

THE FEDERAL RESERVE BANK of KANSAS CITY
RESEARCH WORKING PAPERS

Financial Crises, Unconventional Monetary Policy Exit Strategies, and Agents' Expectations

Andrew T. Foerster

August 2011; Revised May 2015

RWP 11-04

Financial Crises, Unconventional Monetary Policy Exit Strategies, and Agents' Expectations*

Andrew T. Foerster[†]

May 5, 2015

Abstract

This paper studies the implications of exit strategies from unconventional monetary policy. Using a Markov switching DSGE model with financial frictions, agents in the model have rational expectations about the probability of financial crises, the probability of an unconventional response to crises, and the exit strategy used. Selling off assets quickly produces a double-dip recession; in contrast, a slow unwind generates a smooth recovery. Expectations about the exit strategy matter for the initial effectiveness of intervention. Increasing the probability of an unconventional response to crises creates distortions in pre-crisis variables that depend upon the exit strategy. The welfare benefits of increasing the probability of unconventional policy may differ *ex-ante* versus *ex-post*, as can the preferred exit strategy.

Keywords: unconventional monetary policy, exit strategy, Markov switching

JEL Codes: E52, G01, E61

*The views expressed herein are solely those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City or the Federal Reserve System. I thank Juan Rubio-Ramirez, Francesco Bianchi, Craig Burnside, Nobuhiro Kiyotaki, Eric Swanson, Stephanie Schmitt-Grohe, Giuseppe Ferrero, and Jonas Fisher for helpful comments, as well as seminar participants at Duke, the Federal Reserve Banks of Richmond, Kansas City, and San Francisco, the Federal Reserve Board, Boston College, Michigan State, McMaster, the Bank of Canada, the 2011 International Conference on Computing in Economics and Finance, the 2011 North American Summer Meetings of the Econometric Society, the Conference on Zero Bound on Interest Rates and New Directions in Monetary Policy, and the 2012 American Economic Association Meetings.

[†]Research Department, Federal Reserve Bank of Kansas City, 1 Memorial Drive, Kansas City, MO 64198, www.econ.duke.edu/~atf5/, andrew.foerster@kc.frb.org

1 Introduction

The expansion of the Federal Reserve's balance sheet during and after the financial crisis was an attempt to inject capital into the economy and bolster real activity. However, many of the asset purchase programs, designed to alleviate pressure on financial institutions by boosting the value of the assets on their balance sheets, may not be permanent; as a consequence, it remains to be seen how the Federal Reserve will unwind its balance sheet, and how this unwinding may affect the macroeconomy.

This paper addresses the issue of a central bank exiting from its large balance sheet, and how the strategy for unwinding its asset position affects the balance sheets of the financial sector and ultimately the macroeconomy. Using a Markov-switching dynamic stochastic general equilibrium (MS-DSGE) model with a financial sector, it gives agents in the model rational expectations about the probability of a financial crisis occurring, the central bank's decision to intervene with asset purchases or not, and, if asset purchases occur, how they will eventually be unwound.

Using this framework, the paper first considers the effect of exit strategies during and after financial crises occur. After a shock that erodes the value of assets on financial firms' balance sheets, the central bank purchasing assets helps boost their value and limits the damage to the macroeconomy. However, after a crisis has ended, if the central bank sells off assets quickly, a double-dip recession ensues; in contrast, a slow unwind of assets generates a smooth recovery. Through expectations, the exit strategy that the central bank will implement in turn influences the initial effectiveness of intervention.

In addition, expectations about whether a central bank will purchase assets during a financial crisis and how those purchases will be unwound alter the pre-crisis economy. If the central bank always purchases assets when crises occur, knowledge of this policy distorts incentives, leading to changes in the level and volatility of economic activity in periods when there are no crises.

Finally, since a central bank intervening in financial markets lessens the impact of a crisis shock but also distorts the pre-crisis economy, the paper considers the welfare implications of asset purchases and the exit strategy. An issue similar to time inconsistency represents an

important factor in this welfare analysis: *ex-ante*, or before a crisis occurs, making intervention more likely could decrease welfare, but *ex-post*, when a crisis occurs, making intervention more likely could improve welfare. In addition, the welfare benefits of a the exit strategy may depend upon the timing of the decision.

There are many channels through which central bank asset purchases can affect the macroeconomy. Following Gertler and Karadi (2010) and Gertler, et al. (2012), among others, this paper focuses on the balance sheet channel. In this channel, the central bank purchasing assets boosts the price and supply of assets in the economy, which boosts the asset position of financial firms, who are then able to intermediate assets to a greater extent. However, while many papers consider the initial effectiveness of asset purchases when crises occur, this paper's focus is on the exit strategy following purchases.

Models of the Federal Reserve's exit strategy often look at detailed projections of assets, but not within a DSGE framework (Carpenter, et al. (2015), Greenlaw, et al. (2013)). Angeloni, et al. (2011) analyze exit strategies from fiscal and monetary stimulus, but without considering the Markov-switching environment considered in this paper. Markov-switching allows financial crises to be rare events as in Rietz (1988) and Barro (2006), so agents expect them and form expectations over the central bank's decision to intervene conditional upon that rare event occurring. Given the presence of expectations, intervention policy and exit strategies can influence economic behavior prior to crises occurring. These expectations also allow the comparison of *ex-ante* and *ex-post* welfare.

The paper proceeds as follows. Section 2 discusses the model, Section 3 details how the parameters of the economy change according to a Markov Process and details the transitions between regimes. Section 4 discusses the response of the economy to crises with and without intervention, as well as the effects of different exit strategies. Section 5 analyzes the effects of expectations of crisis policies on the pre-crisis economy. Section 6 discusses the welfare implications of policy announcements, and Section 7 concludes.

2 Model

This section describes the model, which is similar to that developed in Gertler and Karadi (2010).¹ The following subsections describe the households, financial intermediaries and government purchase of assets, non-financial firms, and government policy. Section 3 discusses the regime switching and equilibrium in detail.

2.1 Household

A household of unit measure consists of $(1 - \vartheta)$ workers that supply labor and ϑ bankers that own financial intermediaries. Bankers become workers with probability $(1 - \theta)$ in which case they return their accumulated earnings to the household; new bankers receive initial funds from the household. Perfect consumption insurance exists within the household. The household chooses C_t , labor L_t , and bonds B_t to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left\{ \log(C_t - hC_{t-1}) - \frac{\varkappa}{1 + \varphi} L_t^{1+\varphi} \right\}, \quad (1)$$

subject to the budget constraint

$$C_t + B_t + T_t = W_t L_t + \Delta_t + R_{t-1} B_{t-1}. \quad (2)$$

Households earn income from workers earning a wage W_t , receive an amount Δ_t of net profits from financial and non-financial firms, and pay lump sum taxes to the government T_t . Bonds purchased from either financial intermediaries or the government pay a gross real return of R_t in period $t + 1$. In equilibrium, both sources of bonds have no risk, so the household views them as identical with risk-free rate R_t . Households thus have income $R_{t-1} B_{t-1}$ from bonds purchased the previous period.

¹Since the model is well-known, the discussion will be brief. A full set of derivations for the model is available upon request.

2.2 Financial Intermediaries

Financial intermediaries are indexed by j , and use their net worth $N_{j,t}$ along with deposits from households $B_{j,t}$ to purchase claims on non-financial firms $S_{j,t}$ at relative price Q_t

$$Q_t S_{j,t} = N_{j,t} + B_{j,t}. \quad (3)$$

Deposits pay a risk-free rate R_t in period $t + 1$, while claims on non-financial firms pay a stochastic return of $R_{k,t+1}$. As a result, net worth follows

$$N_{j,t+1} = (R_{k,t+1} - R_t) Q_t S_{j,t} + R_t N_{j,t}. \quad (4)$$

Intermediaries can grow their net worth faster than the risk-free rate through either higher realized interest rate spreads $R_{k,t+1} - R_t$ or an expansion of assets $Q_t S_{j,t}$.

Bankers' participation constraint requires a positive expected discounted spread

$$\mathbb{E}_t \beta^i \frac{\varrho_{t+1+i}}{\varrho_t} (R_{k,t+1+i} - R_{t+i}) \geq 0, \text{ for } i \geq 0, \quad (5)$$

where $\beta^i \frac{\varrho_{t+1+i}}{\varrho_t}$ denotes the stochastic discount factor applied to returns in period $t + 1 + i$. They also have expected terminal net worth given by

$$V_{j,t} = \mathbb{E}_t (1 - \theta) \beta \sum_{i=0}^{\infty} \beta^i \theta^i \frac{\varrho_{t+1+i}}{\varrho_t} N_{j,t+1+i} \quad (6)$$

which, along with (4) shows the value of being a financial intermediary increases with expected future interest rate spreads, $(R_{k,t+1+i} - R_{t+i})$, future asset levels $Q_{t+i} S_{j,t+i}$, and the risk-free return on net worth. Financial intermediaries can divert a fraction λ of its assets back to the household every period, which produces an incentive constraint that requires

$$V_{j,t} \geq \lambda Q_t S_{j,t}. \quad (7)$$

Since this constraint always binds for reasonable parameterizations, and all financial intermediaries face this same constraint, total private intermediary assets is given by

$$Q_t S_{p,t} = \phi_t N_t. \quad (8)$$

where ϕ_t denotes the leverage ratio, which in turn is dependent upon expected future interest rate spreads and risk free rates. In the setup of this paper, interest rate spreads and the risk free rate in the future depend upon realizations of financial crises and government policy. New bankers receive start-up funds equal to a fraction $\frac{\omega}{1-\theta}$ of the assets of exiting bankers.

The government possibly serves as a financial intermediary by issuing debt to households and purchasing claims $S_{g,t}$. Therefore, the total value of all assets in the economy equals private assets plus government assets

$$Q_t S_t = Q_t S_{p,t} + Q_t S_{g,t}. \quad (9)$$

The government does not face a moral hazard problem and consequently the central bank faces no constraints on its balance sheet, but it does pay a resource cost of τ for every unit of assets that the central bank owns. This resource cost captures any possible inefficiencies from government intervention. The government's policy rule, discussed in Section 2.4, sets a fraction ψ_t of total intermediated assets, so

$$Q_t S_{g,t} = \psi_t Q_t S_t. \quad (10)$$

As a result, total funds then depends on intermediary net worth by

$$Q_t S_t = \phi_{c,t} N_t, \quad (11)$$

where $\phi_{c,t} = \frac{\phi_t}{1-\psi_t}$ denotes the total leverage ratio for the economy. By setting ψ_t , the central bank manipulates the total leverage ratio $\phi_{c,t}$.

Government intervention through ψ_t has both direct and indirect effects on the economy. A direct effect is that, by purchasing assets, it expands the amount of credit available to non-financial firms. This source of funds used to purchases claims on non-financial firms helps to boost the price of those claims, which indirectly boosts the net worth of private financial intermediaries and allows them to intermediate more assets.

2.3 Non-Financial Firms

There are three types of non-financial firms: intermediate goods producers, capital producers, and retail firms.

2.3.1 Intermediate Goods Producers

Intermediate goods firms operate in a competitive environment, producing using capital and labor according to

$$Y_{m,t} = A_t (U_t \xi_t K_{t-1})^\alpha L_t^{1-\alpha} \quad (12)$$

where A_t denotes total factor productivity, U_t the capital utilization rate, and ξ_t is capital quality. Total factor productivity follows the process

$$\log A_t = \rho_a \log A_{t-1} + \sigma_a \varepsilon_{a,t}, \quad (13)$$

where $\varepsilon_{a,t} \sim N(0, 1)$ denotes the TFP shock. Firms also face changes in capital quality ξ_t , which evolves according to the process

$$\log \xi_t = (1 - \rho_\xi(s_t)) \log \xi_m(s_t) + \rho_\xi(s_t) \log \xi_{t-1} + \sigma_\xi \varepsilon_{\xi,t}, \quad (14)$$

where $\varepsilon_{\xi,t} \sim N(0, 1)$ denotes a capital quality shock and s_t indicates the regime at time t . This regime changes according to a Markov process, and affects the mean of the process $\log \xi_m(s_t)$, and persistence around the mean $\rho_\xi(s_t)$. The capital quality measure ξ_t alters the effective capital stock of the economy $\xi_t K_{t-1}$ and thereby exogenously changes the value of capital in the economy. Section 3 contains a more detailed description of the Markov switching process for s_t .

Firms purchase capital by issuing claims S_t to purchase capital at price Q_t and hire labor at wage W_t . Capital depreciates depending on the utilization rate $\delta(U_t)$, which has an elasticity of ζ . The return on capital is given by

$$R_{k,t} = \frac{\left[P_{m,t} \alpha \frac{Y_{m,t}}{\xi_t K_{t-1}} + Q_t - \delta(U_t) \right] \xi_t}{Q_{t-1}}, \quad (15)$$

which highlights how changes in the capital quality measure ξ_t produce exogenous changes in the return on capital. Financial crises are events in which ξ_t drops precipitously, leading to large declines in the return to capital. Since financial intermediaries are the owners of the claims, big declines in $R_{k,t}$ reduce intermediaries' net worth accumulation, as seen in equation (4). Lower intermediary net worth lowers output in part by decreasing the capital stock. Government intervention through ψ_t on the other hand, helps to raise the price of claims Q_t , which boosts the return on capital. With a higher return on capital, financial firms increase the quantity of capital, which raises output.

2.3.2 Capital Producers

Competitive capital producers buy used capital from intermediate goods firms, repair depreciated capital, build new capital, and sell it to the intermediate goods firms. Gross investment I_t equals the total change in capital

$$I_t = K_t - (1 - \delta(U_t)) \xi_t K_{t-1}. \quad (16)$$

These firms face a quadratic adjustment cost on net investment, defined as gross investment less depreciation, with ι denoting the inverse of the elasticity of net investment to the capital price. Since the capital quality measure ξ_t shows up in the capital accumulation equation (16), financial crises effectively destroy a portion of the capital stock. Government purchases of assets helps to increase demand for assets by financial intermediaries, which channels through intermediate goods produces and ultimately encourages capital producers to increase their production.

2.3.3 Retail Firms

Retail firms, are indexed by $f \in [0, 1]$, and repackage intermediate output $Y_{m,t}$ into differentiated products $Y_{f,t}$ which they sell at price $P_{f,t}$. Firms set their price according to Calvo pricing with indexation to lagged inflation: a firm can re-optimize each period with probability $(1 - \gamma)$, and with probability γ sets $P_{f,t} = \Pi_{t-1}^\mu P_{f,t-1}$, where $\mu \in [0, 1]$ denotes the degree of price indexation. Final output equals a CES aggregate of retail firm goods with elasticity of substitution ε .

When financial crises occur, lower intermediate output and lower demand from households cause retail firms that can re-optimize to lower their prices, so inflation will fall. However, the presence of sticky prices will mean that firms that cannot re-optimize will have prices that are relatively high, leading to lower demand and hence aggregate output will be lower than if prices were fully flexible.

2.4 Government Policy

Government policy has three aspects: a standard monetary policy rule, an unconventional policy rule, and fiscal policy. Conventional monetary policy sets the nominal interest rate r_t according to a Taylor rule

$$\left(\frac{r_t}{r_{ss}}\right) = \Pi_t^{\kappa_\pi} \left(\frac{Y_t}{Y_t^*}\right)^{\kappa_y} \exp(\sigma_r \varepsilon_{r,t}), \quad (17)$$

where r_{ss} denotes the steady state nominal rate, κ_π and κ_y control responses to inflation and to deviations of output from its flexible-price counterpart Y_t^* , respectively, and $\varepsilon_{r,t} \sim N(0, 1)$ denotes an interest rate shock.² Nominal bonds are held in zero net supply, and households have a no-arbitrage condition between deposits and nominal bonds.

The government sets its unconventional asset holding ψ_t according to

$$\psi_t = \nu(s_t) ((\mathbb{E}_t R_{k,t+1} - R_t) - (R_{k,ss} - R_{ss})) + \rho_\psi(s_t) \psi_{t-1} \quad (18)$$

where the response to the expected interest rate spread $\nu(s_t)$ and an autoregressive term $\rho_\psi(s_t)$ change according to a Markov process to be discussed in Section 3. This rule allows the government to purchase assets depending upon the interest rate spread; higher spreads accompany worse financial crises, since an erosion in private intermediaries' net worth means they cannot purchase as many assets due to their leverage constraints, and hence there are arbitrage opportunities they cannot fully exploit. In this way, a shift in the rule for ψ_t when a shift in

²Given the model's complexity, this paper ignores the effects of the zero lower bound on interest rates. Including the zero bound can impact the effectiveness of asset purchases (see Gertler and Karadi (2010), Curdia and Woodford (2011), and Chen, et al. (2012)).

capital quality ξ_t occurs is a policy response directed at the heart of the economy's problems: rather than adjusting the nominal interest rate, the asset purchases directly improve the balance sheets of financial intermediaries by increasing asset prices, and this policy response will tend to produce more muted downturns after a crisis occurs.

Finally, the government has a fixed amount of spending G equal to a fraction \bar{g} of steady state output Y_{ss} , plus it must pay a resource cost τ on its assets. It finances these via lump-sum taxes T_t and the return from its previously held assets. Consequently, the government's budget constraint requires

$$G + \tau\psi_t Q_t K_t = T_t + (R_{k,t} - R_{t-1}) B_{g,t-1}. \quad (19)$$

3 Regime Switching and Equilibrium

This section embeds the core model into a regime switching framework. Parameters in two equations switch according to a Markov process: the exogenous process for capital quality (14) and the unconventional policy rule (18). The next two subsections discuss the switching in these equations, Section 3.3 summarizes the switching, Section 3.4 covers the calibration and solution method, and Section 3.5 relates the model to recent US history.

3.1 Markov Switching in the Capital Quality Process

The first switching equation governs the exogenous process for capital quality (14):

$$\log \xi_t = (1 - \rho_\xi(s_t)) \log \xi_m(s_t) + \rho_\xi(s_t) \log \xi_{t-1} + \sigma_\xi \varepsilon_{\xi,t}. \quad (20)$$

The functional form allows for changes in the mean of the process through the term $\xi_m(s_t)$, and changes in the persistence $\rho_\xi(s_t)$, where s_t denotes the state of the Markov Process. Allowing for changes in the mean and the persistence captures a wide variety of possible switching dynamics. As mentioned, changes in capital quality drive exogenous fluctuations in the value of capital, and significant declines generate a financial crisis.

The two switching parameters $\xi_m(s_t)$ and $\rho_\xi(s_t)$ each take on two values, and these values depend upon a common Markov process. Specifically, the values depend upon whether or not the economy experiences a financial crisis. In non-financial crisis normal times, then the process has mean $\xi_m^n = 1$, and persistence $\rho_\xi^n \in (0, 1)$, where the superscript n denotes "no crisis." With probability p_c , the economy experiences a financial crisis, and the mean of the process switches to a lower level $\xi_m^c < 1$, where the superscript c indicates "crisis" and the persistence switches to $\rho_\xi^c = 0$. With probability p_e , the economy exits the crisis and returns to the "no crisis" mean and persistence.

The dual changes in parameters between non-crisis and crisis have two effects. First, when the economy enters a crisis, the crisis mean $\xi_m^c < 1$ implies that average capital quality decreases. The crisis persistence $\rho_\xi^c = 0$ implies that capital quality jumps downward to this lower mean. Second, when the economy leaves a crisis, the mean $\xi_m^n = 1$ implies that average capital quality returns to its original level, but the persistence $\rho_\xi^n \in (0, 1)$ implies a gradual reversion to this higher mean. These two features capture the typically rapid entry into financial crises, with a quick transition to a low capital quality, while after the crisis ends the economy takes time to return back to its pre-crisis level.

The transition probabilities also assume an asymmetry between entering into and exiting out of financial crises. Independent parameters govern the probability of entering (p_c) and exiting (p_e), which incorporates a wide variety of timing assumptions. In, for example, Gertler and Karadi (2010) or Gertler and Kiyotaki (2010), crises represent zero probability events ($p_c = 0$) that only last one period ($p_e = 1$). On the other hand, in Gertler, et al. (2012), crises occur as independent events ($p_c = 1 - p_e$). Most importantly, in this paper, the probabilities allow agents to expect that crises can occur, and, if a crisis does occur, it can last several quarters.

3.2 Markov Switching in Unconventional Policy

The second switching equation governs unconventional policy (18):

$$\psi_t = \nu(s_t) ((\mathbb{E}_t R_{k,t+1} - R_t) - (R_{k,ss} - R_{ss})) + \rho_\psi(s_t) \psi_{t-1} \quad (21)$$

where the Markov switching affects the response to the expected interest rate spread $\nu(s_t)$ and an autoregressive term $\rho_\psi(s_t)$. While an independent Markov process controls the exogenous process for capital quality, the Markov process for the unconventional policy rule depends on the realization of the capital quality process. This feature captures the fact that, when a crisis occurs, the central bank may or may not intervene, but the onset of a crisis triggers the decision to intervene or not. In other words, the central bank will never begin intervention without a crisis. In addition, the central bank may continue to intervene beyond the end of the crisis.

Prior to a crisis, the central bank does not intervene, so it sets the spread response parameter to $\nu^n = 0$ and the persistence to $\rho_\psi^n \in [0, 1)$, where the superscript n denotes "no intervention." When a crisis occurs, which happens with probability p_e , the central bank intervenes with unconventional policy with probability p_i , where i denotes "intervention." If it does not intervene, then the spread response and persistence remain $\nu^n = 0$ and ρ_ψ^n , respectively. If the central bank does intervene, it responds to the interest rate spread with magnitude $\nu^i > 0$ and sets the persistence parameter at $\rho_\psi^i = 0$. These dual switches in parameters imply that the central bank increases its intervention with the magnitude of the crisis, since the intervention increases with the expected spread. After the crisis ends, which occurs with probability p_e , the central bank enters a period where it doesn't adjust its degree of intervention, so it sets the response to $\nu^h = 0$ but $\rho_\psi^h = 1$, where h denotes "holding;" this form implies that $\psi_t = \psi_{t-1}$. Finally, with probability p_s the central bank stops intervening, in which case $\nu^n = 0$ and $\rho_\psi^n \in [0, 1)$, meaning that intervention eventually returns to zero.

The Markov switching specification implies that when the central bank intervenes, it does so by purchasing assets depending on the expected spread $\mathbb{E}_t R_{k,t+1} - R_t$. When it does not intervene, it sets $\nu^i = 0$, and the persistence remains $\rho_\psi^n \in [0, 1)$. These values imply two features about the no intervention case. First, if $\psi_{t-1} = 0$, meaning the central bank previously had no assets, then it will continue to have no assets. Second, if it does have assets, so $\psi_{t-1} > 0$, then it may continue to hold assets, but will be decreasing its balance sheet size. Consequently, the parameter ρ_ψ^n captures the exit strategy after a crisis. If $\rho_\psi^n \in (0, 1)$, when the rule switches from intervention to no intervention, there will be an unwind of the accumulated assets. On the

other hand, if $\rho_\psi^n = 0$, then when the rule switches to no intervention, then instantly $\psi_t = 0$, meaning the central bank exits the asset market with an immediate sell-off.

The policy rule (18) captures several important aspects of policy. First, the magnitude of the intervention depends upon the size of the crisis, in that larger spreads lead to a higher degree of intervention. Second, it allows the central bank to maintain its degree of intervention for an extended period after the crisis ends and spreads begin returning to their non-crisis level; since slow recoveries often follow financial crises (Reinhart and Rogoff (2009)), the central bank may desire to keep intervention in place for several years after the initial shock. Third, the stopping probability p_s captures some uncertainty in markets about exactly how long the central bank will maintain its intervention. Finally, the exit strategies consider two divergent cases: with slowly unwound intervention, or a fast sell-off caused by a desire to get out of capital markets quickly.

3.3 Regime Switching Summary

Based on the preceding discussion of the switching in the capital quality process and the unconventional policy equation, the model has four total regimes. The first regime, "normal times," has high average capital quality and the central bank either holds no assets or unwinds its assets. The second regime, called "crisis without intervention," has low average capital quality and the central bank holds no assets or unwinds. The third regime, "crisis with intervention," has low average capital quality and the central bank intervenes according to the expected spread. The fourth regime, "post-crisis with intervention," has high average capital quality and the central bank maintaining its intervention level. Table 1 summarizes the switching parameters across these regimes.

Given the transition probabilities between the regimes, the transition matrix has elements $\mathcal{P}_{i,j} = \Pr(s_{t+1} = j | s_t = i)$:

$$\mathcal{P} = \begin{bmatrix} 1 - p_c & p_c(1 - p_i) & p_c p_i & 0 \\ p_e & 1 - p_e & 0 & 0 \\ p_e p_s & 0 & 1 - p_e & p_e(1 - p_s) \\ (1 - p_c)p_s & 0 & p_c & (1 - p_c)(1 - p_s) \end{bmatrix}. \quad (22)$$

3.4 Calibration and Model Solution

Given that the recent financial crisis and the Federal Reserve's unconventional monetary policy response represent unique events in history, the results in this paper rely on calibrated rather than estimated parameters. However, the results explore the implications of key parameters. The unit of time equals a quarter. Table 2 shows the baseline calibration of parameters for preferences and production, which follow Gertler and Karadi (2010).

Table 3 shows the calibration of the regime switching parameters. The transition probabilities and switching parameters introduced in this paper capture various aspects of the recent financial crisis. First, following Barro (2006), the probability of crises equals $p_c = 0.005$, implying a roughly two percent chance of a crisis per year. Motivated by US history in which interest rate spreads spiked to above 5% for seven months, the probability of exiting a crisis equals $p_e = 0.5$, implying an expected duration of two quarters. The probability of intervention p_i and of intervention stopping p_s will vary, but the baseline calibration has $p_s = 1/18$, which, along with the expected crisis duration, implies a total duration of intervention of 20 quarters before exit begins. Alternatively, the central bank could have a shorter or longer expected holding duration of either 12 or 28 total quarters, in which case $p_s = 1/10$ or $1/26$, respectively.

Following Gertler and Karadi (2010) and Gertler and Kiyotaki (2010), among others, during a financial crisis the effective capital stock declines by five percent, roughly matching the shock to the housing market, so $\xi_m^c = 0.95$. Capital quality has persistence $\rho_\psi^n = 0.66$ in normal times. The central bank responds to the expected interest rate spread by a factor $\nu^i = 3$, which produces degrees of intervention comparable to Gertler and Karadi (2010) and the experience of the US economy. The non-intervention regimes have persistence of intervention of $\rho_\psi^n = 0.99$, which means that when intervention ends, the central bank unwinds its assets very slowly. An

alternate calibration considers the effects of $\rho_{\psi}^n = 0$, in which the central bank sells its stock of assets off all at once. Lastly, the benchmark calibration sets the resource cost of intervention at $\tau = 0$, which implies no loss of output generated by the central bank holding assets. The welfare calculations of Section 7 consider $\tau = 0.0005$ and $\tau = 0.003$.

To solve the described Markov switching DSGE model, this paper uses the perturbation approach of Foerster, et al. (2013). Given the switching nature of the parameters, standard perturbation techniques such as those developed in Schmitt-Grohe and Uribe (2004) do not apply. The iterative procedure developed by Foerster, et al. (2013) finds an approximation to the solution to the economy by guessing a set of approximations under each regime; given a guess, each regime's approximation follows from standard perturbation techniques, and the iterative algorithm stops when obtained approximations equal the guesses. This perturbation approach has two major advantages. First, the method introduces Markov switching from first principles, which in turn allows for a flexible environment that includes switching that affects the steady state of the economy. Given that the switching equations involve switching means, the economy's regime-specific steady states will differ, a feature perturbation handles easily. In addition, perturbation allows for second-order approximations, which improve the ability to capture the effects of expectations and for welfare calculations. For tractability, the impulse responses in Section 4 use first-order approximations, while the effects of expectations and welfare in Sections 5 and 6 use second-order approximations.

3.5 Model and the U.S. Experience

As noted, due to the fact that the Federal Reserve's recent asset purchases are a unique experience in history, it is beyond the scope of this paper to estimate the model. However, before proceeding to the results, it is worth considering the parallels between the US experience and the model. First, many of the private sector parameters are based upon estimates of Primiceri, et al. (2006) in order to match features of the non-crises data. Second, while there are many other types of financial frictions such as collateral constraints, the framework used in this paper is based upon the financial accelerator model developed by Bernanke et al. (1999). Brzoza-

Brzeina and Kolasa (2013) find this financial accelerator setup fits US data, although Adrian et al. (2012) note this model may not account for shifts in credit and bonds. In addition, asset purchases work through several channels; by stressing the balance sheet channel and not, for example, the risk-taking-channel, the model may miss on risk premia and bank equity.³ The asset purchase response is broadly consistent with the Federal Reserve's several asset purchase programs, where it expanded its balance sheet into non-Treasury securities, held these assets for many quarters beyond the crisis, and has considered how to unwind its holdings of these assets.

In terms of aggregates, a financial crisis shock produces many of the responses seen in 2008-2009 in the US: declines in output, consumption, and investment along with widening credit spreads. As the following results demonstrate, the response of the macroeconomy depends upon agents expectations and the exit strategy used. Since the parameters governing expectations would be nearly impossible to pin down based upon a single crisis and response, the results explore the effects of different parameters.

4 Crisis Responses and Exit Strategies

Having discussed the basic model and the nature of regime switching, this section considers financial crises, the effects of intervention, and exit strategies. Given the Markov switching transitions, each regime has uncertain duration; the following results describe a "typical" crisis. In these experiments, agents know the probabilities $\{p_c, p_e, p_i, p_s\}$ that dictate the transitions in the economy. In a typical crisis, the realized durations equal the expected durations: the crisis lasts $1/p_e$ periods, and unwinding of intervention begins $1/p_s$ periods after the crisis ends. Given the baseline parameterization of $p_e = 0.5$ and $p_s = 1/18$, if the typical crisis begins at $t = 0$, it ends in $t = 2$, and unwinding begins in $t = 20$. Across models, as Section 5 shows, the pre-crisis level of the economy changes depending upon the expectations about policy during crises. In order to focus on the effects of crises and intervention, the following results therefore simply consider deviations from the pre-crisis average level of the economy produced by the

³See, for example, Angeloni and Faia (2013), Adrian and Boyarchenko (2012), Meh and Moran (2010).

typical crisis for a given model.

4.1 Intervention Versus No Intervention

First, consider the effects of a crisis under a guarantee of intervention ($p_i = 1$) versus one of no intervention ($p_i = 0$). With guaranteed no intervention, the economy remains in the "normal times" regime for $t < 0$, experiences a crisis in period $t = 0$ when it automatically moves to the "crisis without intervention" regime, and then at $t = 2$ the crisis ends and the economy moves back to the "normal times" regime. With guaranteed intervention, the economy remains in the "normal times" regime for $t < 0$, and when a crisis occurs at $t = 0$ it moves automatically to the "crisis with intervention" regime. Then, at $t = 2$, the crisis ends and the economy moves to the "post-crisis with intervention" regime, where it stays until $t = 20$, at which time it switches to the "normal times" regime and the central bank unwinds its intervention.

Figure 1 depicts the responses to the typical crisis with $p_i = 0$ and $p_i = 1$. When a crisis occurs, capital quality drops five percent for the duration of the crisis – two periods in this case – and then returns to its pre-crisis levels. When $p_i = 0$, the level of intervention remains at zero. The shock to capital quality reduces banker net worth, driving the leverage ratio up, and causing a drop in the price of capital, which creates a financial accelerator effect of further diminishing banker net worth. Since the financial intermediaries have less net worth, they borrow less due to their leverage constraint, which drives interest rates down and spreads up, and capital declines with less investment. The increase in spreads lasts two quarters before declining, roughly corresponding to the recent crisis in the US, and in contrast to the one-period spike in spreads generated by a one-period shock. In total, the drop in output from its pre-crisis level nears 8%.

When $p_i = 1$, the central bank intervenes on impact of the crisis, purchasing around six percent of assets and holding them for 20 quarters. The additional demand for capital provided by the central bank in this circumstance works against the financial accelerator effect: the price of capital drops slightly less, leading banker net worth to drop slightly less, and the leverage ratio and interest rate spreads to increase less than without intervention. The increased ability of the

private sector to provide capital, as well as that provided by the central bank, yields a trough in output around 7% of below its pre-crisis level – intervention lessens the recession by about 1 percentage point. At $t = 20$, the central bank begins to unwind, and does so very gradually, since $\rho_\psi^n = 0.99$ in this case, leading to a smooth, albeit slow, transition of the economy back to its pre-crisis levels.

4.2 Exit Strategies

Now suppose the central bank guarantees intervention ($p_i = 1$), but that the unwind rate after intervention ends differs. With a slow unwind ($\rho_\psi^n = 0.99$), the results comparing intervention versus no intervention showed that the economy transitions slowly but smoothly back to its pre-crisis level. Figure 2 shows the effects of this slow unwind contrasted with the case of a sell-off ($\rho_\psi^n = 0$). With both the slow unwind and sell-off, the intervention rule remains the same, but at $t = 20$, when the economy switches back to the "normal times" regime, the central bank immediately unloads its asset holdings rather than unwinding them over an extended period. This sell-off has two main implications: contemporaneous to the sell-off and beforehand through expectations.

When the sell-off occurs at $t = 20$, the central bank unloading its assets immediately serves as a fire sale of assets, which depresses the price of capital. The decline in the price of capital diminishes the net worth of bankers, leading to a decline in interest rates, and a jump in the private leverage ratio and the interest rate spread. Since the central bank no longer provides capital and the loss in net worth decreases the private sector's ability to do so, the rebound in capital slows from a loss in investment, and output drops again, by approximately one percentage point. Importantly, all of these responses mimic what occurred during the initial crisis, except that at $t = 20$ capital quality has fully recovered. In other words, the sell-off creates a second financial crisis, generating a double-dip recession due exclusively to policy.

The expectations of a sell-off also create differences in policy effectiveness even before the sell-off occurs. The slow unwind and sell-off have identical policy rules in that they respond equally to changes in the interest rate spread ($\nu^i = 3$), they only differ in how quickly unwind

occurs. When agents in the economy expect a sell-off to occur at some future date, they must worry about the crisis but also the double-dip recession. In fact, given household consumption smoothing through habits, they have a strong incentive to provide more labor and save to smooth consumption through the ensuing double dip. As a consequence, a policy of a sell-off leads to a smaller initial recession, a lower spread, and hence less intervention than the case with a slow unwind. The sell-off policy produces an initial drop in output about half a percentage point less than with an unwind.

Consequently, the sell-off exit strategy represents an interesting trade-off. It creates a double-dip recession when exit occurs, but because agents in the economy expect this sell-off, they take actions to smooth consumption. These expectations make policy actually more effective during the initial stage of the crisis, since output falls by a lower amount for a smaller degree of intervention.

4.3 Holding Duration

The previous results discussed the fact that an exit strategy of an immediate sell-off produces a slightly better outcome through the expectations channel but creates a double-dip recession when the sell-off occurs. Now consider different holding durations as well as the possibility for a moderate unwind strategy. Figure 3 shows the responses of output to the baseline stopping probability $p_s = 1/18$ versus the alternatives of a shorter or longer holding time, at $p_s = 1/10$ or $p_s = 1/26$ quarters, respectively. For each duration, the figure shows the responses to both the slow unwind ($\rho_\psi^n = 0.99$) and the sell-off ($\rho_\psi^n = 0$) previously considered, but also a moderate unwind ($\rho_\psi^n = 0.50$).

Changing the expected holding duration produces similar responses to the baseline duration. The slow unwind produces a gradual recovery in output to its pre-crisis level, with similar responses for all durations. Similarly, for all durations the sell-off policy produces a smaller initial drop but generates a double-dip recession when the central bank exits from its asset position. The size of the double-dip recession decreases with duration: selling-off assets soon after the crisis with a still-weak economy leads to larger negative effects. The sell-off after the

longer holding duration still produces the double-dip recession, however.

In addition to the change in duration, the moderate unwind case represents a mixture between the slow unwind and the sell-off cases. With the moderate unwind, the central bank exits quickly, leading to a less-immediate but similarly sized double-dip as that experienced with a sell-off. Consequently, a moderate unwind still produces a double-dip recession, but a more gradual one that simply delays the recovery.

5 Pre-Crisis: Effects of Expectations

The previous section focused on the effects of intervention and exit strategies during crises, this section examines how expectations of intervention and exit strategies affect non-crisis times. The Markov switching framework established in Section 3 gives agents expectations that crises can occur, as well as expectations about the probability of intervention by the central bank, and the duration of intervention and exit strategy if intervention does occur. These expectations affect prices and quantities before crises occur. Consequently, this section examines how the stochastic steady state of the economy associated with the "normal times" regime changes as the probability of intervention conditional on a crisis increases from $p_i = 0$ to $p_i = 1$, and the implications of the expected exit strategy.

To characterize the effects of expectations, these results use a second-order approximation, which Foerster, et al. (2013) show improves accuracy and captures lack of certainty equivalence. In other words, expectations about future regimes will affect the means and standard deviations in the normal times regime. The ergodic distribution in the normal regime results from a long simulation in which agents in the economy expect regime switches, but *ex-post* along the simulated path, no switches occur.

5.1 Pre-Crisis Stochastic Steady State Mean

In the baseline parameterization, agents perceive crises occur with probability $p_c = 0.005$, crises end with probability $p_e = 0.5$, and any intervention stops with probability $p_s = 1/18$. Figure 4

shows the percent change in the mean of the "normal times" stochastic steady state relative to a benchmark economy where $p_c = 0$, which implies agents do not expect crises, and hence do not expect intervention.

Consider the baseline parameterization with $p_c = 0.005$ but $p_i = 0$, which makes the exit strategy irrelevant. Moving from an economy where agents do not expect crises ($p_c = 0$), to one where they expect crises without intervention has two main implications. Households, on the one hand, have an incentive to precautionary save in order to smooth consumption during times of crises. In the stochastic steady state, this incentive increases household savings, boosting up capital accumulation and raising output and consumption. On the other hand, crises bring poor interest rate realizations for bankers, who will supply more net worth, have lower leverage, and consequently create a lower amount of capital for the economy, leading to lower output and consumption.

Figure 4 shows that in aggregate, the latter of these effects dominates: the economy with crises and $p_i = 0$ has lower capital, lower labor, and output, and higher banker net worth and capital prices. As p_i increases from 0 to 1, agents expect intervention with a higher probability, and so the exit strategy matters. Since intervention dampens the effects of crises, increasing the probability of intervention tends to erode households' precautionary incentive, lowering the capital stock. However, intervention provides more favorable interest rate conditions for bankers, which will tend to increase their leverage ratio, and increase the capital stock. On net, which effect dominates depends on the exit strategy. Banker net worth increases, the spread between interest rates increases, and the real rate declines. When the exit strategy is the fast sell-off, the capital stock increases slightly from the precautionary channel, leading to a only slight decline in output; when the exit strategy is the slow unwind output falls more dramatically as p_i increases.

These results highlight that even in normal times, the low probability event of financial crises has effects on the level of economic activity, and the expectation of asset purchases or not plus the exit strategy matters.

5.2 Pre-Crisis Stochastic Steady State Standard Deviations

Expectations affect not only the means of economic variables in the "normal times" regime, but their standard deviations as well. Figure 5 shows the percentage change in the "normal times" stochastic steady state's standard deviations relative to the no crises benchmark economy ($p_c = 0$). When agents expect crises but not intervention, households have a precautionary motive and bankers have higher net worth. These effects lead to lower volatility of the economy in the normal times regime, as all variables have lower standard deviations when agents expect crises but expect no intervention. As the intervention probability p_i increases, intervention erodes precautionary behavior and requires bankers to have more net worth. These factors raise the volatility of banker net worth and the interest rate spread. The effects on the volatilities of output, consumption, and capital differ depending on the exit strategy. A sell-off exit strategy led to higher output, consumption, and capital than under the unwind strategy, and this fact translates to lower volatility in each of these variables.

While the previous results showed the level of activity changed based upon expectations and exit strategies, Figure 5 shows this extends to the volatility of the economy in normal times.

6 Welfare Calculations

Having considered the effects of policy announcements and expectations during and before crises, this section turns to evaluating the overall welfare gains or losses from different policy announcements. In particular, Section 4 discussed the fact that guaranteed intervention had benefits relative to no intervention during crises, since intervention helps bolster the economy and alleviate the crisis. However, a slight trade-off depended upon the exit strategy: the immediate sell-off case produced a slightly lower drop in output and consumption, but upon exit, the economy experienced a double-dip recession. In addition, in Section 5, the effects of increasing the probability of intervention created pre-crisis distortions.

Importantly, in addition to the probability of intervention and the exit strategy considered, two factors affect the welfare costs. First, the resource cost τ of central bank intermediation

matters for welfare, since a high cost implies a larger loss of output from intervention, which may lower welfare. Second, the timing of the calculation matters for welfare costs. Specifically, the household's gain or loss in welfare from different policies depends upon whether they experience a crisis or not. The *ex-ante* welfare costs measure the willingness to pay for intervention before a crisis, while the *ex-post* welfare costs measure willingness to pay when a crisis occurs, but before realization of the intervention outcome.

Adding the value function formulation of household preferences (1) to the equilibrium conditions, along with a second-order approximation, gives an accurate measure of the value function given the state, regime, intervention probability p_i , stopping probability p_s , exit strategy ρ_ψ^n , and resource cost τ .

The welfare measure equals the percentage increase in expected lifetime consumption under guaranteed no intervention that would make households indifferent between the increase in consumption and a policy of a given probability of intervention and exit strategy. Positive welfare measures indicate that intervention increases welfare, since households need additional consumption under the given specification to mimic positive intervention probabilities. Negative welfare measures then imply intervention decreases welfare, with households willing to give up consumption rather than have positive intervention probabilities.

6.1 Welfare and the Resource Cost

Figure 6 depicts the change in the welfare measure as the resource cost τ and the intervention probability p_i change. The top two plots shows the baseline case, when $\tau = 0$, which corresponds to no efficiency loss from intermediation, the middle two plots and bottom two plots show $\tau = 0.0005$ and $\tau = 0.003$, respectively. Each panel shows the welfare measure in consumption units for a given intervention probability, both with the slow unwind ($\rho_\psi^n = 0.99$) and the sell-off ($\rho_\psi^n = 0$) cases. The left hand plots show *ex-ante* welfare, while the right hand plots show *ex-post* welfare, which correspond to when the economy experiences a crisis but before realization of the intervention outcome.

The top two plots show that, even when $\tau = 0$, increases in intervention probability can

have different implications *ex-ante* versus *ex-post*. Before crises, increasing the probability of intervention increases the distortion in the economy, and the magnitude of the distortion depends on whether the unwind or sell-off exit strategy is used. As a result, a sell-off strategy is welfare increasing, but a slow unwind is welfare decreasing. When a crisis hits, increasing the intervention probability is welfare increasing under both exit strategies; however in this case, agents prefer intervention with a slow unwind since this strategy limits the magnitude of the drop from the crisis and avoids a double-dip recession.

The middle two panels change some of the implications of intervention. When $\tau = 0.0005$, the *ex-ante* welfare calculations are unchanged, but *ex-post* the sell-off strategy has higher welfare. In this case, the larger resource cost of intervention diminishes the benefits of intervention, so agents are willing to have a double-dip recession in order to get a faster unwind.

Finally, when $\tau = 0.003$, shown in the bottom two plots, intervention decreases welfare in all cases. In the *ex-post* case, the sell-off exit strategy dominates the slow unwind strategy; the high resource cost makes agents in the economy prefer to experience the quick exit and double-dip recession rather than have the central bank slowly unwind its assets.

These changes between the *ex-ante* versus *ex-post* welfare implications suggest that the central bank may face a time-inconsistency problem in designing optimal policy, and hence may find commitment difficult. For example, in the case of $\tau = 0$, when a crisis occurs a government choosing a welfare maximizing intervention strategy would choose $p_i = 1$ with a slow unwind. However, in non-crisis times, this intervention policy is welfare decreasing, so the government would either choose $p_i = 0$ if it had to stick to an exit strategy, or choose $p_i = 1$ and the fast unwind. In the context of this paper, the government can guarantee its strategy, but a government that could change policies would have to internalize the fact that agents might expect different policies in the future, and these expectations in turn could possibly alter the government's choice of policy.

6.2 Welfare and Holding Duration

Figure 7 shows that the differences between the *ex-ante* and *ex-post* welfare measures when $\tau = 0.0005$ for different intervention stopping probabilities. The probability of stopping varies from $p_s = 1/18$ to the shorter holding duration $p_s = 1/10$ and the longer duration $p_s = 1/26$ considered in the crises responses. When the expected total holding duration changes from 20 to either 12 or 28 quarters, the *ex-ante* welfare implications do not change. In contrast, the *ex-post* welfare, while always positive in this case, differs between the exit strategies. When the expected holding time is short, the slow unwind is the welfare-preferred exit strategy, but when the holding time increases, the sell-off strategy becomes preferred. This result is due to the fact that when $\tau = 0.0005$ there is a slight efficiency loss by asset holdings, and so agents must weigh this cost versus the possibility of a double-dip recession.

Again, these results highlight that a government setting optimal policy may face time-inconsistency. While the previous results showed that the intervention decision and exit strategy vary depending in whether the economy is in a crisis or not, these show that a policymaker may have incentives to change their expected holding duration. For example, a government that chose a short holding duration policy with a sell-off exit strategy may wish to either do a slow unwind exit strategy or extend the expected holding duration before sell-off occurs. In the absence of the guarantee assumed in this paper, the government would have to internalize that agents may expect policies to change over time when deciding an optimal set of intervention policies.

7 Conclusion

This paper used a model of unconventional monetary policy along with regime switching to study the effects of exit strategies and expectations. After intervention, if the central bank exits its unconventional policy with a sell-off, the economy experiences a double-dip recession, whereas a slow unwind produces a gradual recovery. Expectations about exit strategies matter for the initial effectiveness of the intervention, and the sell-off exit strategy produces more

effective intervention and a smaller output drop than a slow unwind. In addition, increasing the probability of intervention during crises causes distortions in pre-crisis activity by altering agents' expectations, the magnitude of this distortion depends upon the exit strategy. Finally, the welfare benefits of increasing the probability of intervention can raise or lower welfare, and that the timing of the welfare calculation matters as well as the type of exit strategy used. These differences imply that the central bank may face a time-inconsistency issue in trying to design optimal policy.

Table 1: Markov Switching Parameters

s_t	$\xi_m(s_t)$	$\rho_\xi(s_t)$	$\nu(s_t)$	$\rho_\psi(s_t)$
1) "Normal"	1	ρ_ξ^n	0	ρ_ψ^n
2) "Crisis without Intervention"	ξ_m^c	0	0	ρ_ψ^n
3) "Crisis with Intervention"	ξ_m^c	0	ν^i	0
4) "Post-Crisis with Intervention"	1	ρ_ξ^n	0	1

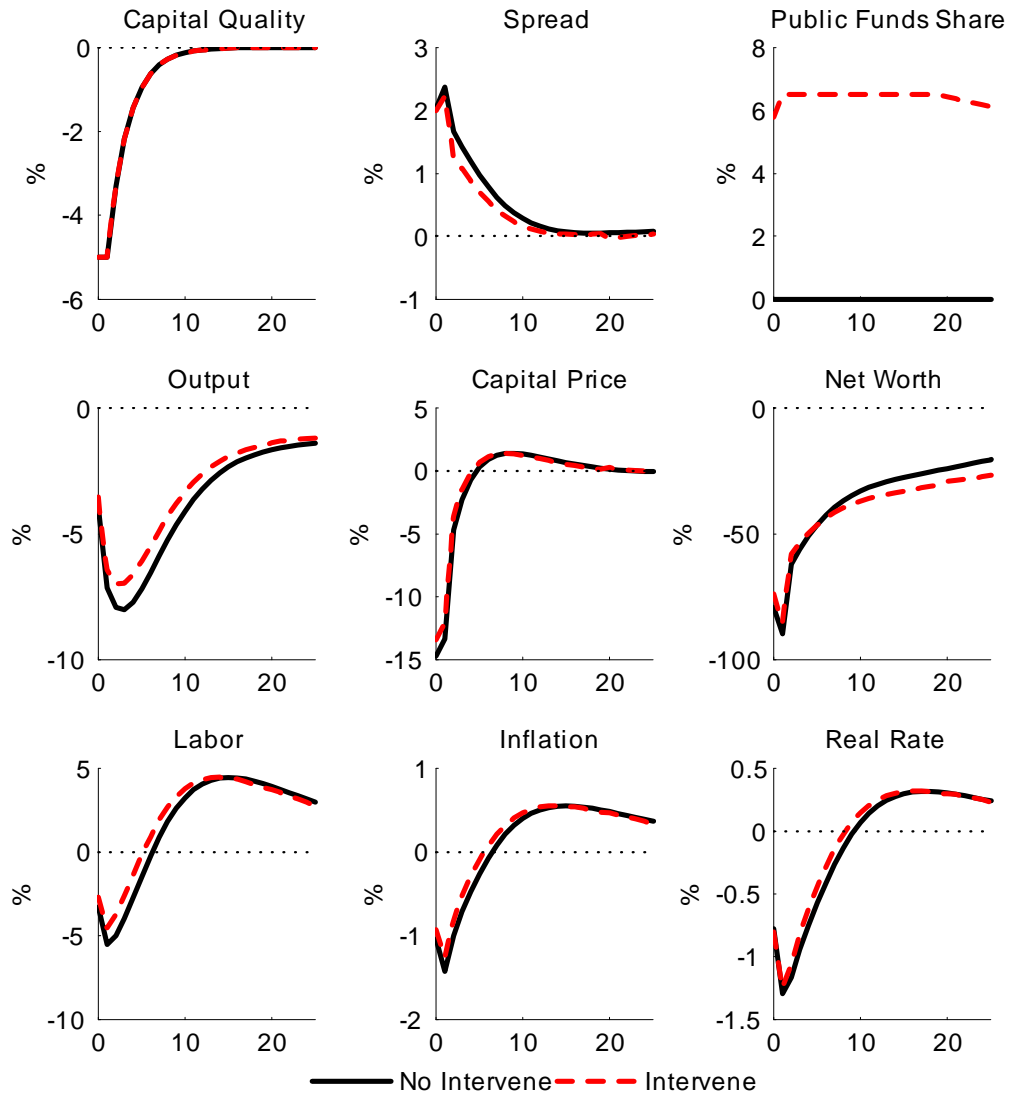
Table 2: Benchmark Parameterization, Preferences and Production

Parameter	Description	Value
β	Discount Factor	0.99
h	Degree of Habit Persistence	0.815
\varkappa	Disutility of Labor	3.409
φ	Inverse Frisch Elasticity of Labor	0.276
λ	Divertable Fraction of Banker Assets	0.381
ω	Transfer to New Bankers	0.002
θ	Survival Rate of Bankers	0.972
α	Capital Share	0.33
U_{ss}	Steady State Capital Utilization	1.00
δ_{ss}	Steady State Depreciation	0.025
ζ	Elasticity of Depreciation to Utilization	7.2
ι	Inverse Elasticity of Net Invest. to Capital Price	1.728
ε	Elasticity of Substitution Between Final Goods	4.167
γ	Probability of No Optimization of Prices	0.779
μ	Degree of Price Indexation	0.241
\bar{g}	Fraction of Steady State Output for Government	0.2
κ_{π}	Response of Interest Rate to Inflation	1.5
κ_y	Response of Interest Rate to Output Gap	0.125
ρ_a	Persistence of TFP Shock	0.95
σ_a	Std Deviation of TFP Shock	0.01
σ_{ξ}	Std Deviation of Capital Quality Shock	0.01
σ_r	Std Deviation of Monetary Policy Shock	0.01

Table 3: Benchmark Parameterization, Key Parameters

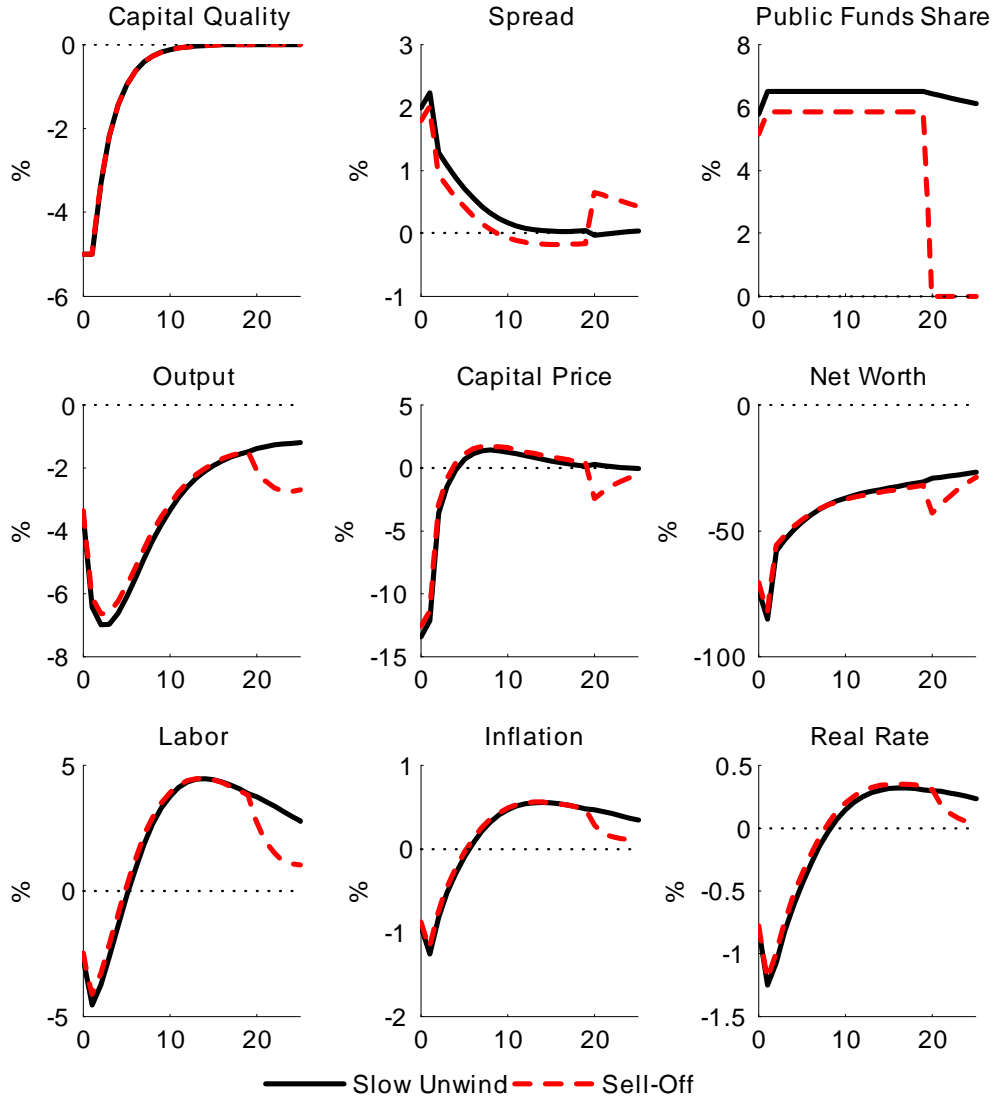
Parameter	Description	Value
p_c	Probability of Crisis	0.005
p_e	Probability of Exiting Crisis	0.5
p_s	Probability of Stopping Intervention	1/18
ν^i	Intervention Response to Spread	3
ξ_m^c	Average Capital Quality in Crisis	0.95
ρ_ξ^n	Capital Quality Persistence in Normal Times	0.66
τ	Resource Cost of Intervention	0

Figure 1: Responses Under a Guarantee, Intervention Versus No Intervention



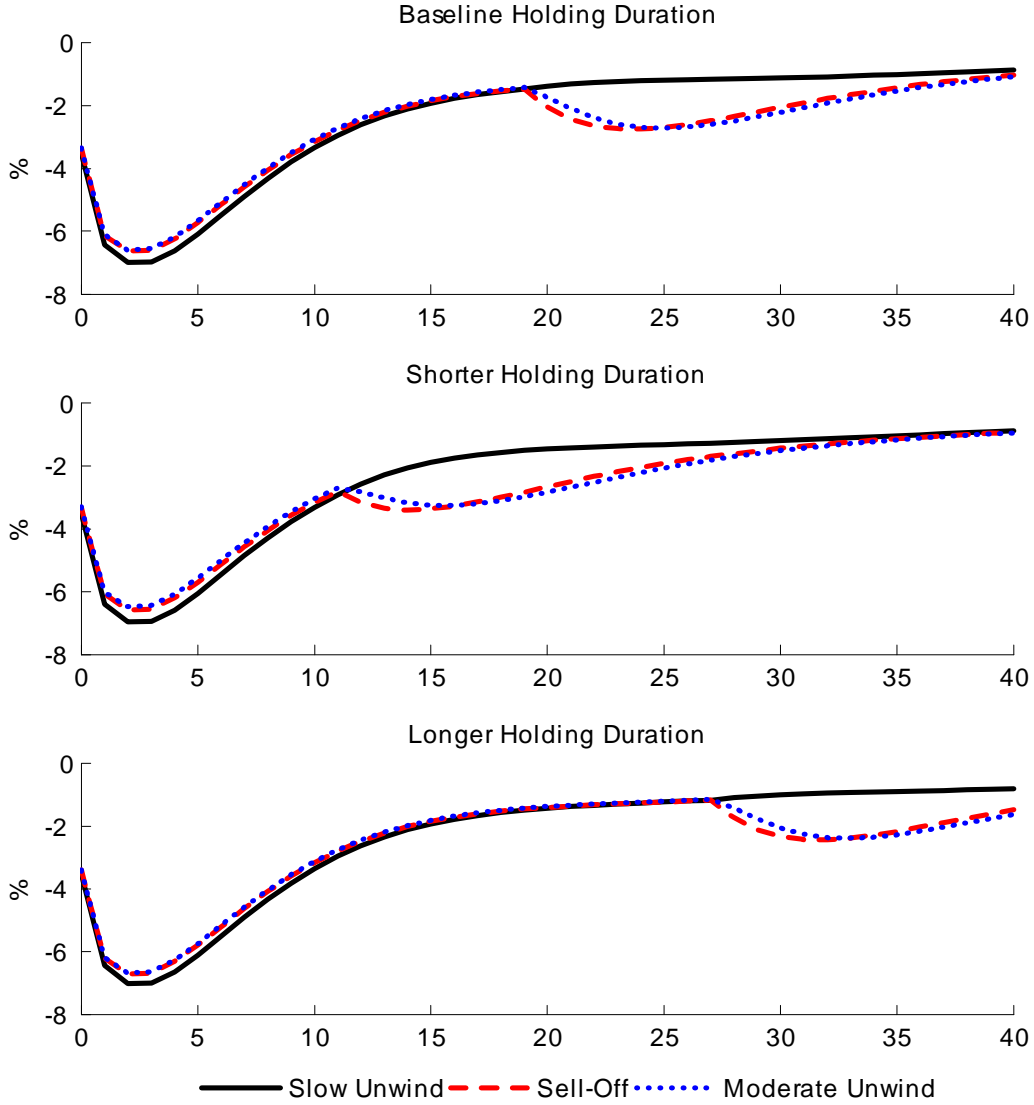
Note: Shows percent deviations from the case where no crisis occurs.

Figure 2: Effects of Expectations, Means



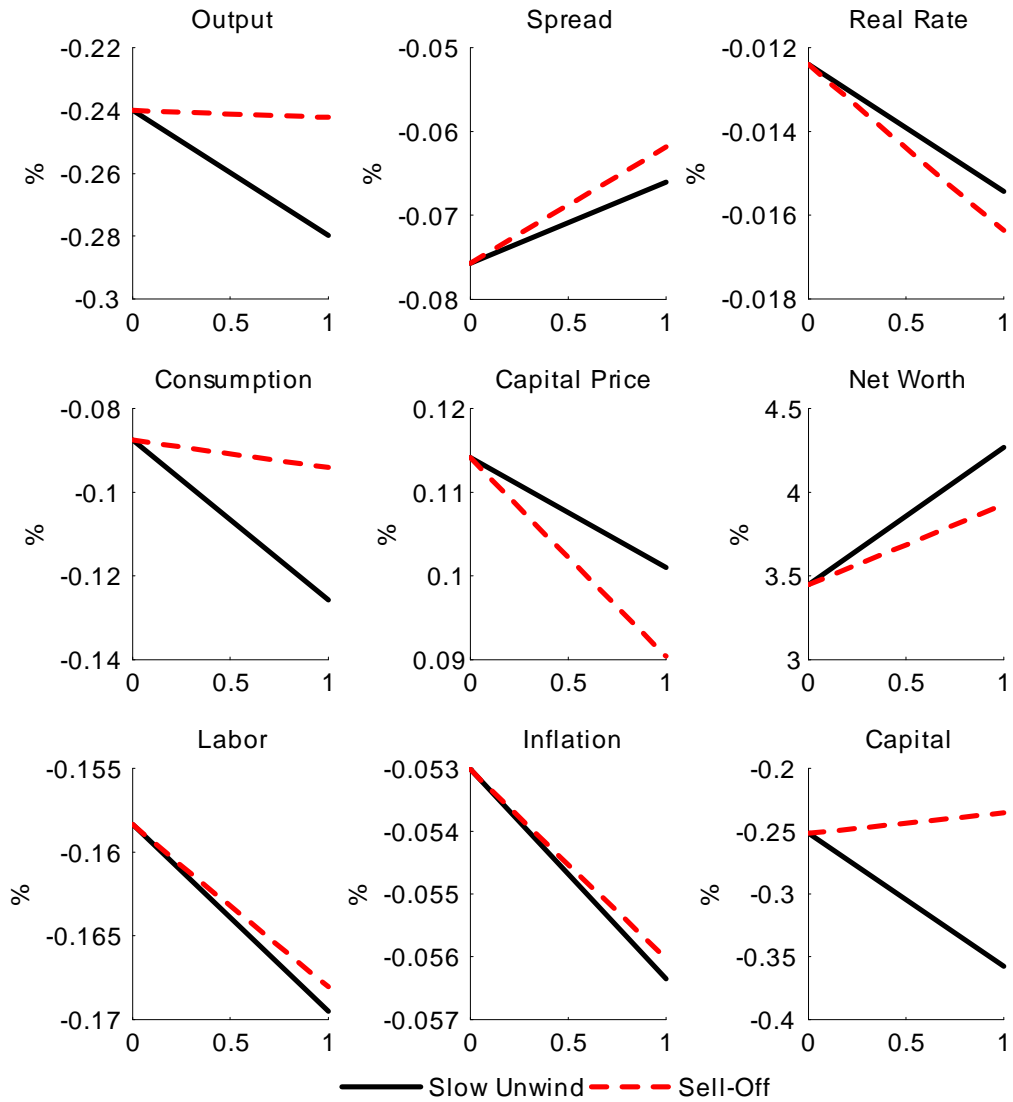
Note: Shows percent deviations from the case where no crisis occurs.

Figure 3: Exit Strategies, Slow Unwind versus Sell-Off



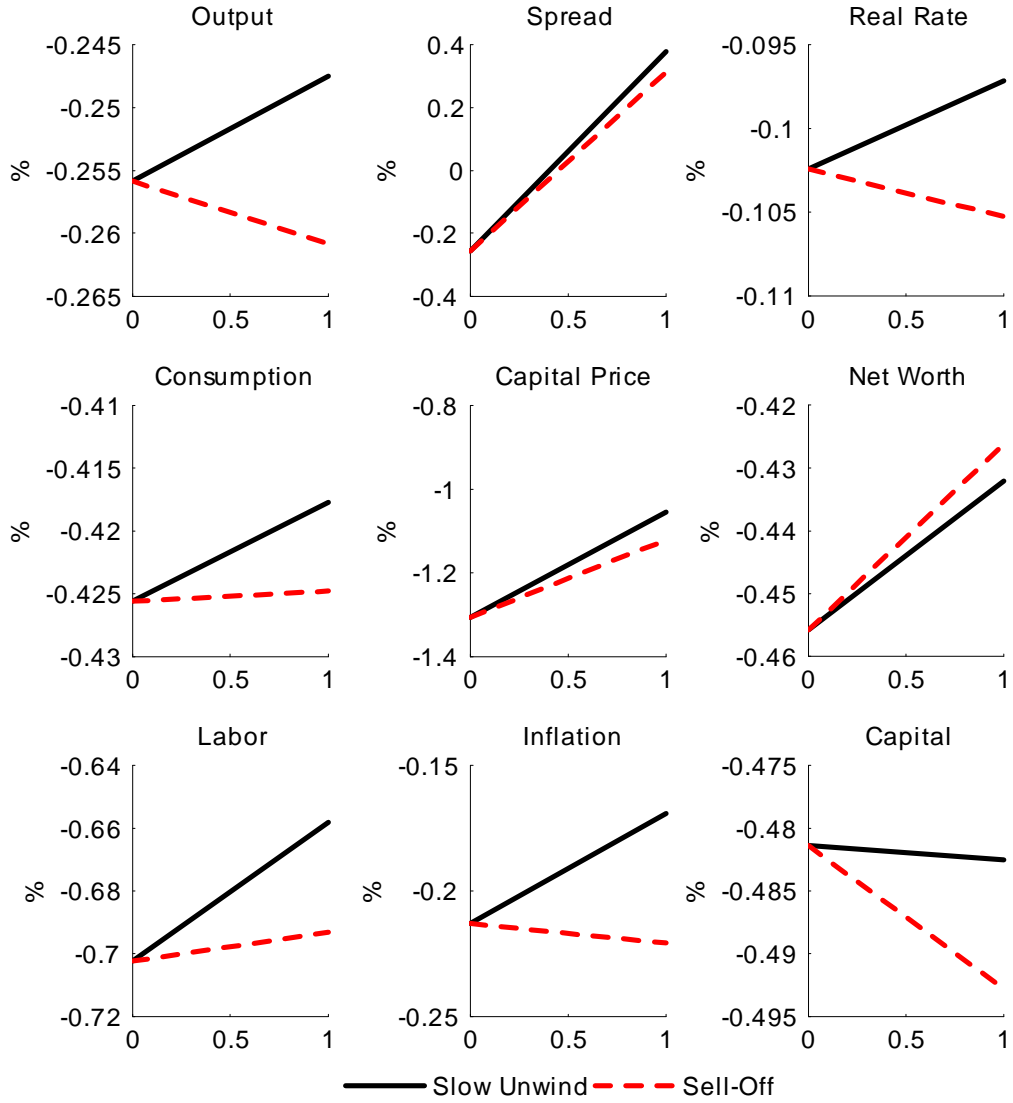
Note: Shows percent deviations from the case where no crisis occurs.

Figure 4: Output Responses to Holding Durations and Exit Strategies



Note: Shows percent deviations from the case where crises never occur.

Figure 5: Effects of Expectations, Std Deviations



Note: Shows percent deviations from the case where crises never occur.

Figure 6: Welfare and the Resource Cost

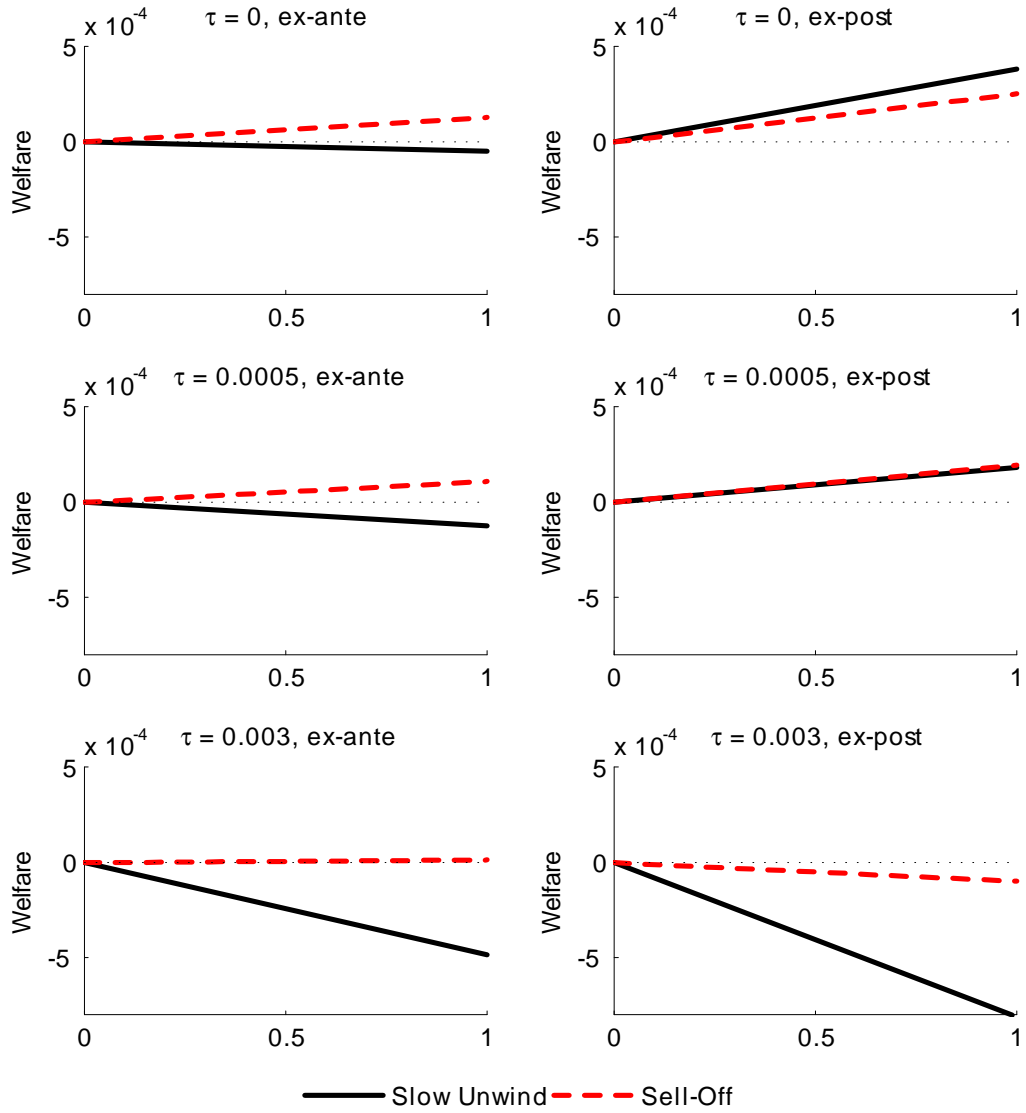


Figure 7: Welfare and Holding Duration, $\tau = 0.0005$

