

The Power of Forward Guidance Revisited

Alisdair McKay
Boston University

Emi Nakamura
Columbia University

Jón Steinsson*
Columbia University

July 14, 2015

Abstract

In recent years, central banks have increasingly turned to “forward guidance” as a central tool of monetary policy, especially as interest rates around the world have hit the zero lower bound. Standard monetary models imply that far future forward guidance is extremely powerful: promises about far future interest rates have huge effects on current economic outcomes, and these effects grow with the horizon of the forward guidance. We show that the power of forward guidance is highly sensitive to the assumption of complete markets. If agents face uninsurable income risk and borrowing constraints, a precautionary savings effect tempers their responses to changes in future interest rates. As a consequence, forward guidance has substantially less power to stimulate the economy. In addition, we show that the business cycle dynamics of our incomplete markets model differ substantially from its complete market counterpart. This contrasts with the well-known results of Krusell and Smith (1998). We present approximate representations that can easily be incorporated into standard business cycle models.

JEL Classification: E40, E50, E21

*We thank Susanto Basu, Gauti Eggertsson, Marc Giannoni, Fatih Guvenen, Nobuhiro Kiyotaki, Benjamin Moll, Michael Peters, Oistein Roisland, Greg Thwaites, Aleh Tsyvinski, Nicolas Werquin and seminar participants at various institutions for valuable comments and discussions. We thank the National Science Foundation (grant SES-1056107) for financial support.

1 Introduction

Forward guidance has become an increasingly important tool of monetary policy in recent years. Gurkaynak, Sack, and Swanson (2005) show that much of the surprise news about monetary policy at the time of FOMC announcements arises from signals about the central bank’s intentions about future monetary policy. In many cases, changes in the current Federal Funds rate are fully expected, and all of the *news* about monetary policy has to do with how the central bank is expected to set interest rates in the future.¹

Promises about future interest rates have been shown to have a powerful effect on the economy in standard monetary models. Eggertsson and Woodford (2003) show that a shock to the natural rate of interest that causes the economy to hit the zero lower bound on nominal interest rates induces a powerful deflationary spiral and a crippling recession. However, the recession can be entirely abated if the central bank commits from the outset to holding interest rates at the zero lower bound for a few additional quarters beyond what is justified by contemporaneous economic conditions.

Recent work argues that the magnitude of the effects of forward guidance in New Keynesian models stretches the limits of credibility. Carlstrom, Fuerst, and Paustian (2012) show that a promise by the central bank to peg interest rates below the natural rate of interest for roughly two years generates explosive dynamics for inflation and output in a workhorse New Keynesian model (the Smets and Wouters (2007) model). Del Negro, Giannoni, and Patterson (2013) refer to this phenomenon as the *forward guidance puzzle*. Along the same lines, consider an experiment whereby the central bank promises a 1 percentage point lower real interest rate *for a single quarter* at some point in the future. We show that in the plain vanilla New Keynesian model, this promise has an eighteen times greater impact on inflation when the promise pertains to interest rates 5 years in the future than when it pertains to the current interest rate.

It may seem unintuitive that an interest rate cut far in the future has a greater effect than a near-term one. To see why this arises in standard models, consider the response of consumption to a decrease in the real interest rate for a single quarter 5 years in the future. The consumption Euler equation dictates that consumption will rise immediately to a higher level and stay constant at that higher level for 5 years before returning to its normal level.² The cumulative response

¹Campbell et al. (2012) reinforce these results using a longer sample period spanning the Great Recession.

²The response is a step function because consumption growth only deviates from normal when the real interest

of consumption to the shock is therefore quite large and gets larger the further in the future the interest rate shock occurs. It is the cumulative response of consumption (with some discounting) that determines the response of current inflation in the basic New Keynesian model. So, the further in the future is the interest rate that the monetary authority announces it will change, the larger is the current response of inflation. At the zero lower bound, this large effect on inflation will lower real rates and thus create a powerful feedback loop on output.

But is it a realistic prediction of the standard model that agents increase their consumption by the same amount in response to an interest rate cut 5 years in the future as they do to a cut in the current interest rate? Many people face some risk of hitting a borrowing constraint over the next five years. This effectively shortens their planning horizon since interest rate changes in states of the world that occur after they hit a borrowing constraint are irrelevant for their current consumption plan. In addition to this, people's desire to maintain a buffer stock of saving for precautionary reasons will temper their response to future interest rate shocks. Taking full advantage of the opportunity for intertemporal substitution presented by the future interest rate change requires people to run down their assets. This is costly since it leaves them more exposed to future income shocks. As a consequence, people will trade off the gains from intertemporal substitution and the costs of running down their buffer stock of savings. As the low interest rate is further in the future, the change in assets needed to take full advantage of intertemporal substitution grows and the countervailing precautionary savings effect therefore grows stronger, tempering the effects of forward guidance.

To investigate the quantitative magnitude of these effects, we consider a general equilibrium model in which agents face uninsurable, idiosyncratic income risk and borrowing constraints. In this model, the effect of forward guidance about future interest rates on current output falls the further out in the future the interest rate change is. For forward guidance about the interest rate 5 years in the future, the effect on output and inflation is roughly 40% as large as in the standard model. For forward guidance about the interest rate 10 years in the future, the effect on current output is essentially zero.

rate deviates from normal and this only occurs in the single quarter in our experiment. Another way to see this is that the forward guidance does not change the relative price of consumption for any two dates before the date of the interest rate change. All these dates must therefore have the same level of consumption. The end-point of consumption is pinned down at the old steady state by the fact that monetary shocks have no effect on real outcomes in the long run.

Our results indicate that forward guidance is a much less effective policy tool at the zero lower bound in a model with a realistic degree of precautionary savings than it is in standard macro models. We consider a shock that lowers the natural rate of interest enough that the zero lower bound binds for 5 years and the initial fall in output is -4% in the absence of forward guidance. If we assume markets are complete (and precautionary savings thus absent), a policy of maintaining interest rates at zero for a little more than three quarters beyond what a strict inflation targeting central bank would do completely eliminates the fall in output. In contrast, in our incomplete markets model with idiosyncratic risk and borrowing constraints, the effect of this amount of forward guidance is substantially smaller and a significant recession remains.

Movements in real interest rates play an important role in the response of the economy to a wide variety of shocks. Moreover, forward guidance has been an important part of monetary policy in “normal times” (i.e., away from the zero lower bound). We next explore how the reduced sensitivity of output to real interest rates in the incomplete markets model alters the response of the economy to a number of different business cycle shocks. Specifically, we consider productivity shocks, preference shocks, and markup shocks, in addition to monetary shocks. We extend our model to include these aggregate shocks and several features typically included in monetary business cycle models. We parameterize the monetary policy rule and monetary shocks to match recent evidence from Nakamura and Steinsson (2015).

To analyze this more complex model, we consider two approximations of the model, both of which allow the model to be analyzed using standard techniques for solving linear rational expectations models (e.g., Sims, 2002). First, we use the method proposed by Reiter (2009) to create a linear representation of the model. Second, we develop a simpler version of the model in which consumption-savings decisions of households can be represented by an Euler equation with discounting. We refer to this version of the model as the “discounted Euler equation model.”³

We show that the business cycle dynamics in our incomplete markets model are substantially different from those of its complete markets counterpart in response to all four of the different shocks we consider. Our result regarding the response to productivity shocks contrasts sharply with the results of Krusell and Smith (1998). They compare the response of an incomplete markets

³Interestingly, such a formulation of the consumption Euler equation has actually been used in policy calculations by the Central Bank of Norway, to combat the forward guidance puzzle. We thank Oistein Roisland for pointing this out to us.

model with flexible prices and its complete markets counterpart to productivity shocks and find that the two models have very similar responses. We show that the difference between our result and theirs arises because of the New Keynesian features we incorporate into the incomplete markets model.

One might ask how our mechanism relates to a model with hand-to-mouth consumers. Including such consumers also reduces the response of output to interest rates. However, this set-up reduces the sensitivity of output to interest rates by the same amount at all horizons—in this respect, it is equivalent to reducing the intertemporal elasticity of substitution. In contrast to the hand-to-mouth model, our incomplete markets model and the discounted Euler equation model reduce the effect of far future interest rate changes relative interest rate changes closer to the present.

Our work builds on recent papers that incorporate market incompleteness and idiosyncratic uncertainty into New Keynesian models starting with Guerrieri and Lorenzoni (2011) and Oh and Reis (2012). The papers closest in spirit to ours are McKay and Reis (2014) who investigate the power of automatic stabilizers at the zero lower bound and Gornemann, Kuester, and Nakajima (2014) who investigate the distributional implications of monetary policy shocks. Several other recent papers suggest “solutions” to the forward guidance puzzle. Del Negro, Giannoni, and Patterson (2013) argue that the experiment that gives rise to the puzzle is, itself, unreasonable. They argue that it is unreasonable to assume that the central bank really can engender substantial changes in long-term interest rates, which are at the core of why the forward guidance puzzle arises. Carlstrom, Fuerst, and Paustian (2012) and Kiley (2014) show that the magnitude of the forward guidance puzzle is substantially reduced in a sticky information (as opposed to a sticky price) model. This is because the sticky information Phillips curve is less forward looking. Our solution instead yields an Euler equation that is less forward looking than in the standard model. Caballero and Farhi (2014) argue that forward guidance is less effective if the reason why the zero lower bound binds is a shortage of safe assets in the economy—a safety trap—as opposed to a deleveraging or patience shock.

The paper proceeds as follows. Section 2 explains why forward guidance is so powerful in standard New Keynesian models. Section 3 presents our incomplete markets model featuring uninsurable idiosyncratic income risk and borrowing constraints. Section 4 describes our results about the reduced power of forward guidance in our incomplete markets model relative to the standard complete markets models. Section 5 extends the model and analyzes business cycle dynamics. It

also explains the two methods we employ to make the model tractable enough to be analyzed using standard techniques for solving linear rational expectations models. Section 6 concludes.

2 Why Is Forward Guidance So Powerful?

It is useful to start with an explanation for why forward guidance is so powerful in standard monetary models. Consider the basic New Keynesian model as developed, e.g., in Woodford (2003) and Gali (2008). The implications of private sector behavior for output and inflation in this model can be described up to a linear approximation by an intertemporal “IS” equation of the form

$$x_t = \mathbb{E}_t x_{t+1} - \sigma(i_t - \mathbb{E}_t \pi_{t+1} - r_t^n), \quad (1)$$

and a Phillips curve of the form

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t. \quad (2)$$

Here, x_t denotes the output gap—i.e., the percentage difference between actual output and the natural rate of output that would prevail if prices were fully flexible— π_t denotes inflation, i_t denotes the nominal, short-term, risk-free interest rate, r_t^n denotes the natural real rate of interest—i.e., the real interest rate that would prevail if prices were fully flexible— σ denotes the intertemporal elasticity of substitution, β denotes the subjective discount factor of households, and κ is the slope of the Phillips curve which is determined by the degree of nominal and real rigidities in the economy. All variables are denoted as percentage deviations from their steady state values.

Suppose for simplicity that the monetary policy of the central bank is given by an exogenous rule for the real interest rate where the real interest rate tracks the natural real rate with some error: $r_t = i_t - \mathbb{E}_t \pi_{t+1} = r_t^n + \epsilon_{t,t-j}$. Here $\epsilon_{t,t-j}$ denotes the shock to the short term real rate in period t that becomes known in period $t - j$.⁴ Absent any monetary shocks, the real interest rate will perfectly track the natural real rate and both the output gap and inflation will be zero. Suppose we start in such a state, but then the monetary authority announces that the real interest rate will be lower by 1% for a single quarter 5 years in the future, but maintained at the natural

⁴Given this specification of monetary policy, the model has a unique solution for which $\lim_{j \rightarrow \infty} \mathbb{E}_t x_{t+j} = 0$ and inflation is bounded. We could alternatively assume that the monetary authority sets the nominal rate according to the following rule $i_t = r_t^n + \phi \pi_t + \tilde{\epsilon}_{t,t-j}$ and $\phi > 1$. In this case, the model has a unique bounded solution (without the additional restriction that $\lim_{j \rightarrow \infty} \mathbb{E}_t x_{t+j} = 0$) and there exists a path for $\tilde{\epsilon}_{t,t-j}$ that gives the same solution as the model with monetary policy given by the exogenous path for the real rate we assume. We prefer to describe the monetary policy as an exogenous rule for the real interest rate because this simplifies our exposition substantially.

real rate of interest in all other quarters (i.e., $\epsilon_{t+20,t} = -0.01$).⁵

Figure 1 plots the response of output to this shock (assuming for simplicity that $\sigma = 1$). Even though the real interest rate does not change until 20 quarters later, output jumps up by a full 1% immediately. Output then stays at this higher level for 20 quarters before falling back to steady state in quarter 21. To understand why output responds in this way, it is important to consider that the shock changes the relative price of consumption between quarters 20 and 21 (since it is the real interest rate in quarter 20 that changes), but leaves the relative price of consumption for any two dates before quarter 20 and any two dates after quarter 20 unchanged. This implies that consumption growth can only deviate from normal in quarter 20. In other words, the response of consumption must be a step function. In addition to this, the level of consumption (and output) is pinned down in the long-run by the fact that monetary shocks have no effect on real outcomes in the long run. This implies that consumption (and output) must rise by 1% immediately, so that they can fall back to steady state in quarter 21.⁶

The step-function shape of the output response in Figure 1 is determined solely by the Euler equation. The level of consumption, and therefore output, is however determined by the intertemporal budget constraint. In general equilibrium, income rises in response to this type of shock because the level of production increases in response to the increase in demand. This increase in income allows households to consume more initially without reducing consumption after period 20. We can compare this general equilibrium case to the response of a single household holding its own income fixed and also holding the actions of all other agents in the economy fixed (call this the partial equilibrium response). Figure A.2 in the appendix plots the partial equilibrium response. The partial equilibrium response is also a step function, since the same Euler equation applies. The difference is that the increase in consumption over the first 20 quarters will cause the household to run down its wealth and imply that consumption going forward after period 20 will be permanently lower (this effect does not occur in general equilibrium due to the offsetting income rise). However, this difference is relatively small, even for a shock 20 quarters out. For this case,

⁵Conventional specifications of monetary shocks affect real rates in more than a single quarter. But in a linear model the effects of such monetary shocks can be “decomposed” into a simple sum of the effects of changes in real rates at each horizon. In this sense one can think of our 20-quarter experiment as one component of a more complex monetary shock that affects real rates in many quarters.

⁶An alternative way to see this is to solve the intertemporal IS equation—equation (1)—forward to get $x_t = -\sigma \sum_{j=0}^{\infty} \mathbb{E}_t(i_{t+j} - \mathbb{E}_{t+j}\pi_{t+j+1} - r_{t+j}^n)$. Notice, that there is no discounting in the sum on the right hand side of this equation. This implies that the output gap will rise immediately by 1% and will stay at that higher level for the next five years and then fall back to zero all at once when the low interest rate period passes.




Figure 1: Response of output to a one-quarter drop in the real interest rate 20 quarters in the future.

the partial equilibrium response of output is 91 basis points rather than the full 100 basis points in general equilibrium (see Figure A.2).

The logic described above for forward guidance 20 quarters in the future, applies for forward guidance at any horizon. As a consequence, the further out in the future the forward guidance is, the larger is the *cumulative* response of output. In the New Keynesian model, it is the the entire cumulative response of the output gap (albeit with some discounting) that determines the current response of inflation to forward guidance. To see this, it is useful to solve the Phillips curve forward to get

$$\pi_t = \kappa \sum_{j=0}^{\infty} \beta^j \mathbb{E}_t x_{t+j}. \quad (3)$$

This equation makes clear that the further in the future is the interest rate that the monetary authority announces it will change, the larger is the current response of inflation. While the response of inflation to a 1% change in the current real rate is $\kappa\sigma$, the response of inflation to a 1% change in the real rate for one quarter in the infinite future is $\kappa\sigma/(1 - \beta)$. If $\beta = 0.99$, the current response of inflation to forward guidance about a single quarter in the infinite future is 100 times larger than the response of inflation to an equally large change in the current real interest rate. Figure 2 plots the response of inflation to forward guidance about interest rates at different




Figure 2: Response of current inflation to forward guidance about interest rates at different horizons relative to response to equally large change in current real interest rate.

horizons relative to the response of inflation to an equally large change in the current real interest rate. We see that the response of inflation to forward guidance about interest rates five years in the future is roughly 18 times larger than the response of inflation to an equally sized change in the current real interest rate.

To build intuition, we have assumed that there is no endogenous feedback from changes in output and inflation back onto real interest rates. Actual monetary policies are more complicated. In normal times, forward guidance about lower real interest rates in the future may be partly undone by higher real interest rates in the intervening period. On the other hand, when monetary policy is constrained by the zero lower bound on short-term nominal interest rates, the higher inflation associated with forward guidance about future interest rates will actually lower current real interest rates and this will in turn raise current output and inflation further. In this case, the outsized effects of forward guidance we describe above will be further reinforced by subsequent endogenous interest rate movements.

3 An Incomplete Markets Model with Nominal Rigidities

Section 2 shows that the huge power of far future forward guidance in standard monetary models depends crucially on the prediction of the model that the current response of output to an expected change in real interest rates in the far future (say 5 years in the future) is equally large as the response of output to a change in the current real interest rate. But is this realistic? With some probability one will hit a borrowing constraint in the next 5 years. This effectively shortens one's planning horizon. Also, increasing consumption today by 1% in anticipation of a 1% change in real interest rates 5 years from today would entail a sizable run down of assets over the 5 years until the interest rate changes. Agents that face uninsurable idiosyncratic income risk and borrowing constraints will trade off the benefits of intertemporal substitution and the costs in terms of reduced ability to smooth consumption over time of having lower buffer stock savings. To analyze these effects we develop a model with uninsurable idiosyncratic shocks to household productivity, borrowing constraints, and nominal rigidities.

3.1 The Environment

The economy is populated by a unit continuum of ex ante identical households with preferences given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_{h,t}^{1-\gamma}}{1-\gamma} - \frac{\ell_{h,t}^{1+\psi}}{1+\psi} \right],$$

where $c_{h,t}$ is consumption of household h at time t and $\ell_{h,t}$ is labor supply of household h at time t . Households are endowed with stochastic idiosyncratic productivity $z_{h,t}$ that generates pre-tax labor income $W_t z_{h,t} \ell_{h,t}$, where W_t is the aggregate real wage. Each household's productivity $z_{h,t}$ follows a Markov chain with transition probabilities $\Pr(z_{h,t+1}|z_{h,t})$. We assume that the initial cross-sectional distribution of idiosyncratic productivities is equal to the ergodic distribution of this Markov chain. As the the Markov chain transition matrix is constant over time, it follows that the cross-sectional distribution of productivities is constant. We use $\Gamma^z(z)$ to denote this distribution.

In this economy, a final good is produced from intermediate inputs according to the production function

$$Y_t = \left(\int_0^1 y_{j,t}^{1/\mu} dj \right)^\mu,$$

where Y_t denotes the quantity of the final good produced at time t and $y_{j,t}$ denotes the quantity of

the intermediate good produced by firm j in period t . The intermediate goods are produced using labor as an input according to the production function

$$y_{j,t} = n_{j,t},$$

where $n_{j,t}$ denotes the amount of labor hired by firm j in period t .

The market structure of this model economy combines elements that are familiar from the standard New Keynesian model with elements that are familiar from the standard incomplete markets model (Bewley, undated; Huggett, 1993; Aiyagari, 1994). While the final good is produced by a representative competitive firm, the intermediate goods are produced by monopolistically competitive firms. The intermediate goods firms face frictions in adjusting their prices that imply that they can only update their prices with probability θ per period as in Calvo (1983). These firms are controlled by a risk-neutral manager who discounts future profits at rate β . Whatever profits are produced are paid out immediately to the households with each household receiving an equal share D_t . Households cannot trade their stakes in the firms.

Households trade a risk-free real bond with real interest r_t between periods t and $t+1$. Borrowing constraints prevent these households from taking negative bond positions. There is a stock of government debt outstanding with real face value B . The government raises tax revenue to finance interest payments on this debt. These taxes are collected by taxing households according to their labor productivity $z_{h,t}$. Let $\tau_t \bar{\tau}(z_{h,t})$ be the tax paid by a household h in period t . By levying the taxes on labor productivity, which is exogenous, the tax does not distort household decisions in the same way that a lump-sum tax does not. At the same time, the dependence of the tax on $z_{h,t}$ allows us to manipulate the cross-sectional correlation of tax payments and earnings.

We assume that the government runs a balanced budget so as to maintain a stable level of debt in each period. The government budget constraint is

$$\frac{B}{1+r_t} + \sum_z \Gamma^z(z) \tau_t \bar{\tau}(z) = B. \quad (4)$$

To illustrate our main results about the power of forward guidance, we will consider several monetary policy experiments involving somewhat different specifications of monetary policy. These are described in Section 4. The relationship between the real interest rate, the nominal interest rate i_t , and inflation π_t is given by the Fisher relation in the usual way

$$1+r_t = \frac{1+i_t}{1+\pi_{t+1}} \quad (5)$$

where $\pi_{t+1} \equiv P_{t+1}/P_t - 1$ and P_t is the aggregate price level.

3.2 Decision Problems

The decision problem faced by the households in the economy is

$$V_t(b_{h,t}, z_{h,t}) = \max_{c_{h,t}, b_{h,t+1}, \ell_{h,t}} \left\{ \frac{c_{h,t}^{1-\gamma}}{1-\gamma} - \frac{\ell_{h,t}^{1+\psi}}{1+\psi} + \beta \sum_{z_{h,t+1}} \Pr(z_{h,t+1}|z_{h,t}) V_{t+1}(b_{h,t+1}, z_{h,t+1}) \right\}$$

subject to

$$\begin{aligned} c_{h,t} + \frac{b_{h,t+1}}{1+r_t} &= b_{h,t} + W_t z_{h,t} \ell_{h,t} - \tau_t \bar{r}(z_{h,t}) + D_t \\ b_{h,t+1} &\geq 0. \end{aligned}$$

Let $c_t(b, z)$ be the decision rule for $c_{h,t}$, $g_t(b, z)$ be the decision rule for household bond holdings $b_{h,t+1}$, and $\ell_t(b, z)$ be the decision rule for $\ell_{h,t}$. These policy rules vary over time in response to aggregate events that affect current or future prices, taxes, or dividends.

The final goods producer's cost minimization problem implies that

$$y_{j,t} = \left(\frac{p_{j,t}}{P_t} \right)^{\mu/(1-\mu)} Y_t, \quad (6)$$

where $p_{j,t}$ is the price charged by firm j in period t and the aggregate price index is given by

$$P_t = \left(\int_0^1 p_{j,t}^{1/(1-\mu)} dj \right)^{1-\mu}.$$

When an intermediate goods producer updates its price it solves

$$\max_{p_t^*, \{y_{j,s}, n_{j,s}\}_{s=t}^{\infty}} \sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} \left(\frac{p_t^*}{P_s} y_{j,s} - W_s n_{j,s} \right)$$

subject to

$$\begin{aligned} y_{j,s} &= \left(\frac{p_t^*}{P_s} \right)^{\mu/(1-\mu)} Y_s, \\ y_{j,s} &= n_{j,s}, \end{aligned}$$

where p_t^* is the price set by firms who are able to update their price at date t .

The solution to this problem satisfies

$$\frac{p_t^*}{P_t} = \frac{\sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} \left(\frac{P_t}{P_s} \right)^{\mu/(1-\mu)} Y_s \mu W_s}{\sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} \left(\frac{P_t}{P_s} \right)^{1/(1-\mu)} Y_s}. \quad (7)$$

3.3 Equilibrium

Let $\Gamma_t(b, z)$ be the distribution of households over idiosyncratic states at date t . This distribution evolves according to

$$\Gamma_{t+1}(\mathcal{B}, z') = \int_{\{(b,z):g_t(b,z)\in\mathcal{B}\}} \Pr(z'|z)d\Gamma_t(b, z) \quad (8)$$

for all sets $\mathcal{B} \subset \mathbb{R}$.

As a result of nominal rigidities, price dispersion will result in some loss of efficiency. Integrating both sides of (6) across firms and using $y_{j,t} = n_{j,t}$ yields an aggregate production function

$$S_t Y_t = \int n_{j,t} dj \equiv N_t, \quad (9)$$

where N_t is aggregate labor demand and $S_t \equiv \int_0^1 \left(\frac{p_{j,t}}{P_t}\right)^{\mu/(1-\mu)} dj$ reflects the efficiency loss due to price dispersion. S_t evolves according to

$$S_{t+1} = (1 - \theta)S_t (1 + \pi_{t+1})^{-\mu/(1-\mu)} + \theta \left(\frac{p_{t+1}^*}{P_{t+1}}\right)^{\mu/(1-\mu)}. \quad (10)$$

Inflation can be written as a function of the relative price selected by firms that update their prices

$$1 + \pi_t = \left(\frac{1 - \theta}{1 - \theta \left(\frac{p_t^*}{P_t}\right)^{1/(1-\mu)}} \right)^{1-\mu}. \quad (11)$$

Aggregate labor supply is given by

$$L_t \equiv \int z \ell_t(b, z) d\Gamma(b, z). \quad (12)$$

and labor market clearing requires

$$L_t = N_t. \quad (13)$$

Bond market clearing requires

$$B = \int g_t(b, z) d\Gamma_t(b, z). \quad (14)$$

The aggregate dividend paid by the intermediate goods firms is

$$D_t = Y_t - W_t N_t. \quad (15)$$

Table 1: Baseline Calibration

Parameter	Description	Value	Target
β	discount factor	0.986	2% annual interest rate
γ	risk aversion	2	
$1/\psi$	Frisch elasticity	1/2	Chetty (2012)
B	supply of assets	$1.4 \times$ annual GDP	aggregate liquid assets (see text)
μ	markup	1.2	Christiano, Eichenbaum, and Rebelo (2011)
θ	price revision rate	0.15	Christiano, Eichenbaum, and Rebelo (2011)

Finally, integrating across the household budget constraints and using the government budget constraint and equation (15) gives

$$C_t = Y_t \tag{16}$$

as the aggregate resource constraint, where $C_t \equiv \int c_t(b, z) d\Gamma_t(b, z)$.

An equilibrium of this economy consists of decision rules and value functions $\{g_t(b, z), \ell_t(b, z), V_t(b, z)\}_{t=0}^{\infty}$ that solve the household's problem, distributions $\{\Gamma_t(b, z)\}_{t=0}^{\infty}$ that evolve according to (8). In addition, an equilibrium involves sequences $\{C_t, L_t, N_t, Y_t, D_t, i_t, W_t, \pi_t, r_t, p_t^*/P_t, S_t, \tau_t\}_{t=0}^{\infty}$, that satisfy the definitions of C_t and L_t and equations (4), (5), (7), (9), (10), (11), (13), (15), (16), and a monetary policy rule as described in section 4.

The main difference between this model and the model discussed in section 2 is the fact that markets are incomplete. If we modified this model to have complete markets and then linearized the equilibrium conditions, we would get the model in section 2. The introduction of incomplete markets yields a role for precautionary savings and it implies that redistribution of wealth across agents can affect the evolution of aggregate output. The fact that the present model is not linearized also implies that price dispersion affects equilibrium outcomes.

3.4 Calibration

Our model period is one quarter and our calibration is summarized in Table 1. We fix the steady state real interest rate at 2% annually and adjust the discount factor to match this.⁷ We set the coefficient of risk aversion to 2. We set the Frisch elasticity of labor supply to 1/2, which is in line with the findings of Chetty (2012). In our baseline calibration we set the supply of government bonds, B , to match the ratio of aggregate liquid assets to GDP. We calculate liquid assets from aggregate household balance sheets reported in the Flow of Funds Accounts and take the average ratio over the period 1970 to 2013.⁸ Our choice to calibrate the aggregate supply of assets to match liquid assets is motivated by the view that much of household net worth is illiquid and therefore not easily used for consumption smoothing and intertemporal substitution.⁹ In a sensitivity analysis we also consider a calibration in which we match aggregate household net worth, which we also calculate from the Flow of Funds Accounts (described below).

For our choices of the desired markup of intermediate firms, μ , and probability of maintaining a fixed price, θ , we follow Christiano, Eichenbaum, and Rebelo (2011) and set $\mu = 1.2$ and $\theta = 0.15$. The implied degree of price stickiness is on the high side of values used in the business cycle literature. This tends to reduce the size of the effects we find on inflation since it makes inflation less sensitive to changes in current marginal costs. In the exercises we do where the zero lower bound on nominal interest rates binds, this also tends to reduce the size of the effects on output since the smaller effects on inflation translate into smaller effects on real interest rates (Werning, 2012).

We calibrate the idiosyncratic wage risk to the persistent component of the estimated wage process in Floden and Lindé (2001).¹⁰ The estimates of Floden and Lindé are for an AR(1) with annual observations of log residual wages after the effects of age, education, and occupation have been removed. Floden and Lindé find an autoregressive coefficient of 0.9136 and an innovation

⁷We use the term “steady state” to refer to the stationary equilibrium in which aggregate quantities and prices are constant and inflation is zero.

⁸We use the same definition of liquid assets as Guerrieri and Lorenzoni (2011). Flow of Funds Table B.100 Lines 10 (deposits), 17 (treasury securities), 18 (agency and GSE securities), 19 (municipal securities), 20 (corporate and foreign bonds), 24 (corporate equities), 25 (mutual fund shares).

⁹Kaplan and Violante (2014) present a lifecycle savings model with liquid and illiquid assets and show that the illiquidity of household net worth leads to stronger and more realistic consumption responses to transitory income fluctuations.

¹⁰While it is common to include a transitory income shock in empirical models of wage dynamics we do not include such transitory shocks in our analysis because their impact on the quantitative results will be small as these shocks are easily smoothed by virtue of being transitory.

variance of 0.0426. We convert these estimates to parameters of a quarterly AR(1) process for log wages by simulating the quarterly process and aggregating to annual observations. We find the parameters of the quarterly process such that estimating an AR(1) on the simulated annual data reproduces the Floden and Lindé estimates, which results in an autoregressive coefficient of 0.966 and an innovation variance of 0.017. We discretize the resulting AR(1) process for log wages to a three-point Markov chain using the Rouwenhorst (1995) method.¹¹

Finally, to capture the fact that the bulk of tax payments are made by those with high earnings we set $\bar{\tau}(z)$ to be positive only for the highest z . Since households are heterogeneous in the incomplete markets model, their MPCs differ widely. Households with little wealth have high MPCs, while households with a great deal of wealth have much lower MPCs. As a consequence, wealth redistribution matters for aggregate consumption dynamics. For example, a reallocation of income from high to low net worth households leads to higher consumption demand, all else equal. One way in which this shows up in our model is that the government levies taxes to finance interest payments on debt. A fall in interest rates, therefore, may lead to a redistribution of wealth due to variation across households in holdings of government debt as well as tax obligations. Our (realistic) assumption that taxes are paid mostly by the rich leads these tax effects to be relatively small.

3.5 Alternative Calibrations

Our baseline calibration potentially implies too little volatility in household earnings. Guvenen, Ozkan, and Song (2014) report the standard deviation of the distribution of five-year earnings growth rates to be 0.73.¹² Our model calibrated as described above implies this standard deviation is only 0.53.

We therefore consider an alternative calibration in which we raise the variance of the productivity shocks in the model so that our model matches this moment of the five-year earnings growth rate distribution. Doing so requires raising the variance of the idiosyncratic productivity innovation from 0.017 to 0.033. With more risk, the larger precautionary savings motive raises the total demand for assets by households. In this calibration we, thus, reduce the discount factor so that

¹¹Kopecky and Suen (2010) prove that the Rouwenhorst method can match the conditional and unconditional mean and variance, and the first-order autocorrelation of any stationary AR(1) process.

¹²This value is the average across years of the values reported in Table A8 of Guvenen, Ozkan, and Song (2014).

the model is again consistent with the total supply of assets and a 2% annual interest rate. This requires a discount factor of 0.978. We refer to this as the High Risk calibration.

We also explore the extent to which our results depend on the average level of assets in the economy. With more assets, households will generally have more self-insurance and therefore will be less concerned with running down their assets. To explore this possibility we consider an alternative calibration in which we raise the supply of government debt, B , so that the average wealth in the economy is equal to the aggregate net worth of the household sector from the Flow of Funds, including both liquid and illiquid wealth.¹³ This yields a ratio of assets to annual GDP of 3.79. With a larger supply of assets, bond market clearing requires that households are more patient so as to increase the demand for assets at a given interest rate. In this calibration, we set the discount factor to 0.992 to be able to match a 2% annual interest rate. We refer to this as the High Asset calibration.

Finally we consider a case where we increase both the supply of assets and the extent of risk that households face. In this case we match a ratio of assets to GDP of 3.79 and the standard deviation of five-year earnings growth rates of 0.73. The discount factor needed to match a 2% annual interest rate is 0.990. We refer to this as the High Risk and Asset calibration.

3.6 Computation

In Section 4, we compute the perfect foresight transition paths of the economy in response to monetary policy and demand shocks. We assume that the economy begins in the steady state and returns to steady state after 250 quarters. We begin by guessing paths for all aggregate quantities and prices. We can then verify whether this guess is an equilibrium by checking that the definition of an equilibrium given above is satisfied. Part of this step involves solving and simulating the households problem at the guessed prices. We solve the household's problem by iterating on the Euler equation backwards through time using the endogenous gridpoint method of Carroll (2006) to compute the policy rules for each period of the transition. We then simulate the population of households forwards through time using a non-stochastic simulation algorithm to compute the distribution of wealth at each date. We can then compute aggregate consumption, labor supply, and bond holdings using the policy rules and distribution of wealth for each date. If our guess is not

¹³Here we use the ratio of household net worth to GDP averaged over 1970 to 2013. Household net worth is taken from Table B.100 Line 42.

an equilibrium we update it to a new guess that is closer to an equilibrium. We generate the new guess of prices and aggregate quantities by making use of an auxiliary model that approximates the aggregate behavior of the incomplete markets households and then solving for an equilibrium under this approximating model. We perform this step using a version of Newton’s method. We provide additional details of the computational methods in Appendix A.

4 Results

Our main result is that the power of forward guidance is substantially muted in the incomplete markets model we present in Section 3 relative to the standard complete-markets New Keynesian model. To illustrate this, we first consider a simple policy experiment: suppose the central bank promises a 50 basis point (i.e., 2% annualized) decrease in the real interest rate for a single quarter 5 years in the future.¹⁴

Figure 3 plots the response of output to this shock in our incomplete markets model and, for comparison, in the complete markets version of this model. As we discuss in section 2, the response of output under complete markets is a step function: Output immediately jumps up by 25 basis points and remains at that elevated level for 20 quarters before returning to steady state (recall that the IES is $1/2$). In contrast, in the incomplete markets model the initial increase in output is only about 40% as large. Output then gradually rises as the interest rate decrease gets closer. But even in the period right before the interest rate decrease, the increase in output is substantially smaller than under complete markets.

There are two main reasons why the response of output is muted under incomplete markets. First, households anticipate that with some probability they will face a sequence of idiosyncratic shocks that leads them to hit the borrowing constraint over the next 20 quarters. Each future state of the world in which they hit the borrowing constraint chops off a branch of their future event tree when it comes to their current intertemporal optimization problem. The possibility of hitting the borrowing constraint therefore shortens the households’ planning horizon and makes them respond less strongly to opportunities to intertemporally substitute that are far in the future. Since the probability that a household will hit the borrowing constraint grows the further they look into the future, households respond less to far future changes in interest rates than to changes in interest

¹⁴As in Section 2, we assume here that the monetary authority sets an exogenous path for the real interest rate.




Figure 3: Response of output to 50 basis point forward guidance about the real interest rate in quarter 20 (with real interest rates in all other quarters unchanged).

rates that will occur sooner.

Second, households are loath to run down their assets for precautionary reasons. Suppose a single household contemplates deviating from the path of consumption (and output) shown in Figure 3 and instead consumes more before quarter 20 and less after quarter 20 so as to take advantage of the opportunity for intertemporal substitution offered by the low future interest rate. This deviation would require the household to run down its wealth, which is costly as it leaves the household more exposed to future idiosyncratic shocks to income. At the equilibrium, such a deviation is not desirable because the marginal benefits of taking advantage of the opportunity for intertemporal substitution exactly equal the marginal cost of having a smaller buffer stock of savings and being more exposed to idiosyncratic shocks.

Under complete markets, it poses no concern for households to run down their wealth, since they are fully insured against all shocks. As a consequence they take full advantage of opportunities to intertemporally substitute. Under incomplete markets, however, households face a trade-off due to the precautionary benefits of maintaining a sizable buffer stock of assets. This precautionary motive leads the incomplete markets household to choose a more conservative consumption response to the change in the interest rate.

Another difference between the complete markets response and the incomplete markets response in Figure 3 is that, after the interest rate change passes, output falls below steady state for some time in the incomplete markets case. The reason for this is that the interest rate shock leads to a redistribution of wealth away from households with high marginal propensities to consume and towards households with low marginal propensities to consume (MPC). This wealth redistribution lowers aggregate demand (and therefore output) for some time until the distribution of wealth has had time to converge back to steady state.¹⁵

Figure 4 plots the response of inflation to this same shock. The five year output boom induced by the forward guidance about real interest rates leads to a large inflation response in the complete markets case. Since the output boom is much smaller in the incomplete markets model, the rise in inflation is also much smaller. The initial response of inflation in the incomplete markets model is again only about 40% as large as in the complete markets model.

Are these responses symmetric for interest rate increases versus decreases? Above, we consider the response of the economy to a decrease in interest rates five years in the future. We get very similar results if we instead consider an equally large increase in interest rates five years in the future. The immediate effect on output is a drop of 12 basis points, while output rises by about 10 basis points in the case of a decrease (Figure 3). The immediate effect on inflation is a drop of 28 basis points, while inflation rises by about 30 basis points in the case of decrease (Figure 4).¹⁶

4.1 Dependence on Horizon of Forward Guidance

The difference between the complete and incomplete markets models grows with the horizon of the interest rate shock. Figure 5 plots the initial response of output to 50 basis point forward guidance about the real interest rate in a single quarter as the horizon of that single quarter changes from

¹⁵There are two reasons why the interest rate shock leads to a redistribution towards households with low MPCs (which tend to be the high productivity households). First, the fall in interest rates implies that the the government needs less tax revenue in period 20 to pay the interest on the debt. Since the high productivity households pay the tax, they benefit from this fall in taxes. If these same households owned all the assets, their asset income would fall by an equal amount. The high productivity households own most of the assets but not all. So, on net, they benefit from the fall in interest rates. Second, the boom in output that the fall in interest rates causes leads wages to rise more than firm dividends. Since a larger fraction of the income of high productivity households is wage income, this also disproportionately benefits the high productivity households.

¹⁶It may seem puzzling that the absolute size of the output response is larger for the interest rate increase, but the ranking is reversed for the inflation response. Both interest rate increases and decreases yield an increase in price dispersion (because they both yield non-zero inflation). This implies that more labor is needed to produce the amount of output demanded, which raises wages in both cases and pushes inflation up, reversing the relative rankings.




Figure 4: Response of inflation to 50 basis point forward guidance about the real interest rate in quarter 20 (with real interest rates in all other quarters unchanged).

one to 40 quarters.¹⁷

In the complete markets model, output always rises by 25 basis points, regardless of the horizon of the forward guidance. In contrast, in the incomplete markets model, the effect is only about 20 basis points for an announcement about the real interest rate next quarter and falls monotonically thereafter. It is roughly 10 basis points for an announcement about the real interest rate 5 years ahead; and essentially zero for an announcement about the real interest rate 10 years ahead.

Intuitively, the probability of hitting a borrowing constraint before the interest rate change occurs rises with the length of time until the interest rate change occurs. For this reason, the benefit of responding falls. In addition, the cost of responding rises since the amount the household would need to run down its assets to take full advantage of the opportunity to intertemporally substitute gets larger and larger. Together, these forces imply that eventually the benefits from intertemporal substitution are simply too small to make it worth it for households to incur the costs associated with running down their buffer-stock savings.

The results for inflation are even starker. Figure 6 plots the initial response of inflation to

¹⁷For example, the points at horizon 20 in Figure 5 are the first points on each line in Figure 3. And the points at horizon 10 in Figure 5 are the initial response of output in the two models if the central bank announces that it will lower the real interest rate by 50 basis points for a single quarter 10 quarters in the future.




Figure 5: Initial response of output to 50 basis point forward guidance about the real interest rate for a single quarter at different horizons.

forward guidance about the real interest rate at different horizons. In the complete markets model, the response of inflation rises explosively with the horizon of the forward guidance.¹⁸ In the incomplete markets model, in contrast, the inflation response is smaller to start out with, grows more slowly, and therefore generates very different results at long horizons.

4.2 Results for Alternative Calibrations

Table 2 presents the results of the forward guidance experiment described above for our baseline incomplete markets model as well as for several alternative calibrations of our incomplete markets model. We also present the results for the the complete markets version of our model, for comparison. In each case, we present the initial response of output and inflation to 50 basis point forward guidance about the real interest rate for a single quarter 5 years in the future.

¹⁸There is a powerful feedback loop in our non-linear model. High inflation leads to an increase in price dispersion, which results in loss of efficiency. Workers must then work harder to produce the output that households demand. To induce workers to work more, wages must rise leading to more inflation. For small amounts of inflation these forces are not important, but if inflation remains elevated for long enough these dynamics lead to a sharp enough increase in wages and relative prices that the model solution can no longer be computed accurately. Interestingly, inflation itself asymptotes to a finite value because, in the non-linear New Keynesian model, no matter how high the current desired price of adjusting firms is, inflation can only asymptote to a certain level, because demand (and the products weight in the price index) falls with the relative price of adjusting firms. See Ascari (2004) for a very related discussion.




Figure 6: Initial response of inflation to 50 basis point forward guidance about the real interest rate for a single quarter at different horizons.

The High Risk calibration features greater uninsurable risk than our baseline calibration. We roughly double the volatility of idiosyncratic productivity shocks relative to our baseline calibration, allowing us to match recent evidence on the volatility of earnings growth from Guvenen, Ozkan, and Song (2014). This boosts the precautionary savings motive and further reduces the power of forward guidance relative to the complete markets benchmark. The response of output in this case is only about 20% of the complete markets benchmark and the response of inflation only about 32% of the complete markets benchmark.

In the High Asset calibration, we set the ratio of assets to GDP in the model to be almost three times higher than in our baseline calibration (3.79 versus 1.4). We do this to match the ratio of total net worth in the economy to GDP (as opposed to total liquid assets as in our baseline calibration). Increasing the quantity of available assets in the economy increases the size of the precautionary savings buffers available to households and thus reduces their reluctance to engage in intertemporal substitution. This change therefore moves the incomplete markets model closer to the complete markets benchmark. The output response rises to 58% of the complete markets benchmark, while the inflation effect rises to 49%.¹⁹

¹⁹One might ask how adding unsecured debt would affect our results. We have considered a case in which the

Table 2: Power of 20 Quarter Ahead Forward Guidance

	Initial Responses of	
	Output	Inflation
Baseline	10.3	29.8
High Risk	4.8	23.8
High Asset	14.5	36.2
High Risk and Asset	11.6	33.8
Complete Markets	25.0	74.3

Initial response of output and inflation (in basis points) to forward guidance that reduces the expected real interest rate 20 quarters ahead by 50 basis point for four different calibrations of our incomplete markets model.

We also consider a High Risk and Asset calibration with both of the above-mentioned alternative parameter values. These two modifications largely offset each other. As a consequence, the results lie between the two calibrations described above and close to the baseline calibration. The response of output in this calibration is 46% of the complete markets benchmark, while the response of inflation is 45% of the complete markets benchmark.

These alternative calibrations demonstrate that, with more risk and less self-insurance, the effects of precautionary savings and credit constraints become more pronounced. There are several reasons to believe that our baseline calibration understates the importance of these forces. In our baseline calibration, the fraction of agents that are constrained in steady state is 13% implying that 87% of households have positive liquid assets. Kaplan, Violante, and Weidner (2014) present estimates from the U.S. Survey of Consumer Finances that 75% of U.S. households have positive liquid assets. On this metric, therefore, our baseline calibration understates the fraction of households with low liquid assets relative to the U.S. data. A related statistic is the average marginal propensity to consume (MPC). In our baseline calibration the average MPC is only 12%, and even in our High Risk calibration it rises only to 14%. In contrast, a substantial amount of empirical

amount of assets in the model is calibrated as in our baseline case but we relax the borrowing constraint by allowing households to borrow up to five times average monthly labor income. Figure 2 of Kaplan, Violante, and Weidner (2014) indicates that very few households have unsecured debt that is larger than five times monthly income. This calibration yields results that are very similar to our baseline case (an initial output response of 10.9bp versus 10.3bp in our baseline case).

evidence suggests larger values for the average MPC, with many studies estimating values close to 20%.²⁰ On the basis of these statistics, one might argue for calibrations in which households are more credit constrained than in the calibrations we have considered. Such a calibration would likely further amplify the effects we emphasize regarding the differences between the complete and incomplete markets models.

4.3 Zero Lower Bound Analysis

In recent years, risk-free nominal interest rates around the world have hit zero. At the zero lower bound (ZLB), forward guidance has become an indispensable policy tool, since it is no longer possible to implement monetary policy via the current policy rate. Eggertsson and Woodford (2003) show how a persistent shock to the natural interest rate that causes the economy to hit the ZLB can provoke a massive recession if the central bank does not engage in unconventional monetary policy. They show, however, that the recession can be fully abated by a relatively modest amount of forward guidance about future interest rates.

Our conclusions above suggest that forward guidance may not be as powerful at the ZLB in our incomplete markets model. To investigate this question, we follow Eggertsson and Woodford (2003) in assuming that the ZLB is brought on by a temporary shock to the subjective discount factor of households in the economy that depresses the natural rate of interest below zero. In other words, we now consider a case where the discount factor can vary over time. The specific shock we consider is an increase in the discount factor that lasts for a known number of quarters and then reverts to normal.²¹ We choose the size and persistence of the shock so that the initial output decline is 4% and the ZLB binds for 20 quarters under a naive monetary policy (described below).²²

²⁰Parker (1999) estimates that household spend 20% of increases in disposable income when they hit the Social Security tax cap. Johnson, Parker, and Souleles (2006) estimate that households spent 20-40% of tax rebate checks they received in 2001, and Parker et al. (2013) estimate that households spent 12-30% of tax rebate checks they received in 2008. These studies as well as most others on this topic consider anticipated changes in income. They therefore provide a lower bound for responses to unanticipated changes in income. See Jappelli and Pistaferri (2010) for an excellent recent survey of the literature on the response of consumption to changes in income.

²¹Our shock differs from the shock considered in Eggertsson and Woodford (2003) in that its persistence is known, implying that agents have perfect foresight about the evolution of the aggregate economy. Eggertsson and Woodford (2003) consider a shock that reverts back to normal with constant probability each period. Clearly, both formulations are approximations. Eggertsson and Woodford's formulation abstracts from time-variation in the probability of the ZLB period ending, while our framework abstracts from uncertainty about when it will end. However, the incomplete markets model is more difficult to solve computationally without the assumption of perfect foresight for aggregate shocks.

²²Hitting these targets requires slightly different calibrations of the discount factor shock in the complete versus the incomplete markets model: it corresponds to a decline in the natural rate of 16.4 basis points in the incomplete markets model, but only 14.8 basis points in the complete markets model. In each case, the duration of the shock is




Figure 7: Response of output to the ZLB shock.

We consider two alternative monetary policies. First, we consider a policy where the central bank sets the nominal interest rate equal to a simplified Taylor rule whenever this yields an interest rate greater than zero, and, otherwise, sets the nominal rate to zero: $i_t = \max[0, \bar{r} + \phi\pi_t]$, where $\phi = 1.5$ and \bar{r} is the steady state real interest rate. We refer to this policy as the “naive” policy. We also consider an “extended” policy whereby the central bank sets the nominal rate to zero for several additional quarters beyond what is implied by the naive policy and then reverts back to the policy rule. We choose the length of the additional monetary stimulus to fully eliminate the initial fall in output in the complete markets model.

Figure 7 shows that forward guidance is substantially less powerful at the ZLB in the incomplete markets model than in the standard New Keynesian model. The bottom two lines show the path of output under the naive monetary policy for the complete and incomplete markets cases; while the top two lines show the effects of the extended monetary policy in these two cases. While the extended monetary policy fully eliminates the recession in the complete markets case, a substantial recession remains in the incomplete markets model. Figure 8 shows the implications for inflation: the extended policy is much more successful in preventing deflation in the complete markets model

33 quarters.




Figure 8: Response of inflation to the ZLB shock.

versus the incomplete markets model. While the initial deflation is only about 30 basis points in the complete markets case, it is more than 100 basis points in the incomplete markets case. The fact that inflation is lower in the incomplete markets case implies that real interest rates are higher (since nominal rates are stuck at zero). This contributes to the larger fall in output.

Figure 9 plots the implications of the naive and extended monetary policies for the nominal interest rate. Under the naive policy the ZLB binds for 20 quarters and then rises gradually to its steady state value of 50 basis points. Under the extended policy, the nominal interest rate remains at zero for 23 quarters (an additional 3 quarters), and interest rates are somewhat lower in quarter 24 than the naive policy implies (this partial stimulus in the 24th quarter is what is needed to exactly eliminate the initial fall in output due to the shock). The difference between the dashed and solid lines, thus, indicates the amount of additional stimulus provided by the extended policy.

5 Business Cycle Analysis

We have argued that market incompleteness affects the response of the economy to changes in real interest rates especially when those changes occur in the future. As the real interest rate is an important price in most macro models, we now explore how these effects of market incompleteness




Figure 9: Response of nominal interest rate to the ZLB shock.

alter the dynamics of the business cycle for a variety of alternative shocks. This analysis is related to the influential work of Krusell and Smith (1998) who ask a related question in the context of a flexible price economy. As we will show, their findings on the limited importance of market incompleteness for aggregate dynamics do not carry over to our New Keynesian environment. We also consider an empirically estimated specification of the monetary policy shock—which illustrates that forward guidance is important even away from the zero lower bound.

5.1 An Extended Model for Business Cycle Analysis

We now extend our model to include a set of aggregate shocks and several features typically included in monetary business cycle models. These features modify the Phillips curve and monetary policy rule of the model to generate more empirically realistic responses to shocks. It is important to note, however, that none of these modifications affect the response of consumers to a shock to the real interest rate—hence, the results of the main experiment in section 4 for output are unchanged.

Specifically, we make the following four changes to the “supply side” of the model presented in section 3. First, we follow Kimball (1995) in adopting a production function for final goods that implies that the elasticity of demand for each intermediate good falls as the relative quantity used

of that intermediate good rises. This makes the pricing decisions of the intermediate goods firms strategic complements. Second, we follow Christiano, Eichenbaum, and Evans (2005) in assuming that intermediate goods prices that are not re-optimized in a particular period are indexed to lagged inflation. This generates inflation inertia, which is needed to fit the empirical response of inflation to monetary shocks. Third, we assume that the production function of intermediate goods is subject to aggregate productivity shocks and therefore given by $y_{j,t} = A_t n_{j,t}$, where $\log A_t$ follows an AR(1) process. Fourth, we assume that the production function for final goods is subject to shocks that lead to variation in the elasticity of demand for each intermediate good even conditional on the relative price of these goods. We refer to these shocks as “markup shocks” and assume that they follow an AR(1) process.

Under these assumptions, we derive the following log-linearized Phillips curve

$$\pi_t - \pi_{t-1} = \beta \mathbb{E}_t [\pi_{t+1} - \pi_t] + \kappa \xi \left(\hat{W}_t - \hat{A}_t \right) + \hat{\mu}_t, \quad (17)$$

where hats denote log deviations from steady state. The parameter $\kappa = (\theta(1 - \beta(1 - \theta)))/(1 - \theta)$ governs the degree of nominal rigidity, while the parameter $\xi = 1/(1 + \mu\Omega/(\mu - 1))$ governs the degree of “real rigidity” (i.e., strategic complementarity in price setting). The parameter Ω governs the rate at which the elasticity of demand for a particular intermediate good changes with the relative price of that good. See appendix B for a detailed derivation of equation (17).

The only change we make to the household’s problem relative to the model presented in section 3 is that we assume that households are affected by an aggregate shock to their time-preference. Specifically, we assume that the discount rate between period t and $t + 1$ is $\beta \times \exp\{q_t\}$ where q_t follows an AR(1) process.

In earlier sections of the paper, we considered rather stylized representations of monetary shocks. We did so for illustrative purposes. A natural question, though, is how incomplete markets affect the response of the economy to the types of monetary shocks that have been estimated in the literature. To answer this question, we now assume that the monetary authority varies the nominal interest rates according to the following relatively standard feedback rule

$$i_t = \rho_i i_{t-1} + (1 - \rho_i)(\bar{r} + \phi\pi_t) + \epsilon_t, \quad (18)$$

where ρ_i is the degree of interest rate smoothing and ϵ_t is a monetary policy shock that follows an AR(1) process.




Figure 10: Response of the real interest rate to a monetary shock.

The calibration of the parameters of the monetary policy rule as well as the degree of real rigidity is designed to match estimates of the response of real interest rates and inflation to monetary policy shocks from Nakamura and Steinsson (2015). Specifically, the autocorrelation of the monetary shocks ϵ_t is set equal to its estimated value of 0.74 from Nakamura and Steinsson (2015). The degree of interest rate smoothing is set to $\rho_i = 0.954$ to match the response of the real rate at a 10-year horizon. We set the real rigidity parameter Ω equal to 56.5 to match the response of inflation at a five-year horizon to a monetary shock estimated by Nakamura and Steinsson (2015). As before, we set ϕ to 1.5. Figure 10 shows that this calibration does a good job of matching the path of the response of the real interest rate to the monetary shocks identified by Nakamura and Steinsson (2015). Finally, we set the autocorrelation of the productivity shocks to 0.95, and the autocorrelation of the preference shocks and markup shocks to 0.8.

5.2 Two Methods for Analyzing Business Cycle Dynamics

There are well-known challenges involved in solving incomplete markets models in the presence of aggregate uncertainty (Krusell and Smith, 1998). We now present two methods that can be used to analyze our extended model. Both of the methods are premised on building an approximate

representation of household behavior that can be linearized and analyzed with standard techniques for linear rational expectations models.

5.2.1 Linearizing the Full Model

First, we use the method proposed by Reiter (2009) to create a linear representation of the dynamics of the incomplete markets model. We approximate the distribution of wealth with a histogram. The mass of households in each bin then becomes a state variable of the model. We also approximate the household decision rules with a discrete approximation (i.e. splines). In this way, the model’s household block is represented with a large, but finite, number of variables and an equal number of equations. These equations replace the Euler equation and the labor supply rule in the representative agent model and we combine them with the other equations of the model such as the Phillips curve and monetary policy rule. We then linearize the model with respect to aggregate states, and solve for the dynamics of the economy as a perturbation around the stationary equilibrium without aggregate shocks using a standard linear rational expectations algorithm (Sims, 2002). The resulting solution is non-linear with respect to the idiosyncratic state variables, but linear with respect to the aggregate states. Appendix A contains more details.

Since this method involves solving the aggregate dynamics using a perturbation approach, it is easy to allow for many aggregate state variables. This implies that, using Reiter’s method, we are able to extend the incomplete markets model to incorporate many of the features found in workhorse monetary models (see also McKay and Reis, 2014).

This method has been shown to provide an accurate approximation to heterogeneous agent models. We have verified that this algorithm yields a solution that is very similar to our fully non-linear solution for the response to anticipated monetary shocks. In addition, appendix A presents an error analysis based on Euler equation errors.

5.2.2 Discounted Euler Equation

We next develop a simpler version of our model that we refer to as the “discounted Euler equation model,” which yields a single equation replacement for the standard consumption Euler equation. This is a version of our main model in which idiosyncratic productivity is i.i.d. and takes one of two values each period, which we refer to as employed ($z = 1$) and unemployed ($z = 0$). The only

income of the unemployed comes from home production in amount $A_t m$, where m is a parameter that determines productivity at home. Furthermore, we assume that there are no assets available ($B = 0$).²³

We show in appendix C that under these assumptions we can derive the following log-linearized consumption Euler equation²⁴

$$\hat{C}_t = \alpha \mathbb{E}_t \hat{C}_{t+1} - \gamma^{-1} (i_t - \mathbb{E}_t \pi_{t+1} - r_t^n). \quad (19)$$

This equation differs from the standard consumption Euler equation because of the parameter $\alpha < 1$. Solving this equation forward yields

$$\hat{C}_t = -\gamma^{-1} \mathbb{E}_t \sum_{j=0}^{\infty} \alpha^j (i_{t+j} - \mathbb{E}_{t+j} \pi_{t+j+1} - r_{t+j}^n). \quad (20)$$

Notice that the effect of interest rates j periods in the future on current consumption is discounted by a factor α^j . For this reason, we refer to equation (19) as the “discounted Euler equation.” The presence of $\alpha < 1$ implies that far future interest rate changes have much smaller effects on current consumption than near term interest rate changes.

In this version of the model, the unemployed households are liquidity constrained, while the employed households are not. The presence of α in equation (19) results from the fact that with some probability the currently employed households will be unemployed next period and next period’s expected marginal utility is therefore partly determined by the exogenous marginal utility in the unemployed state.²⁵

In the incomplete markets model, the response of output to real interest rate changes differs from the response for the complete markets model both because the response declines with the horizon of forward guidance and also because the response to contemporaneous changes in the real rate is smaller. The discounted Euler equation can match the first of these differences with an

²³Krusell, Mukoyama, and Smith (2011) and Ravn and Sterk (2013) have also used this approach to formulate tractable incomplete markets models.

²⁴For simplicity, we present the discounted Euler equation for the case without technology shocks. Appendix C extends the analysis to include those shocks.

²⁵Pierrgallini (2006) and Nistico (2012) provide an alternative micro-foundation for discounting in the Euler equation based on mortality shocks as in Blanchard’s (1985) perpetual-youth model. In their formulation, discounting only arises from aggregation over different generations, and to generate a quantitatively important deviation from the standard Euler equation, these authors assume counter-factually high death rates. Our approach rationalizes why long-lived consumers can have short planning horizons. Moreover, in Pierrgallini and Nistico’s formulation, the discounting in the Euler equation is larger the larger is the amount of financial wealth in the economy and disappears when financial wealth is in zero-net supply. In contrast, in our full model, agents discount the future more when they have little liquid financial wealth to buffer shocks to income.




Figure 11: Response of output to a 50 basis point change in the expected real interest rate in quarter 20 in the incomplete markets model and in the discounted Euler equation model.

appropriate value for α . To match the second of these differences we can change the value of the intertemporal elasticity of substitution, $1/\gamma$, in the discounted Euler equation.

Figure 11 shows that the discounted Euler equation with $\alpha = 0.97$ and $1/\gamma = 3/8$ provides a good approximation to the response of output to a real interest rate shock 20 quarters in the future in the baseline incomplete markets model analyzed in sections 3 and 4. The approximation is nearly perfect up until the time that the interest rate changes. What the discounted Euler equation misses is the fall in consumption after the interest rate shock. This fall is due to redistribution of wealth in the incomplete markets model (from households with high MPCs to households with low MPCs), which the discounted Euler equation does not capture.

In the discounted Euler equation model, the distribution of wealth is degenerate and the equilibrium conditions of the economy can be expressed using a small number of equations that can be analyzed with standard linearization techniques. However, the discounted Euler equation model does not incorporate distributional effects or the dynamics of the distribution of wealth. Nevertheless, we show below that the discounted Euler equation also provides a good approximation to the full incomplete markets model for a range of standard business cycle shocks.




Figure 12: Impulse responses to a monetary shock.

5.3 Results

We first consider the response of the economy to the monetary shock estimated by Nakamura and Steinsson (2015) and discussed above. Figure 12 plots the response of output, inflation, nominal and real interest rates to this monetary shock for three models: (1) the incomplete markets model described in section 5.1, (2) a complete markets version of that same model, and (3) a version of this model where household behavior is governed by the discounted Euler equation. Under complete markets, there is a large decline in output. A substantial part of this decline is driven by changes in real interest rates several years in the future. As the incomplete markets model responds less strongly to all changes in real interest rates and especially those at longer horizons, we find that under incomplete markets the initial decline in output is only 60 percent as large as under complete markets. The smaller decline in output translates into a smaller change in inflation. The figure also shows that the discounted Euler equation model yields a good approximation to the full incomplete markets model for this monetary shock.

Until now, we have focused on analyzing the response of the incomplete markets model to mon-

Table 3: Relative Volatility in Incomplete Markets Model

	Output				Inflation			
	Monetary	Pref.	Markup	Tech.	Monetary	Pref.	Markup	Tech.
Baseline	0.62	0.78	0.64	1.21	0.86	0.79	1.35	1.08
Discounted Euler	0.61	0.69	0.62	1.28	0.92	0.92	1.38	1.00
Flexible prices	–	–	0.92	0.97	–	–	–	–

The table lists the standard deviation of output and inflation in the incomplete markets model as a fraction of the same standard deviation in the corresponding complete markets model. The results in the first row are for the incomplete markets model described in section 5.1, while the results in the second row are for a version of that model where household behavior is governed by the discounted Euler equation. When prices are flexible, demand shocks (monetary and preference) do not affect equilibrium output under complete markets although there are some small effects in the incomplete markets that result from redistribution due to changes in real interest rates.

etary policy shocks. However, the arguments we have made apply to all shocks that yield variation in real interest rates. To illustrate this, Table 3 compares the response of the incomplete markets model and the complete markets model to preference shocks, markup shocks, and technology shocks, in addition to the monetary shocks we consider above. For a given shock, say a technology shock, we simulate a time series for the shock and then compute the response of the incomplete markets model as well as the response of the complete markets version of the model to this series of shocks. We then report the standard deviation of output and inflation for the incomplete markets model relative to the complete markets model. Since our approximate solutions are all linear in the shocks, the scale of the shock is irrelevant to this calculation.

The results of these calculations are reported in the first row of Table 3. Just as one would expect from Figure 12, the volatility of output that arises from monetary shocks is only 62% of its level under complete markets and this implies that the volatility of inflation is only 86% of its level under complete markets. We obtain similar results for the preference shock we consider. Households in the incomplete markets model respond less strongly to changes in the discount rate—recall that our preference shock is a shock to the discount rate—just as they respond less strongly to changes in real interest rates. This implies that the volatility of output and inflation is smaller under incomplete markets than complete markets in response to preference shocks.

Both monetary shocks and discount-rate shocks represent shocks to demand. The remaining

two shocks we consider are supply shocks. There are also considerable differences between the incomplete and complete markets models in response to these shocks. Consider first the markup, or “cost-push,” shock. The direct effect of this shock is to raise inflation. The central bank reacts to this by raising real interest rates. In the incomplete markets model, the central bank’s efforts will be less effective since households will react less strongly to the increase in interest rates. Output will therefore react by less, and a larger share of the initial upward pressure on inflation will remain. This is exactly what Table 3 shows.

In contrast to the other three shocks, when we consider technology shocks the volatility of output is substantially larger under incomplete markets than under complete markets. This results from the fact that the marginal propensity to consume extra income is substantially higher under incomplete markets than complete markets. A positive technology shock results in an increase in household income. The demand of households in the incomplete markets model rises by more in response to this rise in income, which raises wages and results in even higher income, amplifying the response of output to the shock.

The large differences between the incomplete markets and complete markets models that we document in Table 3 stand in stark contrast to the results of Krusell and Smith (1998). Krusell and Smith compare the response of a complete markets and incomplete markets model to a productivity shock and find that the difference is very small. While our model differs in several ways from theirs, we now demonstrate that a crucial difference is that we incorporate nominal and real rigidities. The last row of Table 3 presents results for a flexible price version of our model (we analyze only markup and technology shocks here because monetary and preference shocks have no effect in the complete markets model when prices are flexible).²⁶ Specifically, we compare a flexible price version of our incomplete markets model with a flexible price version of its complete markets counterpart. In this case, the volatility of output is almost identical under incomplete markets and complete markets.

The much larger difference between the complete and incomplete markets models in the sticky price case has to do with the sluggish movements of real interest rates when prices are sticky. The partial equilibrium response of output to both shocks to income and relative prices is quite different

²⁶Recall that there is no way for agents to save on aggregate in our model. For this reason, discount rate shocks have no impact with flexible prices in the complete markets models. These shocks do affect prices and have distributional consequences that lead to some small movements in quantities when markets are incomplete.

under incomplete markets than complete markets. When prices are flexible, these differences are largely undone in general equilibrium by changes in real interest rates. These general equilibrium responses of real interest rates are key to Krusell and Smith's (1998) near equivalence result. However, when prices are sticky, the behavior of real interest rates is governed by monetary policy. If monetary policy follows a conventional Taylor-type rule and the response of inflation is small, the response of real interest rates will be muted.

In the case of the productivity shock, the fact that agents in the incomplete markets model have higher marginal propensities to consume out of increased income than under complete markets implies that demand increases more in the incomplete markets model. If prices are sticky, this increase in demand will translate into an increase in output. If, however, prices are flexible, real interest rates will increase, putting downward pressure on aggregate consumption so that, in the end, output responds almost identically to the complete markets model.

The second row of Table 3 shows that the discounted Euler equation model yields a reasonably good approximation to the full incomplete markets for all four of the business cycle shocks that we have considered. This is useful because it helps us develop intuition for these results. One way in which the discounted Euler equation model differs from the full incomplete markets model is that it does not incorporate redistributive effects. The fact that it yields such a good approximation to the full incomplete markets model sheds some light on the (lack of) importance of redistributive effects for our results.

6 Conclusion

We study the effects of forward guidance about monetary policy. We do this in a standard New Keynesian model augmented with uninsurable income risk and borrowing constraints. Our main finding is that allowing for uninsurable income risk and borrowing constraints substantially decreases the power of forward guidance relative to a New Keynesian model with complete markets.

In the standard New Keynesian model with complete markets, an announcement about interest rates has the same effect on current consumption whether it pertains to the current short rate or the short rate 5 years in the future. In contrast, when markets are incomplete the effect of such an announcement on current consumption declines with the horizon of the announcement. Forward guidance about interest rates 5 years in the future has only about 40% as large an effect

on current consumption as in the complete markets case, while forward guidance about interest rates 10 years in the future has essentially no effect on current consumption. Intuitively, in the incomplete markets model, a precautionary savings effect counteracts the standard intertemporal substitution motive.

We develop two approximations to our model that allow us to analyze it using standard methods for solving linear rational expectations models. We then augment the model with several features typically included in monetary business cycle models and analyze the response of this augmented model to four different types of business cycle shocks. We show that the business cycle dynamics in the incomplete markets model differ substantially from those in its complete markets counterpart. This result contrasts with the influential results of Krusell and Smith (1998). We show that this arises as a consequence of the New Keynesian features we incorporate into the model, which fundamentally change the determination of the real interest rate.

A Computational Methods

A.1 Methods for Section 4

Here we describe the procedure used to find an equilibrium path of the heterogeneous agent model along a perfect foresight transition for the zero-lower-bound episode considered in Section 4.3. The algorithm used to compute the results for a one-time change in the real interest rate is closely related to what we present here.

Writing the firm's first order condition recursively. For the numerical analysis it is convenient to rewrite equation (7) recursively. Define

$$P_t^A \equiv \sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} \left(\frac{P_t}{P_s} \right)^{\mu/(1-\mu)} Y_s \mu W_s \quad (21)$$

$$P_t^B \equiv \sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} \left(\frac{P_t}{P_s} \right)^{1/(1-\mu)} Y_s. \quad (22)$$

then equation (7) becomes

$$\frac{p_t^*}{P_t} = \frac{P_t^A}{P_t^B}. \quad (23)$$

Equations (21) and (22) can be written recursively

$$P_t^A = \mu W_t Y_t + (1-\theta) \beta \mathbb{E}_t (1 + \pi_{t+1})^{-\mu/(1-\mu)} P_{t+1}^A \quad (24)$$

$$P_t^B = Y_t + \beta (1-\theta) \mathbb{E}_t (1 + \pi_{t+1})^{-1/(1-\mu)} P_{t+1}^B. \quad (25)$$

Initial guess. We assume that the economy has returned to steady state after $T = 250$ periods and look for equilibrium values for endogenous variables between dates $t = 0$ to T . In this explanation of our methods we use variables without subscripts to represent sequences from 0 to T . Let X denote a path for all endogenous aggregate variables from date 0 to date T . These variables include aggregate quantities and prices

$$X \equiv \{C_t, L_t, N_t, Y_t, D_t, i_t, W_t, \pi_t, r_t, p_t^*/P_t, S_t, \tau_t, P_t^A, P_t^B\}_{t=0}^T.$$

The dimension of X is given by 14 variables for each date and 251 dates. We require an initial guess X^0 . In most cases we found it sufficient to guess that the economy remains in steady state.

Solving the household’s problem. The household’s decision problem depends on X through r , W , τ , and D . For a given X^i we solve the household’s problem using the endogenous grid point method (Carroll, 2006). We approximate the household consumption function $c(b, z)$ with a shape-preserving cubic spline with 200 unequally-spaced knot points for each value of z with more knots placed at low asset levels where the consumption function exhibits more curvature. Given the consumption function we calculate labor supply from the household’s intratemporal optimality condition and savings from the budget constraint.

Simulating the population of households. We simulate the population of households in order to compute aggregate consumption and aggregate labor supply. We use a non-stochastic simulation method. We approximate the distribution of wealth with a histogram with 1000 unequally-spaced bins for each value of z again placing more bins at low asset levels. We then update the distribution of wealth according to the household savings policies and the exogenous transitions across z . When households choose levels of savings between the center of two bins, we allocate these households to the adjacent bins in a way that preserves total savings. See Young (2010) for a description of non-stochastic simulation in this manner.

Checking the equilibrium conditions. An equilibrium value of X must satisfy equations (4), (5), (9), (10), (11), (13), (15), (16), (23), (24), and (25) and the monetary policy rule $i_t = \max[0, \bar{r} + \phi * \pi_t + \epsilon_t]$, where ϵ_t is the exogenous deviation from the Taylor rule that takes a negative value under our “extended” policy. Call these 12 equations the “analytical” equilibrium conditions. The remaining two equilibrium conditions that pin down X are that C and L are consistent with household optimization and the dynamics of the distribution of wealth given the prices. Call these the “computational” equilibrium conditions.

To check whether a given X represents an equilibrium of the model is straightforward. We can easily verify whether the analytical equilibrium conditions hold at X . In addition, we can solve the household problem and simulate the population of households to verify that aggregated choices for consumption and labor supply of the heterogeneous households match with the values of C and L that appear in X .

Updating X^i The difficult part of the solution method arises when X^i is not an equilibrium. In this case we need to find a new guess X^{i+1} that moves us towards an equilibrium. To do this, we construct an auxiliary model by replacing the computational equilibrium conditions with additional analytical equilibrium conditions that approximate the behavior of the population of heterogeneous households but are easier to analyze. Specifically we use the equations

$$C_t^{-\gamma} = \eta_t^1 \beta (1 + r_t) C_{t+1}^{-\gamma} \quad (26)$$

$$C_t^{-\gamma} W_t = \eta_t^2 L_t^\psi. \quad (27)$$

where η^1 and η^2 are treated as parameters of the auxiliary model. For a given X^i , we have computed C and L from the computational equilibrium conditions. We then calibrate η^1 and η^2 from (26) and (27). We then solve for a new value of X from the 12 analytical equilibrium conditions and (26) and (27). This is a problem of solving for 14 unknowns at each date from 14 non-linear equations at each date for a total of 3514 unknowns and 3514 non-linear equations. We solve this system using the method described by Juillard (1996) for computing perfect foresight transition paths for non-linear models. This method is a variant of Newton's method that exploits the sparsity of the Jacobian matrix. Call this solution $X^{i'}$. We then form X^{i+1} by updating partially from X^i towards $X^{i'}$. We iterate until X^i satisfies the equilibrium conditions within a tolerance of 5×10^{-6} .

A.2 Methods for Section 5

Solving for the household's decision rules without aggregate shocks. For each level of idiosyncratic productivity, z , we solve for the asset level, b , at which the household chooses to leave the period without savings, $b' = 0$, and is therefore just on the threshold of being borrowing constrained. Call this level of assets $\underline{b}(z)$. We create a grid of 100 unevenly spaced points with the first grid point equal to $\underline{b}(z)$. We then solve for the household's savings at each of these points. Constructing the grid in this fashion allows for a more accurate solution because we know the savings policy rule is constant at $b' = 0$ for all $b < \underline{b}(z)$ and we do not interpolate across the kink in the policy rule where the borrowing constraint stops binding. We also fix a grid of 100 points and solve for labor supply at each of these points. Between the grid points we interpolate the savings and labor supply policies linearly. We solve for the household's savings and labor supply policies jointly using Broyden's (1965) method by imposing that the Euler equation and labor supply condition hold with equality at the grid points. Notice that the household policy rule

for savings is parameterized by 300 variables (three levels of z and 100 grid points for each) and similarly the household labor supply policy rule is parameterized by another 300 variables.

Finding the stationary equilibrium. This part of the algorithm is similar in spirit to Aiyagari (1994). We assume the stationary equilibrium interest rate is 2 percent annually and we search for the value of β for which this is an equilibrium. We find it also useful to guess values of equilibrium output and dividends. So we search over values of β , Y , and D such that aggregate asset demand by households matches the supply of government bonds, the labor market clears and the dividend is equal to $Y - WN$. We perform this search using Broyden’s method. To simulate the stationary distribution of wealth, we use non-stochastic simulation as described above.

Solving for aggregate dynamics Here we follow Reiter (2009). Our non-stochastic simulation algorithm tracks the distribution of wealth using a histogram. We consider the mass in each of these bins to be a variable. We have three skill levels and 250 bins for each. In total there are 749 variables because the distribution must sum to one. For a given set of household policy rules, an idiosyncratic shock process, and prices we can formulate an equal number of equations that dictate how the distribution of wealth updates across these bins.

The 300 variables that parameterize the household savings policy and the 300 variables that parameterize the household labor supply policy are unknowns that depend on the aggregate state. As the aggregate state changes, we require that these variables satisfy the Euler equation and labor supply first-order condition on the grids described above, which yields another 600 equations in total.

Next we have a small number of aggregate variables and “aggregate” equations. The variables are Y , L , C , π , W , i , R , τ and the four aggregate shocks. We then impose equations (4), (5), (9), (13), (14), (17), (18), $C_t \equiv \int c_t(b, z) d\Gamma_t(b, z)$, and the four AR(1) stochastic processes for the aggregate shocks. These are 12 additional equations and 12 additional variables.

At this stage, we have a large system of non-linear equations that the discretized model must satisfy. We follow Reiter (2009) in linearizing this system around the stationary equilibrium using automatic differentiation and then solving the linearized system as a linear rational expectations model using the algorithm from Sims (2002). The results take the form of a VAR(1) representation of the discretized model, but crucially, the solution is non-linear with respect to idiosyncratic states.

We use this linear system to compute impulse responses and simulate it to compute covariances of aggregate quantities and prices.

Accuracy check: Euler equation errors. In our calculations for Section 5 there are two sources of errors both of which commonly arise in related algorithms. First, there are errors in the household decision rules between the grid points. These errors are present even in the stationary equilibrium. Away from the stationary equilibrium (the point around which we linearize) there are errors due to non-linear responses to aggregate states as is the case with other applications of perturbation methods. The procedure that follows assesses both of these types of errors by using a finer grid and by checking the error in the non-linear equations away from the stationary equilibrium.

To assess the accuracy of our solution, we calculate unit-free Euler equation errors.²⁷ We use a test grid over asset holdings that is finer than the grid on which we solve for household decision rules.²⁸ For a given state of the economy, \mathcal{X}_t , the distribution of bond holdings at the start of the period and exogenous variables are predetermined. In addition, the interest rate paid on current bond holdings, r_{t-1} , is also predetermined.²⁹ We then use the approximate model solution to determine W_t , R_t , and the household savings and labor supply rules. We then use the non-linear, static relationships, market clearing conditions, and budget constraints to determine τ_t , L_t , N_t , Y_t , D_t , and the distribution of wealth at the start of the next period. In this step, all budget constraints and market-clearing conditions are forced to hold.

From these calculations and a given set of aggregate shocks we can compute the next state of the economy, \mathcal{X}_{t+1} , and repeat these calculations to find W_{t+1} and so on. To compute expectations, we use Gaussian quadrature over the four aggregate shocks using a grid that has seven nodes in each dimension. For a household with particular idiosyncratic states, (b, z) , we can compute the level of consumption implied by the right-hand side of the Euler equation as

$$\hat{c}_t(b, z) \equiv \{\beta(1 + r_t)\mathbb{E}_t [c_{t+1}(b', z')^{-\gamma}]\}^{-\gamma}, \quad (28)$$

where the expectation is over aggregate and idiosyncratic shocks and b' is computed using the

²⁷See Judd (1992) for an explanation of this accuracy check and the interpretation of the errors in terms of bounded rationality.

²⁸Specifically, we use the same 250 point grid that we use to approximate the distribution of wealth.

²⁹In our calculations we track the distribution of the wealth at the start of the period when idiosyncratic shocks have been realized but before interest has been paid.

approximate solution for the policy rule. The unit-free Euler equation error for a given type of household is then $\hat{c}_t(b, z)/c_t(b, z) - 1$, where $c_t(b, z)$ is the level of consumption implied by the approximated decision rules.³⁰ Using the steps above, we can compute the Euler equation error for each type of household.

We can also assess the errors in household labor supply decision rules. Using the approximate solution for the policy rules and the wage we can compute

$$\hat{n}(b, z) \equiv [Wc(b, z)^{-\gamma}]^{1/\psi}$$

We can then compare this value to the approximate solution for the labor supply rule. We express the error in this equation in unit-free terms as $\hat{n}(b, z)/n(b, z) - 1$.

In order to assess the errors due to non-linearities we must take a stand on the magnitude of the shocks hitting the economy. We target a standard deviation of log output of 0.0154, which is the observed value for HP filtered data from 1960 to 2011. We assume that each shock accounts for a quarter of the variance of log output.

The Euler equation and labor supply errors vary over the state space. We randomly draw points in the state space by simulating the model for 100,000 periods and we compute the errors every 1,000 simulated periods. We describe the distribution of errors across the 100 resulting points by reporting the largest absolute error and the mean absolute error for different types of households as defined by assets and productivity. Figure A.1 plots the mean and maximum absolute errors across the 100 aggregate states for both the Euler equation and the labor supply rule. The errors rarely exceed 0.01 and are generally on the order of 0.001 or below. An error of 0.001, means that a household following our approximate policy rule makes an error of \$1 per \$1,000 spent.

B Model with Inflation Inertia and Real Rigidities

B.1 Final Goods Producers

Following Kimball (1995), we adopt a setup where the elasticity of demand for a particular intermediate good j by final goods producer falls as the relative quantity used of that intermediate good

³⁰We approximate the policy rules for savings as opposed to consumption so $c_t(b, z)$ is computed from the budget constraint and depends on the approximate policy rule for labor supply and the market-clearing prices.




Figure A.1: Mean and maximum Euler equation and labor supply equation errors (log base 10). The mean and maximum are taken over 100 randomly selected points in the aggregate state space. The three lines in each figure correspond to idiosyncratic productivities (dark = low-skill, dashed = medium-skill, grey = high-skill). Assets are relative to average quarterly income.

risers. The final goods producer minimizes costs

$$\int_0^1 p_{j,t} y_{j,t} dj$$

subject to the implicit production function

$$\int_0^1 D\left(\frac{y_{j,t}}{Y_t}\right) dj = 1, \quad (29)$$

where following Dotsey and King (2005) we adopt the function form

$$D(x) = \frac{\mu_t}{1+\eta} [(1+\eta)x - \eta]^{1/\mu_t} + 1 - \frac{\mu_t}{1+\eta}.$$

This function is increasing and strictly concave in x and satisfies $D(1) = 1$. Furthermore, with $\eta = 0$ it reduces to the constant elasticity of substitution aggregator we use in section 3. Notice that we allow μ_t to vary over time. We refer to variation in μ_t as markup shocks.

Cost minimization by final producers yields

$$\frac{y_{j,t}}{Y_t} = \frac{1}{1+\eta} \left[\left(\frac{p_{j,t}}{\tilde{P}_t} \right)^{\mu_t/(1-\mu_t)} + \eta \right]. \quad (30)$$

where

$$\tilde{P}_t = \left[\int_0^1 p_{j,t}^{1/(1-\mu_t)} dj \right]^{1-\mu_t}.$$

Starting from the definition of the price index, we can use equation (30) to see that

$$\begin{aligned} 1 &= \int_0^1 \frac{p_{j,t} y_{j,t}}{P_t Y_t} dj \\ &= \int_0^1 \frac{p_{j,t}}{P_t} \frac{\left(\frac{p_{j,t} P_t}{\tilde{P}_t} \right)^{\mu_t/(1-\mu_t)} + \eta}{1 + \eta} dj \\ &= \frac{1}{1 + \eta} \left(\frac{P_t}{\tilde{P}_t} \right)^{\mu_t/(1-\mu_t)} \int_0^1 \left(\frac{p_{j,t}}{P_t} \right)^{1/(1-\mu_t)} dj + \frac{\eta}{1 + \eta} \int_0^1 \frac{p_{j,t}}{P_t} dj \\ &= \frac{1}{1 + \eta} \left(\frac{P_t}{\tilde{P}_t} \right)^{\mu_t/(1-\mu_t)} \left(\frac{\tilde{P}_t}{P_t} \right)^{1/(1-\mu_t)} + \frac{\eta}{1 + \eta} \int_0^1 \frac{p_{j,t}}{P_t} dj \\ &= \frac{1}{1 + \eta} \frac{\tilde{P}_t}{P_t} + \frac{\eta}{1 + \eta} \int_0^1 \frac{p_{j,t}}{P_t} dj \\ P_t &= \frac{1}{1 + \eta} \tilde{P}_t + \frac{\eta}{1 + \eta} \bar{P}_t, \end{aligned} \tag{31}$$

where we have defined $\bar{P}_t \equiv \int_0^1 p_{j,t} dj$. In other words, the price index is a convex combination of a CES aggregator and a linear aggregator.

B.2 Intermediate Goods Producers

We assume prices that are not re-optimized are adjusted with lagged inflation. The evolution of inflation can be defined in terms of the change in the CES and linear aggregators. The two price aggregators evolve according to

$$\left(\frac{\tilde{P}_t}{\tilde{P}_{t-1}} \right)^{1/(1-\mu_t)} = \frac{(1 - \theta)(1 + \pi_{t-1})^{1/(1-\mu_t)}}{1 - \theta \left(\frac{p_t^*}{\tilde{P}_t} \right)^{1/(1-\mu_t)}} \tag{32}$$

$$\bar{P}_t = (1 - \theta)\bar{P}_{t-1}(1 + \pi_{t-1}) + \theta p_t^*, \tag{33}$$

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

We assume the intermediate goods firm produces according to

$$y_{j,t} = A_t n_{j,t} \tag{34}$$

so real marginal cost is $M_t = W_t/A_t$, where W_t is the wage per efficiency unit of labor. The price setting problem is then

$$\max_{p_t^*} \sum_{s=t}^{\infty} \beta^{s-t} (1 - \theta)^{s-t} \left[\frac{p_t^* P_{s-1}}{P_s P_{t-1}} - M_s \right] y_{j,s}$$

where $y_{j,s}$ is given by

$$y_{j,s} = \frac{\left(\frac{p_{j,t}}{\bar{P}_s} \frac{P_{s-1}}{P_{t-1}}\right)^{\mu_t/(1-\mu_t)} + \eta}{1 + \eta} Y_s.$$

In the objective function, the term P_{s-1}/P_{t-1} reflects the indexation to lagged inflation. For the sake of the algebra that follows we call this ratio Z_s . The first order condition is

$$\sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} Y_s \left\{ \frac{1}{\bar{P}_s} Z_s \frac{\left(\frac{p_t^*}{\bar{P}_s} Z_s\right)^{\mu_t/(1-\mu_t)} + \eta}{1 + \eta} + \left[\frac{p_t^*}{\bar{P}_s} Z_s - M_s \right] \frac{\mu_t}{1 - \mu_t} \frac{\left(\frac{p_t^*}{\bar{P}_s} Z_s\right)^{\mu_t/(1-\mu_t)-1}}{1 + \eta} \frac{1}{\bar{P}_s} Z_s \right\} = 0.$$

Some manipulation of this equation yields

$$\sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} Y_s \frac{p_t^*}{\bar{P}_s} Z_s \left\{ \left(\frac{p_t^*}{\bar{P}_s} Z_s\right)^{\mu_t/(1-\mu_t)} + \eta(1 - \mu_t) \right\} = \sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} Y_s M_s \mu_t \left(\frac{p_t^*}{\bar{P}_s} Z_s\right)^{\mu_t/(1-\mu_t)}$$

To calculate the efficiency loss from price dispersion, combine the demand curve for intermediate products—equation (30)—the production function—equation (34)—and $N_t \equiv \int_0^1 n_{j,t} dj$ to arrive at

$$A_t N_t = \frac{1}{1 + \eta} \left[\int_0^1 \left(\frac{p_{j,t}}{\bar{P}_t}\right)^{\mu_t/(1-\mu_t)} dj + \eta \right] Y_t.$$

Defining $S_t \equiv \int_0^1 \left(\frac{p_{j,t}}{\bar{P}_t}\right)^{\mu_t/(1-\mu_t)} dj$ we have

$$A_t N_t = \frac{1}{1 + \eta} [S_t + \eta] Y_t$$

and S_t evolves according to

$$S_t = \theta \left(\frac{p_t^*}{\bar{P}_t}\right)^{\mu_t/(1-\mu_t)} + (1 - \theta) S_{t-1} (1 + \pi_{t-1})^{\mu_t/(1-\mu_t)} \left(\frac{\bar{P}_{t-1}}{\bar{P}_t}\right)^{\mu_t/(1-\mu_t)} \quad (35)$$

B.3 Log-Linearization

We use hats to denote log deviations. We log-linearize around a steady state in which $P = \tilde{P} = \bar{P} = p^*$ and $\pi = 0$. Log-linearizing (31), (32), and (33) yields

$$\hat{P}_t = \frac{1}{1 + \eta} \hat{P}_t + \frac{\eta}{1 + \eta} \hat{\bar{P}}_t \quad (36)$$

$$\hat{\bar{P}}_t - \hat{\bar{P}}_{t-1} = \frac{\theta}{1 - \theta} (\hat{p}_t^* - \hat{\bar{P}}_t) + \pi_{t-1} \quad (37)$$

$$\hat{\bar{P}}_t - \hat{\bar{P}}_{t-1} = \frac{\theta}{1 - \theta} (\hat{p}_t^* - \hat{\bar{P}}_t) + \pi_{t-1}, \quad (38)$$

where we have abused notation to write $\log(1 + \pi_{t-1}) = \pi_{t-1}$. Combining these yields

$$\pi_t = \frac{\theta}{1-\theta} \left(\hat{p}_t^* - \hat{P}_t \right) + \pi_{t-1} \quad (39)$$

(37)-(39) show that to a first-order approximation \tilde{P}_t , \bar{P}_t , and P_t evolve in the same way.

Log-linearization of the firm's price-setting first order condition yields

$$\hat{p}_t^* = (1 - \beta(1 - \theta)) \sum_{s=t}^{\infty} \beta^{s-t} (1 - \theta)^{s-t} \left[\frac{1 + (1 - \mu)\eta}{1 + \eta} \hat{M}_s + \hat{\mu}_t + \hat{P}_s - \hat{Z}_s \right] \quad (40)$$

where we use the fact that in steady state $\mu M = 1 + (1 - \mu)\eta$ and the definition of $\hat{Z}_s = \hat{P}_{s-1} - \hat{P}_{t-1}$.

We also define

$$\hat{\mu}_t = \left[\frac{1 + \eta + \frac{\mu^2 \eta}{(1-\mu)^2}}{1 + (1 - \mu)\eta} \right] \log \left(\frac{\mu_t}{\mu} \right).$$

We can now rewrite this last expression recursively as

$$\begin{aligned} \hat{p}_t^* &= (1 - \beta(1 - \theta)) \left[\frac{1 + (1 - \mu)\eta}{1 + \eta} \hat{M}_t + \hat{\mu}_t + \hat{P}_t \right] + \beta(1 - \theta) \left(\mathbb{E}_t [\hat{p}_{t+1}^*] - (\hat{P}_t - \hat{P}_{t-1}) \right) \\ \hat{p}_t^* - \hat{P}_t &= (1 - \beta(1 - \theta)) \frac{1 + (1 - \mu)\eta}{1 + \eta} \hat{M}_t + \hat{\mu}_t + \beta(1 - \theta) \mathbb{E}_t \left[\hat{p}_{t+1}^* - \hat{P}_t - (\hat{P}_t - \hat{P}_{t-1}) \right] \\ \hat{p}_t^* - \hat{P}_t &= (1 - \beta(1 - \theta)) \frac{1 + (1 - \mu)\eta}{1 + \eta} \hat{M}_t + \hat{\mu}_t + \beta(1 - \theta) \mathbb{E}_t \left[\hat{p}_{t+1}^* - \hat{P}_{t+1} + \pi_{t+1} - \pi_t \right]. \end{aligned}$$

Combining this last expression with equation (39) yields

$$(\pi_t - \pi_{t-1}) = \frac{\theta(1 - \beta(1 - \theta))}{1 - \theta} \frac{1 + (1 - \mu)\eta}{1 + \eta} \hat{M}_t + \hat{\mu}_t + \beta \mathbb{E}_t [\pi_{t+1} - \pi_t]. \quad (41)$$

If we define $\Omega = (\mu - 1)\eta / (1 - (\mu - 1)\eta)$, and notice that $\hat{M}_t = \hat{W}_t - \hat{A}_t$ up to a first order approximation, we have equation (17) in the text.

Just as in the standard case, there is no efficiency loss from price dispersion to a first-order approximation. To see this, log-linearize (35)

$$\hat{S}_t = \theta \frac{\mu}{1 - \mu} \left(\hat{p}_t^* - \hat{P}_t \right) + (1 - \theta) \left[\hat{S}_{t-1} + \frac{\mu}{1 - \mu} \left(\pi_{t-1} + \hat{P}_{t-1} - \hat{P}_t \right) \right]$$

and substitute for $\hat{P}_{t-1} - \hat{P}_t$ using (37) to arrive at

$$\hat{S}_t = (1 - \theta) \hat{S}_{t-1}.$$

C A Simple Incomplete Markets Model

The discounted Euler equation (19) can be micro-founded with a simplified version of the model we present in Section 3 in which:

1. The idiosyncratic productivity shock takes just two values, which we will call employed ($z = 1$) and unemployed ($z = 0$).
2. Idiosyncratic productivity is i.i.d. across time: $\Pr(z'|z) = \Pr(z')$.
3. The supply of government debt is zero: $B = 0$.
4. Unemployed agents have access to a home production technology that produces final goods in amount $A_t m$.
5. Firm dividends D_t are distributed only to the employed households.

This version of our model is analytically tractable because there is no wealth in the economy and there is a strict borrowing constraint, $b' \geq 0$. As the gross supply of assets is zero, there is no possibility of saving in equilibrium and so the distribution of wealth remains degenerate at zero.

As individual assets are constant at zero, it follows that all households of a given employment status must choose the same consumption. Let $c_{e,t}$ be the consumption of the employed and $c_{u,t}$ be the consumption of the unemployed agents. Moreover, the absence of opportunities to borrow and save implies that consumption must be equal to income for all individuals. The unemployed will therefore consume $m A_t$ and the employed will consume

$$c_{e,t} = W_t \ell_t + D_t / (1 - u),$$

where ℓ_t is labor supply per employed household, D_t is per capita dividends paid by firms, and u is the fraction of the population that is unemployed.

The Euler equation for employed agents is

$$c_{e,t}^{-\gamma} \geq \beta e^{qt} R_t \mathbb{E}_t \left[(1 - u) c_{e,t+1}^{-\gamma} + u (m A_{t+1})^{-\gamma} \right] \quad (42)$$

and the Euler equation for an unemployed agent is

$$(m A_t)^{-\gamma} \geq \beta e^{qt} R_t \mathbb{E}_t \left[(1 - u) c_{e,t+1}^{-\gamma} + u (m A_{t+1})^{-\gamma} \right].$$

Notice that the right hand side of the Euler equation is independent of employment status as employment is i.i.d. across time. Therefore, if $mA_t < c_{e,t}$ the unemployed households must be constrained and their Euler equation will not hold with equality. We assume that m is low enough that $mA_t < c_{e,t}$ for all t .³¹

Following Krusell, Mukoyama, and Smith (2011) and Ravn and Sterk (2013), we will focus on the equilibrium of this economy in which the Euler equation of the employed holds with equality in all periods. In this equilibrium, the employed households choose zero savings and are therefore up against their constraint. But the constraint does not bind in the sense that they would not strictly prefer to borrow more if allowed. There are other equilibria of this economy in which the Euler equation for the employed household holds with inequality (implying that the employed households would strictly prefer to borrow if allowed). We focus on the equilibrium in which the employed households' Euler equation holds with equality for the following reason. In the more realistic case in which the supply of government debt is positive ($B > 0$) there is a unique equilibrium in which the employed households' Euler equation holds with equality. Consider a sequence of such economies with smaller and smaller amounts of government debt. The equilibrium we focus on is the unique equilibrium of the limiting economy for which the supply of government bonds goes to zero.³²

We assume that home production is not measured as part of GDP. Therefore measured aggregate consumption is

$$C_t = (1 - u)c_{e,t} \tag{43}$$

and measured aggregate output is

$$Y_t = (1 - u)A_t\ell_t.$$

³¹It is easy to relax the assumption that idiosyncratic productivity is i.i.d., however, with i.i.d. productivity it is especially easy to see that the unemployed will be constrained if m is low enough.

³²Suppose the outstanding supply of government debt is positive ($B = \varepsilon > 0$). In this case, bond market clearing requires some household to hold this debt at the prevailing interest rate (which is set by the monetary authority). This household must be indifferent at the margin regarding how much they hold of the debt (i.e., their Euler equation must hold with equality). If instead the Euler equation of all households held with strict inequality, no household would want to hold the outstanding bonds. There would be excess supply of bonds and excess demand for consumption goods. Since firms must meet demand at posted prices, they would produce more, which would increase income and consumption and lead marginal utility of consumption to fall. This process would continue until the agents with the lowest marginal utility of consumption—the employed agents—were willing to hold the bonds. Thus, for any positive ε the Euler equation of the employed households will hold with equality in equilibrium and the equilibrium we focus on is the limit of these equilibria as ε goes to zero.

From the budget constraint of the employed we have $C_t = Y_t$. Substituting into the employed Euler equation we have

$$(C_t/(1-u))^{-\gamma} = \beta e^{qt} R_t \mathbb{E}_t [(1-u)(C_{t+1}/(1-u))^{-\gamma} + u(mA_{t+1})^{-\gamma}] \quad (44)$$

Log-linearizing this equation and using the Fisher equation yields

$$\hat{C}_t = -\frac{1}{\gamma} [i_t - \mathbb{E}\pi_{t+1} - r_t^n] + \alpha \mathbb{E} [\hat{C}_{t+1}] + (1-\alpha) \mathbb{E} [\hat{A}_{t+1}], \quad (45)$$

where

$$\alpha \equiv \frac{1}{1 + \frac{u}{1-u} \left(\frac{\bar{c}_e}{\bar{m}A}\right)^\gamma}, \quad (46)$$

$$r_t^n \equiv -qt, \quad (47)$$

and bars denote steady state values and hats denote log deviations from steady state. Notice that α is decreasing in the probability of becoming unemployed, u , and the consumption impact of unemployment, $\bar{c}_e/(\bar{A}m)$.

The labor supply condition for the employed is

$$c_{e,t}^{-\gamma} W_t = \ell_t^\psi.$$

Using the fact that aggregate labor supply is $L_t = (1-u)\ell_t$ and aggregate consumption is given by equation (43), this last expression can be rewritten as

$$(1-u)^{\gamma+\psi} C_t^{-\gamma} W_t = L_t^\psi.$$

Log-linearizing this equation and using $C_t = Y_t$ and $L_t = Y_t/A_t$ yields

$$\hat{W}_t = (\psi + \gamma)\hat{Y}_t - (\psi + 1)\hat{A}_t.$$

Finally, market incompleteness does not change the relationship between inflation and marginal cost so we can write the Phillips curve as

$$\pi_t - \pi_{t-1} = \beta \mathbb{E}_t [\pi_{t+1} - \pi_t] + \kappa \xi \left[(\psi + \gamma)\hat{Y}_t - (\psi + 1)\hat{A}_t \right] + \hat{\mu}_t, \quad (48)$$

Figure 11 shows that $\alpha = 0.97$ and $1/\gamma = 3/8$ provides a close fit to the response of the incomplete markets model in the 20-quarter-ahead forward guidance experiment. As our focus is on the Euler equation, we recalibrate ψ and Ω , which determines ξ , so the coefficients on \hat{Y} and \hat{A} in the Phillips curve are unaffected when we change γ . This requires $\psi = 4$ and $\Omega = 94.4$.




Figure A.2: General and partial equilibrium response to a one percentage point reduction in real interest rates in period 20 with a unit intertemporal elasticity of substitution.

References

- AIYAGARI, S. R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109(3), 659–684.
- ASCARI, G. (2004): “Staggered Prices and Trend Inflation: Some Nuisances,” *Review of Economic Dynamics*, 7, 642–667.
- BEWLEY, T. (undated): *Interest Bearing Money and the Equilibrium Stock of Capital*. Manuscript.
- BLANCHARD, O. J. (1985): “Debt, Deficits, and Finite Horizons,” *Journal of Political Economy*, 93(2), 223–247.
- BROYDEN, C. G. (1965): “A Class of Methods for Solving Nonlinear Simultaneous Equations,” *Mathematics of Computation*, 19(92), pp. 577–593.
- CABALLERO, R., AND E. FARHI (2014): “The Safety Trap,” Working Paper, MIT.
- CALVO, G. A. (1983): “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, 383–398.
- CAMPBELL, J. R., C. L. EVANS, J. D. FISHER, AND A. JUSTINIANO (2012): “Macroeconomic Effects of Federal Reserve Forward Guidance,” *Brookings Papers on Economic Activity*, 2012(1), 1–80.
- CARLSTROM, C. T., T. S. FUERST, AND M. PAUSTIAN (2012): “Inflation and Output in New Keynesian Model a Transient Interest Rate Peg,” Working Paper, Federal Reserve Bank of Cleveland.
- CARROLL, C. D. (2006): “The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems,” *Economics Letters*, 91(3), 312 – 320.
- CHETTY, R. (2012): “Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply,” *Econometrica*, 80(3), 969–1018.
- CHRISTIANO, L., M. EICHENBAUM, AND S. REBELO (2011): “When Is the Government Spending Multiplier Large?,” *Journal of Political Economy*, 119(1), 78–121.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 115, 1–45.
- DEL NEGRO, M., M. GIANNONI, AND C. PATTERSON (2013): “The Forward Guidance Puzzle,” Working Paper, Federal Reserve Bank of New York.
- DOTSEY, M., AND R. G. KING (2005): “Implications of State-dependent Pricing for Dynamic Macroeconomic Models,” *Journal of Monetary Economics*, 52, 213–242.
- EGGERTSSON, G. B., AND M. WOODFORD (2003): “The Zero Bound on Interest Rates and Optimal Monetary Policy,” *Brookings Papers on Economic Activity*, 2003(1), 139–211.
- FLODEN, M., AND J. LINDÉ (2001): “Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?,” *Review of Economic Dynamics*, 4(2), 406–437.
- GALL, J. (2008): *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press, Princeton, NJ.
- GORNEMANN, N., K. KUESTER, AND M. NAKAJIMA (2014): “Doves for the Rich, Hawks for the

- Poor? Distributional Consequences of Monetary Policy,” Working Paper, Federal Reserve Bank of Philadelphia.
- GUERRIERI, V., AND G. LORENZONI (2011): “Credit Crises, Precautionary Savings, and the Liquidity Trap,” Working Paper, University of Chicago.
- GURKAYNAK, R. S., B. SACK, AND E. T. SWANSON (2005): “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements,” *International Journal of Central Banking*, 1, 55–93.
- GUVENEN, F., S. OZKAN, AND J. SONG (2014): “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 122(3), pp. 621–660.
- HUGGETT, M. (1993): “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies,” *Journal of Economic Dynamics and Control*, 17(5-6), 953–969.
- JAPPELLI, T., AND L. PISTAFERRI (2010): “The Consumption Response to Income Changes,” *Annual Review of Economics*, 2, 479–506.
- JOHNSON, D. S., J. A. PARKER, AND N. S. SOULELES (2006): “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 96(5), 1589–1610.
- JUDD, K. L. (1992): “Projection methods for solving aggregate growth models,” *Journal of Economic Theory*, 58(2), 410–452.
- JUILLARD, M. (1996): *Dynare: A program for the resolution and simulation of dynamic models with forward variables through the use of a relaxation algorithm*, vol. 9602. CEPREMAP Paris.
- KAPLAN, G., AND G. VIOLANTE (2014): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82(4), 1199–1239.
- KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): “The Wealthy Hand-to-Mouth,” *Brookings Papers on Economic Activity*, 2014(Spring), 77–138.
- KILEY, M. T. (2014): “Policy Paradoxes in the New Keynesian Model,” Finance and Economics Discussion Series, Federal Reserve Board.
- KIMBALL, M. (1995): “The Quantitative Analytics of the Basic Neomonetarist Model,” *Journal of Money, Credit and Banking*, 27(4), 1241–1277.
- KOPECKY, K. A., AND R. M. SUEN (2010): “Finite state Markov-chain approximations to highly persistent processes,” *Review of Economic Dynamics*, 13(3), 701 – 714.
- KRUSELL, P., T. MUKOYAMA, AND A. A. SMITH (2011): “Asset Prices in a Huggett Economy,” *Journal of Economic Theory*, 146(3), 812–844.
- KRUSELL, P., AND A. A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.
- MCKAY, A., AND R. REIS (2014): “The Role of Automatic Stabilizers in the US Business Cycle,” Working Paper, Boston University.
- NAKAMURA, E., AND J. STEINSSON (2015): “High Frequency Identification of Monetary Non-Neutrality,” Working Paper, Columbia University.
- NISTICO, S. (2012): “Monetary Policy and Stock-Price Dynamics in a DSGE Framework,” *Journal of Macroeconomics*, 34, 126–146.

- OH, H., AND R. REIS (2012): “Targeted transfers and the fiscal response to the great recession,” *Journal of Monetary Economics*, 59(S), S50–S64.
- PARKER, J., N. SOULELES, D. JOHNSON, AND R. MCCLELLAND (2013): “Consumer Spending and the Economic Stimulus Payments of 2008,” *American Economic Review*, 103(6), 2530–53.
- PARKER, J. A. (1999): “The Reaction of Household Consumption to Predictable Changes in Social Security Taxes,” *The American Economic Review*, 89(4), pp. 959–973.
- PIERGALLINI, A. (2006): “Real Balance Effects and Monetary Policy,” *Economic Inquiry*, 44(3), 497–511.
- RAVN, M., AND V. STERK (2013): “Job Uncertainty and Deep Recessions,” UCL Working Paper.
- REITER, M. (2009): “Solving Heterogeneous-agent Models by Projection and Perturbation,” *Journal of Economic Dynamics and Control*, 33(3), 649–665.
- ROUWENHORST, K. G. (1995): “Asset pricing implications of equilibrium business cycle models,” in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, pp. 294–330, Princeton, N.J. Princeton University Press.
- SIMS, C. A. (2002): “Solving Linear Rational Expectations Models,” *Computational Economics*, 20(1-2), 1–20.
- SMETS, F., AND R. WOUTERS (2007): “Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97(3), 586–606.
- WERNING, I. (2012): “Managing a Liquidity Trap: Monetary and Fiscal Policy,” Working Paper, MIT.
- WOODFORD, M. (2003): *Interest and Prices*. Princeton University Press, Princeton, NJ.
- YOUNG, E. R. (2010): “Solving the Incomplete Markets Model with Aggregate Uncertainty using the Krusell-Smith Algorithm and Non-Stochastic Simulations,” *Journal of Economic Dynamics and Control*, 34(1), 36 – 41.