogenous values of q_l . The driving factor behind these results is the volatility of the wage bill, which includes $Var(Y_t^\sigma U_t)$. Hedging firms are unable to adjust their output in response to preference shocks. When q_l is low, the economy is mostly hedgers and Y_t thus does not respond much to U_t . τ in (24) is small and the wage bill is relatively volatile. Because hedging firms are harmed by this risk, learning is preferable to hedging when q_l is low. As q_l increases, output has an increasingly strong negative relationship with U_t because learners choose to reduce their output when U_t is high. As $q_l \rightarrow 1$, $\tau \rightarrow -1$, and the wage bill becomes more stable. Higher values of

¹⁵The solution represented by Equation (24) is not necessarily unique, but instead represents the model's minimum state variable solution. We are able to rule out alternate solutions were \tilde{y}_t may be written as an ARMA process. We cannot, however, completely rule out the possibility of more exotic, alternate solutions.

q_l thus make learning less preferable. An equilibrium generally exists for an interior value of q_l where learners and hedgers expected profits are equal.

Our calibration sets $\beta = 0.99$. We follow Woodford (2003) by setting $\epsilon = 7.67$, which implies a markup over marginal cost of 16%. We set $\phi = 1$ so that the mean and variance of profits are of equal importance. We set $\sigma = 1$ so that utility is logarithmic in consumption. We set $c = 1$, and $\gamma = 1$. Increasing either of these parameters reduces steady state output.¹⁶ We set $\eta = 0.2$, implying that preference shocks have a variance equal to 0.0033.¹⁷ For our baseline case, we set $\kappa = 1/3$, $\iota = 0.02$, and $\nu = 0.01$. Because ϕ , η , κ , ι , and ν do not have well established calibrations, we later consider the effects of different values for these parameters.

For this calibration, Figure 1 shows if q_l is endogenous, then a unique equilibrium exists where $q_l = 14.4\%$.

Figure 1: Expected Learning Profits less Hedging Profits, Exogenous q_l




Figure 2 shows welfare, production, and the financial sector. The bulk of the financial sector is learning rather than hedging. Its size thus increases with q_l . As will typically be the case, welfare is closely connected with minimizing the scope of the financial sector which requires labor but does not directly affect consumption.¹⁸ As a result, optimality occurs when there is no learning. Production increases slowly, increasing by 0.5% as q_l increases from zero to one. Because production is slightly increasing in q_l and output is simply the sum of production and the financial sector, optimality therefore minimizes output.

¹⁶If the input requirements of risk management are adjusted in proportion to the change in steady state output, then our results are qualitatively unaffected by changing either c or γ .

¹⁷Supply shocks where c is stochastic have almost the same effect as our demand shocks.

¹⁸Individually, households are better off supplying labor for financial activities. The reduction in profits, which are returned to households, however makes all other households worse off. All households would be better off if they could

Figure 2: Output, Welfare, and the Financial Sector, Exogenous q_t





Figure 3 shows the behavior of hedging firms. Because unresolved risk provides an incentive to reduce their production, hedgers produce less than learners. Notably, the amount of production that hedgers hedge is largely independent of q_t . This is because both hedgers and hedging suppliers are equally risk averse. They respond similarly to changes in risk, thus keeping the amount of hedging fairly constant.


Figure 3: Hedging Sector, Exogenous q_t



Finally, Figure 4 reports the dynamics around the steady state. As q_t increases, price volatility (measured by $\frac{-(\tau+1)}{\sigma}$) decreases while output volatility (measured by τ) increases. This is because hedging firms, unable to adjust their output in response to U_t , instead must allow their prices to respond. Ω_2 , which represents the volatility of the wage bill, also stabilizes as more firms learn.

coordinate not to supply financial labor.

Figure 4: Dynamics Around the Steady State, Exogenous q_l



3.1 Effects of Varying the Model's Parameters

We now consider the impact on the model of varying some of its parameters. Throughout this subsection, we assume that q_l is endogenous, determined by (16). We first formalize the supply side, learning, and financialization effects mentioned in the Introduction. To see the intuition behind each mechanism, suppose that a parameter or policy change results in more stable profits.

The supply side effect exists because, when profits are more stable: (1) hedging firms face less risk, and (2) lower volatility reduces the price of hedging slightly, so they face lower hedging costs. Both effects incentivize hedging firms to increase production. We quantify the supply side effect as the change in production that results from a parameter change, holding q_l constant at its initial level.

When we allow an endogenous learning-hedging choice, we also find a learning effect. The learning effect occurs because more stable profits reduce the incentive to pay the labor cost of learning and fewer firms thus learn. But because learning eliminates all uncertainty, learning firms produce more on average than hedging firms and the learning effect tends to decrease average aggregate production as profits stabilize. We quantify the learning effect by calculating the change in production that occurs when q_l is endogenous less the change in production that occurs when q_l is exogenous. The learning and supply side effects thus sum to the total change in production.


In addition to including firms' production, aggregate output also includes hedging and learning. These financial activities do not directly yield utility but do directly result in labor disutility. We define the financialization effect as the change in financial activities that occurs as a result of a parameter change. As profits become more stable, the incentive to use financial services becomes smaller. The size of the financial sector closely tracks q_l .

As we alter the model's parameters a few results stand out across all of the exercises that follow.

First, changes that either add risk or increase risk aversion incentivize more learning. Second, the representative household's expected utility closely tracks the size of the financial sector, in an inverse way, and is maximized when learning is minimized. Third, production exhibits a stronger response to preference shocks when there is more learning.

We begin by varying ϕ , the amount of risk aversion exhibited by firms:

Figure 5: Variable ϕ




The left panel shows the steady state values as ϕ increases while the right panel shows the dynamics around the steady state. Unsurprisingly, as firms become more risk averse, there is more learning. The learning effect thus leads to more production. The firms that hedge, however, are more risk averse and produce less, leading to a negative supply side effect. Overall, production falls by 0.1% as ϕ rises from 0 to 6. Higher values of ϕ always reduce welfare by inducing more learning and a larger financial sector. Recall that ϕ measures firm risk aversion beyond that of the representative household. It is thus intuitive that households do best when there is no additional risk aversion.

We now vary η , the support of the preference shocks:

Higher values of η have similar effects as higher values of ϕ . As risk increases, more firms learn, leading to a positive learning effect. Hedging firms, however, produce less, leading to a negative supply-side effect.¹⁹ Welfare, as expected, is maximized when U_t is a constant. In most macroeconomic models, agents do best with small shocks. Here, however, the cause is not, as is usually

¹⁹Production is minimized when $\eta = 0.229$. Here, it is 0.6% lower than its values either when $\eta = 0$ or $\eta = 0.43$, the largest value we report.


Figure 6: Variable η



the case, based on stabilization. Rather, small shocks induce firms not to invest in costly financial activities.

Now, we alter σ , the inverse intertemporal elasticity of substitution. Because this term appears in the utility function, we do not report the effects on welfare.

Figure 7: Variable σ




The variance of the wage bill includes $Var(Y_t^\sigma U_t)$. Low values of σ cause Y_t^σ to be more stable, and because Y_t is a decreasing function of U_t , this causes the overall wage bill to be more volatile. As a result, as σ increases, firms respond to the resulting stabilization by learning less, which boosts

welfare. Because price volatility is measured by $\frac{-(\tau+1)}{\sigma}$, its dynamics are much more exotic than when $\sigma = 1$.

We refrain from reporting the effects of varying other parameters, including ι , and v because they do not have large effects. Because hedging suppliers are as risk averse as hedging firms, they choose to purchase only a small share of hedgers output for all calibrations that we consider. Varying the labor requirements of hedging thus has only small effects.

Because this version of the model does not include any nominal rigidity or related distortion, the effects of nominal interest rate policy are limited to the price level. There is, however, an important role for non-conventional monetary policy, interpreted here as the taxation or subsidization of financial instruments (learning and hedging). We now analyze taxes on either hedging or learning under the assumption that revenue is returned to the representative household in the form of a lump sum transfer. Subsidies are likewise paid for using a lump sum tax on households. We begin with a tax on learning.²⁰

Figure 8: Learning tax



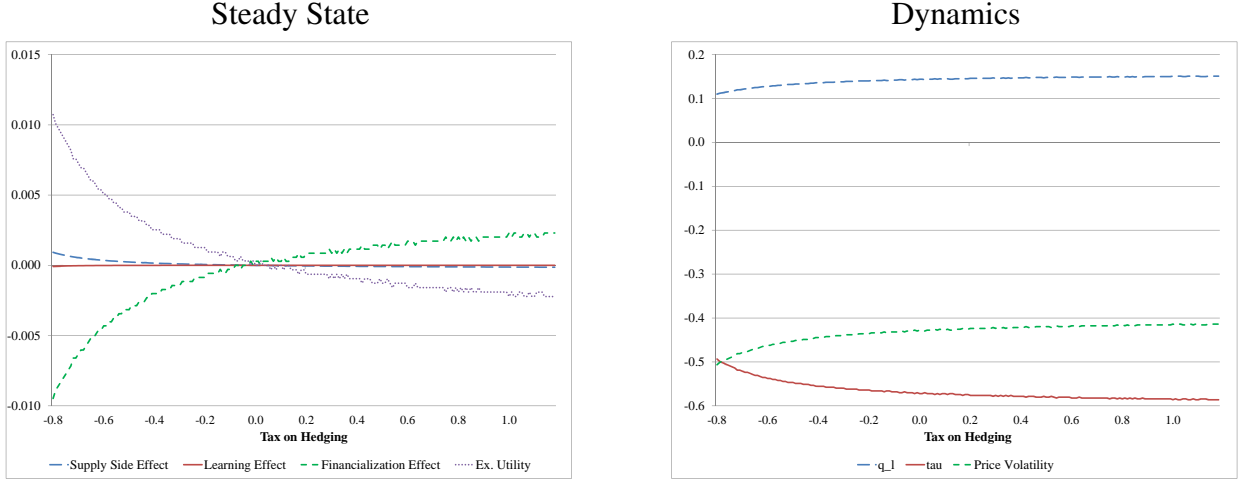
A large tax on learning is desirable because it discourages learning and thus minimizes the size of the financial sector. If this tax is greater than 25%, then there is no learning and welfare is maximized. This is despite a small decline in production as the learning tax increases.

We now report the results for a tax on hedging.

Subsidizing hedging also improves welfare by incentivizing firms to switch from learning to hedging. It is not, however, as effective. Even a 80% subsidy is only able to reduce q_l from its baseline

²⁰The effects of increasing κ , the cost of learning are similar and we thus do not report them.

Figure 9: Hedging tax



value of 14.4% to 11.0%. Optimal monetary policy is thus achieved by minimizing the size of the financial sector by harshly taxing learning by at least 25%.²¹

3.2 Nominal Profit Risk Aversion and Monetary Policy

Because we have assumed that firms are averse to the risk of real profits, interest rate policy has had no real effects in the model. Here, we show that if firms are instead averse to the risk associated with nominal profits, then interest rate policy has important implications in this framework. We replace (8), hedging firms' objective function, with:

$$Max_{Y_{i,t}, Z_t} E_{t-1} \left[\frac{\Pi_t^h}{P_t} - \phi \frac{Var(\Pi_t^h | I_{h,t})}{P_t^2} \right] \quad (27)$$

Equation (27) captures the notion that firms care about inflation risk. Hall and Liebman (1998) argue that stock options cause firm managers to exhibit excess risk aversion. These stock options usually are not indexed to broader stock indices and they rarely, if ever, are indexed to inflation. We thus might expect firms to be risk averse over nominal profits, as in (27), instead of real profits.²²

²¹The nature of optimal policy is the same for all of the alternate calibrations considered in this section. Optimal policy taxes learning in order to drive q_l to zero.

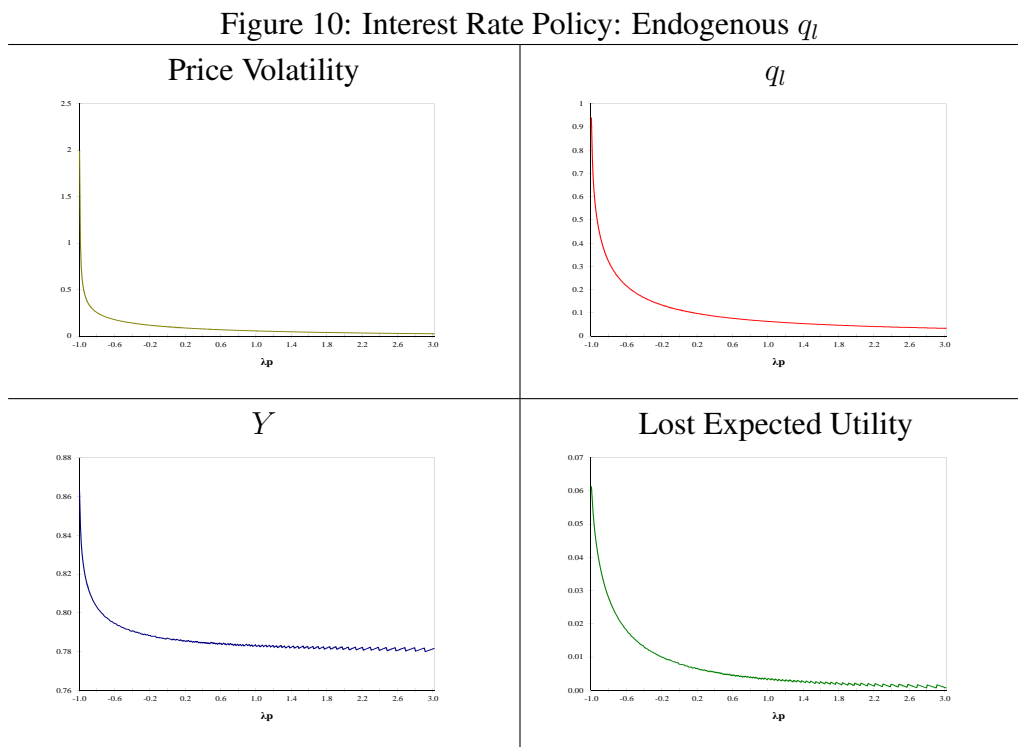
²²By relying on the the ratio of the variance to the squared price level instead of $Var(\Pi_t^h | I_{h,t})/P_t$, we ensure that no real variables depend on the steady state price level, which is indeterminate in the model. When calibrating the model, we normalize the steady state price level to one. The effects of instead using the ratio of the variance to the squared steady state price level are very similar.

Because nominal profits are more volatile than real profits, we reduce the support of the preference shocks so that $\eta = 0.1$ and we reduce risk aversion so that $\phi = 0.5$. The remaining parameters are unchanged from their baseline values. We assume that the monetary authority uses the following interest rate rule:

$$\tilde{R}_t = \lambda_p \tilde{P}_t \quad (28)$$

Because the model exhibits no serial correlation, a policy rule that targets inflation, and thus depends on the lagged price level, only adds noise. We thus assume that interest rates depend only on the price level. Details on solving this version of the model are provided in Appendix 2.


We now report the effects of varying λ_p between -1 and 3 . The bottom-right panel reports the utility loss relative to the optimum ($\lambda_p = 3$), measured as the percentage increase in steady state consumption needed to set expected utility equal to its maximum.



As λ_p increases, the price level stabilizes as shown in the top-left panel. As prices stabilize, fewer firms choose to learn as shown in the top-right panel. Because hedging firms produce less than learners, low values of q_l result in less production. Production thus decreases significantly along with λ_p , as shown in the bottom-left panel. Optimality, however, again depends on minimizing the size of the financial sector. As shown in the bottom-right panel, expected utility converged to its maximum

as learning is minimized as λ_p grows larger.²³ Optimality thus entails minimizing both production and output in order to avoid wasting too many resources on financial activities.²⁴ This can also be seen by breaking down the supply side, learning, and financialization effects. Here, each is measured relative to $\lambda_p = 3$.²⁵

Figure 11: Supply Side, Learning, and Financialization Effects



4 Noisy Signals and Other Alternate Assumptions

This section discusses several alternate modeling approaches. Throughout, our main results, including that volatility increases the size of the financial sector which then reduces welfare, are preserved. We present detailed results for one alternate approach; the case where learning provides only a noisy signal of the demand shock.

#1 Noisy Signals

In Section 3.2, increased price volatility induces more learning, which then results in both more production and a larger financial sector. This result is not inconsistent with the empirical evidence

²³Increasing λ_p above 3 leads to increasingly small increases in utility and decreases in production.

²⁴This result is the same for all of the calibrations discussed in Section 3.1. Optimal interest rate policy minimizes learning by stabilizing prices.

²⁵The oscillatory pattern seen in some of these panels is a remnant of the computational algorithm that optimizes over a discrete set of values for q_l .

for low volatility economies which fails to show a consistent relationship between price volatility and output.²⁶ For high volatility regimes, however, the data show that higher inflation volatility does significantly reduce output.²⁷ We now show that our model can yield this result by making the reasonable assumption that learning yields a signal of U_t , but does not eliminate all uncertainty.

We make several changes to the model, some in the interests of tractability. First, we greatly simplify the model's stochastic structure. We assume that U_t , with probability $\frac{1}{2}$, equals either U^+ or U^- where $U^+ > 1 > U^- > 0$. We set $U^- = \frac{U^+}{2U^+-1}$, which makes average output independent of U^+ . Second, we assume that learners receive a signal that indicates the correct value of U_t with probability 0.9, and the incorrect value with probability 0.1. Third, we assume that monetray policy observes the correct value of U_t and that it reponds directly to the demand shock by setting $R_t = \beta^{-1} + \lambda_u(1 - U^+)$ if $U_t = U^+$ and $R_t = \beta^{-1} - \lambda_u(1 - U^+)$ if $U_t = U^-$. Fourth, we analyze the non-linear model instead of a linear approximation in order to consider large deviations from the model's steady state.

In this version, learners, as well as hedgers, reduce their production as the conditional variance of profits increases. There are four possible equilibrium values of production, corresponding to the two values of U_t and the two signals. Table 1 reports the equilibrium properties for different policies where $U^+ = 1.05$.²⁸

²⁶See Barro (1995) and (1996).

²⁷Temple (2000) discusses the theoretical relationship between price volatility and output. He also discusses the strong correlation between inflation volatility and average inflation that makes them it statistically very difficult to separate their effects.

²⁸ \bar{Y} now reports the average value of Y instead of its non stochastic steady state value.

Table 1: Effects of Monetary Policy with Noisy Signals

λ_u	q_l	\bar{Y}	Y^+	Y^-	Y_h	$Var(P_t)$
0	0.32	0.869	0.837	0.901	0.867	0.007
-0.5	0.32	0.869	0.837	0.901	0.866	0.012
-1	0.32	0.868	0.836	0.900	0.863	0.018
-2	0.34	0.865	0.833	0.898	0.850	0.035
-3	0.37	0.860	0.826	0.894	0.820	0.056
-4	0.44	0.853	0.817	0.890	0.763	0.079
-5	0.54	0.848	0.808	0.887	0.675	0.104
-6	0.68	0.845	0.802	0.888	0.456	0.159
-7	0.75	0.832	0.799	0.868	0.416	0.167
-8	0.89	0.827	0.797	0.858	0.380	0.181
-9	0.98	0.800	0.791	0.809	0.303	0.206
-10	0.99	0.775	0.775	0.774	0.240	0.235
-11	0.99	0.753	0.752	0.754	0.190	0.262
-12	0.99	0.727	0.721	0.734	0.150	0.287
-13	0.99	0.686	0.679	0.694	0.118	0.315
-14	0.99	0.628	0.623	0.634	0.091	0.341
-15	0.99	0.562	0.550	0.574	0.068	0.365
-16	0.99	0.465	0.455	0.475	0.049	0.392
-17	0.99	0.346	0.335	0.356	0.032	0.416

The main effect of this modification is to strengthen the supply side effect. Learners can no longer fully insulate themselves from inflation risk and as prices become more volatile, both hedgers and learners reduce their average production. As λ_u decreases, prices de-stabilize. The supply side effect now dominates the learning effect, and average production thus decreases. The magnitude of the effect, however, is relatively small as long as q_l is below its maximum value.²⁹ Once learning is maximized, however, the learning effect is exhausted and additional volatility begins to rapidly reduce production through the strengthened supply side effect. Whereas increasing the variance of the price level from 0.012 to 0.206 reduces production from 0.869 to 0.800 and learning from $q_l = 0.32$ to $q_l = 0.98$, additional uncertainty results in almost no additional learning. An approximately equal subsequent increase in price volatility now results in a dramatic reduction in average production;

²⁹In this section, expected utility tracks production.

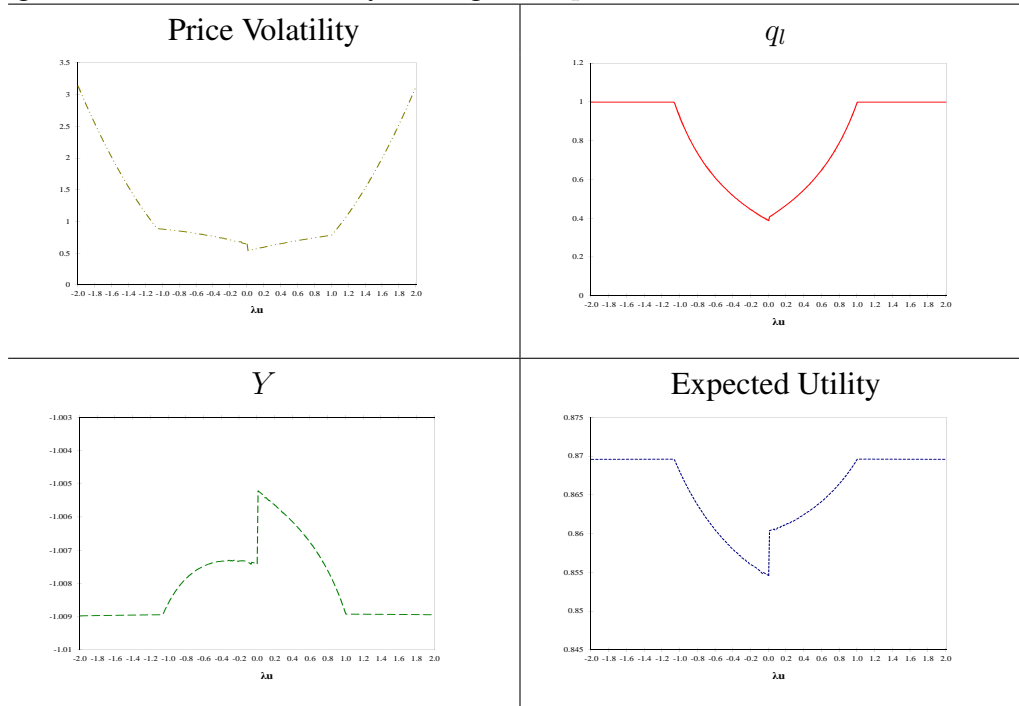
when the variance of the price level equals 0.416, average production only equals 0.346.

The overall pattern thus fits the empirical results. Taken as a whole, the results of this paper provide theoretical ambiguity that helps explain the lack of a strong empirical relationship between inflation volatility and output for low volatility regimes. The non-linear simulation of this section, however, demonstrates a strong negative relationship for high volatility regimes where producers are unable to ever fully insulate themselves from inflation risk.

#2 Firms Choose Prices and Output Adjusts

We also consider a version of the model where firms choose their prices and their output then endogenously adjusts to clear the market.³⁰ The only major change in this version is that hedging, which involves contracting to sell output at its expected price, is no longer well defined. In a version of the model without hedging, however, the results are very similar to Section 3. Parameter changes that either increase risk aversion or risk again induce more learning which reduces welfare. If firms are risk averse over nominal profits, then optimal policy again minimizes price volatility in order to minimize the scope of financial activities. Figure 12 shows these results.³¹

Figure 12: Interest Rate Policy: Endogenous q_t : Model where Firms Choose Prices



#3 Alternate Informational Structures

³⁰This is standard in the New Keynesian literature. See Woodford (2003a).

³¹Here, we assume that interest rates respond directly to the demand shock: $\tilde{R}_t = \lambda_u \tilde{u}_t$.

Our ordering of production, where hedgers producers and hedging suppliers must make their decisions before observing U_t eliminates opportunities for firms to extract information based on other firms' behavior. The primary appeal of this assumption is admittedly its convenience. We now briefly discuss a more complex version of the model that allows for a more plausible informational structure. Suppose that the demand for any firm's output depends both on an observable aggregate demand shock, and a firm specific demand shock. Our risk management framework may then be extended to how firms manage the risk associated with the firm specific shock. The main effect of this modification is that hedging suppliers are able to reduce their risk by purchasing the output of a diversified set of firms. As a result, the equilibrium hedging price is reduced. The effects of this change is to make hedging somewhat cheaper. This has only very small effects on our major results.

#4 Considering Households' Marginal Utility of Consumption

We also consider the model where firm preferences incorporate households' (who own firms), marginal utility of consumption. This replaces (8) with

$$Max_{Y_{i,t}, Z_t} E_{t-1} \left[\frac{\Pi_t^h}{P_t C_t^\sigma} - \phi Var\left(\frac{\Pi_t^h}{P_t C_t^\sigma} | I_{h,t}\right) \right] \quad (29)$$

This modification does not qualitatively affect our results. Learning firms again produce more output than hedging firms. Policy can again improve welfare by taxing learning or subsidizing hedging. And monetary policy maximizes household utility by stabilizing prices in order to minimize the scope of the financial sector.

#5 Cost Shocks

We also consider a version of the model with cost shocks instead of preference shocks. This entails setting $U_t = 1 \forall t$ and assuming a distribution for c_t , the, now stochastic, labor required to produce one unit. Notably, c_t simply replaces U_t in both the representative household's first order conditions, (3) and (4), and that for learners, (15). Inserting (4) into (9) shows that the second term on the right hand side of (11) changes from $\gamma c U_t Y_{i,t} C_t^\sigma P_t$ to $\gamma c_t Y_{i,t} C_t^\sigma P_t$. From the perspective of a hedging firm, this makes no difference. Their wage bill is equally uncertain regardless if the uncertainty comes from a demand shock or a cost shock. Continuing forward, τ has the same solution as in (24) and c_t replaces U_t in the definitions of Ω_1 , Ω_2 , and Ω_3 . Thus, any distribution of U_t produces very similar results for a version of the model where that distribution instead governs c_t .³² It is thus reasonable to think of the variance of U_t as approximately coming from both preference and cost shocks.

³²This change does have small effects when integrating over households' utility function in order to compute welfare. But these effects do not alter any of our major findings.

5 Conclusion

Our paper shows that adding risk management into a business cycle model adds an important motivation where policy makers must consider the size of the financial sector. Increased risk aversion, or changes that add volatility, now have important effects on the steady state levels of production, output, and consumption. Crucially, by encouraging more financial activities (specifically learning), they reduce welfare. Policy makers can improve welfare most effectively by taxing learning or, less effectively, taxing hedging. If firms are risk averse over nominal profits, then interest rate policy is also effective by minimizing price volatility and therefore minimizing the financial sector, even though this also reduces production and consumption.

Our paper focuses on only a few type of financial activities and a few sources of uncertainty. Given the significance of our results for central bankers, it is worthwhile to extend our framework to many of these extensions. We conclude by discussing two. First, if we extend the role of the financial sector so that it acts as an intermediary by extending credit, is it still true that optimal policy minimizes the scope of the financial sector? Second, we assume rational expectations where agents know the model's reduced form solution. This minimizes the amount of uncertainty in the economy which may thus reduce the role of our financial sector. The adaptive learning literature instead assumes that agents revise their beliefs over time using standard econometric techniques. Adaptive learning introduces additional volatility into a model and it is of interest to examine how this extra uncertainty affects risk averse agents.³³ It would be worthwhile to see if our results extend to a version of the model with adaptive learning.

References

- Amihud, Y. and B. Lev. 1981. "Risk Reduction as Managerial Motive for Conglomerate Mergers." *Bell Journal of Economics*, 12: 605-617.
- Barro, R. 1995, "Inflation and Economic Growth." *Bank of England Quarterly Bulletin*, May 1995.
- Barro, R. 1996, "Inflation and Growth." *Federal Reserve Bank of St. Louis Review*, May/June 1996.
- Benigno, P. and L. Paciello. 2014. "Monetary Policy, Doubts, and Asset Prices." *Journal of Monetary Economics*, 34: 85-98.

³³For a general discussion of adaptive learning, see Evans and Honkapohja (2001). For an application of adaptive learning to monetary policy, see Orphanides and Williams (2007).

- Bodnar, Gordon M., Gregory S. Hayt, and Richard C. Marston. 1998. "1998 Wharton Survey of Financial Risk Management by U.S. Non-Financial Firms." *Financial Management* 27: 70-91.
- Branch W.A., and G.W. Evans. 2007. "Model Uncertainty and Endogenous Volatility." *Review of Economic Dynamics* 10: 207-237.
- Branch W.A., and G.W. Evans. 2011. "Learning about Risk and Return: A Simple Model of Bubbles and Crashes." *American Economic Journal: Macroeconomics*, 3(3): 159-191.
- Choudhary, A. and P. Levine. 2010. "Risk Averse Firms and Employment Dynamics." *Oxford Economics Papers*, 62(3): 578-602.
- Christiano, L., Eichenbaum, M., and C. Evans. 1997. "Sticky Price and Limited Participation Models of Money: a Comparison." *European Economic Review*, 41: 1201-1249.
- Clarida, R., Gali, J., and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature*, 37: 1661-1707.
- Cooley, T., and K. Nam. 1998. "Asymmetric Information, Financial Intermediation and Business Cycles." *Economic Theory*, 12: 599-620.
- Evans, G, and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.
- Gali, J. 2008. *Monetary Policy, Inflation, and the Business Cycle*. Princeton. Princeton University Press.
- Geczy, C., Minton, B., and C. Schrand. 1997. "Why Firms Use Currency Derivatives." *Journal of Finance*, 52: 1323-1254.
- Guay, W. 1999. "The Sensitivity of CEO Wealth to Equity Risk: An Analysis of the Magnitude of and Determinates." *Journal of Financial Economics*, 53: 43-71.
- Hall, B. and J. Liebman. 1998. "Are CEOs Really Paid Like Bureaucrats?" *The Quarterly Journal of Economics*, 113(3): 653-690.
- Huizinga, J. 1993. "Inflation Uncertainty, Relative Price Uncertainty, and Investment in U.S. Manufacturing." *The Journal of Money, Credit, and Banking: Inflation Uncertainty: A Conference Sponsored by the Federal Reserve Bank of Cleveland*, 25(3): 521-549.

Karanasos, M., and J. Kim. 2005. “The Inflation-Output Variability Relationship in the G3: A Bivariate GARCH (BEKK) Approach.” *Risk Letters* 1(2): 17-22.

May, D. 1995. “Do Managerial Motives Influence Firm Risk Reduction Strategies.” *Journal of Finance*, 50(4): 1291-1308.

Orphanides, A. and J. Williams. 2007. “Robust Monetary Policy with Imperfect Knowledge.” *Journal of Monetary Economics*, 54: 1407-1435.

Parke, W. and G. Waters. 2007. “An Evolutionary Game Theory Explanation of ARCH.” *Journal of Economic Dynamics and Control*, 31(7), 2234-2262.

Taylor, J. 1980. “Aggregate Dynamics and Staggered Contracts.” *Journal of Political Economy*, 88(1): 1-23.

Temple, J. 2000. “Inflation and Growth: Stories Short and Tall.” *Journal of Economic Surveys*, 14(4): 395-426.

Tufano, P., 1996. “Who Manages Risk? An Empirical Examination of Risk Management Practices in the Gold Mining Industry.” *Journal of Finance*, 51: 1097-1137.

Woodford, M. 2003. *Interest and Prices*. Princeton, NJ and Oxford: Princeton University Press.

Appendix 1: Hedging Suppliers

Using Equations (4) and (7), the representative hedging supplier’s problem may be re-written in terms of Ω_1 , Ω_2 , and Ω_3 .³⁴

$$\text{Max}_{Z_{i,t}} E_{t-1} \left[\frac{h_t Z_{i,t}}{P_t} - \gamma(\iota + v Z_{i,t}) Y_t U_t - \phi(Y_{h,t}^{-\frac{2}{\epsilon}} Z_{i,t}^2 \Omega_1 + \gamma^2 (\iota + v Z_{i,t})^2 \Omega_2 - 2\gamma(\iota + v Z_{i,t}) Y_{h,t}^{-\frac{1}{\epsilon}} Z_{i,t} \Omega_3) \right]. \quad (30)$$

Differentiating with respect to $Z_{i,t}$ yields a first order condition for the representative hedging supplier:

$$Z_{i,t} = E_{t-1} \left[\frac{\frac{h_t}{P_t} - \gamma v Y_t U_t - 2\phi \gamma \iota (\gamma v \Omega_2 - Y_{h,t}^{-\frac{1}{\epsilon}} \Omega_3)}{2\phi(Y_{h,t}^{-\frac{2}{\epsilon}} \Omega_1 + \gamma^2 v^2 \Omega_2 - 2\gamma v Y_{h,t}^{-\frac{1}{\epsilon}} \Omega_3)} \right]. \quad (31)$$

³⁴We treat the output of hedging firms, $Y_{h,t}$, as observable to hedging suppliers.

We impose a zero-profits condition to determine the equilibrium value of $Z_{i,t}$, h_t , and corresponding number of hedging suppliers.

$$E_{t-1}\left[\frac{h}{P_t}\right] = E_{t-1}\left[\gamma v Y_t U_t + 2\phi\gamma\iota(\gamma v\Omega_2 - Y_{h,t}^{\frac{-1}{\epsilon}}\Omega_3)\dots\right. \\ \left.+ 2\sqrt{\phi\iota\gamma(Y_{h,t}^{\frac{-2}{\epsilon}}\Omega_1 + \gamma^2 v^2\Omega_2 - 2\gamma v Y_{h,t}^{\frac{-1}{\epsilon}}\Omega_3)(Y_t U_t + \iota\gamma\phi\Omega_2)}\right] \quad (32)$$

$$Z_{i,t} = E_{t-1}\left[\sqrt{\frac{\iota\gamma(Y_t U_t + \iota\gamma\phi\Omega_2)}{\phi(Y_{h,t}^{\frac{-2}{\epsilon}}\Omega_1 + \gamma^2 v^2\Omega_2 - 2\gamma v Y_{h,t}^{\frac{-1}{\epsilon}}\Omega_3)}}\right]. \quad (33)$$

Appendix 2: The Model with Risk Aversion to Nominal Profits

This appendix provides additional detail on the version of the model from Section 3.2 where firms are risk averse to nominal profits instead of real profits. Equations (1)-(7), which describe the behavior of households, are unchanged as are (14)-(16), which describe the behavior of learners.

Hedging firms face the following optimization problem:

$$Max_{Y_{i,t}, Z_t} E_{t-1}\left[\frac{\Pi_t^h}{P_t} - \phi\frac{Var(\Pi_t^h|I_{h,t})}{P_t^2}\right] \quad (34)$$

$$\Pi_t^h = Y_{i,t}^{\frac{\epsilon-1}{\epsilon}} P_t Y_t^{1/\epsilon} - c Y_{i,t} W_t - Z_t(h_t - E_{t-1}[P_{i,t}] + P_{i,t}) \quad (35)$$

$$Y_{i,t} = \frac{N_{i,t}}{c} \quad (36)$$

Decomposing the variance allows us to rewrite the hedger's optimization problem.

$$Max_{Y_{i,t}, Z_t} E_{t-1}\left[\frac{\Pi_t^h}{P_t} - \phi Y_{i,t}^{-2/\epsilon}(Y_{i,t} - Z_t)^2\Omega_1 - \phi c^2 \gamma^2 Y_{i,t}^2\Omega_2 + 2\phi c \gamma Y_{i,t}^{\frac{\epsilon-1}{\epsilon}}(Y_{i,t} - Z_t)\Omega_3\right] \quad (37)$$

where $\Omega_1 = \frac{Var(P_t Y_t^{1/\epsilon}|I_{h,t})}{P_t^2}$, $\Omega_2 = \frac{Var(P_t Y_t^\sigma U_t|I_{h,t})}{P_t^2}$, and $\Omega_3 = \frac{Covar(P_t Y_t^{1/\epsilon}, P_t Y_t^\sigma U_t|I_{h,t})}{P_t^2}$.

Due to stationarity, Ω_1 , Ω_2 , and Ω_3 are again constants conditional on a hedging firm's information set. Differentiating with respect to $Y_{i,t}$ and Z_t yields two first-order conditions:

$$\frac{\epsilon-1}{\epsilon} Y_{h,t}^{\frac{-1}{\epsilon}} E_{t-1}\left[Y_t^{\frac{1}{\epsilon}}\right] - c \gamma E_{t-1}[Y_t U_t] - \phi\left(\frac{2(\epsilon-1)}{\epsilon} Y_{h,t}^{\frac{\epsilon-2}{\epsilon}} - \frac{2(\epsilon-2)}{\epsilon} Y_{h,t}^{\frac{-2}{\epsilon}} Z_t - \frac{2}{\epsilon} Y_{h,t}^{\frac{-2-\epsilon}{\epsilon}} Z_t^2\right)\Omega_1\dots$$

$$-2\phi c^2 \gamma^2 Y_{h,t} \Omega_2 + 2\phi c \gamma \left(\frac{2\epsilon - 1}{\epsilon} Y_{h,t}^{\frac{\epsilon-1}{\epsilon}} - \frac{\epsilon - 1}{\epsilon} Y_{h,t}^{-1/\epsilon} Z_t \right) \Omega_3 = 0 \quad (38)$$

$$\frac{h_t}{2P_t} - \phi \Omega_1 (Y_{h,t}^{\frac{\epsilon-2}{\epsilon}} - Y_{h,t}^{\frac{-2}{\epsilon}} Z_t) + \phi c \gamma Y_{h,t}^{\frac{\epsilon-1}{\epsilon}} \Omega_3 = 0. \quad (39)$$

Apart from the new formulae for Ω_1 , Ω_2 , and Ω_3 , hedging suppliers face the same problem as in Appendix 1. Solving the model now involves assuming an interest rate rule and using the log-linearized Euler Equation:

$$\tilde{R}_t = \lambda_p \tilde{P}_t \quad (40)$$

$$\sigma \tilde{Y}_t = E_t[\sigma \tilde{Y}_{t+1} + \tilde{\pi}_{t+1}] - \tilde{R}_t - \tilde{u}_t. \quad (41)$$

Combining these, and taking expectations yields:

$$\tilde{P}_t = \frac{-(\tau\sigma + 1)}{1 + \lambda_p} \quad (42)$$

Equation (24) continues to determine the equilibrium value of production. Solving the model again consists of jointly solving for Ω_1 , Ω_2 , and Ω_3 , τ , and the model's steady state.