

Empirical Properties of Inflation Expectations and the Zero Lower Bound*

Mirko Wiederholt
Sciences Po and CEPR

This version: May 2019

Abstract

Survey inflation expectations adjust slowly after shocks and are heterogeneous. The paper therefore solves a New Keynesian model with dispersed information on the household side. The information friction changes results regarding shock propagation and policy effectiveness at the zero lower bound on nominal interest rates. The deflationary spiral in bad states of the world takes off slowly. Communication about the aggregate state affects consumption (and the sign of this effect reverses when the zero lower bound binds). Forward guidance and fiscal policy are less powerful. The effects of dispersed information are larger in states with a smaller prior probability.

Keywords: zero lower bound, monetary policy, fiscal policy, information frictions. (*JEL:* D83, E31, E32, E52).

*E-mail: Mirko.Wiederholt@gmail.com. I thank Christian Hellwig, Lena Körber, Ricardo Reis, Chris Sims, and seminar and conference participants at Banque de France, Berkeley, Bilkent, Bocconi, Boston University, Bundesbank, CEMFI, CERGE-EI, Chicago Fed, Duke, ECB, LSE, Minneapolis Fed, NBERSI 2015, New York Fed, NYU, Norges Bank, Paris School of Economics, SED 2012, Stanford, Toulouse School of Economics, UCL, University of Amsterdam, University of Mannheim, and University of Minnesota for helpful comments. This paper extends and subsumes a 2012 working paper that circulated under the title “Dispersed Inflation Expectations and the Zero Lower Bound.” Sebastian Weiske provided excellent research assistance.

1 Introduction

The main policy rate of the central bank is currently at zero (or close to zero) in 22 out of 34 OECD countries.¹ Moreover, unless central banks significantly raise their inflation targets, policy rates are likely to be at zero again in the future. This raises two important questions. How are shocks propagated when the zero lower bound on the nominal interest rate is binding? What policies are effective under these circumstances?

In New Keynesian models with a binding zero lower bound, movements in household inflation expectations are of great importance for shock propagation and policy effectiveness. Consumption must satisfy the Euler equation for a short-term bond. Furthermore, when monetary policy is constrained by the zero lower bound, the short-term nominal interest rate is constant. Suppose that the zero lower bound is known to bind for at least T periods and consider the log-linearized Euler equation for the short-term bond. Solving this Euler equation forward implies that current consumption depends on a constant, expected inflation over the next T periods, and expected consumption in period $T+1$. Hence, all the amplification of the shock comes from changes in expected inflation over the next T periods and changes in expected consumption in period $T+1$. In addition, government policies like forward guidance or increases in government purchases affect current consumption by changing households' inflation expectations and/or expected consumption in period $T+1$.

It is therefore desirable to model inflation expectations in a way that is consistent with data. Recent papers studying survey data on inflation expectations find that agents' average inflation expectation responds slowly to realized shocks to future inflation (Coibion and Gorodnichenko, 2012) and that agents have heterogeneous inflation expectations (see, e.g., Armantier et al., 2011). By contrast, in any model with perfect information and rational expectations, agents' inflation expectation responds instantly and one-for-one to any realized shock to future inflation, because the shock is in the information set of the agents and agents know how the shock affects future inflation. Furthermore, all agents have the same expectation of aggregate inflation, because all agents have the same information set and the same perceived law of motion for inflation.

Motivated by the frequent use of New Keynesian models with a zero lower bound in policy analysis, the importance of household inflation expectations in those models, and the tension be-

¹The main policy rate is weakly smaller than 25 basis points in 22 out of 34 OECD countries.

tween model properties and data properties of inflation expectations, this paper studies a New Keynesian model with a zero lower bound and dispersed information on the household side. The assumption that households have less than perfect information and different pieces of information yields the slow adjustment and the heterogeneity of inflation expectations. It turns out that this model with sluggish and dispersed household inflation expectations has quite different implications for shock propagation and policy effectiveness.

First of all, the deflationary spiral in bad states of the world is less severe than under perfect information. The slow adjustment of household inflation expectations implies that consumption falls slowly. This effect is amplified by the fact that consumption choices of different households are strategic complements when the zero lower bound is binding.² As a result, even a small information friction on the household side can generate large deviations from the perfect-information outcome.

Second, central bank communication about the state of the economy (without any change in current or future policy) affects aggregate consumption, and the sign of this effect changes when the zero lower bound is binding. The reason for the sign change is that downward movements in household inflation expectations are stabilizing when the zero lower bound is not binding, while they are destabilizing when the zero lower bound is binding. For comparison, in a standard New Keynesian model, central bank communication about the current state has no effect on consumption, because the current state is common knowledge.

Third, a central bank commitment to holding the policy rate at zero for longer beyond what is justified by contemporaneous economic conditions (“forward guidance”) has smaller effects on current consumption than under perfect information (and can even reduce current consumption). The reason is twofold. Households that do not update their inflation expectations do not change their consumption. Households that do update their inflation expectations experience a positive and a negative effect on their inflation expectations. Future monetary policy is expected to be more expansionary (if the announcement is credible), which raises inflation expectations. But today’s commitment to such a policy reveals that the economy is in a bad state, which reduces inflation expectations.

Fourth, the government spending multiplier is smaller than under perfect information. The reason is again twofold. Households that do not update their inflation expectations do not change

²Consumption choices of different households are strategic substitutes when the zero lower bound is not binding.

their consumption. Households that do update inflation expectations in response to the policy announcement experience a positive effect and a negative effect on their inflation expectations. State-contingent expansionary fiscal policy raises inflation in a given state and reveals the state.

Fifth, if individual beliefs about inflation cover regions where the zero lower bound is binding and regions where the zero lower bound is non-binding, individual consumption depends on several moments of the conditional distribution of inflation. The reason is the kink in the monetary policy rule due to the zero lower bound on the nominal interest rate. As a result, aggregate consumption is no longer a function of only the average inflation expectation.

Finally, the deviations from the perfect-information equilibrium are larger in states of the world that have a smaller prior probability. I therefore believe the results listed above are particularly relevant for thinking about the Great Recession in the U.S. and the sovereign debt crisis in Europe.

In a calibrated version of the model, the response of consumption on impact of a large, negative shock equals about 1/3 of the value under perfect information (while the cumulative response of consumption equals about 1/2 of the value under perfect information). Here the parameter controlling the speed at which households update inflation expectations over time is chosen so as to match the speed at which the gap between inflation and inflation expectations closes in the data, according to the estimates by Coibion and Gorodnichenko (2012).

Literature review: My work builds on the existing literature on the implications of the zero lower bound for shock propagation and policy effectiveness. In contrast to all the existing literature on the zero lower bound, households' short-term inflation expectations are sluggish and dispersed.³ In terms of modelling, I follow closely the seminal contributions on the zero lower bound (e.g., Krugman, 1998, Eggertsson and Woodford, 2003) to isolate the implications of slow adjustment and heterogeneity of household inflation expectations.

The part of the paper on forward guidance is related to the emerging literature on the forward guidance puzzle. This literature argues that the effects of forward guidance in benchmark

³Andrade et al. (2015) present a model with heterogeneous *medium-term* inflation expectations. All agents agree that the nominal interest rate will equal zero for exactly N periods, but agents disagree about the reason. Pessimists believe that fundamentals revert back to normal after N periods and the central bank is of the no-commitment type. Optimists believe that fundamentals revert back to normal after $N' < N$ periods and the central bank is of the commitment type. In that model, all agents have the same, correct beliefs about the path of the economy up to period N' . The two groups only have different beliefs about the path of the economy in periods N' to N .

New Keynesian models seem unreasonably large and offers explanations for weaker effects of forward guidance (Carlstrom, Fuerst, and Paustian, 2012, Del Negro, Giannoni, and Patterson, 2012, McKay, Nakamura, and Steinsson, 2015).

The part of the paper on fiscal policy is related to the literature on the government spending multiplier at the zero lower bound (e.g., Christiano, Eichenbaum, and Rebelo, 2011, Woodford, 2011), especially the work arguing that the government spending multiplier at the zero lower bound may not be as large as predicted by the benchmark New Keynesian model with a zero lower bound, e.g., because of distortionary taxation (Uhlig and Drautzburg, 2013) or a non-fundamental liquidity trap (Mertens and Ravn, 2014).

The paper is related to the literature on business cycle models with information frictions on the household side (Mankiw and Reis, 2006, Lorenzoni, 2009, Angeletos and La'O, 2013, Maćkowiak and Wiederholt, 2015). In contrast to the existing work on this topic, I study the implications of the zero lower bound. This difference turns out to be important, e.g., because at the zero lower bound, movements in household inflation expectations are destabilizing (rather than stabilizing), actions of different households are strategic complements (rather than strategic substitutes), and the monetary policy rule has a kink.

My work also builds on the empirical literature on inflation expectations (e.g., Coibion and Gorodnichenko, 2012 and 2015, Armantier et al., 2011), which motivated this paper.

The rest of the paper is organized as follows. Section 2 explains why the modeling of inflation expectations matters. Section 3 presents the model. Section 4 states the optimality conditions of households and firms. Section 5 solves a special case of the model analytically. Section 6 solves a calibrated version of the model numerically and presents quantitative results. Section 7 discusses forward guidance. Section 8 discusses the government spending multiplier. Section 9 turns to the missing deflation puzzle. Section 10 concludes.

2 Why the modeling of inflation expectations matters

2.1 Firms

Consider the New Keynesian Phillips curve

$$\pi_t = \kappa \hat{y}_t + \beta E_t [\pi_{t+1}], \tag{1}$$

where π_t denotes quarterly inflation, \hat{y}_t denotes the output gap, and $E_t[\pi_{t+1}]$ denotes the expectation of next quarter's inflation. Suppose inflation follows an AR(1) process and thus $E_t[\pi_{t+1}] = \rho\pi_t$. Then

$$\pi_t = \frac{1}{1 - \beta\rho} \kappa \hat{y}_t. \quad (2)$$

With a persistence parameter of $\rho = 0.95$ the term $\frac{1}{1 - \beta\rho}$ equals 20.⁴ With a persistence parameter of $\rho = 0.75$ the term $\frac{1}{1 - \beta\rho}$ falls from 20 to 4. Hence, perceived persistence of inflation matters a great deal for the inflation response for a given output gap. The Great Recession and the subsequent slump were highly persistent, implying that a perfect foresight model will have a high value of $\frac{1}{1 - \beta\rho}$; but the slump was believed to be highly transitory in the first couple of years and thus a low value for the perceived persistence of inflation seems more appropriate.

Similarly, suppose that the nowcast of inflation moves less than one-for-one with inflation and thus $E_t[\pi_{t+1}] = \rho E_t[\pi_t] = \rho K \pi_t$, where $K = 1$ indicates that the nowcast of inflation moves one-for-one with inflation and $K < 1$ indicates that the nowcast of inflation underreacts to inflation. Then

$$\pi_t = \frac{1}{1 - \beta\rho K} \kappa \hat{y}_t.$$

With an actual and perceived persistence parameter of $\rho = 0.95$ and $K = 1$ (no underreaction of the nowcast of inflation to inflation) the term $\frac{1}{1 - \beta\rho K}$ equals 20. With $K = 0.8$ (underreaction of the nowcast of inflation by 20%) the term $\frac{1}{1 - \beta\rho K}$ falls again from 20 to 4. Hence, underreaction of the nowcast of inflation to inflation also generates much smaller inflation responses.

The missing deflation puzzle can be illustrated with equation (2). Real GDP fell by 4 percent between the fourth quarter of 2007 and the second quarter of 2009. Without a simultaneous drop in productivity, this implies that the output gap fell by four percentage points. Equation (2) predicts a fall in inflation. By how much? A standard calibration of a New Keynesian model yields $\kappa = 0.045$. With $\kappa = 0.045$, $\beta = 0.99$ and $\rho = 0.95$ equation (2) becomes $\pi_t = (0.76) \hat{y}_t$. A fall in the output gap by four percentage points should lead to a fall in quarterly non-annualized inflation by three percentage points. A fall in quarterly *annualized* inflation by twelve percentage points certainly did not happen.⁵ Authors have proposed various solutions to the missing deflation puzzle: a much lower value for κ (Christiano, Eichenbaum and Rebelo, 2011), a small output gap (Christiano,

⁴For ease of exposition, β is set equal to one. Setting $\beta = 0.99$ does not affect the main point of this paragraph.

⁵See Section 6.1 for a standard calibration of a New Keynesian model that yields $\kappa = 0.045$.

Eichenbaum and Trabandt, 2015), new channels that push inflation up (e.g., Gilchrist, Schoenle, Sim and Zakrajšek, 2017), and a non-linear New Keynesian Phillips curve in combination with a Kimball aggregator (Lindé and Trabandt, 2019). The explanation I propose instead is a modeling of expectations that is consistent with survey data, along the lines of the previous two paragraphs.

Figure 6 illustrates that this is going to work. The blue diamonds show data: quarterly core PCE inflation before, during, and after the Great Recession at a quarterly frequency. The red squares show quarterly inflation implied by the New Keynesian Phillips curve: the right-hand side of equation (1), where parameter values are the same as before ($\kappa = 0.045$ and $\beta = 0.99$), the measure of the output gap is taken from Fernald, Hall, Stock, and Watson (2017), and the measure of expected inflation is the average expectation of next quarter’s core PCE inflation from the Survey of Professional Forecasters.⁶ The red squares track the blue diamonds remarkably well, even though the New Keynesian Phillips curve is completely standard and the values of κ and β are standard. Hence, if a completely standard New Keynesian model matches the measured output gap and measured inflation expectations during the Great Recession, the model also matches measured inflation during the Great Recession.

2.2 Consumers

Inflation expectations are also important for consumption-savings decisions. The reason is simple. The cost of saving is felt today (forgone consumption). The benefit of saving is enjoyed in the future. In many macroeconomic models, there is a consumption Euler equation for a nominal bond, which with CRRA preferences reads

$$C_t^{-\gamma} = E_t \left[\beta \frac{R_t}{\Pi_{t+1}} C_{t+1}^{-\gamma} \right],$$

where $C_t^{-\gamma}$ is the marginal utility of consumption, β is the discount factor, R_t is the gross nominal interest rate on bond holdings between periods t and $t + 1$, and Π_{t+1} is the inflation factor. In log-linearized form the equation becomes

$$c_t = E_t \left[-\frac{1}{\gamma} (r_t - \pi_{t+1}) + c_{t+1} \right],$$

⁶The red squares are simply 0.045 times the time series from Fernald, Hall, Stock, and Watson (2017) plus 0.99 times the time series from the Survey of Professional Forecasters.

where small letters denote log-deviations from the non-stochastic steady state. On the left-hand side we see consumption. On the right-hand side we see expectations. In a model of this form, the modeling of expectations obviously matters for consumption.

When the effective lower bound on the nominal interest rate is binding, expected inflation is particularly important because it is the only variable that moves the perceived real return on saving in a nominal bond. Suppose that a consumer expects the lower bound for the nominal interest rate to be binding for exactly N periods. Solving the last equation forward and using the law of iterated expectations yields

$$c_t = \frac{N}{\gamma} (-r) + \frac{1}{\gamma} \sum_{j=1}^N E_t [\pi_{t+j}] + E_t [c_{t+N}].$$

The term $(-r) > 0$ is the log difference between the gross nominal interest rate in the non-stochastic steady state and the gross nominal interest rate at the effective lower bound. Importantly, the term is strictly positive. The term $E_t [c_{t+N}]$ may be negative but has to be small because otherwise the consumer would not expect the lower bound to stop binding in period $t + N$. Hence, any large, temporary drop in consumption in a model of this form has to come from the expected inflation term. To evaluate such a model quantitatively, it seems crucial to model inflation expectations in a way that is consistent with data and to model the link between consumption-savings decisions and inflation expectations in a way that is consistent with data.

3 Model

The economy is populated by households, firms, and a government. The model setup is close to a standard New Keynesian model with a zero lower bound (in particular, Eggertsson and Woodford, 2003), but in contrast to all the existing literature on the zero lower bound households' short-term inflation expectations adjust slowly to shocks and are heterogeneous.

Households. The economy is populated by a continuum of households of mass one. Households are indexed by $i \in [0, 1]$. The preferences of household i are given by

$$E_0^i \left[\sum_{t=0}^{\infty} \beta^t e^{\xi_{i,t}} \left(\frac{C_{i,t}^{1-\gamma} - 1}{1-\gamma} - N_{i,t} \right) \right],$$

where $C_{i,t}$ is consumption of the household in period t , $N_{i,t}$ is labor supply of the household in period t , and $\xi_{i,t}$ is a preference shock. Here E_0^i is the expectation operator conditioned on the

information of the household in period zero. The parameter $\beta \in (0, 1)$ is a discount factor and $\gamma > 0$ is the inverse of the intertemporal elasticity of substitution.

Following Eggertsson and Woodford (2003), I study the response of the economy to a temporary increase in households' desire to save. In period zero, each household is hit by a preference shock $\xi_{i,0} \in \{\xi_L, \xi_H\}$ with $\xi_L < \xi_H < 0$. To obtain a closed-form solution of the model, I initially assume stochastic decay of the preference shock. That is, in every period $t \geq 1$, $\xi_{i,t} = \xi_{i,t-1}$ with probability μ and $\xi_{i,t}$ returns permanently to its normal value of zero with probability $1 - \mu$. The return to the normal value of zero occurs at the same time for all households. Once the main properties of the model have been shown formally, I also solve the model numerically with deterministic decay. That is, in periods $t \geq 1$, $\xi_{i,t} = \rho \xi_{i,t-1}$ with $\rho \in (0, 1)$.

In contrast to the existing literature, there are two possible aggregate exogenous states in period zero and there is heterogeneity across households in the value of the preference shock. Let λ denote the mass of households with $\xi_{i,0} = \xi_H$. Let $1 - \lambda$ denote the mass of households with $\xi_{i,0} = \xi_L$. The two possible aggregate exogenous states in period zero differ in terms of the mass of households who experience the high realization of the preference shock: $\lambda \in \{\lambda_{bad}, \lambda_{good}\}$ with $0 < \lambda_{bad} < \lambda_{good} < 1$. To study the implications of imperfect information about the aggregate state, one has to introduce at least two possible aggregate states. To ensure that the own preference shock does not perfectly reveal the aggregate state, one has to assume that each realization of the preference shock is possible in both aggregate states. In the following, think of the good state as an aggregate shock that would cause a severe recession under perfect information. Think of the bad state as an aggregate shock that would cause the worst recession in a century under perfect information. Let $\theta \in (0, 1)$ denote the prior probability of the good state.⁷

Households can save or borrow by holding (positive or negative amounts of) nominal government bonds. Let $B_{i,t}$ denote the bond holdings of household i between periods t and $t + 1$. The evolution of the bond holdings of household i is given by

$$B_{i,t} = R_{t-1}B_{i,t-1} + W_{i,t}N_{i,t} + D_{i,t} - P_t C_{i,t} + Z_{i,t}.$$

Here R_{t-1} denotes the gross nominal interest rate on bond holdings between periods $t - 1$ and t , $W_{i,t}$ is the nominal wage rate for labor supplied by household i in period t , and $D_{i,t}$ denotes the

⁷Throughout the paper, the recession is more severe in the state with the lower λ . Furthermore, the efficient allocation features no recession. I therefore refer to the state with the lower λ as the bad state.

difference between dividends received by the household in period t and nominal lump-sum taxes paid by the household in period t . The term $P_t C_{i,t}$ is the household's consumption expenditure, where P_t denotes the price of the final good in period t . The term $Z_{i,t}$ is a net transfer that is specified below. The household can save or borrow (i.e., bond holdings can be positive or negative), but the household cannot run a Ponzi scheme. All households have the same initial bond holdings in period minus one.

For simplicity, I assume that households can trade state-contingent claims with one another in period minus one (i.e., when all households are still ex-ante identical). Recall that each household is hit by a preference shock in period zero and preference shocks of all households revert permanently back to zero in a stochastic period. The contingent claims are settled in that period, denoted $T \geq 1$. A state-contingent claim specifies a payment to the household who purchased the claim that is contingent on the individual history of the household and the aggregate history of the economy (e.g., the claim is contingent on $\xi_{i,0}$, λ , and T). The term $Z_{i,t}$ in the flow budget constraint is the net transfer associated with these state-contingent claims. This term equals zero in all periods apart from period T . The fact that agents can trade these state-contingent claims in period minus one implies that in equilibrium all households will have the same post-transfer wealth in period T . This simplifies the analysis: to solve for consumption of each household one does not have to keep track of the dynamics of the wealth distribution in periods $0 \leq t \leq T - 1$. A similar assumption is made in Lucas (1990), Lorenzoni (2010), and Curdia and Woodford (2011).

Finally, let us turn to expectation formation by households. Households have correct prior beliefs about the probability of the good state and of the bad state, the distribution of types in the two states, and the dynamics of inflation (and other endogenous variables) in the two states. In period zero, each household observes the realization of the own preference shock and updates his or her beliefs about the aggregate shock and future inflation using Bayes' rule. Thereafter, there is slow updating of beliefs, as in Mankiw and Reis (2002, 2006). In every period $0 \leq t < T$, a constant fraction $\omega \in [0, 1]$ of randomly selected households learns the size of the aggregate shock that hit the economy in period zero and moves to perfect-information rational expectations about inflation. The remaining households do not update their beliefs about the aggregate shock and inflation.

The idea is that processing information is costly and therefore in every period only a fraction of households update their beliefs about the aggregate state of the economy and future inflation by

processing information contained in the news media or prices. The standard assumption is $\omega = 1$, i.e., all households move instantly to perfect-information rational expectations about inflation. However, this assumption is inconsistent with both the slow adjustment and the heterogeneity of household inflation expectations in the data.

In the calibrated version of the model, I set the value of ω so as to match the empirical speed at which the gap between future inflation and the average expectation of future inflation closes after shocks, according to the estimates by Coibion and Gorodnichenko (2012).

Firms. There are final good firms and intermediate good firms. To illustrate as clearly as possible the effects of household imperfect information, firms are assumed to have perfect information.⁸ The final good is produced by competitive firms using the technology

$$Y_t = \left(\int_0^1 Y_{j,t}^{\frac{\psi-1}{\psi}} dj \right)^{\frac{\psi}{\psi-1}},$$

where Y_t denotes output of the final good, $Y_{j,t}$ is input of intermediate good j , and $\psi > 1$ is the elasticity of substitution between intermediate goods. Final good firms have fully flexible prices. Profit maximization of final good firms implies the following demand function for good j

$$Y_{j,t} = \left(\frac{P_{j,t}}{P_t} \right)^{-\psi} Y_t,$$

where $P_{j,t}$ denotes the price of intermediate good j and P_t is the price of the final good. Furthermore, the zero profit condition of final good firms implies

$$P_t = \left(\int_0^1 P_{j,t}^{1-\psi} dj \right)^{\frac{1}{1-\psi}}.$$

The intermediate good j is produced by a monopolist using the technology

$$Y_{j,t} = N_{j,t}^\varrho \quad \text{with} \quad N_{j,t} = \left(\int_0^1 N_{i,j,t}^{\frac{\eta-1}{\eta}} di \right)^{\frac{\eta}{\eta-1}}.$$

Here $Y_{j,t}$ is output, $N_{j,t}$ is composite labor input, and $N_{i,j,t}$ is type i labor input of monopolist j . Type i labor is labor supplied by household i . The parameter $\varrho \in (0, 1]$ is the elasticity of output with respect to composite labor and $\eta > 1$ is the elasticity of substitution between types of labor. Cost minimization implies that the demand for type i labor of monopolist j is given by

$$N_{i,j,t} = \left(\frac{W_{i,t}}{W_t} \right)^{-\eta} N_{j,t},$$

⁸I use the term “perfect information” in the usual way: an agent has perfect information if in every period t the agent knows the entire history of the economy up to and including period t .

where

$$W_t = \left(\int_0^1 W_{i,t}^{1-\eta} di \right)^{\frac{1}{1-\eta}}.$$

Furthermore, cost minimization implies that the wage bill of monopolist j in period t equals $W_t N_{j,t}$. Monopolists producing intermediate goods are subject to a price-setting friction, as in Calvo (1983). Each monopolist can optimize its price with probability $1 - \alpha$ in any given period. With probability α the monopolist producing good j sets the price

$$P_{j,t} = P_{j,t-1}.$$

How monopolists value profit in different states of the world is determined by the ownership structure. I assume that each monopolist is owned by a single household and takes the household's marginal utility of consumption as given, because the household also owns many other firms.⁹

Monetary policy. The monetary authority sets the gross nominal interest rate according to the rule

$$R_t = \max \left\{ 1, R \Pi_t^\phi \right\},$$

where $R = (1/\beta)$ denotes the nominal interest rate in the non-stochastic steady state with zero inflation, $\Pi_t = (P_t/P_{t-1})$ is the inflation rate, and $\phi > 1$ is a parameter. According to the last equation, the monetary authority follows a Taylor rule as long as the implied net nominal interest rate is non-negative and the monetary authority sets the net nominal interest rate to zero otherwise.

Fiscal policy. The fiscal authority can purchase units of the final good and can finance these purchases with current lump-sum taxes or future lump-sum taxes. The government flow budget constraint in period t reads

$$T_t + B_t = R_{t-1} B_{t-1} + P_t G_t.$$

The government has to finance maturing nominal government bonds and any purchases of the final good, denoted G_t . The government can collect lump-sum taxes, denoted T_t , or issue new bonds.

⁹To ensure that households are ex-ante identical in period minus one, ownership is assigned randomly in period zero. The individual history of a household in period T consists of the realization of the preference shock in period zero, the realization of ownership, and the date at which the household moves to perfect-information rational expectations. Since the state-contingent claims specify a payment that is contingent on the individual history, in equilibrium all households have the same post-transfer wealth in period T .

Until Section 7, $G_t = 0$ in every period. In Section 7, there are government purchases and I study the size of the government spending multiplier. A change in the path of government purchases is assumed to imply a change in the path of lump-sum taxes so as to maintain intertemporal government solvency.

4 Household and firm optimality

This section states the optimality conditions of households and firms and derives the New Keynesian Phillips curve for this economy. The New Keynesian Phillips curve is derived under different assumptions about wage setting: households set nominal wage rates and households set real wage rates. The sole purpose of the latter assumption is to obtain analytical solutions of the model. Section 4 presents closed-form solutions under the assumption that households set real wage rates. Section 5 shows that results get amplified when households set nominal wage rates.

Households. If households set nominal wage rates, the first-order conditions for consumption and the nominal wage rate read

$$C_{i,t}^{-\gamma} = E_t^i \left[\beta \frac{e^{\xi_{i,t+1}}}{e^{\xi_{i,t}}} \frac{R_t}{\Pi_{t+1}} C_{i,t+1}^{-\gamma} \right],$$

and

$$E_t^i \left[\frac{W_{i,t}}{P_t} \right] = \frac{\eta}{\eta - 1} C_{i,t}^\gamma.$$

Let $\tilde{W}_{i,t} = (W_{i,t}/P_t)$ denote the real wage rate for type i labor. If households set real wage rates, the first-order condition for consumption remains unchanged and the first-order condition for the real wage rate reads

$$\tilde{W}_{i,t} = \frac{\eta}{\eta - 1} C_{i,t}^\gamma.$$

Let small letters denote log-deviations from the non-stochastic steady state with zero inflation. Log-linearizing the consumption Euler equation around the non-stochastic steady state yields

$$c_{i,t} = E_t^i \left[-\frac{1}{\gamma} (\xi_{i,t+1} - \xi_{i,t} + r_t - \pi_{t+1}) + c_{i,t+1} \right]. \quad (3)$$

Log-linearizing the two wage setting equations around the non-stochastic steady state yields

$$w_{i,t} = \gamma c_{i,t} + E_t^i [p_t], \quad (4)$$

and

$$\tilde{w}_{i,t} = \gamma c_{i,t}. \quad (5)$$

Firms. An intermediate good firm j that can adjust its price in period t and is owned by household i sets the price

$$X_{j,t}^i = \arg \max_{P_{j,t} \in \mathbb{R}_{++}} E_t \left[\sum_{s=t}^{\infty} (\alpha\beta)^{s-t} \left(\frac{e^{\xi_{i,s}} C_{i,s}^{-\gamma} P_t}{e^{\xi_{i,t}} C_{i,t}^{-\gamma} P_s} \right) \left(P_{j,t} \left(\frac{P_{j,t}}{P_s} \right)^{-\psi} Y_s - W_s \left(\left(\frac{P_{j,t}}{P_s} \right)^{-\psi} Y_s \right)^{\frac{1}{\varrho}} \right) \right].$$

Log-linearizing the first-order condition for the adjustment price around the non-stochastic steady state with zero inflation yields

$$x_{j,t}^i = (1 - \alpha\beta) E_t \left[\sum_{s=t}^{\infty} (\alpha\beta)^{s-t} \left(p_s + \frac{1}{1 + \frac{1-\varrho}{\varrho}\psi} (w_s - p_s) + \frac{\frac{1-\varrho}{\varrho}}{1 + \frac{1-\varrho}{\varrho}\psi} y_s \right) \right].$$

Note that the log-linearized adjustment price is independent of who owns the firm and is the same for all adjusting firms. Therefore, one can drop the superscript i and the subscript j . Furthermore, the last equation can be stated in recursive form as

$$x_t = (1 - \alpha\beta) \left(p_t + \frac{1}{1 + \frac{1-\varrho}{\varrho}\psi} (w_t - p_t) + \frac{\frac{1-\varrho}{\varrho}}{1 + \frac{1-\varrho}{\varrho}\psi} y_t \right) + \alpha\beta E_t [x_{t+1}].$$

New Keynesian Phillips curve: Log-linearizing the equation for the price of the final good given in Section 2 and using the fact that adjusting firms are selected randomly and the log-linearized adjustment price is the same for all firms yields

$$p_t = \int_0^1 p_{j,t} dj = \alpha p_{t-1} + (1 - \alpha) x_t.$$

Using the last equation to substitute for the adjustment prices x_t and x_{t+1} in the previous equation and rearranging yields

$$\pi_t = \frac{(1 - \alpha)(1 - \alpha\beta)}{\alpha} \left(\frac{1}{1 + \frac{1-\varrho}{\varrho}\psi} (w_t - p_t) + \frac{\frac{1-\varrho}{\varrho}}{1 + \frac{1-\varrho}{\varrho}\psi} y_t \right) + \beta E_t [\pi_{t+1}]. \quad (6)$$

Finally, log-linearizing the equation for the wage index presented in Section 2 yields

$$w_t = \int_0^1 w_{i,t} di.$$

Substituting the log-linearized wage index and the wage setting equation (4) into equation (6) and using $y_t = c_t$, where c_t denotes aggregate consumption of the final good, yields a modified version

of the New Keynesian Phillips curve

$$\pi_t = \frac{(1 - \alpha)(1 - \alpha\beta)}{\alpha} \left(\frac{\gamma + \frac{1-\varrho}{\varrho}}{1 + \frac{1-\varrho}{\varrho}\psi} c_t + \frac{1}{1 + \frac{1-\varrho}{\varrho}\psi} (\bar{E}_t[p_t] - p_t) \right) + \beta E_t[\pi_{t+1}], \quad (7)$$

where $\bar{E}_t[p_t] = \int_0^1 E_t^i[p_t] di$ denotes households' average expectation of the price level. The new term in the modified New Keynesian Phillips curve reflects the following effect: when households' expectations of the price level are above the price level, households set nominal wage rates that are too high, which raises marginal costs and inflation. Using instead the wage setting equation (5) yields the standard New Keynesian Phillips curve

$$\pi_t = \frac{(1 - \alpha)(1 - \alpha\beta)}{\alpha} \frac{\gamma + \frac{1-\varrho}{\varrho}}{1 + \frac{1-\varrho}{\varrho}\psi} c_t + \beta E_t[\pi_{t+1}]. \quad (8)$$

5 Analytical solutions

This section presents analytical solutions of the model. I first solve the model under the assumption of perfect information to illustrate that movements in household inflation expectations are crucial for the propagation of shocks.

5.1 Perfect information

Perfect information is a special case of the model. When $\omega = 1$, all households learn instantly the exact size of the aggregate shock that hit the economy in period zero. Building on the work by Eggertsson and Woodford (2003), Christiano, Eichenbaum, and Rebelo (2011), and many others, I initially assume that preference shocks decay stochastically and I consider equilibria with the following properties: consumption, inflation, and the nominal interest rate are constant from period zero until preference shocks revert permanently back to zero, and the economy is in the non-stochastic steady state with zero inflation thereafter.

It is an equilibrium that all households have the same consumption level once preference shocks revert back to zero (i.e., in period T) because all households have the same post-transfer wealth in period T due to the trade in state-contingent claims in period minus one. Furthermore, $c_{i,t} = c_t = \pi_t = r_t = 0$ in every period $t \geq T$ satisfies the consumption Euler equation (3) with $\xi_{i,t} = \xi_{i,t+1} = 0$, the New Keynesian Phillips curve (7) with $\bar{E}_t[p_t] = p_t$, and the monetary policy rule.

Let us turn to consumption, inflation, and the nominal interest rate in periods $0 \leq t \leq T - 1$. Since consumption, inflation, and the nominal rate are constant across those periods but depend on the size of the aggregate shock, I replace the time subscript t by the state subscript $s \in \{good, bad\}$. Furthermore, since preference shocks do not change in the next period with probability μ , revert back to zero in the next period with probability $1 - \mu$, and households have perfect information, the consumption Euler equation (3) reduces to

$$c_{i,s} = -\frac{1}{\gamma} [(\mu - 1) \xi_{i,0} + r_s - \mu \pi_s] + \mu c_{i,s}.$$

Let $\bar{\xi}_s = \lambda_s \xi_H + (1 - \lambda_s) \xi_L$ denote the cross-sectional mean of the preference shock in state s . Integrating across households yields aggregate consumption in state s

$$c_s = -\frac{1}{\gamma} [(\mu - 1) \bar{\xi}_s + r_s - \mu \pi_s] + \mu c_s. \quad (9)$$

The New Keynesian Phillips curve (7) under perfect information reduces to

$$\pi_s = \kappa c_s + \beta \mu \pi_s, \quad (10)$$

with

$$\kappa = \frac{(1 - \alpha)(1 - \alpha\beta)}{\alpha} \frac{\gamma + \frac{1-\varrho}{\varrho}}{1 + \frac{1-\varrho}{\varrho}\psi}.$$

The monetary policy rule reads

$$r_s = \max \{-\ln(R), \phi \pi_s\}. \quad (11)$$

If the zero lower bound on the nominal interest rate is binding, $\max \{-\ln(R), \phi \pi_s\} = -\ln(R)$. Substituting $r_s = -\ln(R)$ and the New Keynesian Phillips curve (10) into the aggregated consumption Euler equation (9) and solving for consumption yields¹⁰

$$c_s = \frac{\frac{1}{\gamma} \bar{\xi}_s + \frac{1}{1-\mu} \ln(R)}{1 - \frac{\frac{1}{\gamma} \mu}{1-\mu} \frac{\kappa}{1-\beta\mu}}. \quad (12)$$

If the zero lower bound on the nominal interest rate is not binding, $\max \{-\ln(R), \phi \pi_s\} = \phi \pi_s$. Substituting $r_s = \phi \pi_s$ and the New Keynesian Phillips curve (10) into the aggregated consumption Euler equation (9) and solving for consumption yields

$$c_s = \frac{\frac{1}{\gamma} \bar{\xi}_s}{1 + \frac{\frac{1}{\gamma}(\phi - \mu)}{1-\mu} \frac{\kappa}{1-\beta\mu}}. \quad (13)$$

¹⁰Following common practice in the literature on the zero lower bound, I assume that parameters are such that the denominator in equation (12) is positive. See, e.g., Woodford (2011), Section IV.A.

Finally, the zero lower bound on the nominal interest rate is binding in state s if the cross-sectional mean of the preference shocks is sufficiently negative, that is,¹¹

$$\bar{\xi}_s < \bar{\xi}_{crit} = -\frac{1 + \frac{\frac{1}{\gamma}(\phi-\mu)}{1-\mu} \frac{\kappa}{1-\beta\mu}}{\frac{\frac{1}{\gamma}\phi}{1-\mu} \frac{\kappa}{1-\beta\mu}} \frac{1}{1-\mu} \ln(R). \quad (14)$$

An important insight in the literature on the zero lower bound is that the fall in consumption can be very large when the zero lower bound is binding. The condition $\bar{\xi}_s < \bar{\xi}_{crit}$ in combination with a positive denominator on the right-hand side of equation (12) implies a negative numerator on the right-hand side of equation (12). Moreover, the denominator on the right-hand side of equation (12) is a difference between two positive numbers that can be arbitrarily small in absolute value. Thus, even if the zero lower bound is only marginally binding in state s , consumption can fall by a very large amount in state s .

To understand why the fall in consumption can be so large, I propose the following decomposition. The aggregated consumption Euler equation (9) can be written as

$$c_s = \frac{1}{\gamma} \bar{\xi}_s - \frac{1}{1-\mu} r_s + \frac{1}{1-\mu} \mu \pi_s.$$

Consumption in state s equals the sum of three terms: the first term is the direct effect of the preference shock on consumption, the second term is the effect of the nominal interest rate on consumption, and the third term is the effect of expected inflation on consumption. Substituting in the equilibrium nominal interest rate when the zero lower bound is binding (i.e., $r_s = -\ln(R)$) and equilibrium inflation when the zero lower bound is binding yields

$$c_s = \frac{1}{\gamma} \bar{\xi}_s + \frac{1}{1-\mu} \ln(R) + \underbrace{\frac{1}{1-\mu} \mu \frac{\kappa}{1-\beta\mu} \frac{\frac{1}{\gamma} \bar{\xi}_s + \frac{1}{1-\mu} \ln(R)}{1 - \frac{\frac{1}{\gamma}\mu}{1-\mu} \frac{\kappa}{1-\beta\mu}}}_{\text{expected inflation}}.$$

The reason why the fall in consumption can be so large for a given size of the shock is the third term. The amplification of the shock comes from movements in household inflation expectations. If household inflation expectations did not move in response to the shock, consumption would be given by the sum of the first two terms. The second term is actually positive because the central bank can

¹¹Formally, substituting equations (10) and (12) (or equations (10) and (13)) into $-\ln(R) = \phi\pi_s$ and solving for $\bar{\xi}_s$ yields equation (14).

lower the nominal interest rate to some extent before the zero lower bound becomes binding. Hence, how one models household inflation expectations seems crucial for results concerning dynamics at the zero lower bound.

5.2 Imperfect information

To understand the implications of slow adjustment and heterogeneity of household inflation expectations in an economy with a zero lower bound on the nominal interest rate, let us turn to the model with dispersed information on the household side (i.e., $\omega < 1$). I first solve the model analytically in the following special case: households form beliefs about the aggregate state of the economy based only on their own local conditions (i.e., $\omega = 0$) and households set real wage rates. This special case of the model can be solved analytically. In Section 5, I relax these two assumptions. The main results remain unchanged.

As in the previous subsection, I consider equilibria of the following form: consumption, inflation, and the nominal interest rate are constant over time in periods $0 \leq t \leq T-1$, and the economy is in the non-stochastic steady state with zero inflation in periods $t \geq T$. The latter is an equilibrium for the same reasons as in the previous subsection: all households have the same post-transfer wealth in period T due to the trade in state-contingent claims in period minus one, and $c_{i,t} = c_t = \pi_t = r_t = 0$ in every period $t \geq T$ satisfies the consumption Euler equation (3) with $\xi_{i,t} = \xi_{i,t+1} = 0$, the New Keynesian Phillips curve (8), and the monetary policy rule.

Let us turn to consumption, inflation, and the nominal interest rate in periods $0 \leq t \leq T-1$. Since consumption, inflation, and the nominal interest rate are constant over time but depend on the size of the aggregate shock, I replace the time subscript t by the state subscript $s \in \{good, bad\}$. Since the random period T arrives with probability μ , the consumption Euler equation (3) can be written as

$$c_i = -\frac{1}{\gamma} [(\mu - 1) \xi_{i,0} + E^i [r_S - \mu\pi_S]] + \mu c_i. \quad (15)$$

Here $E^i [r_S - \mu\pi_S]$ is household i 's expectation of the real interest rate. The household's expectation has no time subscript because it is constant over time due to the assumption that $\omega = 0$. (This assumption is relaxed in the following section.) The household's expectation has no state subscript, because the household's expectation depends only on the household's type ($\xi_{i,0} = \xi_H$ or $\xi_{i,0} = \xi_L$). As a result, consumption depends only on the household's type.

Let p_H^{good} denote the probability that a high type assigns to the good state. Let p_L^{good} denote the probability that a low type assigns to the good state. Let $\bar{p}_s^{good} = \lambda_s p_H^{good} + (1 - \lambda_s) p_L^{good}$ denote the average probability that households assign to the good state when the economy is in state s . Integrating equation (15) across households yields aggregate consumption in state s

$$c_s = -\frac{1}{\gamma} [(\mu - 1) \bar{\xi}_s + \bar{E}_s [r_S - \mu\pi_S]] + \mu c_s, \quad (16)$$

where $\bar{\xi}_s = \lambda_s \xi_H + (1 - \lambda_s) \xi_L$ is the cross-sectional mean of the preference shock in state s and $\bar{E}_s [r_S - \mu\pi_S] = \int_0^1 E^i [r_S - \mu\pi_S] di$ is the average expectation of the real interest rate in state s . This average expectation equals

$$\bar{E}_s [r_S - \mu\pi_S] = \bar{p}_s^{good} (r_{good} - \mu\pi_{good}) + \bar{p}_s^{bad} (r_{bad} - \mu\pi_{bad}), \quad (17)$$

where $\bar{p}_s^{bad} = 1 - \bar{p}_s^{good}$. Finally, the New Keynesian Phillips curve and the monetary policy rule are again given by equations (10) and (11).

An important difference to the case of perfect information is that aggregate consumption in each state depends on the real interest rate in both aggregate states through equations (16)-(17). For this reason, one cannot solve the model state by state. Instead, one has to distinguish three cases: the zero lower bound is binding in both states, the zero lower bound is binding in no state, and the zero lower bound is binding only in the bad state.

For the moment, I assume that the set of states with a binding zero lower bound under perfect information equals the set of states with a binding zero lower bound under imperfect information. This assumption is relaxed below.

Zero lower bound binding in both states. When the zero lower bound is binding in both states, imperfect information on the household side increases consumption in the bad state. Formally, consumption in the good state equals

$$c_{good} = \frac{\frac{1}{\gamma} \bar{\xi}_{good} + \frac{1}{1-\mu} \ln(R)}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} - \bar{p}_{good}^{bad} \frac{\frac{1}{\gamma} \frac{\mu\kappa}{1-\beta\mu}}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} (c_{good} - c_{bad}), \quad (18)$$

and consumption in the bad state equals

$$c_{bad} = \frac{\frac{1}{\gamma} \bar{\xi}_{bad} + \frac{1}{1-\mu} \ln(R)}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} + \bar{p}_{bad}^{good} \frac{\frac{1}{\gamma} \frac{\mu\kappa}{1-\beta\mu}}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} (c_{good} - c_{bad}), \quad (19)$$

with

$$c_{good} - c_{bad} = \frac{\frac{1}{\gamma}(\bar{\xi}_{good} - \bar{\xi}_{bad})}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} > 0. \quad (20)$$

$$1 + \left(\bar{p}_{bad}^{good} + \bar{p}_{good}^{bad} \right) \frac{\frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}}$$

To obtain these equations, state equation (16) for the good state and the bad state after using equation (17) to substitute for the average expectation of the real interest rate, equation (10) to substitute for inflation, and $r_{good} = r_{bad} = -\ln(R)$ to substitute for the nominal interest rate. One obtains a system of two equations in two unknowns, c_{good} and c_{bad} . Rearranging and using $\bar{p}_{bad}^{good} = 1 - \bar{p}_{bad}^{bad}$ and $\bar{p}_{good}^{bad} = 1 - \bar{p}_{good}^{good}$ yields equations (18)-(20).

To interpret these equations, note that the first term on the right-hand side of equation (18) equals consumption in the good state under perfect information, the first term on the right-hand side of equation (19) equals consumption in the bad state under perfect information, and the numerator on the right-hand side of equation (20) equals the difference between the two. Furthermore, $\bar{p}_{bad}^{good} > 0$ is the average probability assigned to the good state when the economy is in the bad state. Similarly, $\bar{p}_{good}^{bad} > 0$ is the average probability assigned to the bad state when the economy is in the good state.

Dispersed information on the household side increases consumption in the bad state. The reason is that downward movements in household inflation expectations are *destabilizing* at the zero lower bound and the fact that households assign some probability to the wrong state keeps the average inflation expectation high in the bad state.

To understand the *magnitude* of the effect of dispersed information, I follow Angeletos and La'O (2013) by drawing an analogy between the equilibrium of the economy and the perfect Bayesian equilibrium of a fictitious game. Consider the following abstract game. There is a continuum of agents and each agent i chooses an action $c_i \in \mathbb{R}$. There are two types of agents and the aggregate state is the cross-sectional distribution of types. Let c_s denote the average action in the population. Payoff functions are such that the action of agent i equals a linear combination of an agent-specific fundamental, ϖ_i , and the agent's expectation of the average action, $E^i[c_S]$:

$$c_i = (1 - \varsigma) \varpi_i + \varsigma E^i[c_S]. \quad (21)$$

When $\varsigma > 0$, actions are said to be *strategic complements*. It is well known from the literature

on dispersed information that dispersed information has larger effects when the degree of strategic complementarity in actions, $\varsigma \in (0, 1)$, is larger.

Let us return to the economy presented before. Substituting $r_s = -\ln(R)$ and $\pi_s = \frac{\kappa}{1-\beta\mu}c_s$ into equation (15) yields an equation of the form (21) with

$$\varpi_i = \frac{\frac{1}{\gamma}\xi_{i,0} + \frac{1}{1-\mu}\ln(R)}{1 - \frac{1}{1-\mu}\frac{\mu\kappa}{1-\beta\mu}} < 0 \quad \text{and} \quad \varsigma = \frac{1}{1-\mu}\frac{\mu\kappa}{1-\beta\mu} > 0.$$

Furthermore, returning to equations (18)-(20), note that for a given size of the information friction (i.e., for given values of \bar{p}_{bad}^{good} and \bar{p}_{good}^{bad}) and for a given difference between consumption in the good state under perfect information and consumption in the bad state under perfect information (i.e., for a given value of the numerator in equation (20)), the effect of dispersed information is larger when ς is larger and the effect becomes arbitrarily large as ς converges to one.¹² In sum, at the zero lower bound, dispersed information on the household side can have very large effects, because consumption choices of different households are strategic complements.

Zero lower bound binding in no state. When the zero lower bound is binding in no state, imperfect information on the household side decreases consumption in the bad state. Formally, consumption in the good state equals

$$c_{good} = \frac{\frac{1}{\gamma}\bar{\xi}_{good}}{1 + \frac{1}{1-\mu}\frac{(\phi-\mu)\kappa}{1-\beta\mu}} + \bar{p}_{good}^{bad} \frac{\frac{1}{\gamma}\frac{(\phi-\mu)\kappa}{1-\mu}\frac{(\phi-\mu)\kappa}{1-\beta\mu}}{1 + \frac{1}{1-\mu}\frac{(\phi-\mu)\kappa}{1-\beta\mu}} (c_{good} - c_{bad}), \quad (22)$$

and consumption in the bad state equals

$$c_{bad} = \frac{\frac{1}{\gamma}\bar{\xi}_{bad}}{1 + \frac{1}{1-\mu}\frac{(\phi-\mu)\kappa}{1-\beta\mu}} - \bar{p}_{bad}^{good} \frac{\frac{1}{\gamma}\frac{(\phi-\mu)\kappa}{1-\mu}\frac{(\phi-\mu)\kappa}{1-\beta\mu}}{1 + \frac{1}{1-\mu}\frac{(\phi-\mu)\kappa}{1-\beta\mu}} (c_{good} - c_{bad}), \quad (23)$$

with

$$c_{good} - c_{bad} = \frac{\frac{1}{\gamma}(\bar{\xi}_{good} - \bar{\xi}_{bad})}{1 + \frac{1}{1-\mu}\frac{(\phi-\mu)\kappa}{1-\beta\mu}} \frac{1}{1 - \left(\bar{p}_{good}^{bad} + \bar{p}_{bad}^{good} \right) \frac{\frac{1}{\gamma}\frac{(\phi-\mu)\kappa}{1-\mu}\frac{(\phi-\mu)\kappa}{1-\beta\mu}}{1 + \frac{1}{1-\mu}\frac{(\phi-\mu)\kappa}{1-\beta\mu}}} > 0. \quad (24)$$

To obtain these equations, follow the same procedure as before but use $r_s = \phi\pi_s$ instead of $r_s = -\ln(R)$ to substitute for the nominal interest rate. To interpret these equations, note that the first

¹²One can change the value of $\frac{1}{1-\mu}\frac{\mu\kappa}{1-\beta\mu}$ and hold the value of the numerator in equation (20) constant by adjusting the value of $\bar{\xi}_{good} - \bar{\xi}_{bad}$. Finally, recall that the literature assumes that $\frac{1}{1-\mu}\frac{\mu\kappa}{1-\beta\mu}$ is smaller than one (Footnote 7).

term on the right-hand side of equation (22) equals consumption in the good state under perfect information, the first term on the right-hand side of equation (23) equals consumption in the bad state under perfect information, and the numerator on the right-hand side of equation (24) equals the difference between the two.

Dispersed information on the household side now decreases consumption in the bad state. The reason is that downward movements in household inflation expectations are *stabilizing* when the Taylor principle is satisfied (i.e., when the central bank lowers the nominal interest rate more than one-for-one with inflation).

However, consumption choices of different households are strategic substitutes when the Taylor principle is satisfied (substituting $r_s = \phi\pi_s$ and $\pi_s = \frac{\kappa}{1-\beta\mu}c_s$ into equation (15) yields an equation of the form (21) with $\varsigma < 0$), which limits the *magnitude* of the effect of household dispersed information on consumption (the second term on the right-hand side of equation (23)).

Zero lower bound binding only in the bad state. When the zero lower bound is binding only in the bad state, the sign of the effect of household imperfect information depends on the sign of the difference between the real rate in the good state and the real rate in the bad state. Formally, consumption in the good state equals

$$c_{good} = \frac{\frac{1}{\gamma}\bar{\xi}_{good}}{1 + \frac{1}{1-\mu} \frac{(\phi-\mu)\kappa}{1-\beta\mu}} + \bar{p}_{good}^{bad} \frac{\frac{1}{1-\mu}}{1 + \frac{1}{1-\mu} \frac{(\phi-\mu)\kappa}{1-\beta\mu}} [(r_{good} - \mu\pi_{good}) - (r_{bad} - \mu\pi_{bad})], \quad (25)$$

and consumption in the bad state equals

$$c_{bad} = \frac{\frac{1}{\gamma}\bar{\xi}_{bad} + \frac{1}{1-\mu} \ln(R)}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} - \bar{p}_{bad}^{good} \frac{\frac{1}{1-\mu}}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} [(r_{good} - \mu\pi_{good}) - (r_{bad} - \mu\pi_{bad})], \quad (26)$$

with

$$(r_{good} - \mu\pi_{good}) - (r_{bad} - \mu\pi_{bad}) = \frac{\frac{(\phi-\mu)\kappa}{1-\beta\mu} \frac{\frac{1}{\gamma}\bar{\xi}_{good}}{1 + \frac{1}{1-\mu} \frac{(\phi-\mu)\kappa}{1-\beta\mu}} - \left[-\ln(R) - \frac{\mu\kappa}{1-\beta\mu} \frac{\frac{1}{\gamma}\bar{\xi}_{bad} + \frac{1}{1-\mu} \ln(R)}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} \right]}{1 - \bar{p}_{good}^{bad} \frac{\frac{1}{1-\mu} \frac{(\phi-\mu)\kappa}{1-\beta\mu}}{1 + \frac{1}{1-\mu} \frac{(\phi-\mu)\kappa}{1-\beta\mu}} - \bar{p}_{bad}^{good} \frac{-\frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}}}. \quad (27)$$

The derivation is almost identical to the derivation in the previous two cases. The only difference is that one has to use $r_{good} = \phi\pi_{good}$ and $r_{bad} = -\ln(R)$ to substitute for the nominal interest rate. To interpret these equations, note that the first term on the right-hand side of equation (25) equals

consumption in the good state under perfect information, the first term on the right-hand side of equation (26) equals consumption in the bad state under perfect information, and the numerator on the right-hand side of equation (27) equals the difference between the real rate in the good state under perfect information and the real rate in the bad state under perfect information.

When the real rate is lower in the good state than in the bad state under perfect information, the same is true under imperfect information and dispersed information on the household side increases consumption in the bad state.¹³

Due to the kink in the monetary policy rule, the real interest rate under perfect information can be lower in the good state, higher in the good state, or the same in the two states. In the second case, dispersed information on the household side decreases consumption in the bad state. In the third case, dispersed information on the household side has no effect on consumption.¹⁴

Furthermore, note that in the case when the zero lower bound is binding only in the bad state, individual consumption depends on several moments of the conditional distribution of inflation. For comparison, when the zero lower bound is binding in both states, individual consumption equals

$$c_i = \frac{1}{\gamma} \xi_{i,0} + \frac{1}{1-\mu} \ln(R) + \frac{1}{1-\mu} \mu E^i[\pi_S],$$

and aggregate consumption equals

$$c_s = \frac{1}{\gamma} \bar{\xi}_s + \frac{1}{1-\mu} \ln(R) + \frac{1}{1-\mu} \mu \int_0^1 E^i[\pi_S] di.$$

When the zero lower bound is binding in no state, individual consumption equals

$$c_i = \frac{1}{\gamma} \xi_{i,0} - \frac{1}{1-\mu} (\phi - \mu) E^i[\pi_S],$$

and aggregate consumption equals

$$c_s = \frac{1}{\gamma} \bar{\xi}_s - \frac{1}{1-\mu} (\phi - \mu) \int_0^1 E^i[\pi_S] di.$$

In both cases, individual consumption depends only on the conditional mean of inflation and aggregate consumption depends only on the average inflation expectation. By contrast, when the zero

¹³The denominator on the right-hand side of equation (27) is positive.

¹⁴Dispersed information on the household side still implies that household inflation expectations are too high in the bad state, because inflation is higher in the good state than in the bad state and households assign some probability to the wrong state. Dispersed information on the household side just does not affect consumption.

lower bound binds only in the bad state, individual consumption equals

$$c_i = \frac{1}{\gamma} \xi_{i,0} - \frac{\frac{1}{\gamma}}{1-\mu} (\phi - \mu) E^i [\pi_S] - \frac{\frac{1}{\gamma}}{1-\mu} p_i^{bad} (-\ln(R) - \phi \pi_{bad}),$$

and aggregate consumption equals

$$c_s = \frac{1}{\gamma} \bar{\xi}_s - \frac{\frac{1}{\gamma}}{1-\mu} (\phi - \mu) \int_0^1 E^i [\pi_S] di - \frac{\frac{1}{\gamma}}{1-\mu} \bar{p}_s^{bad} (-\ln(R) - \phi \pi_{bad}).$$

Individual consumption now depends on the conditional mean of inflation, the probability assigned to the bad state, and inflation in the bad state. The reason is the kink in the monetary policy rule. The term in the second parentheses is the difference between the actual nominal rate and the nominal rate prescribed by a Taylor rule without a zero lower bound. (In a model with more than two states, inflation in the bad state is replaced by the mean of a truncated distribution: the mean of inflation given a binding zero lower bound. Here this mean simply equals inflation in the bad state.) Aggregate consumption depends on the average inflation expectation, the average probability assigned to the bad state, and inflation in the bad state. Nevertheless, one can solve the model analytically. See equations (25)-(27).

Transitions. In the interpretation (not the derivation) of equations (18)-(20), (22)-(24), and (25)-(27), I assumed up to now that the set of states with a binding zero lower bound under perfect information equals the set of states with a binding zero lower bound under imperfect information. These sets may differ. In particular, when the zero lower bound is binding in no state under perfect information, the zero lower bound may be binding in the bad state under imperfect information. The reason is that imperfect information decreases consumption in the *bad* state.¹⁵ Furthermore, when the zero lower bound is binding only in the bad state under perfect information (and the real rate is lower in the good state than in the bad state), the zero lower bound may be binding in both states under imperfect information. The reason is that imperfect information decreases consumption in the *good* state.¹⁶ However, these transitions do not affect the earlier statements about how imperfect information changes consumption in the two states.

Conditional probabilities. So far we used no properties of the variables \bar{p}_s^{good} and \bar{p}_s^{bad} , apart from the fact that they are non-negative and sum to one. Next, let us compute the average

¹⁵Consumption in the bad state is then given by equation (13) under perfect information and by equation (26) under imperfect information.

¹⁶Consumption in the good state is then given by equation (13) under perfect information and by equation (18) under imperfect information.

probability that households assign to the good state and to the bad state when the economy is in state s .

The probability that a high type assigns to the good state by Bayes' rule equals

$$p_H^{good} = \frac{\lambda_{good}\theta}{\lambda_{good}\theta + \lambda_{bad}(1-\theta)}, \quad (28)$$

while the probability that a low type assigns to the good state by Bayes' rule equals

$$p_L^{good} = \frac{(1-\lambda_{good})\theta}{(1-\lambda_{good})\theta + (1-\lambda_{bad})(1-\theta)}. \quad (29)$$

The average probability assigned to the good state when the economy is in the bad state equals

$$\bar{p}_{bad}^{good} = \lambda_{bad}p_H^{good} + (1-\lambda_{bad})p_L^{good}. \quad (30)$$

The average probability assigned to the bad state when the economy is in the good state equals

$$\bar{p}_{good}^{bad} = \lambda_{good}p_H^{bad} + (1-\lambda_{good})p_L^{bad}. \quad (31)$$

Of course, the variable \bar{p}_{bad}^{good} is an increasing function of the prior probability of the good state, θ , whereas the variable \bar{p}_{good}^{bad} is a decreasing function of the prior probability of the good state. Hence, when the good state has a high prior probability and the zero lower bound is binding in both states, household dispersed information has only a small negative effect on consumption in the good state (equation (18)) but a large positive effect on consumption in the bad state (equation (19)). This result will be important in the following section, where I think of the bad state as a rare event.

6 Numerical solutions

This section relaxes the simplifying assumptions of the previous section and presents quantitative results. I first discuss the parameter values that serve as benchmark parameter values throughout the section.

6.1 Calibration

The main ideas underlying the parameter choices are the following. The preference, technology, price stickiness, and policy parameters are set to their most standard values. The information friction parameter ω is chosen so as to match the speed at which the gap between inflation and

inflation expectations closes in survey data. Finally, the shock parameters are chosen so that under perfect information $\bar{\xi}_{good}$ would create a severe recession and $\bar{\xi}_{bad}$ would create the worst recession in almost a century.

One period corresponds to one quarter. I assume a long-run annual real interest rate of 4% and set $\beta = 0.99$. The intertemporal elasticity of substitution is $(1/\gamma) = 1$ and the elasticity of output with respect to labor is $\varrho = (2/3)$. These are the most common values in the business cycle literature. The elasticity of substitution between intermediate goods is $\psi = 10$, which implies a long-run markup of 11%, a common target in the New Keynesian literature. The probability that a firm cannot adjust its price in a given quarter is $\alpha = 0.66$, implying that one third of prices change per quarter, a value consistent with micro evidence on prices once sales prices have been removed. See Nakamura and Steinsson (2008). For these parameters, the slope of the New Keynesian Phillips curve is $\kappa = 0.045$. I set $\phi = 1.5$, which is the most standard value for the coefficient on inflation in a Taylor rule.

Coibion and Gorodnichenko (2012) estimate impulse responses of inflation and inflation expectations to shocks. Under the null of full-information rational expectations, inflation expectations should adjust to a realized shock by the same amount as the conditional mean of future inflation. Full information implies that the shock is in the information set of the agents. Rational expectations implies that agents understand how the shock affects future inflation. By contrast, Coibion and Gorodnichenko (2012) find that the responses of inflation expectations to shocks are dampened and delayed relative to the responses of future inflation to the same shocks. After an inflationary (disinflationary) shock, inflation expectations rise (fall) by less than future inflation and this difference becomes smaller over time and eventually converges to zero. This result is obtained for all four types of shocks they consider (technology shocks, news shocks, oil shocks, unidentified shocks), inflationary and disinflationary shocks, and different types of agents.¹⁷ As a next step, Coibion and Gorodnichenko (2012) estimate the degree of information rigidity that matches the empirical speed of response of inflation expectations to shocks. In the context of a sticky information model, their estimated degree of information rigidity corresponds to the fraction of agents that do not update

¹⁷For professional forecasters, the inflation forecasts are from the Survey of Professional Forecasters and the time sample is 1976-2007. For households, the inflation forecasts are from the Michigan Survey of Consumers and the time sample is 1976-2007. For each group, Coibion and Gorodnichenko (2012) study the response of the average forecast.

their inflation expectations in a given quarter. On page 143 they write: “This procedure yields estimates between 0.86 and 0.89 for technology, news, and oil price shocks as well as for unidentified shocks.” Based on these estimates, I set $\omega = 1 - 0.875$.¹⁸

Finally, let us turn to the shock parameters. I set the persistence of the preference shocks to $\mu = 0.8$, which is a common value in the New Keynesian literature on the zero lower bound. I set $\xi_H = -0.05$ and $\xi_L = -0.075$, which implies that the shock term $(1 - \mu)\xi_{i,0}$ in the consumption Euler equation equals -1% for a high type and -1.5% for a low type. I set the fraction of high types in the good state to $\lambda_{good} = (3/4)$ and the fraction of high types in the bad state to $\lambda_{bad} = (1/4)$. The shock term $(1 - \mu)\bar{\xi}_s$ in the aggregated Euler equation thus equals -1.125% in the good state and -1.375% in the bad state. Under perfect information, the zero lower bound is marginally binding in the good state and clearly binding in the bad state (for comparison, $(1 - \mu)\bar{\xi}_{crit} = -1.09\%$). Furthermore, under perfect information, consumption drops by 4% in the good state and by 13% in the bad state. Hence, I think of the good state as a shock that would create a serious recession under perfect information, while I think of the bad state as a shock that would create the worst recession since World War II under perfect information. Finally, I set $\theta = 0.9$. That is, the prior probability of the bad state equals 10%. This seems a reasonable value given that recessions with a 13 percent fall in consumption are rare.

6.2 Slow updating of inflation expectations

This subsection presents the solution of the model when households update their inflation expectations slowly over time (i.e., $0 < \omega < 1$). The next three paragraphs describe how the model is solved. The following paragraph presents the solution for the benchmark parameter values.

There are four types of households at each point in time t : households with a high realization of the preference shock in period zero ($\xi_{i,0} = \xi_H$) and perfect-information rational expectations in period t (“high, informed types”), households with a low realization of the preference shock ($\xi_{i,0} = \xi_L$) and perfect-information rational expectations in period t (“low, informed types”),

¹⁸These estimates are for professional forecasters. Coibion and Gorodnichenko (2012) obtain slightly lower estimates of the degree of information rigidity for households (see their Table 4). In the benchmark calibration, I use the estimates that Coibion and Gorodnichenko (2012) obtain for inflation expectations data from the Survey of Professional Forecasters. In Section 5.5, I show that results are essentially unchanged when one uses the estimates that Coibion and Gorodnichenko (2012) obtain for inflation expectations data from the Michigan Survey of Consumers.

households with a high realization of the preference shock in period zero and imperfect information in period t (“high, uninformed types”), and households with a low realization of the preference shock in period zero and imperfect information in period t (“low, uninformed types”).

The consumption Euler equation of a high, informed type in period t in state $s \in \{good, bad\}$ reads

$$c_{s,t}^{hi} = -\frac{1}{\gamma} ((\mu - 1) \xi_H + r_{s,t} - \mu\pi_{s,t+1}) + \mu c_{s,t+1}^{hi}. \quad (32)$$

The consumption Euler equation of a low, informed type in period t in state $s \in \{good, bad\}$ reads

$$c_{s,t}^{li} = -\frac{1}{\gamma} ((\mu - 1) \xi_L + r_{s,t} - \mu\pi_{s,t+1}) + \mu c_{s,t+1}^{li}. \quad (33)$$

The consumption Euler equation of a high, uninformed type in period t reads

$$\begin{aligned} c_t^{hu} &= -\frac{1}{\gamma} \left((\mu - 1) \xi_H + p_H^{good} (r_{good,t} - \mu\pi_{good,t+1}) + p_H^{bad} (r_{bad,t} - \mu\pi_{bad,t+1}) \right) \\ &+ \omega \left(p_H^{good} \mu c_{good,t+1}^{hi} + p_H^{bad} \mu c_{bad,t+1}^{hi} \right) + (1 - \omega) \mu c_{t+1}^{hu}. \end{aligned} \quad (34)$$

The parameter ω reflects the fact that an uninformed household becomes informed in the next period with probability ω . Furthermore, p_H^{good} and p_H^{bad} are the probabilities that a high type assigns to the good state and the bad state. The consumption Euler equation of a low, uninformed type in period t reads

$$\begin{aligned} c_t^{lu} &= -\frac{1}{\gamma} \left((\mu - 1) \xi_L + p_L^{good} (r_{good,t} - \mu\pi_{good,t+1}) + p_L^{bad} (r_{bad,t} - \mu\pi_{bad,t+1}) \right) \\ &+ \omega \left(p_L^{good} \mu c_{good,t+1}^{li} + p_L^{bad} \mu c_{bad,t+1}^{li} \right) + (1 - \omega) \mu c_{t+1}^{lu}. \end{aligned} \quad (35)$$

Aggregate consumption is a weighted average of the consumption of the four types of households. The weight on consumption of a particular type equals the mass of this type. Hence, aggregate consumption in the good state equals

$$c_{good,t} = \left[1 - (1 - \omega)^{t+1} \right] \left[\lambda_{good} c_{good,t}^{hi} + (1 - \lambda_{good}) c_{good,t}^{li} \right] + (1 - \omega)^{t+1} \left[\lambda_{good} c_t^{hu} + (1 - \lambda_{good}) c_t^{lu} \right], \quad (36)$$

and aggregate consumption in the bad state equals

$$c_{bad,t} = \left[1 - (1 - \omega)^{t+1} \right] \left[\lambda_{bad} c_{bad,t}^{hi} + (1 - \lambda_{bad}) c_{bad,t}^{li} \right] + (1 - \omega)^{t+1} \left[\lambda_{bad} c_t^{hu} + (1 - \lambda_{bad}) c_t^{lu} \right]. \quad (37)$$

The mass of uninformed households equals $(1 - \omega)$ in period zero, $(1 - \omega)^2$ in period one, and so on. The mass of informed households equals one minus the mass of uninformed households. Finally, the New Keynesian Phillips curve in the good state reads

$$\pi_{good,t} = \kappa c_{good,t} + \beta \mu \pi_{good,t+1}, \quad (38)$$

and the New Keynesian Phillips curve in the bad state reads

$$\pi_{bad,t} = \kappa c_{bad,t} + \beta \mu \pi_{bad,t+1}. \quad (39)$$

For the benchmark parameter values, the zero lower bound is binding in each period in both states. The system of equations can then be written as a linear difference equation

$$A_t x_t = b + B x_{t+1}, \quad (40)$$

with

$$x_t = \begin{pmatrix} c_{good,t}^{hi} \\ c_{bad,t}^{hi} \\ c_{good,t}^{li} \\ c_{bad,t}^{li} \\ c_t^{hu} \\ c_t^{lu} \\ c_{good,t} \\ c_{bad,t} \\ \pi_{good,t} \\ \pi_{bad,t} \end{pmatrix}. \quad (41)$$

The matrix A_t is a function of t , because aggregate consumption depends on the mass of households with perfect-information rational expectations and this mass increases over time. Let $A = \lim_{t \rightarrow \infty} A_t$ denote the limit of the matrix A_t as t goes to infinity. The matrix A also equals the value of the matrix A_t at each point in time in the special case of the model where all households have perfect-information rational expectations already in period zero (i.e., $\omega = 1$). I first solve for the steady state of the linear difference equation (40) after replacing A_t by A . The solution $x = (A - B)^{-1} b$ equals the equilibrium under perfect information presented in Section 4.1. Thereafter, I compute the earlier x_t from equation (40) and the endpoint $x_{10,000} = x$. That is, I assume the following:

if preference shocks have not yet reverted back, 10,000 periods after the shock the economy has converged to the solution under perfect information.

Figure 1 shows the solution for the benchmark parameter values. Each line depicts consumption in periods $0 \leq t \leq T - 1$. Recall that the preference shocks revert back to zero with probability $1 - \mu = 0.2$ in any given period and the economy is in the non-stochastic steady state with zero inflation thereafter. The thin black line shows consumption in the good state. The thick black line shows consumption in the bad state. For comparison, the dotted upper line shows consumption in the good state under perfect information. The dotted lower line shows consumption in the bad state under perfect information. Each dot corresponds to one quarter. The slow adjustment of household inflation expectations increases consumption in the bad state. The reason is that downward movements in household inflation expectations are destabilizing. The effect of dispersed information is large and persistent. The reason is that Coibion and Gorodnichenko (2012) estimate a speed of updating of inflation expectations that is far away from $\omega = 1$ and consumption choices of different households are strategic complements when the zero lower bound is binding. Finally, dispersed information has only a small negative effect on consumption in the good state and a large positive effect on consumption in the bad state. The reason is the small prior probability of the bad state.

6.3 Deterministic decay

Let us introduce deterministic decay as a next step towards relaxing the simplifying assumptions of Section 4. Formally, $\xi_{i,t} = \rho \xi_{i,t-1}$ in periods $1 \leq t \leq T - 1$. Fundamentals now change deterministically in addition to the stochastic decay. In the consumption Euler equations (32)-(35), the term $(\mu - 1) \xi_{i,0}$ with $\xi_{i,0} \in \{\xi_L, \xi_H\}$ then becomes $(\mu\rho - 1) \rho^t \xi_{i,0}$ and the nominal interest rate is given by $r_{s,t} = \max\{-\ln(R), \phi\pi_{s,t}\}$. To solve the model, I make a guess regarding the number of periods for which the zero lower bound is binding in the good state and in the bad state. If the guess turns out to be incorrect, I update the guess until a fixed point is reached.

Figure 2 shows aggregate consumption (upper panel) and the nominal interest rate (lower panel) in periods $0 \leq t \leq T - 1$ for the benchmark parameter values, $\rho = 0.99$, $\mu = 0.95$, $\xi_H = -0.05 * 1.4$, and $\xi_L = -0.075 * 1.4$. I assume some stochastic decay to match the observation that in December

of 2008 the Federal Reserve expected the zero lower bound episode to be shorter than seven years.¹⁹ Furthermore, I scale the preference shocks in order to keep roughly constant the *perfect-information* response of consumption on impact in the two states across subsections. The good state is again an aggregate shock that would create a serious recession under perfect information. The bad state is again an aggregate shock that would create the worst recession since World War II under perfect information.

The effect of household dispersed information is large and persistent. The thin black lines show consumption and the nominal interest rate in the good state. The thick black lines show consumption and the nominal interest rate in the bad state. For comparison, the dotted upper line in the upper panel shows consumption in the good state under perfect information. The dotted lower line in the upper panel shows consumption in the bad state under perfect information. The first difference to Figure 1 is that consumption in the bad state is almost flat over time for many quarters and then converges slowly back to the non-stochastic steady state. The reason why consumption is almost flat over time in the first couple of years is that there are two forces working in opposite directions. On the one hand, the fundamentals are improving, which drives aggregate consumption up. On the other hand, more and more households have updated their expectations since the period of the shock, which drives aggregate consumption down. In the first couple of years, these two effects roughly cancel. The second difference to Figure 1 is that the economy can now transit through all three regions discussed in Section 4.2. In the first two years, the zero lower bound is binding in both states. In the following 5 years, the zero lower bound is binding only in the bad state. Thereafter, the zero lower bound is binding in no state. Finally, turning to magnitudes, in the bad state the response of consumption on impact of the shock equals about 1/3 of the value under perfect information.

6.4 Households set nominal wage rates

Finally, let us relax the assumption that households set real wage rates. When households set nominal wage rates and household inflation expectations adjust slowly over time, the New Keynesian Phillips curve changes. The New Keynesian Phillips curve is given by equation (7). As a result, there is a time- and state-varying intercept in equations (38)-(39) equal to $\frac{(1-\alpha)(1-\alpha\beta)}{\alpha} \frac{1}{1+\frac{1-\epsilon}{\epsilon}\psi} (\bar{E}_t [p_t] - p_t)$.

¹⁹ Another explanation is that even the Fed assigned some probability to the wrong state in December of 2008.

To solve the model, I make a guess regarding the value of this term in every period in each state. I solve the difference equation (40), which now has a time-varying vector b_t , using the same endpoint as before. Afterwards, I compute the actual path of $(\bar{E}_t [p_t] - p_t)$ in each state. If the initial guess turns out to be incorrect, I update the guess until a fixed point is reached.

Figure 3 shows aggregate consumption (upper panel) and the nominal interest rate (lower panel) for the parameter values of the previous subsection when households set nominal wage rates. The line styles in Figure 3 equal the line styles in Figure 2. The main difference to Figure 2 is that consumption falls even less in the bad state. The reason is simple. When households' expectations of the price level are above the price level, households set nominal wage rates that are too high, which raises marginal costs and inflation. This attenuates the deflationary spiral and increases aggregate consumption. Turning to magnitudes, in the bad state the response of consumption on impact of the shock equals about 1/4 of the value under perfect information.

The path in the bad state matches basic features of the U.S. economy during the Great Recession. Real nondurable and services consumption fell by several percentage points in 2008 and was almost flat over time for several years (relative to a log-linear trend). The Federal Reserve has maintained a 0 to 1/4 percent target range for the federal funds rate for almost seven years. The cross-sectional dispersion of short-term inflation expectations increased in 2008 and then declined almost monotonically (Andrade et al. (2015) document this feature for data from the Survey of Professional Forecasters (SPF), and this feature is also present in data from the Michigan Survey of Consumers (MSC)). Household inflation expectations have been above realized inflation during the Great Recession and in the following years (Armantier et al. (2011) document this feature for data from the survey of consumers that is conducted by the Federal Reserve Bank of New York, and Coibion and Gorodnichenko (2015b) document this feature for data from the MSC).²⁰

6.5 Faster updating of inflation expectations

Figure 4 shows the solution for $\omega = 1 - 0.77 = 0.23$ instead of $\omega = 1 - 0.875 = 0.125$. The value 0.77 is the average estimate of the degree of information rigidity that Coibion and Gorodnichenko (2012) obtain when they use inflation forecasts from the Michigan Survey of Consumers (see their Table

²⁰Armantier et al. (2011) find in addition that survey respondents act on their reported beliefs in a financially incentivized investment experiment (designed such that future inflation affects payoffs).

4). Otherwise the model setup and parameter values are the same as in the previous subsection. In the bad state, the response of consumption on impact of the shock is somewhat larger than in Figure 3 (about $1/3$ instead of $1/4$ of the value under perfect information).

6.6 Lower prior probability of the good state

Figure 5 shows the solution for a higher prior probability of the bad state. Here $1 - \theta = 0.5$ instead of $1 - \theta = 0.1$. (Otherwise model setup and parameters are the same as in Section 5.4). The response of consumption on impact of the shock becomes more negative in both aggregate states.

Recall that households learn the aggregate state over time. Suppose that households use this information to update beliefs about the probability of the bad state in the future.²¹ According to the model, consumption will then drop more strongly in both aggregate states in the next recession. It therefore may be useful to communicate to households that the Great Recession has been a very unusual recession.

7 Monetary policy

7.1 Central bank communication about the current state

Central bankers frequently make statements about the current state of the economy. In fact, the Federal Reserve's official statement after a regular Federal Open Market Committee (FOMC) meeting *starts* with a paragraph on the current state of the economy.

In most models, central bank communication about the current state is irrelevant, because the current state is common knowledge and central bank communication also does not affect equilibrium selection. In the model studied here, central bank communication about the current state of the economy affects aggregate consumption, because the current state is not common knowledge.

To show this result as clearly as possible, let us return to the closed-form solutions of Section 4. Suppose that in period zero the central bank makes a correct statement about the aggregate state of the economy, and this statement reaches a fraction $\zeta \in [0, 1]$ of randomly selected households.²²

²¹For a model of Bayesian learning of the probability of a rare event, see Maćkowiak and Wiederholt (2014), Section 3.5.

²²The individual history of a household in period T now includes whether the central bank statement has reached the household. Since the state-contingent claims specify a payment that is contingent on the individual history, in

Then, the probabilities \bar{p}_{bad}^{good} and \bar{p}_{good}^{bad} given by equations (30)-(31) have to be multiplied by a factor of $1 - \zeta$, because a fraction ζ of households is reached by the central bank communication and these households hold correct beliefs about the aggregate state of the economy.

The effect on consumption can be seen directly from equations (18)-(20), (22)-(24), and (25)-(27). When the zero lower bound is binding in both states (or more generally the real interest rate is lower in the good state than in the bad state), the communication *decreases* consumption in the bad state. In the bad state, households' inflation expectations are above the true conditional mean of future inflation. Therefore, the central bank communication reduces the inflation expectations of the households that are reached. Furthermore, at the zero lower bound, reductions in households' inflation expectations are destabilizing. Hence, aggregate consumption falls. By contrast, when the zero lower bound is binding in no state (or more generally the real interest rate is higher in the good state than in the bad state), the same communication *increases* consumption in the bad state. In the bad state, households' inflation expectations are once again above the true conditional mean of future inflation and the central bank communication reduces the inflation expectations of the households that are reached. The Taylor principle implies that downward movements in households' inflation expectations are stabilizing. Hence, aggregate consumption increases. In sum, central bank communication about the aggregate state of the economy affects aggregate consumption and the sign of this effect reverses when the zero lower bound is binding.

These results imply that central banks probably face a new trade-off when inflation expectations adjust slowly downwards and the zero lower bound is binding. On the one hand, communicating low inflation risk to the public at the zero lower bound is destabilizing. On the other hand, the central bank probably would like to communicate truthfully the current state to the public in order to maintain a good reputation. The following excerpt from the September 4, 2014, press conference with the President of the European Central Bank (ECB) after that day's meeting of the ECB Governing Council suggests that this trade-off is present in practice:²³

“Question: Mr Draghi, isn't there a risk that with the ECB emphasising so much the risk of low inflation that this itself could trigger a de-anchoring of expectations?”

Draghi: Well, you see, this question is actually a question we also asked ourselves. But

equilibrium all households have the same post-transfer wealth in period T .

²³See www.ecb.europa.eu/press/pressconf/2014/html/is140904.en.html

the answer to this question is: would the truth be a risk? In other words, do we really think that telling people things other than the truth would affect their behaviour? And the answer is no. We think that we ought to state things as they are. We don't see deflation. We have seen, as a matter of fact, low inflation for a long time. As I've said several times, the longer the period of low inflation the higher the risks of de-anchoring."

7.2 Central bank communication about future policy

Let us turn to forward guidance. Following Bassetto (2015), I define forward guidance as a direct statement by the central bank about the future path of its policy tools. These statements can take different forms. Campbell et al. (2012) distinguish between Odyssean forward guidance, which publicly commits the central bank to a future action, and Delphic forward guidance, which merely forecasts macroeconomic performance and likely monetary policy actions. Eggertsson and Woodford (2003) demonstrate that in a standard New Keynesian model a commitment by the central bank to create inflation in the future is a powerful way of stimulating the economy when the zero lower bound is binding. I therefore first study the effects of Odyssean forward guidance and then turn to Delphic forward guidance.

Consider again the closed-form solutions of Section 4. Suppose that in period zero the central bank publicly commits itself to a future path of the nominal interest rate and this statement reaches a fraction $\zeta \in [0, 1]$ of randomly selected households. In the good state, the central bank announces that it will set the nominal interest rate in periods $t \geq T$ so as to maintain a long-run inflation target of zero. In the bad state, the central bank instead announces that it will set the nominal interest rate in periods $t \geq T$ so as to achieve a long-run inflation target of $\bar{\pi} > 0$ in order to raise inflation expectations in periods $t < T$.

Suppose that the zero lower bound is binding in both states without and with forward guidance. Without forward guidance, consumption in the good state and in the bad state is given by equations (18)-(20) and (30)-(31). With forward guidance, consumption in the good state equals

$$c_{good} = \frac{\frac{1}{\gamma} \bar{s}_{good} + \frac{1}{1-\mu} \ln(R) + \bar{p}_{good}^{bad} \left(\frac{1}{1-\beta\mu} \frac{1}{\gamma} \bar{\pi} + \bar{c} \right)}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} - \bar{p}_{good}^{bad} \frac{\frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} (c_{good} - c_{bad}), \quad (42)$$

and consumption in the bad state equals

$$c_{bad} = \frac{\frac{1}{\gamma} \bar{\xi}_{bad} + \frac{1}{1-\mu} \ln(R) + \left(1 - \bar{p}_{bad}^{good}\right) \left(\frac{1}{1-\beta\mu} \frac{1}{\gamma} \bar{\pi} + \bar{c}\right)}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} + \bar{p}_{bad}^{good} \frac{\frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}}{1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}} (c_{good} - c_{bad}), \quad (43)$$

with

$$c_{good} - c_{bad} = \frac{\frac{1}{\gamma} (\bar{\xi}_{good} - \bar{\xi}_{bad}) - \left(1 - \bar{p}_{good}^{bad} - \bar{p}_{bad}^{good}\right) \left(\frac{1}{1-\beta\mu} \frac{1}{\gamma} \bar{\pi} + \bar{c}\right)}{1 - \left(1 - \bar{p}_{good}^{bad} - \bar{p}_{bad}^{good}\right) \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}}. \quad (44)$$

The probabilities \bar{p}_{bad}^{good} and \bar{p}_{good}^{bad} in equations (42)-(44) are given by the expressions on the right-hand side of equations (30)-(31) multiplied by a factor of $1 - \zeta$, because the central bank's communication reaches a fraction ζ of randomly selected households and the central bank's statement reveals the aggregate state to those households. Finally, the variable $\bar{c} \geq 0$ in equations (42)-(44) denotes consumption in the non-stochastic steady state with inflation rate $\bar{\pi} > 0$. The derivation of equations (42)-(44) is in Appendix A.

Forward guidance is less powerful than under perfect information. The commitment to create inflation in the future can even reduce current consumption in the bad state. The reason is twofold. Households that are not reached by the communication do not update their inflation expectations. Households that are reached by the communication experience a positive and a negative effect on their inflation expectations. Formally, under perfect information, consumption without forward guidance is given by equation (12) and consumption in the bad state with forward guidance is given by equation (43) with $\bar{p}_{bad}^{good} = 0$. Hence, forward guidance increases consumption in the bad state by $\left(\frac{1}{1-\beta\mu} \frac{1}{\gamma} \bar{\pi} + \bar{c}\right) / \left(1 - \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}\right)$. With dispersed information on the household side, this term is multiplied by $1 - \bar{p}_{bad}^{good}$, because not all households are aware of the fact that the commitment to create inflation in the future is currently in place. Furthermore, forward guidance reduces the probability \bar{p}_{bad}^{good} by a factor of $1 - \zeta$, because the households that are reached by the central bank communication assign probability one to the bad state. Households that are reached by the communication experience a positive and a negative effect on their inflation expectations. They learn that the central bank has committed itself to create inflation in the future and they learn that the economy is in the bad state, which is the state with the lower inflation rate.

Due to the second effect, forward guidance can decrease aggregate consumption in the bad state. To see this as clearly as possible, suppose that all households are reached by the communication

($\zeta = 1$). Consumption in the bad state without forward guidance is given by equations (19)-(20) and (30)-(31). Consumption in the bad state with forward guidance is given by equation (43) with $\bar{p}_{bad}^{good} = 0$. Hence, forward guidance increases aggregate consumption in the bad state if and only if

$$\frac{1}{1 - \beta\mu} \frac{1}{\gamma} \bar{\pi} + \bar{c} > \bar{p}_{bad}^{good} \frac{\frac{1}{\gamma} (\bar{\xi}_{good} - \bar{\xi}_{bad}) \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}}{1 - \left(1 - \bar{p}_{good}^{bad} - \bar{p}_{bad}^{good}\right) \frac{1}{1-\mu} \frac{\mu\kappa}{1-\beta\mu}},$$

where \bar{p}_{bad}^{good} and \bar{p}_{good}^{bad} are given by equations (30)-(31).

Next, consider Delphic forward guidance, which merely forecasts macroeconomic performance and likely monetary policy actions. For this purpose, let us return to Section 5.4 and Figure 3. Suppose that in period zero the central bank makes a correct statement about the number of periods that the zero lower bound is expected to bind and this statement reaches all households ($\zeta = 1$). In the bad state this number of periods is much larger than in the good state. Therefore, the statement reveals the aggregate state to households, and aggregate consumption in the bad state is now given by the dotted lower line instead of the thick black line in Figure 3. Hence, Delphic forward guidance unambiguously decreases aggregate consumption in the bad state. In practice, the effect of Delphic forward guidance is probably much weaker, because the official statement by the central bank does not reach all households ($\zeta \in [0, 1)$).

The results in this subsection are related to two literatures. First, the literature on signaling effects of monetary policy (e.g., Melosi, 2015, and Tang, 2015) studies environments where the current policy rate provides information about current fundamentals to the private sector. Note that this effect is muted when the zero lower bound is binding. I instead study a model where forward guidance provides information about the current state to households. Second, the emerging literature on the forward guidance puzzle (Carlstrom, Fuerst, and Paustian, 2012, Del Negro, Giannoni, and Patterson, 2012, McKay, Nakamura, and Steinsson, 2015) argues that the effects of forward guidance in benchmark New Keynesian models seem unreasonably large and offers explanations for weaker effects of forward guidance. I offer complementary explanations for weaker effects of forward guidance (not all households are reached and those that are reached may revise their inflation expectations downwards).

The negative effects originating from central bank communication about the current state in the bad state can probably be reduced by designing communication in the right way. In particular, according to the model, it would be desirable to separate communication about future inflation from

communication about the current state. This seems feasible. For example, the central bank could say: “We raise the long-run inflation target, because it is a good idea in general.” Alternatively, the central bank could say: “We keep the interest rate at zero until inflation equals a certain value (below the long-run target),” without specifying how long the interest rate would have to be kept at zero in expectation. One would still need to reach households to make this a powerful policy, but at least the negative effects of the policy could be reduced.

8 Fiscal policy

Shortly after the Federal Reserve lowered the target for the federal funds rate to 0-0.25 percent in December 2008, the U.S. Congress passed a major fiscal stimulus package - the American Recovery and Reinvestment Act of 2009. An interesting question is how expansionary fiscal policy in a bad state affects consumption in a world where household inflation expectations adjust slowly.

To address this question formally, consider the closed-form solutions of Section 4. Suppose that the level of government spending equals G in the long run (i.e., in periods $t \geq T$). Let $g_t = \ln(G_t/G)$ denote the log-deviation of current government spending from long-run government spending. In period zero, the fiscal authority makes an announcement about the level of government spending during the recession (i.e., in periods $0 \leq t < T$) and this announcement reaches a fraction $\zeta \in [0, 1]$ of randomly selected households. In the good state, the government announces $g_t = g_{good} = 0$. In the bad state, the government announces $g_t = g_{bad} = g > 0$.

Government spending does not change the consumption Euler equation (3), but it does change the New Keynesian Phillips curve (8). Substituting the log-linearized wage index, the wage setting equation (5), and the log-linearized resource constraint $y_t = (C/Y)c_t + (G/Y)g_t$ into equation (6) yields the modified New Keynesian Phillips curve

$$\pi_t = \kappa_c c_t + \kappa_g g_t + \beta E_t [\pi_{t+1}],$$

with

$$\kappa_c = \frac{(1-\alpha)(1-\alpha\beta)\gamma + \frac{1-\varrho}{\varrho} \frac{C}{Y}}{\alpha + \frac{1-\varrho}{\varrho} \psi} \quad \text{and} \quad \kappa_g = \frac{(1-\alpha)(1-\alpha\beta)\frac{1-\varrho}{\varrho} \frac{G}{Y}}{\alpha + \frac{1-\varrho}{\varrho} \psi}.$$

It is then straightforward to derive aggregate consumption in the good state and aggregate consumption in the bad state. When the zero lower bound is binding in both states, consumption in

the good state equals

$$c_{good} = \frac{\frac{1}{\gamma} \bar{\xi}_{good} + \frac{1}{1-\mu} \ln(R) + \bar{p}_{good}^{bad} \frac{1}{1-\mu} \frac{\mu \kappa_g}{1-\beta \mu} g}{1 - \frac{1}{1-\mu} \frac{\mu \kappa_c}{1-\beta \mu}} - \bar{p}_{good}^{bad} \frac{\frac{1}{1-\mu} \frac{\mu \kappa_c}{1-\beta \mu}}{1 - \frac{1}{1-\mu} \frac{\mu \kappa_c}{1-\beta \mu}} (c_{good} - c_{bad}), \quad (45)$$

and consumption in the bad state equals

$$c_{bad} = \frac{\frac{1}{\gamma} \bar{\xi}_{bad} + \frac{1}{1-\mu} \ln(R) + \left(1 - \bar{p}_{bad}^{good}\right) \frac{1}{1-\mu} \frac{\mu \kappa_g}{1-\beta \mu} g}{1 - \frac{1}{1-\mu} \frac{\mu \kappa_c}{1-\beta \mu}} + \bar{p}_{bad}^{good} \frac{\frac{1}{1-\mu} \frac{\mu \kappa_c}{1-\beta \mu}}{1 - \frac{1}{1-\mu} \frac{\mu \kappa_c}{1-\beta \mu}} (c_{good} - c_{bad}), \quad (46)$$

with

$$c_{good} - c_{bad} = \frac{\frac{1}{\gamma} (\bar{\xi}_{good} - \bar{\xi}_{bad}) - \left(1 - \bar{p}_{good}^{bad} - \bar{p}_{bad}^{good}\right) \frac{1}{1-\mu} \frac{\mu \kappa_g}{1-\beta \mu} g}{1 - \left(1 - \bar{p}_{good}^{bad} - \bar{p}_{bad}^{good}\right) \frac{1}{1-\mu} \frac{\mu \kappa_c}{1-\beta \mu}}. \quad (47)$$

The probabilities \bar{p}_{bad}^{good} and \bar{p}_{good}^{bad} are again given by the expressions on the right-hand side of equations (30)-(31) multiplied by the factor $1 - \zeta$, because the government's state-contingent statement reveals the aggregate state to the households that are reached by the communication. The derivation of equations (45)-(47) is in Appendix B.

There are two differences to the case of perfect information. First, the expression $\frac{1}{1-\mu} \frac{\mu \kappa_g}{1-\beta \mu} g$ in the first term on the right-hand side of equation (46) is multiplied by a factor of $1 - \bar{p}_{bad}^{good}$, because the government's announcement does not reach all households (or equivalently, not all households update their inflation expectations based on the announcement). Second, the probability \bar{p}_{bad}^{good} in the second term on the right-hand side of equation (46) is multiplied by a factor of $1 - \zeta$, because households that are reached by the communication assign probability one to the bad state.

Due to the second effect, the state-contingent fiscal policy can reduce aggregate consumption in the bad state. To see this as clearly as possible, suppose that the government's announcement reaches all households ($\zeta = 1$). Then, consumption in the bad state with state-contingent fiscal policy is given by equation (46) with $\bar{p}_{bad}^{good} = 0$. Consumption in the bad state without state-contingent fiscal policy is given by equations (19)-(20) and (30)-(31). The policy increases aggregate consumption in the bad state if and only if

$$\kappa_g g > \bar{p}_{bad}^{good} \kappa (c_{good} - c_{bad}),$$

where $c_{good} - c_{bad}$ is given by equation (20) and \bar{p}_{bad}^{good} is given by equation (30). Substituting in the

expressions for κ_g and κ yields

$$\frac{G}{Y} \frac{\frac{1-\varrho}{\varrho}}{\gamma + \frac{1-\varrho}{\varrho}} g > \bar{p}_{bad}^{good} (c_{good} - c_{bad}).$$

The results in this section are related to other work arguing that the government spending multiplier at the zero lower bound may not be as large as predicted by the benchmark New Keynesian model with a zero lower bound, e.g., because of distortionary taxes (Uhlig and Drautzburg, 2013) or a non-fundamental liquidity trap (Mertens and Ravn, 2014). I offer complementary arguments (not all households revise their inflation expectations directly after the government’s announcement of the fiscal package and those that do may revise their inflation expectations downwards).

9 The missing deflation puzzle

The missing deflation puzzle was explained in Section 2.1. I now show that Bayesian learning about persistence resolves the missing deflation puzzle.

To this end, I return to the model of Section 6.3. Parameters are the same as before. The discount factor is $\beta = 0.99$. The slope of the New Keynesian Phillips curve is $\kappa = 0.045$. The two persistence parameters are $\rho = 0.99$ and $\mu = 0.95$. Households who experience a large shock experience $\xi_L = -0.075 * 1.4$ and households who experience a small shock experience $\xi_H = -0.05 * 1.4$. In the good state, one quarter of households experience the large shock. In the bad state, three quarters of households experience the large shock. The prior probability of the good state is $\theta = 0.9$. The information diffusion parameter is $\omega = 1 - 0.875$. The intertemporal elasticity of substitution is $1/\gamma = 1$. The Taylor rule coefficient is $\phi = 1.5$. The solution was depicted in Figure 2.

I now introduce Bayesian learning about persistence. Recall that the slump persists with probability μ . In every period $t \geq 1$, the economy switches to the steady state with probability $1 - \mu$ and does not switch to the steady state with probability μ . Up to now the value of μ has been common knowledge. This has the implication that agents in the model expect a highly persistent slump from the beginning. However, agents did not expect the slump to be highly persistent at the beginning of the Great Recession.²⁴ I therefore introduce Bayesian learning about μ . Agents’ prior

²⁴See, for example, Figure 4 in Swanson and Williams (2014).

for $1 - \mu$ is a beta distribution with parameters $\hat{\alpha} > 0$ and $\hat{\beta} > 0$. In every period $t > 0$, agents observe whether the economy has switched back to the steady state and revise their beliefs about μ using Bayes' rule. Agents' posterior for $1 - \mu$ in period $t > 0$ is then a beta distribution with parameters $\hat{\alpha}$ and $\hat{\beta} + t$, if the economy has not switched back yet in period $t > 0$. The posterior mean of $1 - \mu$ in period t is given by

$$E[1 - \mu] = \frac{\hat{\alpha}}{\hat{\alpha} + \hat{\beta} + t}.$$

Hence, as the slump persists, agents revise downwards their beliefs about the probability of a quick recovery and they revise upwards their beliefs about persistence. Figure 7 shows the solution when $\frac{\hat{\alpha}}{\hat{\alpha} + \hat{\beta}} = 0.25$ and $\hat{\alpha}$ is large. That is, consistent with survey data, agents initially expect the slump to persist for about a year and learning about persistence is slow. Agents are assumed to be smart. They take into account the uncertainty about μ in their decisions and they anticipate how they will revise their beliefs about μ in the future if the slump persists.

I assume that Bayesian learning about persistence happens only on the side of firms. The only purpose of this assumption is to demonstrate that firms' perceived persistence is crucial for inflation as suggested by Section 2.1. I therefore assume that households know the true value of μ . Finally, the low initial perceived persistence on the firm side will reduce the firms' inflation response, which will reduce the households' consumption response. To make Figures 2 and 7 comparable in terms of consumption, I slightly increase the value of β from 0.99 to 0.9933. A larger value of β implies a smaller steady state interest rate. Hence, the zero lower bound on the nominal interest rate binds earlier and the consumption response becomes larger. The consumption response in Figure 7 with $\beta = 0.9933$ and Bayesian learning about persistence on the firm side is comparable to the consumption response in Figure 2 with $\beta = 0.99$ and common knowledge about μ . The different lines refer to the same scenarios as before. The thin black line shows consumption in the good state. The thick black line shows consumption in the bad state. For comparison, the dotted upper line shows consumption in the good state under perfect information (households know they are in the good state). The dotted lower line shows consumption in the bad state under perfect information (households know they are in the bad state). In all scenarios, firms learn about μ .

There is no missing deflation puzzle. The very large drop in output (6 percent in the bad state) leads only to a modest fall in quarterly non-annualized inflation (0.8 percentage points in the bad state). Scaling the shock down by 1/3 yields the numbers observed during the Great Recession.

Output fell by four percent and quarterly non-annualized inflation fell by 0.5 percentage points. See Figure 6. The fact that there is no missing deflation puzzle is entirely due to the modeling of firm inflation expectations. The low perceived persistence of inflation in the initial periods mutes the inflation response by a large amount. See Section 2.1. The New Keynesian Phillips curve is not particularly flat, there is no simultaneous drop in productivity to obtain a small output gap (in fact, the output gap equals minus six percent in period zero in the bad state), and there is no new channel that pushes inflation up.

10 Conclusion

In New Keynesian models, movements in household inflation expectations are of great importance for the propagation of shocks and the effectiveness of policy, especially when the nominal interest rate is at zero. It is therefore desirable to model household inflation expectations in a way that is consistent with data. To this end, I assume that households have different pieces of information and update their inflation expectations slowly over time. The resulting model has properties that are quite different from a model with full-information rational expectations on the household side. The deflationary spiral takes off slowly in bad states of the world. Central bank communication about the aggregate state affects aggregate consumption, and the sign of this effect reverses when the zero lower bound is binding. Forward guidance is less powerful than under perfect information, and the effect on consumption can even have the opposite sign compared to perfect information. The government spending multiplier is smaller than under perfect information. All these results are more pronounced when the economy ends up in an aggregate state with a small prior probability, such as the financial crisis in the U.S. and the sovereign debt crisis in Europe.

The negative effects originating from central bank communication about the current state in the bad state can probably be reduced by designing communication in the right way. In particular, according to the model, it would be desirable to separate communication about future inflation from communication about the current state.

One interesting extension would be to introduce slow updating of inflation expectations on the firm side. I conjecture that this modification would reinforce the main conclusions.²⁵ Another

²⁵Kiley (2014) studies the effects of various policies in a model with sticky information instead of sticky prices on the firm side. Coibion and Gorodnichenko (2015b) show that a standard Phillips curve can match the absence of

interesting extension would be to move from a model with an exogenous information structure to a rational inattention model, as in Sims (2003). One could then study how policy-makers' action and communication strategies affect households' attention to fundamentals (ω) and households' attention to official statements by policy-makers (ζ). As a second step, one could derive optimal action and communication strategies. I leave these extensions to future research.

a persistent decline in inflation during the Great Recession once model inflation expectations on the firm side are replaced by survey data.

A Derivation of equations (42)-(44)

First, state the consumption Euler equation of the four types of households: high, informed types (i.e., households with $\xi_{i,0} = \xi_H$ that are reached by the communication), low, informed types (i.e., households with $\xi_{i,0} = \xi_L$ that are reached by the communication), high, uninformed types (i.e., households with $\xi_{i,0} = \xi_H$ that are not reached by the communication), and low, uninformed types.

The consumption Euler equation of a high, informed type in the good state reads

$$c_{good}^{hi} = -\frac{1}{\gamma} [(\mu - 1) \xi_H - \ln(R) - \mu\pi_{good}] + \mu c_{good}^{hi}.$$

In the bad state, it reads

$$c_{bad}^{hi} = -\frac{1}{\gamma} [(\mu - 1) \xi_H - \ln(R) - (\mu\pi_{bad} + (1 - \mu) \bar{\pi})] + \mu c_{bad}^{hi} + (1 - \mu) \bar{c}.$$

The consumption Euler equation of a low, informed type in the good state reads

$$c_{good}^{li} = -\frac{1}{\gamma} [(\mu - 1) \xi_L - \ln(R) - \mu\pi_{good}] + \mu c_{good}^{li}.$$

In the bad state, it reads

$$c_{bad}^{li} = -\frac{1}{\gamma} [(\mu - 1) \xi_L - \ln(R) - (\mu\pi_{bad} + (1 - \mu) \bar{\pi})] + \mu c_{bad}^{li} + (1 - \mu) \bar{c}.$$

The consumption Euler equation of a high, uninformed type reads

$$\begin{aligned} c^{hu} &= -\frac{1}{\gamma} \left[(\mu - 1) \xi_H - \ln(R) - \left(p_H^{good} \mu\pi_{good} + p_H^{bad} (\mu\pi_{bad} + (1 - \mu) \bar{\pi}) \right) \right] \\ &\quad + p_H^{good} \mu c^{hu} + p_H^{bad} \left(\mu c^{hu} + (1 - \mu) \bar{c} \right). \end{aligned}$$

The consumption Euler equation of a low, uninformed type reads

$$\begin{aligned} c^{lu} &= -\frac{1}{\gamma} \left[(\mu - 1) \xi_L - \ln(R) - \left(p_L^{good} \mu\pi_{good} + p_L^{bad} (\mu\pi_{bad} + (1 - \mu) \bar{\pi}) \right) \right] \\ &\quad + p_L^{good} \mu c^{lu} + p_L^{bad} \left(\mu c^{lu} + (1 - \mu) \bar{c} \right). \end{aligned}$$

Hence, aggregate consumption in the good state equals

$$\begin{aligned} c_{good} &= \zeta \left[\lambda_{good} c_{good}^{hi} + (1 - \lambda_{good}) c_{good}^{li} \right] + (1 - \zeta) \left[\lambda_{good} c^{hu} + (1 - \lambda_{good}) c^{lu} \right] \\ &= \frac{1}{\gamma} (1 - \mu) \bar{\xi}_{good} + \frac{1}{\gamma} \ln(R) + \frac{1}{\gamma} \mu\pi_{good} + \mu c_{good} \\ &\quad + \bar{p}_{good}^{bad} (1 - \mu) \left(\frac{1}{\gamma} \bar{\pi} + \bar{c} \right) - \bar{p}_{good}^{bad} \frac{1}{\gamma} \mu (\pi_{good} - \pi_{bad}). \end{aligned}$$

Aggregate consumption in the bad state equals

$$\begin{aligned}
c_{bad} &= \zeta \left[\lambda_{bad} c_{bad}^{hi} + (1 - \lambda_{bad}) c_{bad}^{li} \right] + (1 - \zeta) \left[\lambda_{bad} c^{hu} + (1 - \lambda_{bad}) c^{lu} \right] \\
&= \frac{1}{\gamma} (1 - \mu) \bar{\xi}_{bad} + \frac{1}{\gamma} \ln(R) + \frac{1}{\gamma} \mu \pi_{bad} + \mu c_{bad} \\
&\quad + \left(1 - \bar{p}_{bad}^{good} \right) (1 - \mu) \left(\frac{1}{\gamma} \bar{\pi} + \bar{c} \right) + \bar{p}_{bad}^{good} \frac{1}{\gamma} \mu (\pi_{good} - \pi_{bad}).
\end{aligned}$$

The probabilities \bar{p}_{bad}^{good} and \bar{p}_{good}^{bad} are given by the expressions on the right-hand sides of equations (30)-(31) multiplied by the factor $(1 - \zeta)$.

Second, state the New Keynesian Phillips curve for the good state and the bad state. In the good state, the New Keynesian Phillips curve reads

$$\pi_{good} = \kappa c_{good} + \beta \mu \pi_{good}.$$

In the bad state, the New Keynesian Phillips curve reads

$$\pi_{bad} = \kappa c_{bad} + \beta (\mu \pi_{bad} + (1 - \mu) \bar{\pi}).$$

Finally, using the last two equations to substitute for π_{good} and π_{bad} in the previous two equations and rearranging yields equations (42)-(44).

B Derivation of equations (45)-(47)

First, state the consumption Euler equation of the four types of households: high, informed types (i.e., households with $\xi_{i,0} = \xi_H$ that are reached by the announcement), low, informed types (i.e., households with $\xi_{i,0} = \xi_L$ that are reached by the announcement), high, uninformed types (i.e., households with $\xi_{i,0} = \xi_H$ that are not reached by the announcement), and low, uninformed types.

The consumption Euler equation of a high, informed type in state $s \in \{good, bad\}$ reads

$$c_s^{hi} = -\frac{1}{\gamma} [(\mu - 1) \xi_H - \ln(R) - \mu \pi_s] + \mu c_s^{hi}.$$

The consumption Euler equation of a low, informed type in state $s \in \{good, bad\}$ reads

$$c_s^{li} = -\frac{1}{\gamma} [(\mu - 1) \xi_L - \ln(R) - \mu \pi_s] + \mu c_s^{li}.$$

The consumption Euler equation of a high, uninformed type reads

$$c^{hu} = -\frac{1}{\gamma} \left[(\mu - 1) \xi_H - \ln(R) - \left(p_H^{good} \mu \pi_{good} + p_H^{bad} \mu \pi_{bad} \right) \right] + p_H^{good} \mu c^{hu} + p_H^{bad} \mu c^{hu}.$$

The consumption Euler equation of a low, uninformed type reads

$$c^{lu} = -\frac{1}{\gamma} \left[(\mu - 1) \xi_L - \ln(R) - \left(p_L^{good} \mu \pi_{good} + p_L^{bad} \mu \pi_{bad} \right) \right] + p_L^{good} \mu c^{lu} + p_L^{bad} \mu c^{lu}.$$

Hence, aggregate consumption in the good state equals

$$\begin{aligned} c_{good} &= \zeta \left[\lambda_{good} c_{good}^{hi} + (1 - \lambda_{good}) c_{good}^{li} \right] + (1 - \zeta) \left[\lambda_{good} c^{hu} + (1 - \lambda_{good}) c^{lu} \right] \\ &= \frac{1}{\gamma} (1 - \mu) \bar{\xi}_{good} + \frac{1}{\gamma} \ln(R) + \frac{1}{\gamma} \mu \pi_{good} + \mu c_{good} - \bar{p}_{good}^{bad} \frac{1}{\gamma} \mu (\pi_{good} - \pi_{bad}). \end{aligned}$$

Furthermore, aggregate consumption in the bad state equals

$$\begin{aligned} c_{bad} &= \zeta \left[\lambda_{bad} c_{bad}^{hi} + (1 - \lambda_{bad}) c_{bad}^{li} \right] + (1 - \zeta) \left[\lambda_{bad} c^{hu} + (1 - \lambda_{bad}) c^{lu} \right] \\ &= \frac{1}{\gamma} (1 - \mu) \bar{\xi}_{bad} + \frac{1}{\gamma} \ln(R) + \frac{1}{\gamma} \mu \pi_{bad} + \mu c_{bad} + \bar{p}_{bad}^{good} \frac{1}{\gamma} \mu (\pi_{good} - \pi_{bad}). \end{aligned}$$

The probabilities \bar{p}_{bad}^{good} and \bar{p}_{good}^{bad} are given by the expressions on the right-hand sides of equations (30)-(31) multiplied by the factor $1 - \zeta$.

Second, state the New Keynesian Phillips curve for the good state and the bad state. In the good state, the New Keynesian Phillips curve reads

$$\pi_{good} = \kappa_c c_{good} + \beta \mu \pi_{good}.$$

In the bad state, the New Keynesian Phillips curve reads

$$\pi_{bad} = \kappa_c c_{bad} + \kappa_g g + \beta \mu \pi_{bad}.$$

Finally, using the last two equations to substitute for π_{good} and π_{bad} in the previous two equations and rearranging yields equations (45)-(47).

References

- [1] Andrade, Philippe, Gaetano Gaballo, Eric Mengus, and Benoit Mojon (2015). “Forward Guidance and Heterogeneous Beliefs,” manuscript, Banque de France and HEC Paris.
- [2] Angeletos, George-Marios, and Jennifer La’O (2013). “Sentiments,” *Econometrica*, 81(2), 739-780.
- [3] Armantier, Olivier, Wändi Bruine de Bruin, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar (2011). “Inflation Expectations and Behavior: Do Survey Respondents Act on Their Beliefs?” *Federal Reserve Bank of New York Staff Reports*, no. 509.
- [4] Bassetto, Marco (2015). “Forward Guidance: Communication, Commitment, or Both?” manuscript, Federal Reserve Bank of Chicago and University College London.
- [5] Calvo, Guillermo A. (1983). “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12(3), 383-398.
- [6] Campbell, Jeffrey, Charles Evans, Jonas Fisher, and Alejandro Justiniano (2012). “Macroeconomic Effects of Federal Reserve Forward Guidance,” *Federal Reserve Bank of Chicago Working Papers*, no. 2012-03.
- [7] Carlstrom, Charles, Timothy Fuerst, and Matthias Paustian (2012). “Inflation and Output in New Keynesian Models with a Transient Interest Rate Peg,” *Bank of England Working Papers*, no. 459.
- [8] Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo (2011). “When is the Government Spending Multiplier Large?” *Journal of Political Economy*, 119(1), 78-121.
- [9] Christiano, Lawrence, Martin Eichenbaum, and Mathias Trabandt (2015). “Understanding the Great Recession,” *American Economic Journal: Macroeconomics*, 7(1), 110-167.
- [10] Coibion, Olivier, and Yuriy Gorodnichenko (2012). “What Can Survey Forecasts Tell Us about Information Rigidities?” *Journal of Political Economy*, 120(1), 116-159.

- [11] Coibion, Olivier, and Yuriy Gorodnichenko (2015). “Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts,” *American Economic Review*, 105(8), 2644-2678.
- [12] Coibion, Olivier, and Yuriy Gorodnichenko (2015b). “Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation,” *American Economic Journal: Macroeconomics*, 7(1), 197-232.
- [13] Curdia, Vasco, and Michael Woodford (2011). “The Central-Bank Balance Sheet as an Instrument of Monetary Policy,” *Journal of Monetary Economics*, 58(1), 54-79.
- [14] Del Negro, Marco, Marc Giannoni, and Christina Patterson (2012). “The Forward Guidance Puzzle,” *Federal Reserve Bank of New York Staff Reports*, no. 574.
- [15] Drautzburg, Thorsten, and Harald Uhlig (2013). “Fiscal Stimulus and Distortionary Taxation,” *Federal Reserve Bank of Philadelphia Working Papers*, no. 13-46.
- [16] Eggertsson, Gauti, and Michael Woodford (2003). “The Zero Bound on Interest Rates and Optimal Monetary Policy,” *Brookings Papers on Economic Activity*, 2003(1), 139-211.
- [17] Fernald, John, Robert Hall, James Stock, and Mark Watson (2017). “The Disappointing Recovery of Output after 2009,” *Brookings Papers on Economic Activity*, Spring 2017, 1-58.
- [18] Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek (2017). “Inflation Dynamics during the Financial Crisis,” *American Economic Review*, 107(3), 785-823.
- [19] Kiley, Michael (2014). “Policy Paradoxes in the New Keynesian Model,” Finance and Economics Discussion Series, Federal Reserve Board.
- [20] Krugman, Paul (1998). “It’s Baaack: Japan’s Slump and the Return of the Liquidity Trap,” *Brookings Papers on Economic Activity*, 1998(2), 137-205.
- [21] Lindé, Jesper, and Mathias Trabandt (2019). “Resolving the Missing Deflation Puzzle,” manuscript, Sveriges Riksbank and Freie Universitaet Berlin.
- [22] Lorenzoni, Guido (2009). “A Theory of Demand Shocks,” *American Economic Review*, 99(5), 2050-2084.

- [23] Lorenzoni, Guido (2010). “Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information,” *Review of Economic Studies*, 77(1), 305-338.
- [24] Lucas, Robert E. Jr. (1990). “Liquidity and Interest Rates,” *Journal of Economic Theory*, 50(2), 237-264.
- [25] Maćkowiak, Bartosz, and Mirko Wiederholt (2014). “Inattention to Rare Events,” manuscript, European Central Bank and Goethe University Frankfurt.
- [26] Maćkowiak, Bartosz, and Mirko Wiederholt (2015). “Business Cycle Dynamics under Rational Inattention,” *Review of Economic Studies*, 82(4), 1502-1532.
- [27] Mankiw, N. Gregory, and Ricardo Reis (2002). “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 117(4), 1295-1328.
- [28] Mankiw, N. Gregory, and Ricardo Reis (2006). “Pervasive Stickiness,” *American Economic Review*, 96(2), 164-169.
- [29] McKay, Alisdair, Emi Nakamura, and Jon Steinsson (2015). “The Power of Forward Guidance Revisited,” manuscript, Boston University and Columbia University.
- [30] Mertens, Karel, and Morten Ravn (2014). “Fiscal Policy in an Expectations-Driven Liquidity Trap,” *Review of Economic Studies*, 81(4), 1637-1667.
- [31] Nakamura, Emi, and Jón Steinsson (2008). “Five Facts About Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 123(4), 1415-1464.
- [32] Swanson, Eric, and John Williams (2014). “Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates,” *American Economic Review*, 104(10), 3154-3185.
- [33] Woodford, Michael (2011). “Simple Analytics of the Government Expenditure Multiplier,” *American Economic Journal: Macroeconomics*, 3(1), 1-35.

Figure 1: consumption over time, benchmark




Figure 2: consumption and nominal interest rate, deterministic decay




Figure 3: consumption and nominal interest rate, households set nominal wage rate




Figure 4: faster updating of inflation expectations




Figure 5: lower prior probability of the good state




Figure 6: Inflation before, during, and after the Great Recession




Figure 7: Bayesian learning about persistence

