

NBER WORKING PAPER SERIES

UNCERTAINTY ABOUT GOVERNMENT POLICY AND STOCK PRICES

Lubos Pastor
Pietro Veronesi

Working Paper 16128
<http://www.nber.org/papers/w16128>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2010

Both authors are at the Booth School of Business, University of Chicago, 5807 South Woodlawn Avenue, Chicago, IL 60637, USA. Email: lubos.pastor@ChicagoBooth.edu and pietro.veronesi@ChicagoBooth.edu. We are grateful for comments from Fernando Alvarez, Frederico Belo, Jaroslav Borovika, Pierre Chaigneau, David Chapman, Josh Coval, Max Croce, Steve Davis, Allan Drazen, Francisco Gomes, John Heaton, Monika Piazzesi, Amit Seru, Rob Stambaugh, Milan Švolík, Francesco Trebbi, Rob Vishny, as well as from the conference participants at the 2010 NBER AP meeting, 2010 NBER EFG meeting, 2011 WFA meeting, 2011 SED meeting, 2011 Prague Economic Meeting, 2011 European Summer Symposium in Financial Markets (Gerzensee), 2011 Chicago Booth CFO Forum, and workshop participants at the Bank of England, Chicago, Einaudi Institute (Rome), Erasmus, European Central Bank, Harvard, LBS, LSE, McGill, Northwestern, Notre Dame, Oxford, Vanderbilt, Warwick, and WU Vienna. This research was funded in part by the Initiative on Global Markets at the University of Chicago Booth School of Business. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Lubos Pastor and Pietro Veronesi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Uncertainty about Government Policy and Stock Prices
Lubos Pastor and Pietro Veronesi
NBER Working Paper No. 16128
June 2010, Revised August 2011
JEL No. G01,G12,G14,G18

ABSTRACT

We analyze how changes in government policy affect stock prices. Our general equilibrium model features uncertainty about government policy and a government that has both economic and non-economic motives. The government tends to change its policy after performance downturns in the private sector. Stock prices fall at the announcements of policy changes, on average. The price fall is expected to be large if uncertainty about government policy is large, as well as if the policy change is preceded by a short or shallow downturn. Policy changes increase volatility, risk premia, and correlations among stocks. The jump risk premium associated with policy decisions is positive, on average.

Lubos Pastor
University of Chicago
Booth School of Business
5807 South Woodlawn Ave
Chicago, IL 60637
and NBER
lubos.pastor@chicagobooth.edu

Pietro Veronesi
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
pietro.veronesi@chicagobooth.edu

1. Introduction

Governments shape the environment in which the private sector operates. They affect firms in many ways: they levy taxes, provide subsidies, enforce laws, regulate competition, define environmental policies, etc. In short, governments set the rules of the game.

Governments change these rules from time to time, eliciting price reactions in financial markets. These reactions are weak if the change is widely anticipated, but they can be strong if the markets are caught by surprise. This paper analyzes the effects of changes in government policy on stock prices in a theoretical setting. Naturally, government policy changes are not exogenously given; they are determined by various economic and political forces. To account for this endogeneity, we develop a simple asset pricing model in which the government plays the role of an active decision maker.

We interpret policy changes broadly as government actions that change the economic environment. Recent examples include the health care reform and the ongoing overhaul of financial regulation. Another example is the U.S. government's decision to allow Lehman Brothers to go bankrupt in 2008, which was perceived by many as signaling a shift in the government's implicit too-big-to-fail policy. Numerous other government policy changes also took place during the financial crisis of 2007-2008. Many of those changes were unprecedented, with long-term effects that were difficult to predict in advance.

A key role in our analysis belongs to uncertainty about government policy, which is an inevitable by-product of policymaking. We consider two types of uncertainty. The first type, which we call *policy uncertainty*, relates to the uncertain impact of a given government policy on the profitability of the private sector. The second type, which we call *political uncertainty*, captures the private sector's uncertainty whether the current government policy will change. In other words, there is uncertainty about what the government is going to do, as well as what the effect of its action is going to be. We find that both types of uncertainty affect stock prices in important ways.

We develop a general equilibrium model in which firm profitability follows a stochastic process whose mean is affected by the prevailing government policy. The policy's impact on the mean is uncertain. Both the government and the investors (firm owners) learn about this impact in a Bayesian fashion by observing realized profitability. All agents have the same prior beliefs about the current policy's impact, and the same beliefs apply to any other future government policy. The prior standard deviation is labeled "policy uncertainty."

At a given point in time, the government decides whether or not to change its policy. If a policy change occurs, the agents' beliefs are reset: the posterior beliefs about the old policy's impact are replaced by the prior beliefs about the new policy's impact. When making its policy decision, the government is motivated by both economic and non-economic objectives: it maximizes the investors' welfare, as a social planner would, but it also takes into account the political cost (or benefit) incurred by changing the policy. This cost is unknown to the investors, who therefore cannot fully anticipate whether the policy change will occur. The standard deviation of the political cost is labeled "political uncertainty."

We find that it is optimal for the government to replace its policy by a new one if the old policy's impact on profitability is perceived as sufficiently unfavorable; i.e., if the posterior mean of the impact is below a given threshold. This threshold decreases with policy uncertainty as well as with the political cost. If the government derives an unexpectedly large political benefit from changing its policy, the policy will be replaced even if it worked well in the past. In expectation, however, the threshold is below the prior mean of the policy's impact. To push the posterior mean below the prior mean, Bayes' rule dictates that the signal—the realized profitability in the private sector—must be lower than expected. As a result, policy changes are expected to occur after periods of unexpectedly low realized profitability, which we refer to as "downturns."

We derive the conditions under which stock prices fall at the announcement of a policy change. There are two effects. First, a policy change typically increases the firms' expected future profitability, as a direct result of the government's optimal decision rule described in the previous paragraph. This cash flow effect generally pushes stock prices up. Second, a policy change increases the discount rates because the new policy's impact on profitability is more uncertain; adopting a new policy essentially undoes the gains from learning about the old policy. This discount rate effect pushes stock prices down. We find that the discount rate effect is stronger than the cash flow effect unless the old policy's impact on profitability is perceived as sufficiently negative. That is, stock prices fall at the announcement of a policy change unless the posterior mean of the old policy's impact is below a negative threshold. This threshold is generally below the policy-change-triggering threshold discussed in the previous paragraph. Stock prices rise if the government replaces an old policy so unproductive that its posterior mean falls below both thresholds.

We find that, on average, stock prices fall at the announcement of a policy change. Positive announcement returns are typically small because they tend to occur in states of the world in which the policy change is largely anticipated by investors. Given the government's

economic motive, any policy change that lifts stock prices is mostly expected, so that much of its effect is priced in before the announcement. In contrast, negative announcement returns tend to be larger because they occur when the announcement of a policy change contains a bigger element of surprise. The probability distribution of the announcement returns is left-skewed and, most interesting, its mean is below zero. We prove analytically that the expected value of the stock return at the announcement of a policy change is negative.

We also show numerically that this expected announcement return is more negative when there is more uncertainty about government policy. When policy uncertainty is larger, so is the risk associated with a new policy, and so is the discount rate effect that pushes stock prices down when the new policy is announced. When political uncertainty is larger, so is the element of surprise in the announcement of a policy change.

We relate the stock return at the announcement of a policy change to the length and depth of the preceding downturn. We find that the announcement returns are negative especially after downturns that are short or shallow. The announcement returns can also be positive, mostly after downturns that are long or deep. However, such positive returns tend to be small because after long or deep downturns, policy changes are largely anticipated.

Before the announcement of a policy decision, investors are uncertain whether the policy will change. If it does change, stock prices tend to jump down; if it does not change, prices tend to jump up. The expected jump in stock prices at the announcement of a policy decision is generally nonzero. This expected jump captures the risk premium demanded by investors for facing an uncertain jump in the stochastic discount factor when the policy decision is announced. The conditional jump risk premium can be positive or negative, depending on the posterior mean of the policy impact. The unconditional premium is positive and increasing in both policy uncertainty and political uncertainty.

The volatilities and correlations of stock returns are also affected by changes in government policy. By introducing new policies whose impact is more uncertain, policy changes increase the volatility of the stochastic discount factor. As a result, risk premia go up and stock returns become more volatile and more highly correlated across firms. All of these effects are stronger when policy uncertainty is larger. Before a policy change, stock returns are affected by fluctuations in the investors' beliefs about the current policy as well as in the investors' assessment of the probability of a policy change. The latter fluctuations stop after the government makes its policy decision, typically resulting in lower volatilities and correlations if the decision does not change the prevailing policy.

The government's ability to change its policy has a substantial effect on stock prices. We compare the model-implied stock prices with their counterparts in a hypothetical scenario in which policy changes are precluded. We find that the government's ability to change its policy amplifies the stock price declines around policy changes. In addition, this ability can imply a higher or lower level of stock prices compared to the hypothetical scenario.

In the benchmark version of our model, discussed above, the government can change its policy only at a predetermined point in time. In the first extension of the model, we endogenize the timing of the policy change, by letting the government solve for the optimal time to change the policy. We lose our closed-form solutions, but we find numerically that the key results from our benchmark model continue to hold.

In the benchmark model, the firms' investment decisions are not modeled explicitly. In the second extension of the model, we allow firms to disinvest in response to changes in government policy. In this extension, the firms' investment decisions and the government's policy decision are made simultaneously: the government takes into account the firms' anticipated response, and each firm considers the decisions of the other firms as well as the government. In the resulting Nash equilibrium, a fraction of firms remain invested while the rest hold cash. We show numerically that firms respond to government-induced uncertainty by cutting their investment. We also show that our key results continue to hold.

In the third and final extension of our model, we introduce heterogeneity in the firms' exposures to government policy. We find that firms with more exposure typically have higher expected returns, though the premium for government policy risk is state-dependent and can be negative. The main results from our benchmark model continue to hold.

Prior studies have analyzed the effects of uncertainty, broadly defined, on various aspects of economic activity. For example, it is well known that uncertainty generally reduces firm investment when this investment is at least partially irreversible.¹ The impact of uncertainty about government policy on investment has been analyzed both theoretically and empirically.² The literature has also analyzed the effects of uncertainty about government policy on inflation, capital flows, and welfare.³ However, the literature does not have much to say

¹See, for example, Bernanke (1983), McDonald and Siegel (1986), Pindyck (1988), and Dixit (1989). Bloom (2009) provides a structural analysis of various real effects of macroeconomic uncertainty shocks.

²For example, Rodrik (1991) shows that even moderate amount of policy uncertainty can impose a hefty tax on investment. Hassett and Metcalf (1999) find that the impact of tax policy uncertainty on investment depends on the process followed by the tax policy. Julio and Yook (2011) and Yonce (2009) find that firms reduce their investment in years leading up to major elections. Durnev (2011) finds that corporate investment is less sensitive to stock prices during election years.

³For example, Drazen and Helpman (1990) study how uncertainty about a future fiscal adjustment affects

about how this uncertainty affects asset prices, which is the subject of our study. A rare exception is Sialm (2006), who analyzes the effects of stochastic taxes on asset prices.⁴ Tax uncertainty also features in Croce, Kung, Nguyen, and Schmid (2011), who explore its asset pricing implications in a production economy with recursive preferences. Finally, Ulrich (2011) analyzes how bond yields are affected by Knightian uncertainty about both Ricardian equivalence and the size of the government multiplier. All of these studies are quite different from ours. They analyze fiscal policy, whereas we consider a broader set of government actions. We focus on the stock market reactions to government policy announcements, unlike any of these studies. These studies use very different modeling techniques, and they do not model the government’s policy decision explicitly as we do. None of these studies feature Bayesian learning, which plays an important role here.

Our model is also different from the learning models that were recently proposed in the political economy literature, such as Callander (2008) and Strulovici (2010). In Callander’s model, voters learn about the effects of government policies through repeated elections. In Strulovici’s model, voters learn about their preferences through policy experimentation. Neither study analyzes the asset pricing implications of learning.

The paper is organized as follows. Section 2. presents the benchmark model. Section 3. analyzes the equilibrium stock prices. Section 4. extends the model by endogenizing the time of the policy change. Section 5. extends the model by allowing firms to disinvest in response to government-induced uncertainty. Section 6. introduces heterogeneity in the firms’ exposures to government policy. Section 7. concludes. The Appendix contains technical details and a reference to the Technical Appendix, which contains all the proofs.

2. The Model

We consider an economy with a finite horizon $[0, T]$ and a continuum of firms $i \in [0, 1]$. Let B_t^i denote firm i ’s capital at time t . Firms are financed entirely by equity, so B_t^i can also

the dynamics of inflation. Hermes and Lensink (2001) show that uncertainty about budget deficits, tax payments, government consumption, and inflation is positively related to capital outflows at the country level. Gomes, Kotlikoff, and Viceira (2008) calibrate a life-cycle model to measure the welfare losses resulting from uncertainty about government policies regarding taxes and Social Security. They find that policy uncertainty materially affects the agents’ consumption, saving, labor supply, and portfolio decisions.

⁴Other studies, such as McGrattan and Prescott (2005), Sialm (2009), and Gomes, Michaelides, and Polkovnichenko (2009), relate stock prices to tax rates, without emphasizing tax-related uncertainty. In empirical work, Erb, Harvey, and Viskanta (1996) relate political risk to future country stock returns, and Boutchkova, Doshi, Durnev, and Molchanov (2010) relate political uncertainty to stock volatility.

be viewed as book value of equity. At time 0, all firms employ an equal amount of capital, which we normalize to $B_0^i = 1$. Firm i 's capital is invested in a linear technology whose rate of return is stochastic and denoted by $d\Pi_t^i$. All profits are reinvested, so that firm i 's capital evolves according to $dB_t^i = B_t^i d\Pi_t^i$. Since $d\Pi_t^i$ equals profits over book value, we refer to it as the profitability of firm i . For all $t \in [0, T]$, profitability follows the process

$$d\Pi_t^i = (\mu + g_t) dt + \sigma dZ_t + \sigma_1 dZ_t^i, \quad (1)$$

where (μ, σ, σ_1) are observable constants, Z_t is a Brownian motion, and Z_t^i is an independent Brownian motion that is specific to firm i . The variable g_t denotes the impact of the prevailing government policy on the mean of the profitability process of each firm.⁵ When $g_t = 0$, the government policy is “neutral” in that it has no impact on profitability.

The government policy's impact, g_t , is constant while the same policy is in effect. At an exogenously given time τ , $0 < \tau < T$, the government makes an irreversible decision whether or not to change its policy.⁶ As a result, g_t is a simple step function of time:

$$g_t = \begin{cases} g^{\text{old}} & \text{for } t \leq \tau \\ g^{\text{old}} & \text{for } t > \tau \text{ if there is no policy change} \\ g^{\text{new}} & \text{for } t > \tau \text{ if there is a policy change,} \end{cases} \quad (2)$$

where g^{old} denotes the impact of the government policy prevailing at time 0. A policy change replaces g^{old} by g^{new} , thereby inducing a permanent shift in average profitability. A policy decision becomes effective immediately after its announcement at time τ .

The value of g_t is unknown. This key assumption captures the idea that government policies have an uncertain impact on firm profitability. The prior distributions of both g^{old} and g^{new} at time 0 are normal with mean zero and known variance σ_g^2 :

$$g \sim N(0, \sigma_g^2), \quad (3)$$

for both $g \equiv g^{\text{old}}$ and $g \equiv g^{\text{new}}$. The prior is the same for both policies, the one prevailing at time 0 and the one that might replace it at time τ . Both policies are expected to be neutral a priori.⁷ We refer to σ_g as *policy uncertainty*. The value of g_t is unknown for all $t \in [0, T]$ to all agents—the government as well as the investors who own the firms.

⁵The simplifying assumption that the government policy affects all firms equally is relaxed in Section 6.

⁶The simplifying assumption that τ is exogenous allows us to obtain analytical results. Section 4. shows numerically that all of our key results survive when τ is endogenous, i.e., when the government chooses the optimal time to change its policy.

⁷Any government effect on profitability that is stable across policy changes is included in μ .

The firms are owned by a continuum of identical investors who maximize expected utility derived from terminal wealth. For all $j \in [0, 1]$, investor j 's utility function is given by

$$u(W_T^j) = \frac{(W_T^j)^{1-\gamma}}{1-\gamma}, \quad (4)$$

where W_T^j is investor j 's wealth at time T and $\gamma > 1$ is the coefficient of relative risk aversion. At time 0, all investors are equally endowed with shares of firm stock. Stocks pay liquidating dividends at time T .⁸ Investors observe whether a policy change takes place at time τ .

When making its policy decision at time τ , the government maximizes the same objective function as the investors, except that it also faces a nonpecuniary cost (or benefit) associated with a policy change. The government changes its policy if the expected utility under the new policy is higher than it is under the old policy. Specifically, the government solves

$$\max \left\{ E_\tau \left[\frac{W_T^{1-\gamma}}{1-\gamma} \mid \text{no policy change} \right], E_\tau \left[\frac{CW_T^{1-\gamma}}{1-\gamma} \mid \text{policy change} \right] \right\}, \quad (5)$$

where $W_T = B_T = \int_0^1 B_T^i di$ is the final value of aggregate capital and C is the “political cost” incurred by the government if a new policy is introduced. Values of $C > 1$ represent a cost (e.g., the government must exert effort or burn political capital to implement a new policy), whereas $C < 1$ represents a benefit (e.g., the government makes a transfer to a favored constituency, or it simply wants to be seen doing something).⁹ The value of C is randomly drawn at time τ from a lognormal distribution centered at $C = 1$:

$$c \equiv \log(C) \sim N \left(-\frac{1}{2}\sigma_c^2, \sigma_c^2 \right), \quad (6)$$

where C is independent of the Brownian motions in equation (1). As soon as the value of C is revealed to the agents at time τ , the government uses it to make the policy decision. We refer to σ_c as *political uncertainty*. Political uncertainty introduces an element of surprise into policy changes, resulting in stock price reactions at time τ .

Given its objective function in equation (5), the government is “quasi-benevolent”: it is expected to maximize the investors’ welfare (because $E[C] = 1$), but also to deviate from this objective in a random fashion. The assumption that governments do not behave as fully benevolent social planners is widely accepted in the political economy literature.¹⁰ This

⁸No dividends are paid before time T because the investors’ preferences (equation (4)) do not involve intermediate consumption. Firms in our model reinvest all of their earnings, as mentioned earlier.

⁹We refer to C as a cost because higher values of C translate into lower utility (since $W_T^{1-\gamma}/(1-\gamma) < 0$). The assumption that this cost is nonpecuniary simplifies the analysis. Pecuniary costs would eventually be passed on to the investors, and modeling this transmission would be distracting given our objectives.

¹⁰Drazen (2000) provides a useful overview of this literature.

literature presents various reasons why governments might not maximize aggregate welfare. For example, governments often redistribute wealth.¹¹ Governments tend to be influenced by special interest groups.¹² They might also be susceptible to corruption.¹³ Instead of modeling these political forces explicitly, we adopt a simple reduced-form approach to capturing departures from benevolence. In our model, all aspects of politics—redistribution, corruption, special interests, etc.—are bundled together in the political cost C . The randomness of C reflects the difficulty in predicting the outcome of a political process, which can be complex and non-transparent. For example, it can be hard to predict the outcome of a battle between special interest groups. By modeling politics in such a reduced-form fashion, we are able to focus on the asset pricing implications of government policy changes.

Our modeling of government policies also merits some discussion. We impose the same prior distribution in equation (3) on all policies. In reality, the priors might differ across policies. For example, if a potential new policy has already been tried in the past, the agents may believe they have more prior information about that policy’s impact, so they might assign a lower prior variance to it. The lower variance might not necessarily be warranted, though, because the same policy implemented at a different time in a different environment may have a different impact on profitability. Our assumption that all policies are perceived by agents as identical *ex ante* seems like a natural starting point in an asset pricing analysis. The simplicity of our setup allows us to derive interesting analytical results.

We interpret policy changes as government actions that change the economic landscape, such as major reforms. Political uncertainty, σ_c , is the uncertainty about whether a reform will take place. Policy uncertainty, σ_g , is the uncertainty about the effect of the reform on profitability. We show below that both σ_c and σ_g affect stock prices in interesting ways.

2.1. Learning

The value of g_t is unknown to all agents, investors and the government alike. At time 0, all agents share the same prior beliefs about g_t , summarized by the distribution in equation (3). All agents learn about g_t in the same Bayesian fashion by observing the realized profitabilities of all firms. The learning process is described in the following proposition.

¹¹Redistribution is a major theme in political economy. Prominent studies of redistribution include Alesina and Rodrik (1994) and Persson and Tabellini (1994), among others. Our model is not well suited for analyzing redistribution effects because all of our investors are identical *ex ante*, for simplicity.

¹²See, for example, Grossman and Helpman (1994) and Coate and Morris (1995).

¹³See, for example, Shleifer and Vishny (1993) and Rose-Ackerman (1999).

Proposition 1. *Observing the continuum of signals $d\Pi_t^i$ in equation (1) across all firms $i \in [0, 1]$ is equivalent to observing a single aggregate signal about g_t :*

$$ds_t = (\mu + g_t) dt + \sigma dZ_t. \quad (7)$$

Under the prior in equation (3), the posterior for g_t at any time $t \in [0, T]$ is given by

$$g_t \sim N(\hat{g}_t, \hat{\sigma}_t^2). \quad (8)$$

For all $t \leq \tau$, the mean and the variance of this posterior distribution evolve as

$$d\hat{g}_t = \hat{\sigma}_t^2 \sigma^{-1} d\hat{Z}_t \quad (9)$$

$$\hat{\sigma}_t^2 = \frac{1}{\frac{1}{\sigma_g^2} + \frac{1}{\sigma^2} t}, \quad (10)$$

where the “expectation error” $d\hat{Z}_t$ is given by $d\hat{Z}_t = (ds_t - E_t(ds_t)) / \sigma$ for all $t \in [0, T]$.

If there is no policy change at time τ , then the processes (9) and (10) hold also for $t > \tau$.

If there is a policy change at τ , then \hat{g}_t jumps from \hat{g}_τ to zero right after the policy change, and for $t > \tau$, \hat{g}_t follows the process in equation (9). In addition, for $t > \tau$, $\hat{\sigma}_t$ follows

$$\hat{\sigma}_t^2 = \frac{1}{\frac{1}{\sigma_g^2} + \frac{1}{\sigma^2} (t - \tau)}. \quad (11)$$

Comparing equations (1) and (7), idiosyncratic shocks dZ_t^i wash out upon aggregating infinitely many independent signals. The aggregate signal in equation (7) is the average profitability across firms. When this signal is higher than expected, the agents revise their beliefs about g_t upward, and vice versa (see equation (9)). Uncertainty about g_t declines deterministically over time due to learning (see equations (10) and (11)), except for a discrete jump up at time τ in case of a policy change. A policy change resets the agents’ beliefs about g_t from the posterior $N(\hat{g}_\tau, \hat{\sigma}_\tau^2)$ to the prior $N(0, \sigma_g^2)$, where $\hat{\sigma}_\tau < \sigma_g$ due to learning between times 0 and τ . Before time τ , the agents learn about g^{old} ; after τ , they learn about g^{old} or g^{new} , depending on whether a policy change occurs (see equation (2)).

2.2. Optimal Changes in Government Policy

After a period of learning about g^{old} , the government decides whether or not to change its policy at time τ . If the change occurs, the value of g_t changes from g^{old} to g^{new} and the perceived distribution of g_t changes from the posterior in equation (8) to the prior in equation (3). According to equation (5), the government changes its policy if and only if

$$E_\tau \left[\frac{CB_T^{1-\gamma}}{1-\gamma} \mid \text{policy change} \right] > E_\tau \left[\frac{B_T^{1-\gamma}}{1-\gamma} \mid \text{no policy change} \right]. \quad (12)$$

Since government policy affects future profitability, the two expectations in equation (12) are computed under different stochastic processes for the aggregate capital $B_t = \int_0^1 B_t^i di$. We show in the Technical Appendix that the aggregate capital at time T is given by

$$B_T = B_\tau e^{(\mu+g-\frac{\sigma^2}{2})(T-\tau)+\sigma(Z_T-Z_\tau)}, \quad (13)$$

where $g \equiv g^{\text{old}}$ if there is no policy change and $g \equiv g^{\text{new}}$ if there is one. Evaluating the expectations in equation (12) based on equation (13) yields the following proposition.

Proposition 2. *The government changes its policy at time τ if and only if*

$$\widehat{g}_\tau < \underline{g}(c), \quad (14)$$

where

$$\underline{g}(c) = -\frac{(\sigma_g^2 - \widehat{\sigma}_\tau^2)(\gamma - 1)(T - \tau)}{2} - \frac{c}{(T - \tau)(\gamma - 1)}. \quad (15)$$

The government follows a simple cutoff rule: it changes its policy if the posterior mean of the old policy's impact, \widehat{g}_τ , is below a given threshold. That is, a policy is replaced if its effect on firm profitability is perceived as sufficiently unfavorable. The two terms on the right-hand side of equation (15) reflect the government's economic and political motives. The first term reflects the increase in risk associated with adopting a new policy ($\sigma_g > \widehat{\sigma}_\tau$). Since $\gamma > 1$, higher policy uncertainty σ_g reduces $\underline{g}(c)$, making a policy change less likely. The second term reflects the political cost or benefit incurred by the government if the new policy is adopted. If $c > 0$, the government incurs a cost, the second component is negative, and the new policy is less likely to be adopted. If $c < 0$, the government benefits from a policy change and the new policy is more likely to be adopted. Since the investors do not know c , they cannot fully anticipate whether a policy change will occur. The only exception occurs in the special case of $\sigma_c = 0$; in that case, the investors know the threshold $\underline{g}(c)$, so they can fully predict the outcome of the government's policy decision.

The investors expect c to be close to zero, so their expectation of $\underline{g}(c)$ is close to $\underline{g}(0)$.¹⁴ Since $\sigma_g > \widehat{\sigma}_\tau$, equation (15) implies $\underline{g}(0) < 0$. That is, in expectation, a policy is replaced if its impact is perceived as sufficiently negative. It is not enough for \widehat{g}_τ to be negative; it must be sufficiently negative for the expected gain from a policy change to outweigh the higher risk associated with a new policy.¹⁵ For the posterior mean of g^{old} to be negative at

¹⁴Recall from equation (6) that $E(C) = 1$ implies $E(c) = -\sigma_c^2/2$ rather than zero for $c = \log(C)$.

¹⁵The political economy literature recognizes that uncertainty associated with a policy change may lead to a bias toward status quo, but the literature focuses mostly on the uncertainty about how the gains and losses from reform will be distributed across individuals (e.g., Fernandez and Rodrik (1991)). Our investors are homogeneous; we focus on uncertainty about the aggregate effects of a policy change.

time τ while the prior mean is zero, Bayes' rule implies that the signal received by investors before time τ —realized profitability—must be unexpectedly low. Therefore, Proposition 2 implies that policy changes are expected to occur after periods of unexpectedly low realized profitability.¹⁶ We refer to such periods as “downturns.”

To illustrate this implication, we plot the expected dynamics of profitability conditional on a policy change. We choose a set of plausible parameter values, which are summarized in Table 1. We simulate many samples of shocks in our economy and record the paths of \hat{g}_t and realized profitability in each sample. Realized profitability is the average profitability across all firms, reported in excess of μ so that its unconditional mean is zero. We split the samples into two groups, depending on whether the government changes its policy at time τ .¹⁷ We then plot the average paths of \hat{g}_t and realized profitability across all samples within both groups. Panel A of Figure 1 plots \hat{g}_t (solid line) and profitability (dashed line) conditional on a policy change at time τ . Panel B plots the same quantities conditional on no policy change. Both panels also plot the average value of the cutoff $\underline{g}(c)$ (dotted line).

Panel A of Figure 1 shows that policy changes tend to be preceded by periods of low realized profitability. Between times 0 and $\tau = 10$ years, profitability averages -2.5% per year. This negative profitability is due to ex-post conditioning on $\hat{g}_\tau < \underline{g}(c)$. The average value of $\underline{g}(c)$ is -0.5% per year, and the average value of \hat{g}_t gradually falls from 0 to -1.5% per year between times 0 and τ . For the posterior mean \hat{g}_t to decline in this manner, realized profitability must be below \hat{g}_t , as discussed earlier. When the policy changes at time τ , both \hat{g}_t and profitability jump up to zero, on average, and stay there until time $T = 20$ since there is no more ex-post conditioning after time τ .

Panel B of Figure 1 shows the opposite but milder patterns when there is no policy change. Due to ex-post conditioning on $\hat{g}_\tau > \underline{g}(c)$, the average value of \hat{g}_t rises from 0 to 1% per year at time τ , and average realized profitability is 1.5% until time τ . After time τ , the average \hat{g}_t remains unchanged at 1% because the policy remains the same. Profitability therefore jumps from 1.5% to 1% at time τ , and it stays there until time T .

Figure 2 shows how the expected dynamics of profitability around policy changes depend

¹⁶Similar results have been obtained in the political economy literature by using different mechanisms. Rodrik (1996, p.27) writes: “Reform naturally becomes an issue only when current policies are perceived to be not working.” According to Drazen (2000, p.449), “it is striking how little formal empirical testing there has been of the view that a crisis is necessary for significant policy change.” An early exception is Bruno and Easterly (1996), who find that inflation crises tend to be followed by reforms. Drazen and Easterly (2001) and Alesina, Ardagna, and Trebbi (2006) also find evidence supporting the hypothesis that crises induce reforms, although the former study finds this evidence only for a subset of the crisis indicators.

¹⁷A policy change occurs in about 38% of all simulated samples.

on policy uncertainty and political uncertainty. Each of the four panels is analogous to Panel A of Figure 1. Whereas Figure 1 is constructed using the values from Table 1, $\sigma_g = 2\%$ and $\sigma_c = 10\%$, Figure 2 uses the values of $\sigma_g = 1\%$ and 3% per year, and $\sigma_c = 0$ and 20% .

Figure 2 shows that the downturns preceding policy changes tend to be worse when σ_g is larger. Larger values of σ_g imply more negative expected values of $\underline{g}(c)$, which imply steeper declines in \hat{g}_t conditional on $\hat{g}_\tau < \underline{g}(c)$. To induce steeper declines in \hat{g}_t , realized profitability must be more negative. For $\sigma_g = 3\%$, realized profitability before time τ is as low as -4% per year. The effect of σ_c is weaker. Changing σ_c from 0 to 20% has a negligible effect on the average value of $\underline{g}(c)$, which is close to $\underline{g}(0)$. The main effect of σ_c is that when σ_c is higher, large negative values of c are more likely, so a policy change is more likely to occur even if \hat{g}_t is high. For example, changing σ_c from 0 to 20% increases the likelihood of a policy change from 46% to 50% when $\sigma_g = 1\%$, and from 30% to 31% when $\sigma_g = 3\%$. As a result, the pre- τ declines in \hat{g}_t and profitability are less steep, on average.

Having established the government's optimal policy rule and its key implications for profitability, we are ready to analyze the asset pricing implications. The following section presents our main results on the stock price dynamics around policy changes.

3. Stock Prices

Firm i 's stock represents a claim on the firm's liquidating dividend at time T , which is equal to B_T^i . The investors' total wealth at time T is equal to $B_T = \int_0^1 B_T^i di$. Stock prices adjust to make the investors hold all of the firms' stock. In addition to stocks, there is also a zero coupon bond in zero net supply, which makes a unit payoff at time T with certainty. We use this risk-free bond as the numeraire.¹⁸ Under the assumption of market completeness, standard arguments imply that the state price density is uniquely given by

$$\pi_t = \frac{1}{\lambda} E_t [B_T^{-\gamma}], \quad (16)$$

where λ is the Lagrange multiplier from the utility maximization problem of the representative investor. The market value of stock i is given by the standard pricing formula

$$M_t^i = E_t \left[\frac{\pi_T}{\pi_t} B_T^i \right]. \quad (17)$$

¹⁸This assumption is equivalent to assuming a risk-free rate of zero. Such an assumption is innocuous because without intermediate consumption, there is no intertemporal consumption choice that would pin down the interest rate. This modeling choice ensures that interest rate fluctuations do not drive our results.

3.1. Stock Price Reaction to the Announcement of a Policy Change

When the government announces its policy decision at time τ , stock prices jump. The direction and size of the jump depend on whether the government decides to change or maintain its policy, as well as on the extent to which this decision is unexpected. We derive a closed-form solution for each firm's "announcement return," defined as the instantaneous stock return at time τ conditional on the announcement of a *change* in government policy.¹⁹

Proposition 3. *Each firm's stock return at the announcement of a policy change is given by*

$$R(\hat{g}_\tau) = \frac{(1 - p(\hat{g}_\tau)) F(\hat{g}_\tau) (1 - G(\hat{g}_\tau))}{p(\hat{g}_\tau) + (1 - p(\hat{g}_\tau)) F(\hat{g}_\tau) G(\hat{g}_\tau)}, \quad (18)$$

where

$$F(\hat{g}_\tau) = e^{-\gamma \hat{g}_\tau (T-\tau) - \frac{1}{2} \gamma^2 (T-\tau)^2 (\sigma_g^2 - \hat{\sigma}_\tau^2)} \quad (19)$$

$$G(\hat{g}_\tau) = e^{\hat{g}_\tau (T-\tau) - \frac{1}{2} (1-2\gamma) (T-\tau)^2 (\sigma_g^2 - \hat{\sigma}_\tau^2)} \quad (20)$$

$$p(\hat{g}_\tau) = N\left(\hat{g}_\tau (1 - \gamma) (T - \tau) - \frac{(1 - \gamma)^2}{2} (T - \tau)^2 (\sigma_g^2 - \hat{\sigma}_\tau^2); -\frac{\sigma_c^2}{2}, \sigma_c^2\right), \quad (21)$$

and $N(x; a, b)$ denotes the c.d.f. of a normal distribution with mean a and variance b .

The function in equation (21) is the probability of a policy change perceived by the investors just before time τ . According to Proposition 2, $p(\hat{g}_\tau) = \text{Prob}(\hat{g}_\tau < \underline{g}(c))$. Since the investors observe \hat{g}_τ but not c , $p(\hat{g}_\tau)$ can be expressed as the probability that c is below a given threshold, implying equation (21). The rest of Proposition 3 follows from the comparison of stock prices right before and right after the policy decision at time τ . Right after time τ , at time $\tau+$, the market value of each firm i takes one of two values:

$$M_{\tau+}^i = \begin{cases} M_{\tau+}^{i,yes} = B_{\tau+}^i e^{(\mu - \gamma \sigma^2)(T-\tau) + \frac{1-2\gamma}{2} (T-\tau)^2 \sigma_g^2} & \text{if policy changes} \\ M_{\tau+}^{i,no} = B_{\tau+}^i e^{(\mu - \gamma \sigma^2 + \hat{g}_\tau)(T-\tau) + \frac{1-2\gamma}{2} (T-\tau)^2 \hat{\sigma}_\tau^2} & \text{if policy does not change} \end{cases} \quad (22)$$

The values of the market-to-book (M/B) ratios, M_t^i/B_t^i , are equal across firms for all t because all firms are identical ex ante. We show that right before time τ , the market value of firm i is given by

$$M_\tau^i = \omega M_{\tau+}^{i,yes} + (1 - \omega) M_{\tau+}^{i,no}, \quad (23)$$

where the weight ω , which is always between 0 and 1, is given by

$$\omega = \frac{p_\tau}{p_\tau + (1 - p_\tau) F(\hat{g}_\tau)}, \quad (24)$$

¹⁹The stock return conditional on the announcement of *no* change in policy is analyzed in Section 3.2.

using the abbreviated notation $p_\tau \equiv p(\hat{g}_\tau)$. As one would expect, ω increases with p_τ , with $p_\tau \rightarrow 0$ implying $\omega \rightarrow 0$ and $p_\tau \rightarrow 1$ implying $\omega \rightarrow 1$. The announcement return $R(\hat{g}_\tau)$ in equation (18) is given by the ratio of the quantities in equations (22) and (23):

$$R(\hat{g}_\tau) = \frac{M_{\tau+}^{i,yes}}{M_\tau^i} - 1. \quad (25)$$

To gain some intuitive insight into the announcement return $R(\hat{g}_\tau)$, recall that a policy change replaces a policy whose impact is perceived to be distributed as $N(\hat{g}_\tau, \hat{\sigma}_\tau^2)$ by a policy whose impact is perceived as $N(0, \sigma_g^2)$. Relative to the old policy, the new policy typically increases the firms' expected future cash flows (because \hat{g}_τ must be below a threshold that is typically negative).²⁰ However, the new policy also increases the discount rates due to its higher uncertainty (because $\sigma_g > \hat{\sigma}_\tau$ due to learning). The cash flow effect pushes stock prices up, whereas the discount rate effect pushes them down. Either effect can win, depending on \hat{g}_τ and the parameter values. The following two corollaries describe the behavior of $R(\hat{g}_\tau)$ when risk aversion $\gamma > 1$ takes extreme values.

Corollary 1. *As risk aversion $\gamma \rightarrow \infty$, the announcement return $R(\hat{g}_\tau) \rightarrow -1$ for any \hat{g}_τ .*

Corollary 2. *As risk aversion $\gamma \rightarrow 1$, the expected value of the announcement return goes to zero ($E\{R(\hat{g}_\tau)\} \rightarrow 0$), where the expectation is computed with respect to \hat{g}_τ as of time 0.*

Both corollaries follow quickly from Proposition 3. When $\gamma \rightarrow \infty$, the discount rate effect discussed above dwarfs the cash flow effect and stocks lose all of their value when the more uncertain policy is installed. In this limiting case, the probability of a policy change goes to zero. By continuity, Corollary 1 implies that if risk aversion is large, policy changes are unlikely but if they do occur, stock prices fall dramatically. When $\gamma \rightarrow 1$, the discount rate effect and the cash flow effect cancel out, on average.

3.1.1. When is the Announcement Return Negative?

In this subsection, we derive the conditions under which the announcement return $R(\hat{g}_\tau) < 0$.

Proposition 4. *The market value of each firm drops at the announcement of a policy change (i.e., $R(\hat{g}_\tau) < 0$) if and only if*

$$\hat{g}_\tau > g^*, \quad (26)$$

²⁰To clarify the word “typically,” exceptions occur if the government derives an unexpectedly large political benefit from changing its policy. If c is sufficiently negative, the cutoff $\underline{g}(c)$ in equation (15) can be positive, and the old policy can be replaced even if it has a positive perceived impact on profitability.

where

$$g^* = -(\sigma_g^2 - \hat{\sigma}_\tau^2)(T - \tau) \left(\gamma - \frac{1}{2} \right). \quad (27)$$

Proposition 4 shows that stock prices drop at the announcement of a new policy unless the old policy is perceived as having a sufficiently negative impact on profitability. The relative importance of the cash flow and discount rate effects depends on \hat{g}_τ . Under (26), the discount rate effect is stronger and the announcement return is negative. Formally, (26) implies $G(\hat{g}_\tau) > 1$, resulting in $R(\hat{g}_\tau) < 0$ (see equations (18) and (20)).

Combining the results in Propositions 2 and 4, stock prices drop at the announcement of a policy change if and only if

$$g^* < \hat{g}_\tau < \underline{g}(c). \quad (28)$$

That is, \hat{g}_τ must be sufficiently low for the policy change to occur (Proposition 2), but it must be sufficiently high for the discount rate effect to overcome the cash flow effect (Proposition 4). Comparing the definitions of $\underline{g}(c)$ and g^* in equations (15) and (27),

$$\underline{g}(c) - g^* = \frac{c}{(T - \tau)(1 - \gamma)} + (\sigma_g^2 - \hat{\sigma}_\tau^2)(T - \tau) \frac{\gamma}{2}. \quad (29)$$

The first term on the right-hand side of equation (29) is expected to be close to zero. The second term is always positive, so $\underline{g}(c) - g^*$ is generally positive, and its magnitude (and thus also the size of the interval for which the condition (28) holds) increases with σ_g .

Since g^* in equation (27) can also be expressed as $g^* = \underline{g}(0) - (\sigma_g^2 - \hat{\sigma}_\tau^2)(T - \tau)\gamma/2$, we have $g^* < \underline{g}(0) < 0$. Figure 3 illustrates four possible locations of \hat{g}_τ relative to g^* , $\underline{g}(0)$, and zero. Also plotted is the distribution of $\underline{g}(c)$, as perceived by investors (who do not observe c) just before time τ . This normal distribution is centered at a value just above $\underline{g}(0)$ (see equations (6) and (15)). This distribution, along with the values of g^* and $\underline{g}(0)$, are computed based on the parameters in Table 1. The shaded area represents the probability of a policy change, as perceived by the investors just before time τ .

In Panel A of Figure 3, \hat{g}_τ is very low, and the probability of a policy change is nearly one. Since $\hat{g}_\tau < g^*$, stock prices rise at the announcement of a policy change. Given the high probability of such a change, the price increase will be small because most of it is already priced in. In contrast, stock prices plunge in the unlikely event of no policy change, which occurs if such a change imposes a very large political cost on the government.

In Panels B through D, $\hat{g}_\tau > g^*$, so stock prices fall if the policy is changed. In Panel B, $g^* < \hat{g}_\tau < \underline{g}(0)$, and the policy change reduces stock prices even though it increases

the investors' expected utility. Stock prices are lower due to higher discount rates, but the expected utility is higher due to higher expected wealth. Expected utility and stock prices need not move in the same direction because stock prices are related to marginal utility rather than the level of utility. In Panel D, $\hat{g}_\tau > 0$, indicating that the prevailing policy is boosting profitability. The probability of a policy change is then small, but should such a change occur (if $c \ll 0$), the stock price reaction will be strongly negative. If the government derives an unexpectedly large political benefit from changing its policy, it replaces even a policy that appears to work well, and stock prices exhibit a large drop as a result.

Figure 3 captures uncertainty about c while conditioning on \hat{g}_τ . Such a perspective is relevant at time τ when \hat{g}_τ is known, but less so before time τ while \hat{g}_τ is uncertain. As of time 0, the perceived distribution of \hat{g}_τ is $N(0, \sigma_g^2 - \hat{\sigma}_\tau^2)$. Therefore, the values of \hat{g}_τ considered in Panel A of Figure 3, for which stock prices rise at the announcement of a policy change, are less likely than those in Panels B through D, for which stock prices fall. The distribution of $R(\hat{g}_\tau)$ that is relevant from the perspective of an econometrician (or an investor at time 0) must incorporate uncertainty about both c and \hat{g}_τ . We integrate out all of that uncertainty in computing the expected value of $R(\hat{g}_\tau)$ in the following subsection.

3.1.2. The Expected Announcement Return (EAR)

We define the expected announcement return (EAR) as the expected value of $R(\hat{g}_\tau)$ after integrating out all uncertainty about c and \hat{g}_τ as of time 0 and conditioning on $\hat{g}_\tau < \underline{g}(c)$:

$$\text{EAR} \equiv \text{E} \{R(\hat{g}_\tau) | \text{Policy Change}\} . \quad (30)$$

Note that EAR is the return that we expect to see at the announcement of a policy change—the expected value of the return at time τ conditional on a policy change at the same time τ . Given the conditioning on a contemporaneous event, EAR does not represent the more traditional expectation of a future return based on information available today. (The latter expected returns are analyzed in Section 3.2.) Instead, EAR is the expected return relevant from an event study perspective. The following proposition presents our main result.

Proposition 5. *The expected announcement return is negative ($\text{EAR} < 0$).*

According to Proposition 5, stock prices are expected to fall at the announcement of a policy change. This result relies on the government's quasi-benevolence. Since the government is expected to maximize the investors' welfare, the government's value-enhancing policy decisions are mostly expected by the market participants. Positive announcement

returns tend to occur when the policy change is widely anticipated ($\hat{g}_\tau < g^*$, see Panel A of Figure 3). The positive effect of the policy change is thus largely priced in before the announcement, resulting in a weak price reaction to the announcement. In contrast, negative announcement returns occur when the policy change comes as a bigger surprise (Panels B through D), so the price reaction is stronger. In short, the positive returns tend to be small and the negative returns tend to be large. In addition, the negative announcement returns are fairly likely since even some utility-increasing policy changes reduce market values, as explained earlier while discussing Panel B of Figure 3.

To understand Proposition 5, it also helps to realize that when investors observe a policy change, they revise their beliefs about the political cost C . Before time τ , they expect C to be one (see equation (6)), but conditional on a policy change at time τ , the expectation drops below one: $E(C|\text{policy change}) < E(C) = 1$. This fact follows from the condition (14) for the optimality of a policy change, which can be rewritten as $c < \xi(\hat{g}_\tau)$, where $\xi(\hat{g}_\tau)$ is a decreasing function of \hat{g}_τ . (The expected value of c conditional on $c < \xi(\hat{g}_\tau)$ is less than the unconditional expected value of c .) As a result, investors expect the government to derive a political benefit from any policy change. Even though C does not affect stock prices directly, it adds noise to the policy decision. Observing a policy change, investors infer that the increase in risk is certain but the increase in cash flows is not, and stock prices typically fall. They fall more when \hat{g}_τ is higher because the policy change is then more likely to be politically motivated. After integrating \hat{g}_τ out, we prove that $\text{EAR} < 0$.

The magnitude of EAR depends on uncertainty about government policy. We relate EAR to σ_g and σ_c in Figure 4. We compute EAR by averaging the announcement returns $R(\hat{g}_\tau)$ across all simulated paths for which a policy change occurs at time τ . The parameter values are from Table 1, as before, except that we vary σ_g and σ_c .

Figure 4 shows that EAR is more negative for larger values of σ_g and σ_c . Fixing $\sigma_c = 20\%$, EAR is only -0.3% for $\sigma_g = 1\%$ but -2% for $\sigma_g = 3\%$. When σ_g is larger, so is the risk associated with a new policy, and so is the discount rate effect that pushes stock prices down. Fixing $\sigma_g = 2\%$, EAR is -0.5% for $\sigma_c = 10\%$ and -1.1% for $\sigma_c = 20\%$. When σ_c is larger, so is the element of surprise in the announcement of a new policy. Also note that when either uncertainty is zero, so is the announcement return. When $\sigma_c = 0$, the policy change is fully anticipated, and all of its effect is priced in before the announcement. When $\sigma_g = 0$, one neutral policy is replaced by another, making no difference to investors.

3.1.3. The Determinants of the Announcement Return

In this subsection, we analyze two key determinants of the announcement return $R(\hat{g}_\tau)$ from equation (18): the length and depth of the downturns that precede policy changes. Recall that a downturn is defined as a period of unexpectedly low realized profitability. We vary the downturns' length and depth in a simple way. We measure the length of a downturn as

$$\text{LENGTH} = \tau - t_0, \quad (31)$$

where we set $\hat{g}_{t_0} = 0$. At time t_0 , the posterior mean of the policy impact is then equal to the prior mean, but after t_0 , \hat{g}_t generally falls, conditional on a policy change at time τ (Proposition 2). We refer to t_0 as the beginning of a downturn.

We measure the depth of a downturn by the number of standard deviations by which \hat{g}_t drops during the downturn. Conditional on $\hat{g}_{t_0} = 0$, the distribution of \hat{g}_τ is normal with mean zero and standard deviation $\text{Std}(\hat{g}_\tau) = \sqrt{\hat{\sigma}_{t_0}^2 - \hat{\sigma}_\tau^2}$. We define

$$\text{DEPTH} = \frac{\hat{g}_\tau}{\text{Std}(\hat{g}_\tau)}. \quad (32)$$

Panel A of Figure 5 plots the announcement return as a function of LENGTH and DEPTH. Panel B plots the corresponding probability of a policy change, as perceived by investors just before time τ (cf. Figure 3). Whenever this probability is close to one, the announcement return is close to zero because the announcement is already priced in. When the probability is smaller than one, a policy change contains an element of surprise, and the announcement return is nonzero. The parameter values are from Table 1, as before.

Figure 5 shows that DEPTH has a large effect on the announcement return. When the downturn is shallow (i.e., DEPTH is not a large negative value), the announcement return is negative. This result makes sense given Proposition 4, since DEPTH is proportional to \hat{g}_τ . The magnitude of the announcement return can be as large as -10% if the downturn is sufficiently shallow. Even larger returns can happen if $c \ll 0$, i.e., if the government derives a huge political benefit from changing its policy. In contrast, when the downturn is sufficiently deep, the announcement return is essentially zero. The return is positive (Proposition 4), but it is small because the policy change is fully anticipated (Panel B).

The announcement return also depends on LENGTH. Figure 5 shows that shorter downturns are followed by more negative announcement returns, holding DEPTH at a constant negative value. The reason is that as LENGTH increases, so does $\text{Std}(\hat{g}_\tau)$ in equation (32),

so does the negative magnitude of \hat{g}_τ required to keep DEPTH constant, and so does the probability of a policy change, resulting in a smaller element of surprise in the announcement of such a change. (This effect reverses when DEPTH is positive because higher LENGTH then implies a higher positive magnitude of \hat{g}_τ and thus a lower probability of a policy change.) Put differently, shorter downturns are less likely to induce a policy change, so if such a change occurs, it comes as a bigger surprise. It is also more likely to be politically motivated, resulting in a more negative announcement return.

The announcement returns in Figure 5 are nonpositive, but they can be positive for other parameter values. Recall from equation (29) that the size of the interval for which prices drop at the announcement increases with σ_g . If σ_g is small, so is the interval, and positive announcement returns are more likely. To see this effect, consider Figure 6, which is analogous to Figure 5, except that the baseline value $\sigma_g = 2\%$ is replaced by $\sigma_g = 1\%$ (left panels) and $\sigma_g = 3\%$ (right panels). The patterns for $\sigma_g = 3\%$ are similar to those in Figure 5 but more pronounced; the announcement return can be -20% after a shallow downturn. However, for $\sigma_g = 1\%$, the announcement return can be positive for downturns that are neither too short nor too shallow. These positive returns are small, less than 0.5% . Slightly larger positive announcement returns obtain when we raise σ_c from 10% to 20% , but even those returns are smaller than 1% (results are not shown to save space).

Figures 5 and 6 show that shorter downturns are followed by more negative announcement returns when we hold DEPTH fixed at a constant negative value. The same result holds in expectation when we integrate out uncertainty about DEPTH. Figure 7 is the counterpart of Figure 4: it plots EAR for two different LENGTHs: 5 years in Panel A and 1 year in Panel B. (In Figure 4, LENGTH = 10 years.) Figure 7 shows that EAR grows more negative as LENGTH shortens. For example, in the benchmark case from Table 1, in which $\sigma_g = 2\%$ and $\sigma_c = 10\%$, EAR goes from -0.5% for LENGTH = 10 (Figure 4) to -1% for LENGTH = 5 to -4% for LENGTH = 1 (Figure 7). The effect of LENGTH is even stronger when there is more government-induced uncertainty: when $\sigma_g = 3\%$ and $\sigma_c = 20\%$, EAR goes from -2% for LENGTH = 10 to -4.6% for LENGTH = 5 to -18% for LENGTH = 1.

3.1.4. The Distribution of Stock Returns on the Announcement Day

In this subsection, we examine the probability distribution of stock returns on the day of the announcement of a policy change, without conditioning on DEPTH. This distribution is relevant from the empirical perspective if the event study of the announcement effects is

conducted on daily returns, which is commonly done.²¹ The announcement day returns have a jump component $R(\hat{g}_\tau)$ pertaining to the instant of the announcement, as well as a diffusion component $M_{\tau+\text{day}}/M_\tau - 1$ covering the rest of the announcement day. Figure 8 plots the distribution of the announcement day returns across all simulated paths for which a policy change occurs at time τ . The randomness in the simulations comes from the randomness in \hat{g}_τ and c . The downturn length is 5 years. The parameter values are from Table 1, as before, except that we vary σ_g and σ_c .

Figure 8 shows that the distribution of the announcement day returns is strongly left-skewed. The mode is close to zero, returns larger than 2% are extremely rare, whereas returns below -5% and even -10% are more common, especially when σ_g and σ_c are large. This skewness is due to the asymmetry discussed earlier. Positive announcement returns tend to occur when the policy change is widely anticipated, whereas negative returns occur when the policy change comes as a bigger surprise. We also see that as we increase σ_g or σ_c , the probability mass in Figure 8 shifts to the left, consistent with our earlier result that EAR is more negative when σ_g or σ_c are larger.

3.2. Stock Price Dynamics

The dynamics of stock prices are closely related to the dynamics of the stochastic discount factor from equation (16), which are described by the following proposition.

Proposition 6. *The stochastic discount factor (SDF) follows the process*

$$\frac{d\pi_t}{\pi_t} = -\sigma_{\pi,t}d\hat{Z}_t + J_\pi 1_{\{t=\tau\}} , \quad (33)$$

where $d\hat{Z}_t$ is the Brownian motion from Proposition 1, $1_{\{t=\tau\}}$ is an indicator function equal to one for $t = \tau$ and zero otherwise, and the jump component J_π is given by

$$J_\pi = \begin{cases} J_\pi^{yes} = \frac{(1-p_\tau)(1-F(\hat{g}_\tau))}{p_\tau+(1-p_\tau)F(\hat{g}_\tau)} & \text{if policy changes} \\ J_\pi^{no} = \frac{p_\tau(F(\hat{g}_\tau)-1)}{p_\tau+(1-p_\tau)F(\hat{g}_\tau)} & \text{if policy does not change} . \end{cases} \quad (34)$$

For $t > \tau$, $\sigma_{\pi,t}$ is given by

$$\sigma_{\pi,t} = \gamma [\sigma + (T - t)\hat{\sigma}_t^2\sigma^{-1}] , \quad (35)$$

²¹For example, Savor and Wilson (2010) use daily returns to analyze the announcement effects of macroeconomic news announcements regarding employment, inflation, and interest rates. If the event study is conducted on tick-by-tick returns instead, the distribution of the instantaneous return $R(\hat{g}_\tau)$ would be more relevant. That distribution is similar to the distribution of the announcement day returns plotted here, except that it looks even more left-skewed and it has less probability mass above zero.

and for $t \leq \tau$, it is given in equation (A2) in the Appendix.

Proposition 6 shows that SDF jumps at time τ when the policy decision is announced. The magnitude of the jump depends on whether the policy is changed as well as on \hat{g}_τ . The jumps J_π^{yes} and J_π^{no} always have the opposite signs. The expected value of the jump, as perceived just before time τ , is zero: $E_\tau(J_\pi) = p_\tau J_\pi^{yes} + (1 - p_\tau) J_\pi^{no} = 0$. It is also easy to show that $J_\pi^{yes} < 0$ (and $J_\pi^{no} > 0$) if and only if $\hat{g}_\tau < g^{**}$, where g^{**} is given by

$$g^{**} = -\frac{\gamma}{2}(T - \tau)(\sigma_g^2 - \hat{\sigma}_\tau^2) . \quad (36)$$

Proposition 7. *The return process for stock i is given by*

$$\frac{dM_t^i}{M_t^i} = \mu_{M,t} dt + \sigma_{M,t} d\hat{Z}_t + \sigma_1 dZ_t^i + J_M 1_{\{t=\tau\}} , \quad (37)$$

where the jump component J_M is given by

$$J_M = \begin{cases} J_M^{yes} = R(\hat{g}_\tau) & \text{if policy changes} \\ J_M^{no} = R(\hat{g}_\tau) G(\hat{g}_\tau) + G(\hat{g}_\tau) - 1 & \text{if policy does not change .} \end{cases} \quad (38)$$

For $t > \tau$, we have

$$\mu_{M,t} = \gamma [\sigma + (T - t) \hat{\sigma}_t^2 \sigma^{-1}]^2 \quad (39)$$

$$\sigma_{M,t} = \sigma + (T - t) \hat{\sigma}_t^2 \sigma^{-1} , \quad (40)$$

and for $t \leq \tau$, $\sigma_{M,t}$ and $\mu_{M,t}$ are given in equations (A4) and (A5) in the Appendix.

Proposition 7 shows that stock returns have two components: a diffusion component and a jump component. We discuss these components in separate subsections.

3.3. Stock Price Jump at the Announcement of a Policy Decision

Stock prices jump at time τ when the government announces its policy decision. If the decision is to change the prevailing policy, the jump J_M^{yes} is equal to $R(\hat{g}_\tau)$, which is given in equation (18) and analyzed extensively in Section 3.1. If the decision is *not* to change the policy, the jump J_M^{no} in equation (38) involves also the function $G(\hat{g}_\tau)$ from equation (20).

Corollary 3. *The market value of each firm increases at the announcement of no policy change (i.e., $J_M^{no} > 0$) if and only if $\hat{g}_\tau > g^*$, where g^* is given in equation (27).*

Comparing Corollary 3 with Proposition 4, we see that the jumps J_M^{yes} and J_M^{no} always have the opposite signs: whenever one is positive, the other is negative. When $\hat{g}_\tau > g^*$, we have $J_M^{yes} < 0$ and $J_M^{no} > 0$. When $\hat{g}_\tau < g^*$, we have $J_M^{yes} > 0$ and $J_M^{no} < 0$.

In the remainder of this subsection, we focus on the expected value of J_M , or the expected jump in stock prices at the announcement of a policy decision. This expectation captures the risk premium that investors demand for facing jumps in SDF at time τ . We consider both conditional and unconditional expectations. The *conditional* expectation of J_M , denoted by $E_\tau(J_M)$, is perceived by agents just before time τ :

$$E_\tau(J_M) = p_\tau J_M^{yes} + (1 - p_\tau) J_M^{no} . \quad (41)$$

This conditional expectation conditions on \hat{g}_τ , which is observable just before the policy decision. We also compute the *unconditional* expectation $E(J_M) = E(E_\tau(J_M))$ by integrating out uncertainty about \hat{g}_τ as of time 0. (Recall that $\hat{g}_\tau \sim N(0, \sigma_g^2 - \hat{\sigma}_\tau^2)$.) This latter expectation matters to an econometrician who averages jumps in stock prices across all announcements of policy decisions, without controlling for \hat{g}_τ . Neither expectation conditions on whether the decision changes the policy or not, so that both $E_\tau(J_M)$ and $E(J_M)$ can be viewed as expected returns (or risk premia) in the traditional forward-looking sense.

3.3.1. The Conditional Jump Risk Premium

Proposition 8. *The conditional expected jump in stock prices at time τ , as perceived by investors just before time τ , is given by*

$$E_\tau(J_M) = - \frac{p_\tau (1 - p_\tau) (1 - F(\hat{g}_\tau)) (1 - G(\hat{g}_\tau))}{p_\tau + (1 - p_\tau) F(\hat{g}_\tau) G(\hat{g}_\tau)} . \quad (42)$$

Corollary 4. *We have $E_\tau(J_M) < 0$ if and only if*

$$g^* < \hat{g}_\tau < g^{**} , \quad (43)$$

where g^* is given in equation (27) and g^{**} is given in equation (36).

Corollary 4 shows that $E_\tau(J_M)$ can be positive or negative, depending on \hat{g}_τ .²² This expected jump, which captures the jump risk premium, is related to the covariance between the jumps in prices and SDF, as perceived by the agents just before the policy decision:

$$E_\tau(J_M) = -\text{Cov}_\tau(J_\tau, J_M) , \quad (44)$$

²²The size of the interval in (43) is always positive because $g^{**} - g^* = -\underline{g}(0) > 0$.

where J_π is given in equation (34). The sign of this covariance depends on whether (43) is satisfied. If the covariance is positive, upward jumps in marginal utility at time τ tend to be accompanied by upward jumps in stock prices, which makes stocks an effective hedge against jumps in SDF. As a result, investors are willing to accept a negative jump risk premium for holding stocks (i.e., $E_\tau(J_M) < 0$). In contrast, if the covariance is negative, stocks are a poor hedge against jumps in marginal utility and the jump risk premium is positive.

Corollary 5. *As risk aversion $\gamma \rightarrow \infty$, $E_\tau(J_M) \rightarrow 0$ from above for any value of \widehat{g}_τ .*

Corollary 6. *As risk aversion $\gamma \rightarrow 1$, $E_\tau(J_M)$ converges to a nonnegative value for any \widehat{g}_τ . It converges to zero if and only if $\widehat{g}_\tau = -\frac{1}{2}(T - \tau)(\sigma_g^2 - \widehat{\sigma}_\tau^2)$.*

As $\gamma \rightarrow \infty$, the probability of a policy change goes to zero. Given the diminishing element of surprise in the policy decision, the jump risk premium diminishes as well. In contrast, the jump risk premium remains positive almost surely as $\gamma \rightarrow 1$.

3.3.2. The Unconditional Jump Risk Premium

The expectation $E(J_M)$, which does not condition on \widehat{g}_τ , represents the unconditional jump risk premium. As noted earlier, $E(J_M)$ is the expected value of the quantity in Proposition 8, where the expectation is taken with respect to \widehat{g}_τ as of time 0.

Corollary 7. *As risk aversion $\gamma \rightarrow \infty$, $E(J_M) \rightarrow 0$ from above.*

Corollary 8. *As risk aversion $\gamma \rightarrow 1$, $E(J_M)$ converges to a positive value.*

According to Corollaries 7 and 8, there exist values $\underline{\gamma}$ and $\overline{\gamma}$ exceeding one such that $E(J_M) > 0$ for every $\gamma \in (1, \overline{\gamma})$ as well as for every $\gamma \in (\underline{\gamma}, \infty)$. In fact, $E(J_M) > 0$ appears to hold for any value of γ . While we are unable to prove $E(J_M) > 0$ in its full generality (the analytical challenges are formidable), our numerical investigation supports this statement. Having evaluated $E(J_M)$ on a large multidimensional grid of parameter values, we have not found any set of parameters for which $E(J_M) > 0$ is violated. This result is not surprising because the condition (43) is unlikely to hold (since $g^* < g^{**} < \underline{g}(0) < 0$, $\widehat{g}_0 = 0$, and \widehat{g}_t is a martingale). We thus conclude that our model implies a positive unconditional premium for the jump risk associated with the announcements of policy decisions.

Panel A of Figure 9 plots $E(J_M)$ as a function of σ_g and σ_c for the parameter values from Table 1. We set the downturn length to 5 years, so that uncertainty about \widehat{g}_τ is

integrated out as of time $t_0 = 5$ conditional on $\widehat{g}_{t_0} = 0$. The figure shows that the jump risk premium $E(J_M)$ increases with both σ_g and σ_c . In the benchmark case of $\sigma_g = 2\%$ and $\sigma_c = 10\%$, $E(J_M) = 5$ basis points; for $\sigma_g = 3\%$, it is 10 basis points. These magnitudes are comparable to the 9.6 basis point jump risk premium estimated by Savor and Wilson (2010). These authors find that stock market returns tend to be higher on days when news about inflation, unemployment, or interest rates is announced. We focus on a different kind of announcements, those of government policy decisions, so Savor and Wilson's estimates are not directly comparable to the magnitudes in our Figure 9. Nonetheless, it seems interesting that the magnitude of the jump risk premium produced by our model is similar to the jump risk premium estimated empirically for a different kind of market-wide announcements.

Panels B and C of Figure 9 plot the unconditional expected jumps $E(J_M^{yes})$ and $E(J_M^{no})$, which are computed analogously to $E(J_M)$ by integrating out uncertainty about \widehat{g}_τ . Panel B, showing $E(J_M^{yes})$, coincides with Panel A of Figure 7; it is included for symmetry. Panel C of Figure 9 shows that $E(J_M^{no})$ is positive and increasing in both σ_g and σ_c , but its magnitude is much smaller than that of $E(J_M^{yes})$. The reason is that the probability-weighted average of the jumps in Panels B and C is close to zero (see Panel A) and the unconditional probability of a policy change is less than 0.5. This probability, which is assessed at the beginning of the downturn, ranges from 11% to 49%, depending on the values of σ_g and σ_c (it is 29% in the benchmark case of $\sigma_g = 2\%$ and $\sigma_c = 10\%$).

3.4. The Diffusion Component of Stock Returns

Whereas the jump components are nonzero only at time τ , the diffusion components in equations (33) and (37) drive the dynamics of returns and SDF for all $t \in (0, T)$. Both components are affected by uncertainty about government policy. The jump component is driven by the resolution of political uncertainty, whereas the diffusion component is driven by policy uncertainty. Equations (35), (39), and (40) show that policy uncertainty increases the volatility of SDF as well as the mean and variance of stock returns. After time τ , $\sigma_{\pi,t}$, $\mu_{M,t}$, and $\sigma_{M,t}$ all vary over time as increasing functions of $\widehat{\sigma}_t$. Similar dependence on uncertainty about g_t is present before time τ . If the policy change occurs at time τ , uncertainty about g_t jumps up (since $\sigma_g > \widehat{\sigma}_\tau$), causing upward jumps in $\sigma_{\pi,t}$, $\mu_{M,t}$, and $\sigma_{M,t}$.

Figure 10 plots the expected dynamics of $\sigma_{\pi,t}$, $\mu_{M,t}$ and $\sigma_{M,t}$ around a policy change. We simulate many samples of shocks in the model, maintaining $\widehat{g}_{t_0} = 0$ for $t_0 = 5$ to keep the downturn length at 5 years. (The results for downturn lengths of 1 and 10 years look very

similar.) We plot the paths of $\sigma_{\pi,t}$, $\mu_{M,t}$ and $\sigma_{M,t}$ averaged across the subset of samples in which a policy change occurs in year $\tau = 10$. The parameters are from Table 1, except for σ_g , which takes three different values, 1%, 2%, and 3% per year.

Figure 10 is dominated by the upward jumps in $\sigma_{\pi,t}$, $\mu_{M,t}$ and $\sigma_{M,t}$ at the time of a policy change. These jumps are induced by increases in uncertainty about g_t , as discussed earlier. For example, for $\sigma_g = 2\%$, $\sigma_{\pi,t}$ jumps from 28% to 65% per year, $\mu_{M,t}$ jumps from 2% to 9% per year, and $\sigma_{M,t}$ jumps from 12% to 16% per year. The magnitudes of all three quantities increase in σ_g : when $\sigma_g = 3\%$, both $\mu_{M,t}$ and $\sigma_{M,t}$ jump to about 25% per year. After time τ , all quantities gradually fall as a result of the learning-induced gradual decline in $\hat{\sigma}_t$.

Stock returns before time τ are affected by multiple forces. Let p_t denote the probability of a policy change at time τ , as perceived by investors at time $t < \tau$. Fluctuations in p_t contribute to volatility: stock prices fluctuate as investors change their minds about what the government is going to do. Conditional on a policy change, p_t grows towards $p_\tau = 1$, and the volatility induced by fluctuations in p_t typically increases as time τ approaches. There are also two opposite effects. First, $\hat{\sigma}_t$ declines over time due to learning, thereby pushing $\sigma_{\pi,t}$, $\mu_{M,t}$, and $\sigma_{M,t}$ down. Second, when p_t increases, the current value of g_t becomes less likely to matter after time τ , making stock prices less sensitive to \hat{g}_t . The rising probability of a policy change makes the current policy less relevant, reducing the sensitivity of stock prices to time-varying beliefs about this policy. Which of these effects prevails depends on the parameter values. In Figure 10, all three quantities fall before τ , suggesting that the effect of growing fluctuations in the probability of a policy change is weaker than the combined effects of learning and the diminished sensitivity of prices to \hat{g}_t .²³

Corollary 9. *The correlation between the returns of any pair of stocks for $t > \tau$ is given by*

$$\rho_t = \frac{[\sigma + (T - t) \hat{\sigma}_t^2 \sigma^{-1}]^2}{[\sigma + (T - t) \hat{\sigma}_t^2 \sigma^{-1}]^2 + \sigma_1^2} . \quad (45)$$

For $t < \tau$, the correlation is given in equation (A6) in the Appendix.

For $t = \tau$, the instantaneous correlation is one.

Equation (45) shows that after time τ , the correlation ρ_t increases with $\hat{\sigma}_t$. Intuitively, uncertainty about g_t increases each stock's sensitivity to the common factor \hat{g}_t , thereby making returns more correlated. Similar dependence of ρ_t on $\hat{\sigma}_t$ is present also before time

²³In unreported results for LENGTH=1 year, both $\mu_{M,t}$ and $\sigma_{M,t}$ fall before τ when $\sigma_g = 1\%$ and 2%, but they rise when $\sigma_g = 3\%$. In the latter case, the effect of growing fluctuations in the probability of a policy change prevails over the combined effects of learning and the diminished sensitivity of stock prices to \hat{g}_t .

τ , when ρ_t depends also on the probability of a policy change. At time τ , $\rho_t = 1$ due to the resolution of political uncertainty, which results in a common jump in stock prices.

Panel D of Figure 10 plots the expected dynamics of ρ_t around a policy change at time $\tau = 10$. The correlation jumps up at time τ , for reasons discussed earlier. This jump is substantial: from 30% to 62% when $\sigma_g = 2\%$, and from 39% to 84% when $\sigma_g = 3\%$. (We do not plot the instantaneous jump in ρ_t to one at time τ .) After time τ , the correlation falls due to learning. Before time τ , the correlation falls due to the previously-discussed effects of learning and the diminished price sensitivity to \hat{g}_t .

Figure 11 is equivalent to Figure 10, except that it focuses on policy decisions that result in no policy change. That is, we plot the paths of $\sigma_{\pi,t}$, $\mu_{M,t}$, $\sigma_{M,t}$, and ρ_t averaged across the subset of simulated samples in which no policy change occurs at time τ . Unlike in Figure 10, all four quantities drop at τ for $\sigma_g = 2\%$ and 3% . The main reason behind the drop is the resolution of political uncertainty at τ . As noted earlier, fluctuations in p_t contribute to stock price volatility before τ . Once the government makes its decision, p_t stops fluctuating and this component of volatility disappears. Under no policy change, there is no opposing increase in uncertainty about g_t (unlike in Figure 10), so the four quantities drop.

The only exception occurs for $\sigma_g = 1\%$, for which the four quantities rise slightly at τ . In this case, the resolution of political uncertainty is outweighed by the increase in stock price sensitivity to \hat{g}_t after time τ . When σ_g is small, both \hat{g}_t and the threshold $\underline{g}(c)$ are expected to be close to zero before τ , so that p_t is typically nontrivial for $t < \tau$ even in those samples in which no policy change occurs at τ . The nontrivial probability of a policy change reduces the pre- τ sensitivity of stock prices to \hat{g}_t , as discussed earlier. This reduction in sensitivity is more pronounced on the downside: when \hat{g}_t falls, p_t increases, raising the probability that the current g_t will not matter after τ and thereby cushioning the stock price drop. However, once p_t jumps to zero at τ , the price sensitivity to \hat{g}_t increases, pushing volatility up.

3.5. Price Dynamics When Policy Changes Are Precluded

In this subsection, we compare the model-implied stock prices with their counterparts in the hypothetical scenario in which policy changes are precluded. This scenario matches our model in all respects except that the government cannot change its policy at time τ (and investors know that). First, we compare the aggregate stock market values at time 0, M_0 , given the parameter values from Table 1. We find that M_0 is smaller when policy changes

are precluded: 2.81 vs 3.25.²⁴ However, this result is not general—while it holds for $\sigma_g = 3\%$ as well, it reverses for $\sigma_g = 1\%$ (4.81 vs 4.67). Second, we conduct similar comparisons while varying the time remaining until the policy decision. We find that if the policy decision is close enough (eight years or less), the threat of systematic risk going up in the event of a policy change reduces stock prices relative to the hypothetical scenario, even for $\sigma_g = 2\%$.²⁵ The government’s ability to change its policy can thus increase or decrease market values, depending on the parameter values and the time remaining until the policy decision.

We also compare the expected dynamics of M_t when policy changes are allowed versus when they are precluded, while conditioning on (14) and the parameters from Table 1. That is, we compare the values of M_t averaged across the subsample of all simulations in which a policy change occurs in our model. These average values are plotted in Panels A and B of Figure 12. Panel A conditions on the downturn length of one year, whereas Panel B conditions on five years. In both panels, M_t falls before time τ whether policy changes are allowed or precluded, due to the ex-post conditioning on (14). (Declines in \hat{g}_τ tend to be accompanied by declines in stock prices.) In addition, the average stock price drop at time τ is clearly visible. Interestingly, in Panel B, M_t around time τ is mostly higher when policy changes are allowed, but in Panel A, M_t is higher when policy changes are precluded. That is, the government’s ability to change its policy reduces market values during short downturns. In both panels, this ability amplifies the stock price decline before time τ .

Panels C and D of Figure 12 plot the expected dynamics of market volatility, $\sigma_{M,t}$, conditional on (14). In both panels, after time τ , volatility is higher when policy changes are allowed. Before time τ , though, the ability to change the policy can increase or decrease volatility compared to the hypothetical scenario: increase when the downturn lasts one year (Panel C), but decrease when it lasts five years (Panel D). On the one hand, the ability to change the current policy decreases pre- τ volatility by reducing the sensitivity of stock prices to news about the impact of the current policy. On the other hand, this ability increases pre- τ volatility by introducing fluctuations in the agents’ assessment of the probability of a policy change. The latter force prevails in Panel C, while the former prevails in Panel D. Overall, we conclude that the government’s ability to change its policy substantially affects the level and volatility of stock prices.

²⁴Market values are measured in units of book value at time 0. Recall that we set $B_0^i = 1$ for all i .

²⁵The results are available upon request. We do not plot them here to save space.

4. Extension: Endogenous Timing of Policy Change

Our benchmark model assumes that the government can change its policy only at a given time τ . This assumption buys us closed-form solutions for many quantities of interest, and it allows us to formally prove our propositions and corollaries. In this section, we extend the model by endogenizing the timing of the policy change. No closed-form solutions are available but we can solve the problem numerically. To give a preview, we find that our main results continue to hold in this setting. In addition, we find that policy changes that happen earlier tend to produce more negative announcement returns.

We assume that the government can change its policy at any time $\tau \in [\tau_1, \tau_2, \dots, \tau_n]$, where $\tau_i = i$ and $n = 19$, so the policy change can occur in any year ($T = 20$). Once the change is made, it is irreversible. At each time τ_i , a new value of the political cost C_i is drawn from the distribution in equation (6). The values of C_i are iid. After observing C_i , the government decides whether to change its policy, maximizing the same objective as before. Let $V(\hat{g}_t, B_t, t)$ denote the value function given no policy change at or before time t :

$$V(\hat{g}_t, B_t, t) = E_t \left\{ \max_{\tau > t} \left\{ E_\tau \left[\frac{B_T^{1-\gamma}}{1-\gamma} \mid \text{No change at } \tau \right], E_\tau \left[C \frac{B_T^{1-\gamma}}{1-\gamma} \mid \text{Change at } \tau \right] \right\} \right\} . \quad (46)$$

This value represents a solution to a partial differential equation, with the final condition $V(\hat{g}_T, B_T, T) = B_T^{1-\gamma}/(1-\gamma)$, as detailed in the Appendix. Conditional on a policy change at a given time τ , the value function at that time is available in closed form:

$$\bar{V}(B_\tau, \tau, c) = \frac{B_\tau^{1-\gamma}}{1-\gamma} e^{c+\mu(1-\gamma)(T-\tau)+\sigma_g^2(T-\tau)^2-\gamma(1-\gamma)\frac{\sigma_c^2}{2}(T-\tau)} . \quad (47)$$

The government changes its policy at time $\tau = \tau_i$ if and only if

$$\bar{V}(B_{\tau_i}, \tau_i, c) > V(\hat{g}_\tau, B_{\tau_i}, \tau_i) . \quad (48)$$

Figure 13 shows that our key results continue to hold when τ is endogenous. For the parameter values from Table 1, Panel A plots EAR as a function of σ_g and σ_c . This panel is analogous to Figure 4, in which τ is exogenous. As in Figure 4, EAR is negative and its magnitude is increasing in both σ_g and σ_c . The magnitudes are generally larger than in Figure 4: for example, in the benchmark case with $\sigma_g = 2\%$ and $\sigma_c = 10\%$, EAR is -1.9%, compared to -0.5% in Figure 4. The reason why the magnitudes are larger when τ is endogenous is difficult to convey precisely because the results are obtained numerically. However, some intuition emerges when we consider Propositions 2 and 4, which apply when τ

is exogenous. These propositions imply that a policy change occurs at time τ if $\hat{g}_\tau < \underline{g}(c)$, and that stock prices drop at the announcement if $g^* < \hat{g}_\tau < \underline{g}(c)$. Under the bold assumption that these results are approximately relevant also when τ is endogenous, the policy change will occur at time τ_i if $\hat{g}_{\tau_i} < \underline{g}(c_i)$ but $\hat{g}_{\tau_{i-1}} > \underline{g}(c_{i-1})$. Since \hat{g}_t has only just fallen below the threshold at time τ_i (but not at time τ_{i-1}), it seems unlikely that \hat{g}_{τ_i} is so low that it is also below g^* , in which case the announcement return would be positive. This intuition is only approximate, not only because Propositions 2 and 4 rely on exogenous τ but also because the threshold $\underline{g}(c_i)$ is time-varying due to time variation in c_i . Nonetheless, it is comforting to see that our main result holds and even strengthens when we endogenize τ .

Panel B of Figure 13 plots the same quantity as Panel A, the expected announcement return conditional on a policy change, but this time as a function of the time τ at which the policy is changed. Recall that τ is optimally chosen by the government from the set $\tau \in [1, 2, \dots, 19]$. Whereas Panel A averages the announcement returns across all τ , Panel B averages them only across those simulations in which the policy change occurs at the given τ . We keep $\sigma_c = 10\%$ as in Table 1 and vary σ_g .²⁶ There are two basic results: (i) the announcement return is negative for all τ , and (ii) its magnitude is larger when τ is smaller. For example, holding σ_g at 2%, the announcement return is -1% for $\tau = 15$ but almost -5% for $\tau = 5$.²⁷ This result is easy to understand. At each point in time, the government weighs the costs and benefits of changing the policy. One important cost is that the irreversible policy change destroys the option to wait (e.g., the political benefit from changing the policy might be higher in the future). This option value of waiting declines as time passes. Early on, while this option value is still high, it takes a large political benefit for the policy change to occur. As a result, policy changes occurring at low τ_i 's tend to be politically motivated and highly unexpected, producing larger negative announcement returns.

Panels C and D of Figure 13 plot the expected dynamics of $\sigma_{M,t}$ and ρ_t , respectively, around a policy change. These dynamics are computed by averaging across those simulated paths in which a policy change occurs at any $\tau \in [1, 2, \dots, 19]$. The results are plotted in event time, with time 0 marking the policy change (time τ). Panels C and D are the counterparts of the same panels in Figure 10 for an exogenously given τ . As in Figure 10, both volatility and correlation jump up at the time of a policy change and decline afterwards. The magnitudes of these effects are almost as large as those in Figure 10.

²⁶We vary σ_g because the effect of σ_c is predictable and weaker than the effect of σ_g .

²⁷The magnitudes are even more negative for $\tau < 5$ but the simulation results then become less precise due to the small probability of policy changes for small values of τ . To keep the number of simulations manageable and the plot smooth rather than jagged, we do not plot the announcement returns for $\tau < 5$.

Finally, we examine the expected dynamics of profitability and \hat{g}_t around a policy change. In unreported results, we find that both profitability and \hat{g}_t decline before the policy change. This pattern indicates that policy changes tend to occur after downturns, as they do when τ is exogenous (cf. Panel A of Figure 1). To summarize, all of our key results remain unchanged when τ is endogenous.

5. Extension: Investment Adjustment

In the benchmark model analyzed in Sections 2. and 3., the firms' decision-making is not modeled explicitly. Any investment decisions taken by firms are assumed to be reflected in the profitability process in equation (1). This process might not adequately capture all investment decisions. For example, prompted by policy uncertainty, a firm might disinvest by shutting down its operations, making the process (1) obsolete. In this section, we extend the model by allowing each firm to make a major investment decision at time τ . To preview our results, we find that firms often reduce their investment in response to government-induced uncertainty. We also find that EAR remains negative, as in the benchmark model.

Between times 0 and τ , each firm is fully invested in its risky technology (equation (1)), as in the benchmark model. At time τ , each firm has the option to disinvest and switch all of its capital into a risk-free storage technology whose constant return is normalized to zero. If the firm decides to disinvest, its capital earns the zero rate of return from time τ until time T ; otherwise its profitability continues to follow equation (1). Partial investment in the two technologies is ruled out, for simplicity.

All firms make their investment decisions at the same time τ at which the government makes its policy decision. This simplifying assumption captures the idea that firms decide on their investment while facing uncertainty about the government's future policy, and the government decides on its policy while taking into account its impact on firm investment. As before, the government and the firms have the same beliefs about government policies (equations (3) and (8)). We solve for the *Nash equilibrium* in which the government's policy choice is optimal given the firms' investment decisions, and each firm's investment decision is optimal given the decisions of all other firms as well as the government's decision rule.

Proposition 9. *In equilibrium at time τ , a randomly-selected fraction $\alpha_\tau \in [0, 1]$ of firms continue investing in their risky technologies, while the remaining firms disinvest and park*

their capital in the risk-free technology. The government changes its policy if and only if

$$\widehat{g}_\tau < \underline{g}(c, \alpha_\tau) , \quad (49)$$

where the threshold $\underline{g}(c, \alpha_\tau)$ and the equilibrium value of α_τ are described below.

Similar to Proposition 2, Proposition 9 shows that a policy change occurs if \widehat{g}_τ is below a threshold, so that a policy is replaced if it is perceived to have a sufficiently unfavorable impact on profitability. A key difference from Proposition 2 is that the threshold $\underline{g}(c, \alpha_\tau)$ depends on the fraction of investing firms, α_τ , which in turn depends on \widehat{g}_τ .

To shed light on Proposition 9, we first take the government's perspective. The government makes its policy decision by taking the firms' investment decisions, α_τ , as given. Recall that a policy change occurs if and only if the condition (12) holds. This condition involves functions of the aggregate capital B_T . Given α_τ , the value of B_T is given by

$$\frac{B_T}{B_\tau} = \alpha_\tau e^{(\mu+g-\frac{\sigma^2}{2})(T-\tau)+\sigma(Z_T-Z_\tau)} + (1-\alpha_\tau) , \quad (50)$$

where $g \equiv g^{\text{old}}$ if there is no policy change and $g \equiv g^{\text{new}}$ if there is a policy change. Using equation (50), we find that the condition (12) holds if and only if

$$e^c E_\tau \left\{ [\alpha_\tau e^{\varepsilon^{yes}(\tau, T)} + (1-\alpha_\tau)]^{1-\gamma} \right\} < E_\tau \left\{ [\alpha_\tau e^{\varepsilon^{no}(\tau, T)} + (1-\alpha_\tau)]^{1-\gamma} \right\} , \quad (51)$$

where

$$\varepsilon^i(\tau, T) \sim N\left((\mu + g_i - \sigma^2/2)(T - \tau), (T - \tau)^2 \sigma_i^2 + \sigma^2(T - \tau)\right) , \quad (52)$$

for $i = yes, no$, and $g_{no} = \widehat{g}_\tau$, $\sigma_{no}^2 = \widehat{\sigma}_\tau^2$, $g_{yes} = 0$, $\sigma_{yes}^2 = \sigma_g^2$. The right-hand side of (51) decreases with \widehat{g}_τ , whereas the left-hand side does not depend on \widehat{g}_τ . Therefore, the condition (51) implies the cutoff rule in Proposition 9, where the cutoff $\underline{g}(c, \alpha_\tau)$ is the value of \widehat{g}_τ for which the left-hand side of (51) equals the right-hand side.

Next, we take the firms' perspective. At the time of their investment decisions, firms do not know whether the government's policy will change. Even though firms observe \widehat{g}_τ , they do not observe $\underline{g}(c, \alpha_\tau)$ (because they do not observe c), so they are unable to fully anticipate the government's action. Firms invest in a way that maximizes their market value. Market value obeys equation (17), where the state price density π_t in equation (16) is computed based on equation (50). If firm i decides to switch to the risk-free technology at time τ , then its market value is $M_\tau^i = B_\tau^i$ (because $B_T^i = B_\tau^i$, so that $E_\tau \left[\frac{\pi_T}{\pi_\tau} B_T^i \right] = E_\tau \left[\frac{\pi_T}{\pi_\tau} \right] B_\tau^i = B_\tau^i$). Therefore, each firm chooses to remain invested in the risky technology if and only if doing so results in a market-to-book ratio greater than one at time τ .

If firm i decides at time τ to remain invested in its risky technology, its market value right after time τ is given by

$$M_{\tau+}^i = \begin{cases} M_{\tau+}^{i,yes} = B_{\tau+}^i \frac{\mathbb{E}_{\tau+} [e^{\varepsilon^{yes}(\tau,T)} [\alpha_{\tau} e^{\varepsilon^{yes}(\tau,T)} + (1-\alpha_{\tau})]^{-\gamma}]}{\mathbb{E}_{\tau+} [[\alpha_{\tau} e^{\varepsilon^{yes}(\tau,T)} + (1-\alpha_{\tau})]^{-\gamma}]} & \text{if policy changes} \\ M_{\tau+}^{i,no} = B_{\tau+}^i \frac{\mathbb{E}_{\tau+} [e^{\varepsilon^{no}(\tau,T)} [\alpha_{\tau} e^{\varepsilon^{no}(\tau,T)} + (1-\alpha_{\tau})]^{-\gamma}]}{\mathbb{E}_{\tau+} [[\alpha_{\tau} e^{\varepsilon^{no}(\tau,T)} + (1-\alpha_{\tau})]^{-\gamma}]} & \text{if policy does not change} \end{cases} \quad (53)$$

Right before time τ , the market value of firm i is given by a weighted average of $M_{\tau+}^{i,yes}$ and $M_{\tau+}^{i,no}$ as in equation (23), except that the weights ω_{τ} are given in equation (A9) in the Appendix. The values of M_{τ}^i/B_{τ}^i are equal across i because all firms are identical ex ante. Therefore, if $M_{\tau}^i/B_{\tau}^i > 1$, all firms invest in their risky technologies, and $\alpha_{\tau} = 1$ in equilibrium. For an interior solution $0 < \alpha_{\tau} < 1$ to occur, firms must be indifferent between the risky and risk-free technologies, so that $M_{\tau}^i/B_{\tau}^i = 1$ for all i . Combining this condition with equation (23), we obtain the condition for equilibrium with $0 < \alpha_{\tau} < 1$.²⁸

$$1 = \omega_{\tau} \left(\frac{M_{\tau+}^i}{B_{\tau+}^i} \right)^{yes} + (1 - \omega_{\tau}) \left(\frac{M_{\tau+}^i}{B_{\tau+}^i} \right)^{no} \quad (54)$$

The right-hand-side of equation (54) depends on α_{τ} , which affects both ω_{τ} and the M/B ratios. If there exists a value of α_{τ} strictly between 0 and 1 for which equation (54) holds, then this is the equilibrium value of α_{τ} in Proposition 9. If instead the right-hand side of equation (54) exceeds one for all $\alpha_{\tau} \in [0, 1]$, then $\alpha_{\tau} = 1$ in equilibrium, and all the results from Sections 2. and 3. hold. If the right-hand side of equation (54) is smaller than one for all $\alpha_{\tau} \in [0, 1]$, then the equilibrium features $\alpha_{\tau} = 0$, an uninteresting case of no investment.

For the parameter values in Table 1, the equilibrium solution is $\alpha_{\tau} = 1$ (no disinvestment), so that all the results from the benchmark model apply also to the problem analyzed here. To analyze the effect of disinvestment, we reduce the value of μ from its benchmark value of 10% to 2%, which is the largest integer percentage value for which $\alpha_{\tau} < 1$ in equilibrium (i.e., for all $\mu \geq 3\%$, the equilibrium has $\alpha_{\tau} = 1$). All other parameter values are in Table 1. We set the downturn length to five years, as before. We solve the problem numerically and plot the results in Figure 14.

Panel A of Figure 14 reports the equilibrium value of α_{τ} as a function of σ_c and σ_g . We see that higher values of σ_c or σ_g imply lower values of α_{τ} . Fixing $\sigma_c = 10\%$ and increasing σ_g from 1% to 2% to 3%, the fraction of firms that remain invested in the risky technology decreases from 1 to 0.87 to 0.77, respectively. The effect of σ_c is weaker: fixing $\sigma_g = 2\%$, α_{τ}

²⁸The same condition is obtained as a first-order condition in an alternative formulation of the problem in which a social planner chooses α_{τ} to maximize the investors' expected utility. See the Technical Appendix.

decreases from 0.88 to 0.85 as we increase σ_c from 0 to 20%. In short, both policy uncertainty and political uncertainty reduce aggregate investment.

Panel B of Figure 14 plots EAR as a function of σ_c and σ_g . The plot is similar to that in Panel A of Figure 7, which corresponds to the benchmark model (in which we force $\alpha_\tau = 1$). For $\sigma_g = 1\%$, EAR is the same as in Figure 7 because $\alpha_\tau = 1$ in equilibrium (see Panel A of Figure 14). For $\sigma_g = 2\%$ and 3% , EAR is more negative than for $\sigma_g = 1\%$ but less negative than its counterparts in Figure 7. For example, for $\sigma_g = 2\%$ and $\sigma_c = 10\%$, EAR is -0.6% compared to -1.1% in Figure 7, and the difference from Figure 7 is even larger for $\sigma_g = 3\%$. The reason is that a substantial fraction of firms disinvest in equilibrium ($\alpha_\tau < 1$), thereby reducing aggregate risk. By reducing their investment, firms essentially undo some of the uncertainty associated with government policy. Nonetheless, EAR remains negative for all values of σ_c and σ_g , as in the benchmark model.

Panels C and D of Figure 14 plot the expected dynamics of return volatility, $\sigma_{M,t}$, and correlation, ρ_t , respectively, around a policy change at time $\tau = 10$. Both quantities exhibit patterns very similar to those in Panels C and D of Figure 10, which correspond to the benchmark model ($\alpha_\tau = 1$). The post- τ increases in $\sigma_{M,t}$ and ρ_t are slightly smaller than in Figure 10, as one would expect as a result of a reduction in aggregate risk, but the conclusions are exactly the same. In short, the key asset pricing results from the benchmark model hold also when we allow for disinvestment.

6. Extension: Different Policy Exposures

Government policies typically affect all firms, but they may affect some firms more than others. In this section, we extend our benchmark model by allowing firms to have different exposures to government policy. We lose many closed-form solutions and propositions, but we gain some additional insight into the cross-section of expected stock returns. In particular, we find that firms with larger government exposures tend to have higher expected returns. The government risk premium can be negative, though, and it is highly state-dependent. We also find that EAR remains negative, as in the benchmark model.

Firms are divided into N sectors, $n = 1, \dots, N$, where N is finite. Each sector n is characterized by a nonnegative “government beta” β^n , which represents a loading on the policy impact. Specifically, the profitability of each firm i in sector n is given by

$$d\Pi_t^i = (\mu + \beta^n g_t) dt + \sigma dZ_t^n + \sigma_1 dZ_t^i, \quad (55)$$

where dZ_t^n is a sector-specific shock, dZ_t^i is a firm-specific shock, and all Brownian motions are independent of each other.²⁹ Equation (55) replaces equation (1) in the benchmark model from Section 2.; the rest of the model remains unchanged. The benchmark model is thus a special case of this model with $\beta^n = 1$ and $N = 1$. Similar to Proposition 1, by observing each firm's profitability, all agents effectively receive N signals

$$ds_t^n = (\mu + \beta^n g_t) dt + \sigma dZ_t^n, \quad n = 1, \dots, N. \quad (56)$$

Let B_t denote the total capital of all firms and B_t^n denote the total capital of all firms in sector n . We denote sector n 's share of total capital by w_t^n , so that

$$w_t^n = \frac{B_t^n}{B_t}, \quad n = 1, \dots, N. \quad (57)$$

The $N \times 1$ vector of these shares, which we denote by \mathbf{w}_t , represents additional state variables introduced by the presence of multiple sectors.

Proposition 10. *The government changes its policy at time τ if and only if*

$$\widehat{g}_\tau < \underline{g}(c, \mathbf{w}_\tau), \quad (58)$$

where the threshold $\underline{g}(c, \mathbf{w}_\tau)$ is described in the Technical Appendix.

Proposition 10 shows that a policy is replaced if its impact on profitability is perceived as sufficiently unfavorable. This result is similar to Proposition 2, except that the threshold for \widehat{g}_τ now depends also on the vector \mathbf{w}_t of the sector shares.

We solve for stock prices and their dynamics before time τ by using numerical methods. The pricing results after time τ are in closed form. Those results highlight the importance of the new state variable \mathbf{w}_t , which affects not only the level of stock prices but also their volatility as well as the expected rates of return. All of the asset pricing results, along with the results on learning, are presented and proved in the Technical Appendix.

Figure 15 plots various results obtained in a two-sector framework (i.e., $N = 2$). We denote the two sectors by H and L . The government betas of the two sectors are β^H and β^L , where $\beta^H \geq \beta^L$. We vary these betas while holding their average equal to one. In the baseline case, explored in Panel A, we set $\beta^H = 1.5$ and $\beta^L = 0.5$. We also vary the relative size of the high-beta sector, which we denote by $w_t = B_t^H / (B_t^H + B_t^L)$. In the baseline case, examined in Panel A, we set $w_0 = 0.5$, so that as of time 0, half the firms have high betas and half have low betas. Unless specified otherwise, the parameters are from Table 1.

²⁹We could make the dZ_t^n shocks correlated at the cost of additional analytical complexity. Given our interest in the cross-sectional differences in stock returns as opposed to their level, such cross-correlation seems to be of second-order importance, so we focus on the simpler case of no correlation.

6.1. The Expected Announcement Return

Panel A of Figure 15 plots EAR, obtained by simulations for the baseline case, as a function of σ_g and σ_c . This panel is the counterpart of Figure 4, which corresponds to the benchmark model. As in Figure 4, EAR is always negative and its magnitude is increasing in both σ_g and σ_c . The magnitudes of EARs are also close to those from Figure 4.

Panel B of Figure 15 plots EAR as a function of $\beta^H - \beta^L$, for three different values of w_0 . The values of σ_g and σ_c are fixed at their baseline values from Table 1. When $\beta^H - \beta^L = 0$, both sectors have the same government beta of one. EAR is then -0.6%, which is close to the corresponding value of -0.5% in the benchmark model (Figure 4). These two values are different even though all firms have a beta of one in both models. The reason is that the speeds of learning in the two models are different. The two-sector model has two signals about g_t (equation (56)), whereas there is only one signal in the benchmark model (equation (7)). As a result, learning in the two-sector model is faster.³⁰ Therefore, the difference between the prior and posterior variances is larger, and the discount rate effect is stronger; hence the slightly more negative EAR. Panel B also shows that EAR is more negative for larger values of w_0 . For example, for $\beta^H - \beta^L = 1$, EAR is -0.4% for $w_0 = 0.2$, -0.6% for $w_0 = 0.5$, and -0.8% for $w_0 = 0.8$. When w_0 is larger, the high-beta firms account for a larger fraction of the aggregate market. As a result, the risk associated with learning about g_t is harder to diversify and shocks to \hat{g}_t are more highly correlated with the stochastic discount factor. The increase in uncertainty at time τ thus generates a larger discount rate effect, leading to a more negative EAR when w_0 is larger. Finally, it is useful to consider the case of $\beta^H - \beta^L = 2$, in which $\beta^H = 2$ and $\beta^L = 0$. In the limiting case $w_0 \rightarrow 0$, all firms have zero beta; as a result, policy changes become irrelevant and $\text{EAR} \rightarrow 0$.

6.2. Government Betas and Expected Stock Returns

Given the government's pervasive effect on the private sector, policy uncertainty is unlikely to be fully diversifiable. It thus seems natural to ask whether firms that are more exposed to policy uncertainty offer higher expected returns. In this subsection, we relate government betas to expected stock returns. The instantaneous expected stock returns are given by the drifts of the stocks' return processes. We denote these drifts for sectors H and L at time t by μ_t^H and μ_t^L , respectively.

³⁰We can prove that learning in the multiple-sector model is faster than in the benchmark model under the condition $\beta' \beta > 1$, where β is an $N \times 1$ vector of β^m 's. This condition is trivially satisfied here.

Panels C and D of Figure 15 plot $\mu_t^H - \mu_t^L$ as a function of $\beta^H - \beta^L$. We consider three different values of w_t and set $t = 5$ years (which is half-way between time 0 and time τ). We set w_t rather than w_0 as in the previous paragraph because our focus is on the conditional expected return at time t rather than EAR as of time 0.

Panels C and D show that as the share of sector H increases, so does $\mu_t^H - \mu_t^L$, holding $\beta^H - \beta^L$ constant. There are two reasons for this result. First, a larger w_t makes the \hat{g}_t shocks more powerful, as explained above. As a result, a higher government beta has a larger effect on the expected return. Second, there is an interesting size effect even when both sectors have equal government betas ($\beta^H - \beta^L = 0$). Both panels show that equal betas imply equal risk premia (i.e., $\mu_t^H - \mu_t^L = 0$) only when both sectors are equally large ($w_t = 0.5$). When $w_t \neq 0.5$, the risk premia are different. For example, $w_t = 0.8$ implies $\mu_t^H - \mu_t^L = 0.7\%$ per year, and $w_t = 0.2$ implies $\mu_t^H - \mu_t^L = -0.7\%$. Firms in the larger sector command a higher risk premium because their capital B_t^i covaries more closely with aggregate capital B_t , and the higher covariance makes the larger-sector firms riskier. This size effect can make $\mu_t^H - \mu_t^L$ negative even when $\beta^H - \beta^L$ is positive if w_t is low; we observe such instances in both panels when $w_t = 0.2$.

To abstract from the size effect, suppose both sectors are equally large ($w_t = 0.5$). In that case, Panels C and D always show $\mu_t^H - \mu_t^L > 0$, indicating that the higher-beta firms earn higher expected returns. In addition, $\mu_t^H - \mu_t^L$ is generally increasing in $\beta^H - \beta^L$. The only exception occurs in Panel D for large values of $\beta^H - \beta^L$. The reason is that in Panel D, we have $\hat{g}_t = -2\%$. With $\hat{g}_t < 0$, the H sector is expected to shrink relative to the L sector (because it loads more heavily on \hat{g}_t), especially when $\beta^H - \beta^L$ is large. As a result, the expected future impact of \hat{g}_t is diminished, reducing the impact of the \hat{g}_t shocks on stock prices when $\hat{g}_t < 0$ and $\beta^H - \beta^L$ is large. With this caveat, we can state that larger differences in government betas generally translate into larger differences in expected returns.

Finally, a comparison of Panels C and D shows that $\mu_t^H - \mu_t^L$ is substantially larger for $\hat{g}_t = 2\%$ (Panel C) than for $\hat{g}_t = -2\%$ (Panel D). The \hat{g}_t shocks are less powerful in Panel D, for two reasons. The first reason, related to the expected shrinkage of the H sector, is explained in the previous paragraph. The second reason is that the probability of a policy change is higher in Panel D. When $\hat{g}_t = 2\%$, the prevailing policy is likely to be retained; as a result, the \hat{g}_t shocks are likely to be permanent. In contrast, the same shocks have a shorter expected duration when $\hat{g}_t = -2\%$ because the policy change is then more likely. Since the \hat{g}_t shocks are expected to be longer-lasting when \hat{g}_t is higher, the risk premium associated with exposure to these shocks is larger as well. Overall, Panels C and D show that

the premium for government beta risk is typically positive but also highly state-dependent.

In related empirical work, Belo, Gala, and Li (2011) find that firms with higher exposures to the government sector (measured by the fraction of sales generated by government spending) earn higher average stock returns, but only during Democratic presidential terms; the relation flips during Republican terms.³¹ The match between their analysis and ours is not perfect. They measure government spending, whereas we focus on government policy more generally. They measure return exposure to government, whereas our government betas are profitability exposures. They study Democrats versus Republicans, whereas our simple model does not distinguish between different types of government. Despite these differences, it seems interesting that they estimate a state-dependent government risk premium. More research into this state dependence is clearly warranted.

7. Conclusions

We conduct a theoretical analysis of the effects of changes in government policy on stock prices. Our simple general equilibrium model makes numerous testable predictions. Stock market returns at the announcements of policy changes should be negative unless the policy being replaced is perceived as sufficiently harmful to profitability. Averaging across all policy changes, the expected announcement return should be negative. The magnitude of this negative return should be large if uncertainty about government policy is large, as well as if the policy change is preceded by a short or shallow downturn. The distribution of stock returns at the announcements of policy changes should be left-skewed. Policy changes should make stock returns more volatile and more highly correlated across firms. The average stock market return at the announcements of policy decisions, without conditioning on whether these decisions change the policy or not, should be positive.

Our key result—that stock prices are expected to fall at the announcement of a policy change—hinges on our assumption that the government is quasi-benevolent. Due to that assumption, policy changes that raise stock prices are largely anticipated. Policy changes that reduce prices contain a larger element of surprise, resulting in larger negative announcement returns. The result might flip around if the government were instead perceived as malevolent because policy changes that benefit investors would then be largely unexpected and thus associated with large positive announcement returns. In reality, the government's

³¹In another empirical study of political cycles, Santa-Clara and Valkanov (2003) find that the average stock market return is higher under Democratic presidencies than under Republican presidencies.

objectives are surely more complex than benevolence or malevolence, but they are likely to have a benevolent tilt due to the mechanics of a democratic process. For example, the government might maximize the probability of reelection, which is imperfectly but positively related to the agents' material well-being.³² Analyzing alternative objective functions for the government in the context of our model is an interesting direction for future research.

Our model makes the simplifying assumption that the prevailing government policy can only be replaced by a single policy that is identical a priori. It would be useful to extend the model to allow the government to choose from multiple policies. The policy decision would then resolve the uncertainty about which policy is chosen. As long as all potential new policies are identical a priori, knowing which policy is chosen is irrelevant, and the multiple-policy setting collapses to our single-policy setting. The convergence to a single-policy setting obtains also at the opposite extreme, if the potential new policies are so different from each other that one of them dominates a priori. In the intermediate cases, though, the results might depend on the prior heterogeneity in interesting ways. We are currently analyzing the multiple-policy setting in a work in progress.

Another interesting extension would endogenize the political cost. We treat this cost as exogenous, in the interest of simplicity, but it would be useful to build the microfoundations for it based on the insights from the political economy literature. Such an extension could establish new theoretical links between the financial markets and various political economy variables. Another theoretical extension could make the government and the investors asymmetrically informed. While our focus is on stocks, future work can also investigate the effects of uncertainty about government policy on the prices of other assets, such as bonds.

There is also need for empirical work. The effects of policy changes on asset prices have been analyzed in various contexts, with mixed results.³³ A broader analysis of policy changes, or reforms, would be beneficial. One challenge for the empirical work is the timing of the policy decision. In some cases, this timing is clear.³⁴ However, most reforms occur

³²See Downs (1957) for an early model in which the government maximizes the probability of reelection. After surveying the empirical literature, Lewis-Beck and Stegmaier (2000) conclude that economic indicators explain most of the variance in government support. A more recent survey by Anderson (2007) finds that the relation between economic outcomes and election outcomes is imperfect and complicated. Brender and Drazen (2008) find that this relation is significant only among the less developed countries.

³³For example, Thorbecke (1997), Rigobon and Sack (2004), and Bernanke and Kuttner (2005) examine the effects of changes in monetary policy on stock prices. Klingebiel, Kroszner, Laeven, and Oijen (2001) examine stock market responses to the announcements of bank restructuring policies by East Asian governments during the 1997-98 crisis. Ait-Sahalia et al (2010) investigate the effects of policy changes during the recent financial crisis on interbank credit and liquidity risk premia.

³⁴One example is the abrupt shift in the too-big-to-fail policy, mentioned in the introduction. The U.S. government's decision to let Lehman Brothers fail was followed by a 4.7% drop in the S&P 500 index on

gradually, clearing multiple hurdles. For example, the ongoing reform of U.S. financial regulation has involved separate passages of related bills by the House and the Senate, a subsequent reconciliation, and a continuing emergence of specific regulatory measures. According to our model, stock prices should respond at each step of the way, with bigger price responses following bigger increases in the probability of a policy change.³⁵ Identifying the pivotal steps of the key reforms is a nontrivial task. For example, Cutler (1988) examines the stock market reaction to the Tax Reform Act of 1986. He identifies two key events in the bill’s progress—the vote by the House for the bill in December 1985, and the vote for a similar bill by the Senate Finance Committee in May 1986—and argues that both events surprised the market.³⁶ With the whole paper in the *American Economic Review* devoted to a single policy change, it seems clear that a careful empirical analysis of multiple policy changes is well beyond the scope of this study. It also seems clear that such analysis would be interesting. Our model provides a simple benchmark for the empirical evidence. Overall, we have much to learn about the role of the government in asset pricing.

September 15, 2008. Another example, fresh as of this writing, is Germany’s surprise announcement of a partial ban on naked short-selling late on May 18, 2010. This announcement was followed by roughly 2% drops in the European share indices in the following morning (e.g., Germany’s DAX dropped 1.9%). This drop is consistent with our model, but not with theories predicting that short-sale constraints lift asset prices.

³⁵A key step in the financial regulation reform took place on Thursday May 20, 2010. At 2.31pm, the Senate voted by the narrowest possible margin (60 to 40) to overcome filibusters and send the bill to a final vote. A CNN Money article on the same afternoon read: “Wall Street reform cleared a crucial test vote on Thursday, all but assuring final Senate passage of the most sweeping regulatory overhaul since the New Deal.” By the end of the day, the S&P 500 index fell 3.9%, with about half of the decline occurring in the last hour. The Senate passed the reform bill at 8.25pm on the same day. The following morning, the U.S. stock market fell about 1% in the first minute of trade, followed by a rebound later in the day.

³⁶Cutler reports that on the day after the late-night favorable House vote, the S&P 500 index fell by 0.40%; on the day after the late-night favorable Senate Finance Committee vote, the index fell by 0.49%.

Table 1
Parameter Choices

This table reports the parameter values used in the simulations. All variables are reported on an annual basis (except for γ , which denotes risk aversion).

σ_g	σ_c	μ	σ	σ_1	T	τ	γ
0.02	0.10	0.10	0.05	0.10	20	10	5




Figure 1. Profitability dynamics and the policy decision. This figure plots the expected dynamics of \hat{g}_t (solid line) and realized profitability (dashed line) under the parameter values from Table 1. Realized profitability is the average profitability across all firms, plotted in excess of μ so that its unconditional mean is zero. The figure also plots the threshold $g(c)$ (dotted line). All three lines are average paths across many simulated samples. The top panel averages across the samples in which a policy change occurred at time $\tau = 10$, whereas the bottom panel conditions on no policy change.

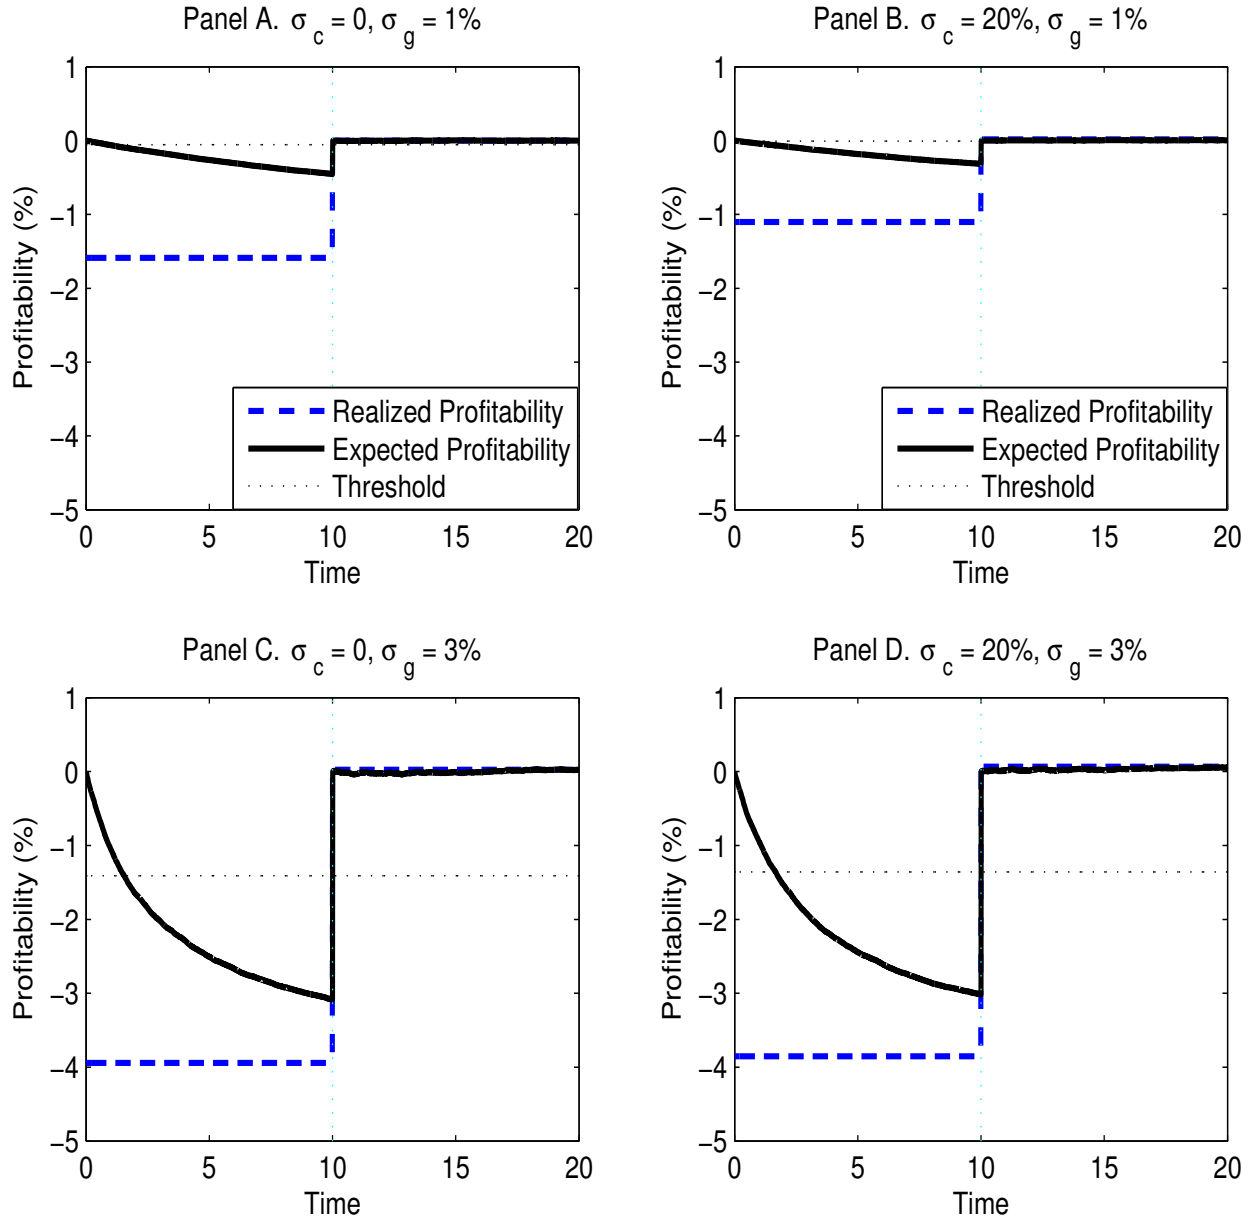


Figure 2. Profitability dynamics conditional on a policy change: The roles of policy uncertainty and political uncertainty. This figure plots the expected dynamics of \hat{g}_t (solid line) and realized profitability (dashed line) conditional on a policy change at time $\tau = 10$. Realized profitability is the average profitability across all firms, plotted in excess of μ so that its unconditional mean is zero. The figure also plots the threshold $\underline{g}(c)$ (dotted line). All three lines are average paths across many simulated samples in which a policy change occurred at time τ . The parameter values are in Table 1 except for policy uncertainty (σ_g) and political uncertainty (σ_c), which vary across the four panels.

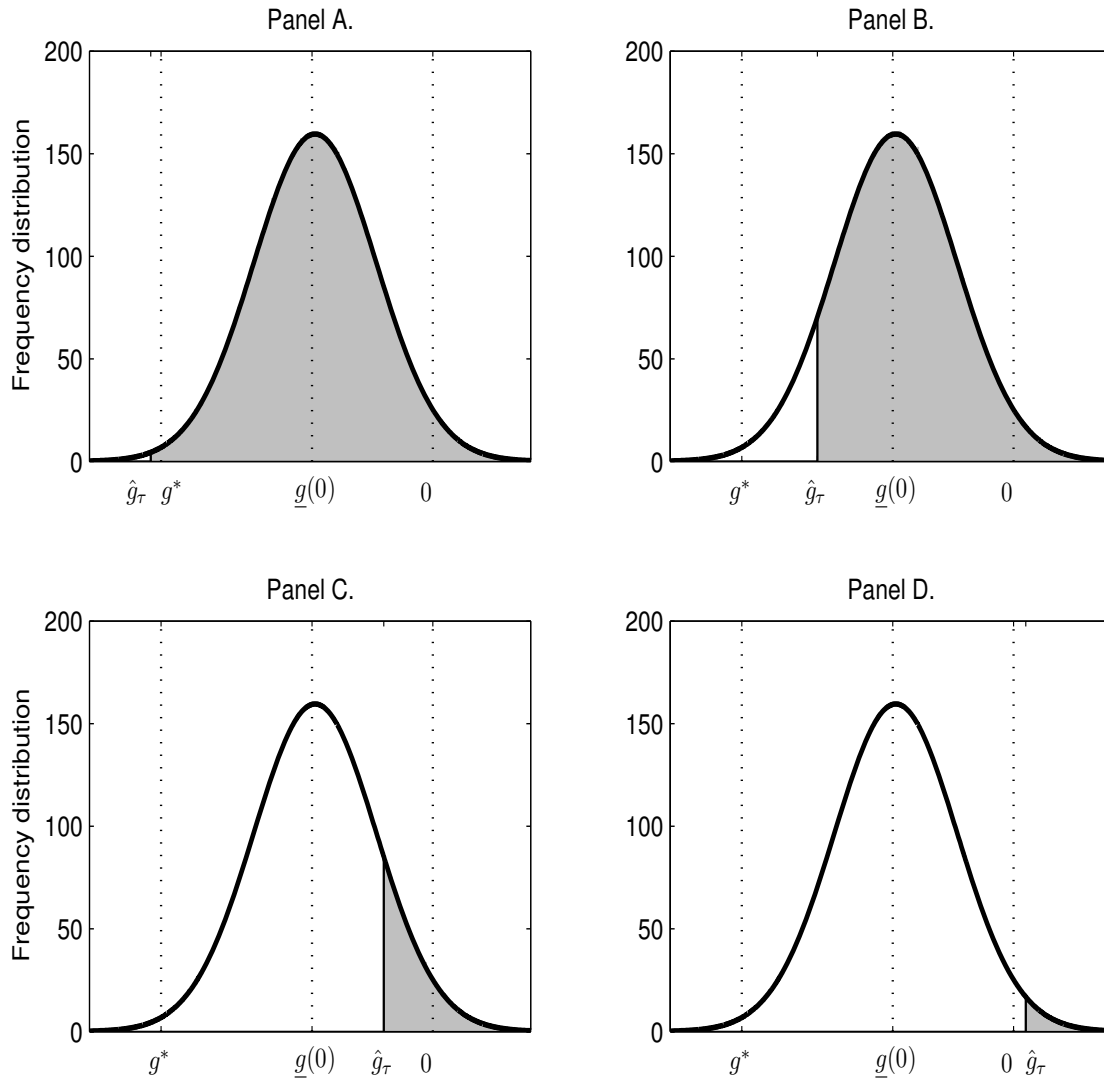


Figure 3. Probability of a policy change. The shaded area represents the probability of a policy change, as perceived by the investors just before time τ . The bell curve represents the normal distribution of the random threshold $g(c)$. The four panels illustrate four possible locations of \hat{g}_τ relative to g^* , $\underline{g}(0)$, and zero. The vertical dotted lines are drawn at g^* , $\underline{g}(0)$, and zero. The normal distribution as well as the values of g^* and $\underline{g}(0)$ are computed based on the parameter values in Table 1.




Figure 4. Expected announcement return. This figure plots the expectation of the announcement return $R(\hat{g}_\tau)$, which is the instantaneous stock return at the announcement of a policy change at time τ . The expectation integrates out uncertainty about \hat{g}_τ as perceived at time 0, as well as uncertainty about c conditional on the policy being changed. We vary σ_g and σ_c , and all other parameters are from Table 1.




Figure 5. Announcement return and the downturn length and depth. Panel A plots the announcement return $R(\hat{g}_\tau)$, the instantaneous stock return at the announcement of a policy change at time τ , as a function of the length and depth of the downturn that precedes the policy change. The downturn length is computed as $\tau - t_0 > 0$, where $\hat{g}_{t_0} = 0$. The downturn depth is given by the number of standard deviations by which \hat{g}_t drops during the downturn. Panel B plots the corresponding probability of a policy change, as perceived by investors just before time τ . The parameters are from Table 1.

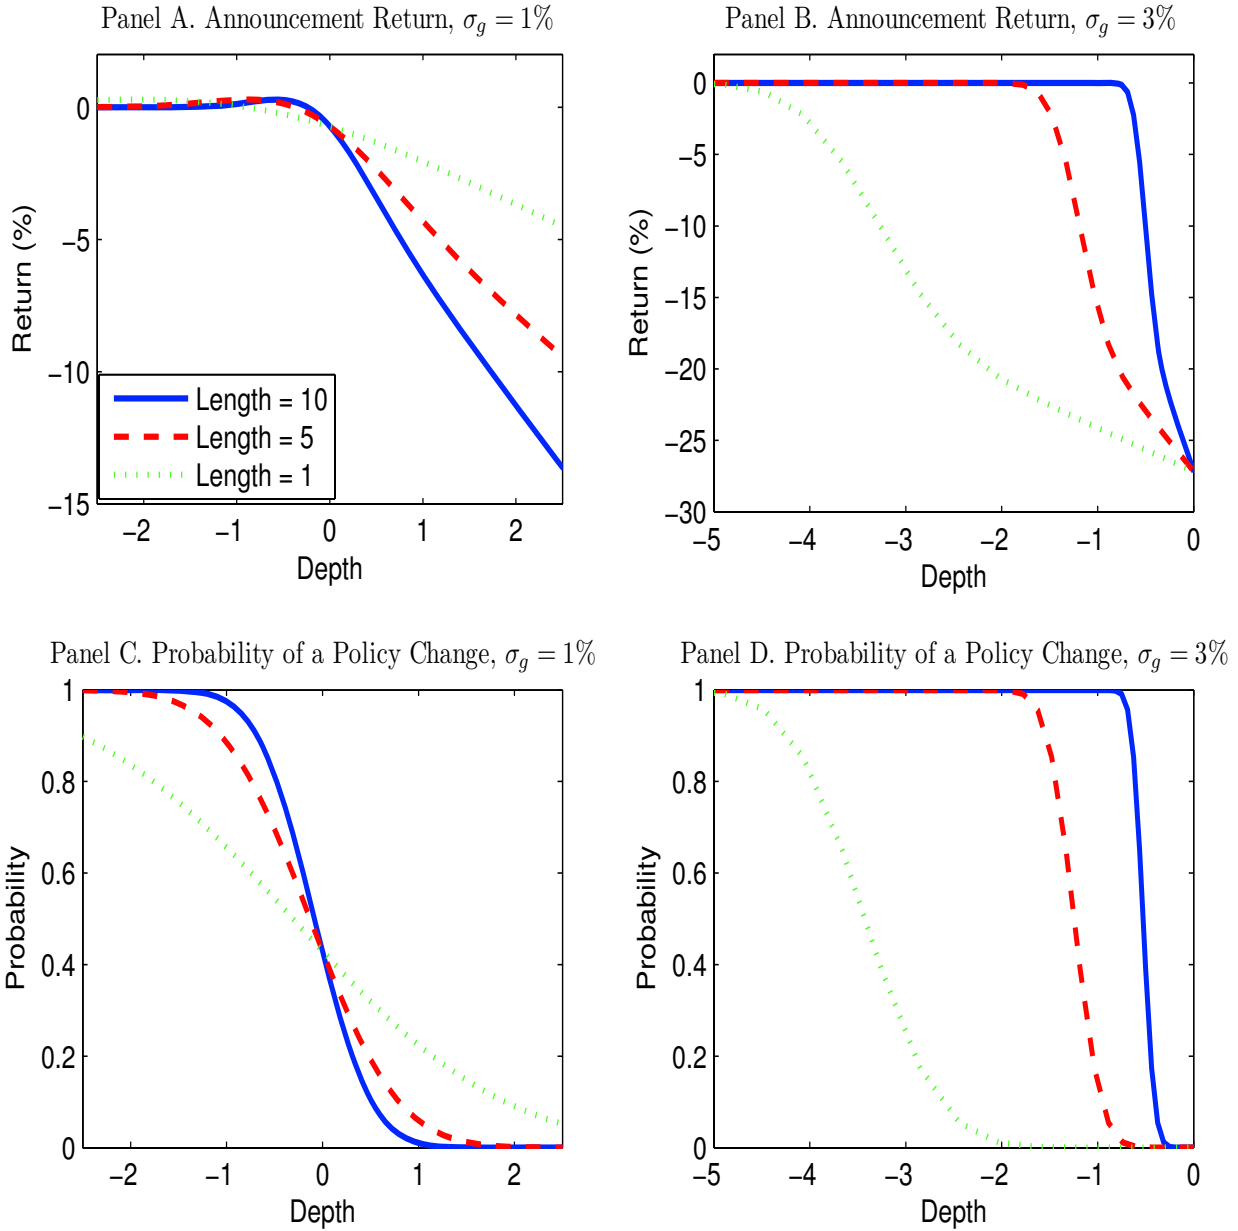


Figure 6. Announcement return and the downturn length and depth: The role of policy uncertainty. Panels A and B plot the announcement return $R(\hat{g}_\tau)$, the instantaneous stock return at the announcement of a policy change at time τ , as a function of the length and depth of the downturn that precedes the policy change. The downturn length is computed as $\tau - t_0 > 0$, where $\hat{g}_{t_0} = 0$. The downturn depth is given by the number of standard deviations by which \hat{g}_t drops during the downturn. Panels C and D plot the corresponding probabilities of a policy change, as perceived by investors just before time τ . The parameters are from Table 1, except that the baseline value $\sigma_g = 2\%$ is replaced by $\sigma_g = 1\%$ (Panels A and C) and $\sigma_g = 3\%$ (Panels B and D).

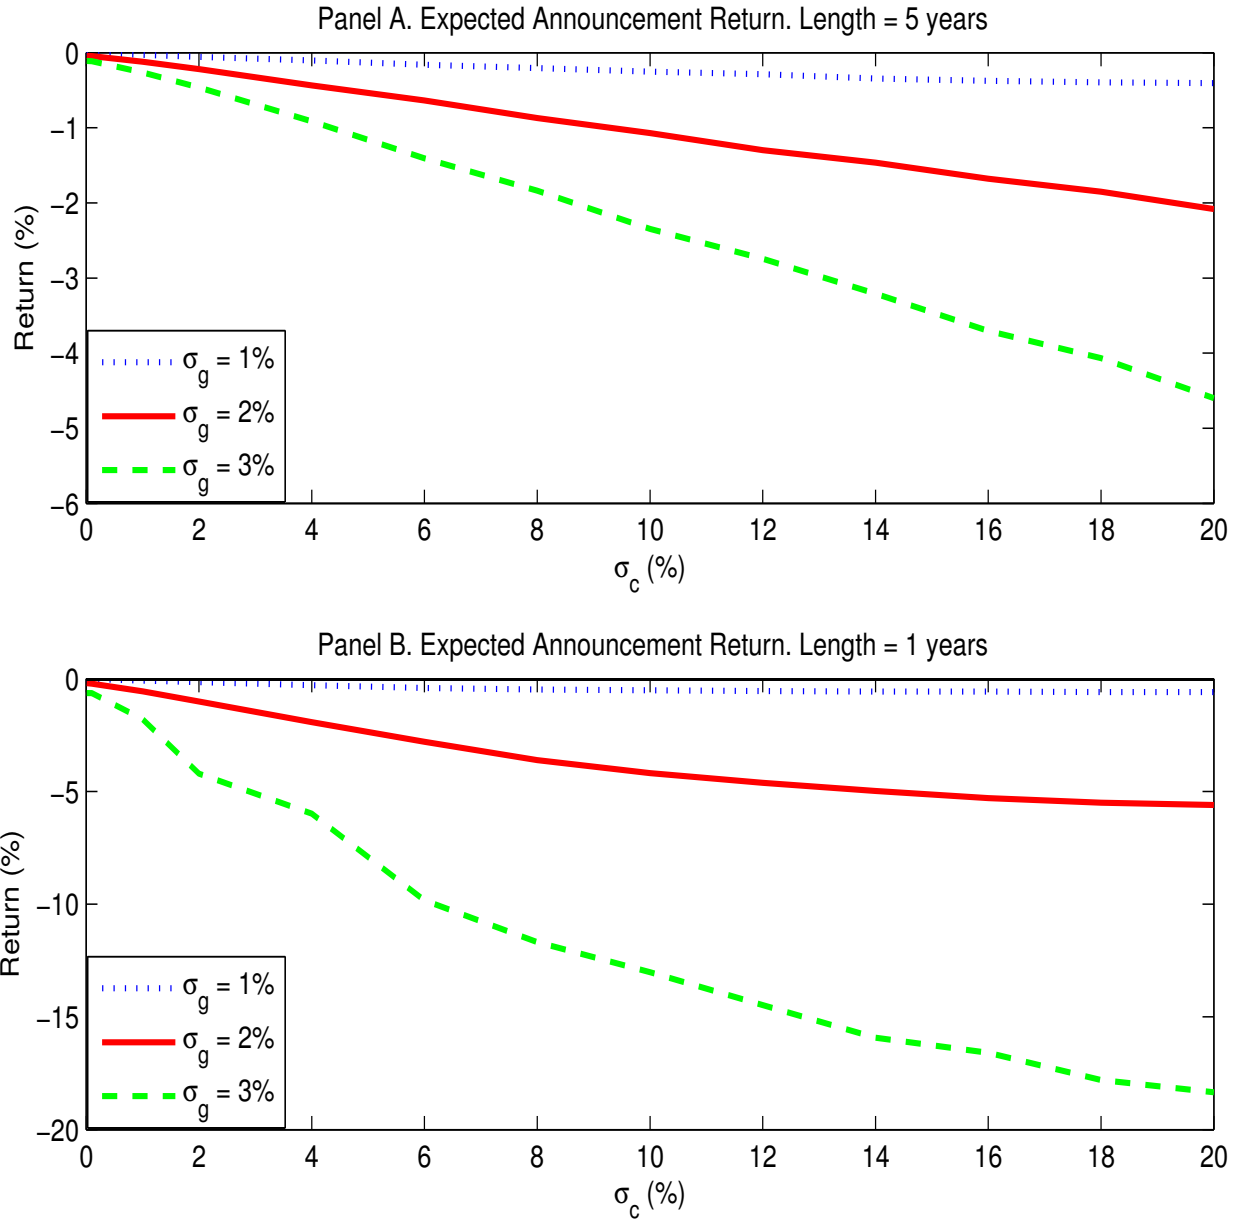


Figure 7. Expected announcement return and the downturn length. This figure plots the expectation of the announcement return $R(\hat{g}_\tau)$, which is the instantaneous stock return at the announcement of a policy change at time τ . The expectation integrates out uncertainty about \hat{g}_τ as perceived at the beginning of the downturn, as well as uncertainty about c conditional on the policy being changed. The downturn length is $\tau - t_0 = 5$ years in Panel A and 1 year in Panel B. We vary σ_g and σ_c , and all other parameters are from Table 1.




Figure 8. Probability distribution of stock returns on the day of the announcement of a policy change. This figure plots the probability distribution of stock returns on the day when a policy change is announced. These per-day stock returns have two components: a jump component $R(\hat{g}_\tau)$ pertaining to the instant of the announcement, and the diffusion component $M_\tau/M_{\tau-\Delta\tau} - 1$ covering the rest of the announcement day ($\Delta\tau = 1/252$ years). This distribution is comparable with an empirical distribution of daily announcement returns. The downturn length is $\tau - t_0 = 5$ years. We vary σ_g and σ_c , and all other parameters are from Table 1.




Figure 9. Expected return at the announcement of a policy decision. Panel A plots the expected value of the jump in stock prices at the announcement of a policy decision, without conditioning on whether the decision is to change the policy or not. This value, $E(J_M)$, represents the unconditional risk premium investors demand for facing a jump in SDF at the announcement. Panel B reports the expected jump conditional on the announcement of a policy change, $E(J_M^{yes})$. This panel coincides with Panel A of Figure 7. Panel C reports the expected jump conditional on the announcement of *no* policy change, $E(J_M^{no})$. All three expectations integrate out uncertainty about \hat{g}_τ as perceived at the beginning of the downturn. The downturn length is $\tau - t_0 = 5$ years. We vary σ_g and σ_c , and all other parameters are from Table 1.

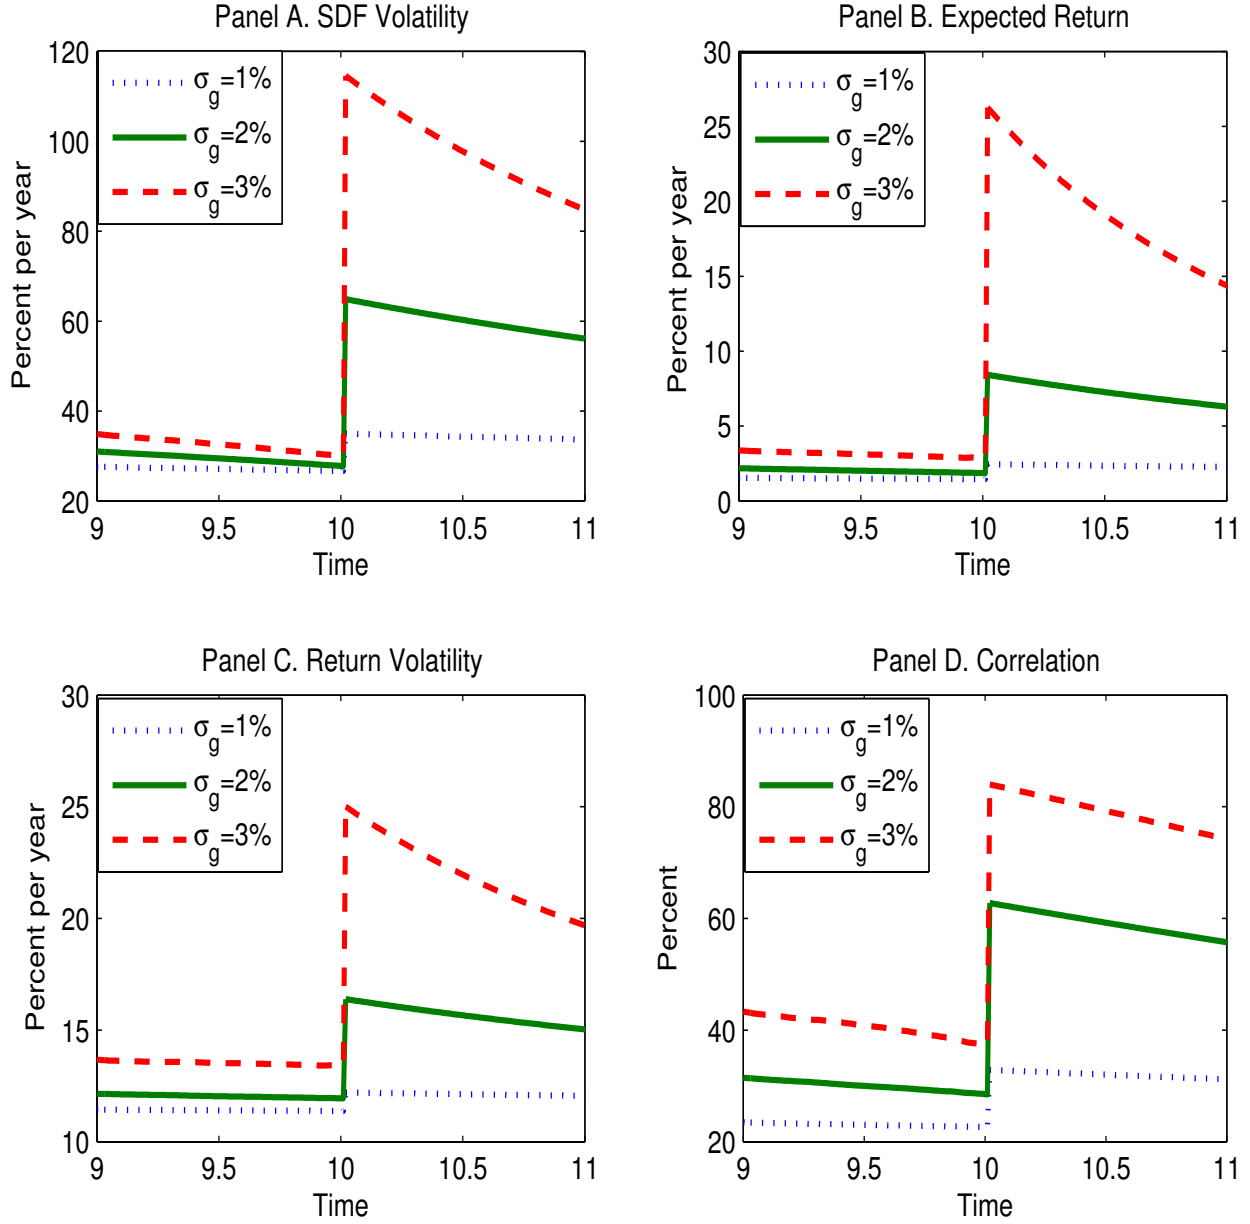


Figure 10. Properties of returns around policy decisions that result in a policy change. This figure plots the expected dynamics of the volatility of the stochastic discount factor, $\sigma_{\pi,t}$ (Panel A), the conditional expected stock market return, $\mu_{M,t}$ (Panel B), the conditional volatility of stock market returns, $\sigma_{M,t}$ (Panel C), and the pairwise correlation between stocks, ρ_t (Panel D), all conditional on a policy change at time $\tau = 10$. The jump-related components at time τ are not plotted. The downturn length is $\tau - t_0 = 5$ years, so that $\hat{g}_5 = 0$. The parameters are in Table 1, except for σ_g , which takes three different values, 1%, 2%, and 3% per year.

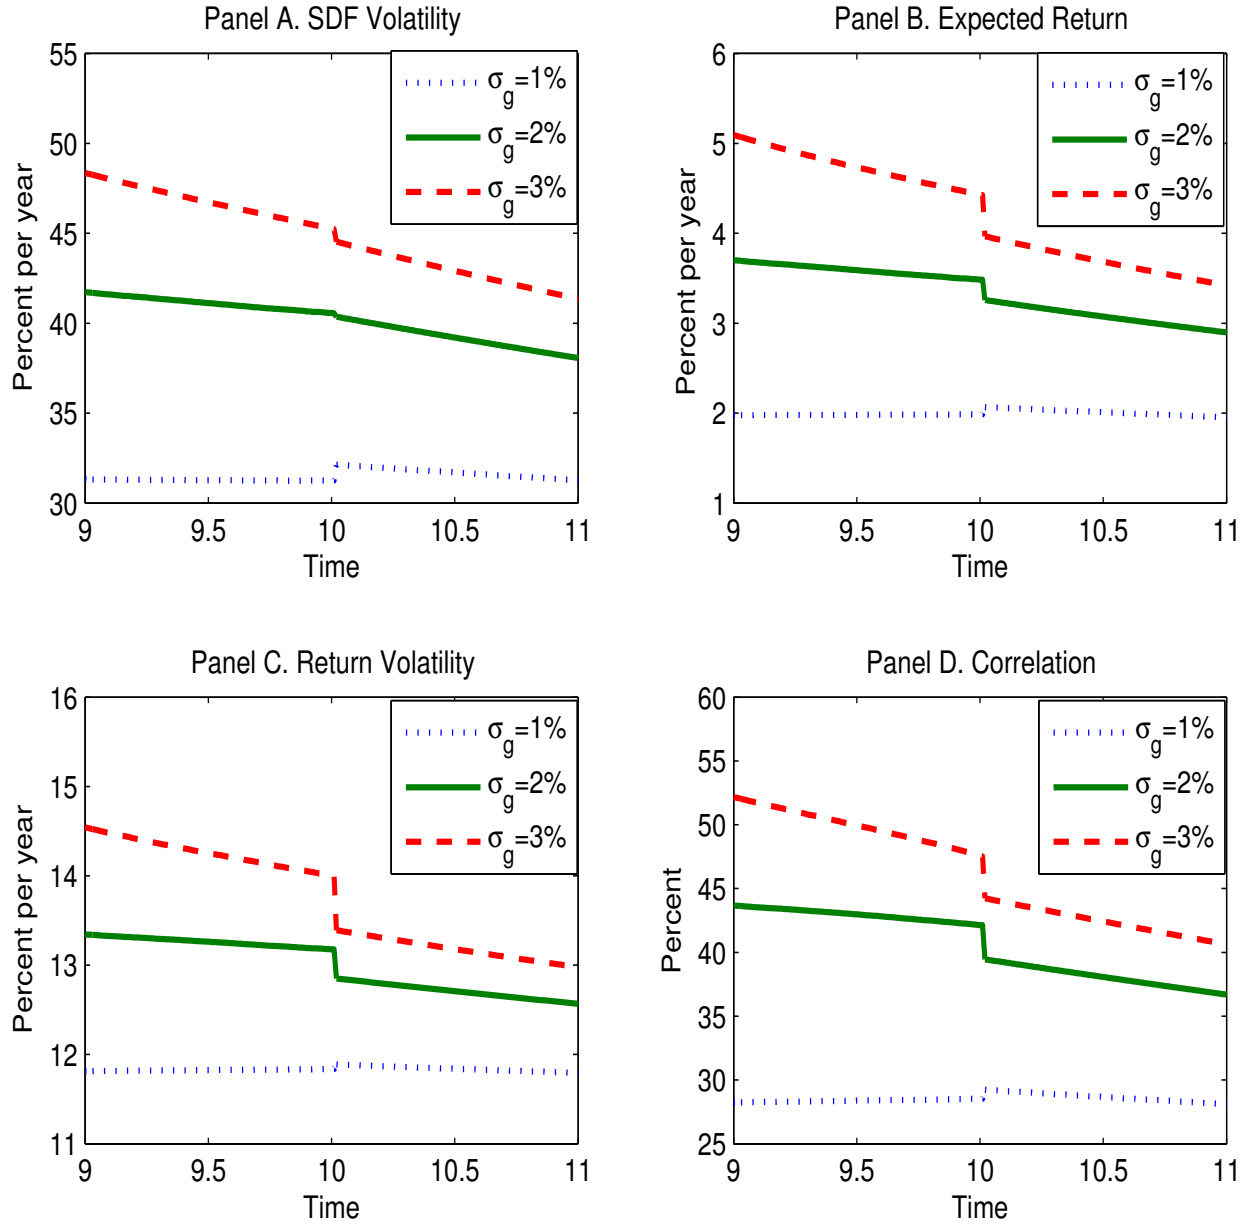


Figure 11. Properties of returns around policy decisions that result in *no* policy change. This figure plots the expected dynamics of the volatility of the stochastic discount factor, $\sigma_{\pi,t}$ (Panel A), the conditional expected stock market return, $\mu_{M,t}$ (Panel B), the conditional volatility of stock market returns, $\sigma_{M,t}$ (Panel C), and the pairwise correlation between stocks, ρ_t (Panel D), all conditional on *no* policy change at time $\tau = 10$. The jump-related components at time τ are not plotted. The downturn length is $\tau - t_0 = 5$ years, so that $\hat{g}_5 = 0$. The parameters are in Table 1, except for σ_g , which takes three different values, 1%, 2%, and 3% per year.




Figure 12. The level and volatility of stock prices around policy changes. This figure plots the dynamics of the level and volatility of stock prices around policy changes at time $\tau = 10$ in two scenarios: (i) when policy changes are allowed (our model—solid line), and (ii) when policy changes are precluded (hypothetical scenario—dashed line). Panels A and B plot the expected dynamics of the stock market value M_t , whereas Panels C and D plot the expected dynamics of stock market volatility $\sigma_{M,t}$. Market value is in units of book value at time 0; volatility is in percent per year. The downturn length is $\tau - t_0 = 1$ year in Panels A and C, but 5 years in Panels B and D. The parameters are in Table 1.




Figure 13. Endogenous timing of a policy change. Panel A plots the expected stock return, EAR, at the endogenously-timed announcement of a policy change. The timing of the policy change is optimally chosen by the government from the set $\tau \in [1, 2, \dots, 19]$. Whereas Panel A averages returns across all τ , Panel B plots EAR as a function of τ . Panels C and D plot the expected dynamics of $\sigma_{M,t}$ and ρ_t , respectively, around a policy change. The results are plotted in event time, with time 0 marking the policy change (time τ). The parameters are in Table 1, except for σ_g , which takes three different values, 1%, 2%, and 3% per year.




Figure 14. Investment adjustment. This figure plots the results from the extension of the benchmark model in which firms can disinvest at time τ . Panel A plots the equilibrium fraction of firms that choose to remain invested in the risky technology after the policy change. Panel B plots the corresponding stock return expected at the announcement of a policy change, or EAR. Panels C and D plot the expected dynamics of $\sigma_{M,t}$ and ρ_t , respectively, around a policy change at time $\tau = 10$. The parameters are in Table 1, except that $\mu = 0.02$ (we reduce μ from its benchmark value to obtain $\alpha < 1$ in equilibrium) and σ_g takes three different values, 1%, 2%, and 3% per year. The downturn length is set to 5 years in all four panels.




Figure 15. Different government policy exposures. This figure plots the results from the extension of the benchmark model in which there are two types of firms with different exposures to government policy, β^H and β^L . Panel A plots the stock return expected at the announcement of a policy change, or EAR, when half of the firms as of time 0 have $\beta^H = 1.5$ and the other half have $\beta^L = 0.5$. Panel B plots EAR as a function of $\beta^H - \beta^L$ while varying w_0 , the fraction of high-beta firms as of time 0. Panels C and D plot the difference in the conditional expected returns of the high-beta and low-beta firms, $\mu^H - \mu^L$, as a function of $\beta^H - \beta^L$ while varying w_t , the fraction of high-beta firms as of time $t = 5$ years. In Panel C, $\hat{g}_t = 2\%$, whereas in Panel D, $\hat{g}_t = -2\%$. In Panels B, C, and D, we vary $\beta^H - \beta^L$ while holding the average of the two betas constant, $(\beta^H + \beta^L)/2 = 1$. Unless specified otherwise, the parameters are from Table 1.

Appendix

The Appendix contains the formulas that are mentioned in the text but omitted for the sake of brevity. The proofs of all results are available in the companion Technical Appendix, which is downloadable from the authors' websites.

Proposition A1. In the benchmark model for $t \leq \tau$, the state price density is given by

$$\pi_t = B_t^{-\gamma} \Omega(\hat{g}_t, t) , \quad (\text{A1})$$

where

$$\begin{aligned} \Omega(\hat{g}_t, t) &= p_t^{yes} G_t^{yes} + (1 - p_t^{no}) G_t^{no} \\ G_t^{yes} &= e^{-\gamma\mu(T-t) - \gamma\hat{g}_t(\tau-t) + \frac{\gamma^2}{2}((T-\tau)^2\hat{\sigma}_g^2 + (\tau-t)^2\hat{\sigma}_t^2) + \gamma(1+\gamma)\frac{\sigma_c^2}{2}(T-t)} \\ G_t^{no} &= e^{-\gamma(\mu+\hat{g}_t)(T-t) + \frac{\gamma^2}{2}(T-t)^2\hat{\sigma}_t^2 + \gamma(1+\gamma)\frac{\sigma_c^2}{2}(T-t)} \end{aligned}$$

and

$$\begin{aligned} p_t^{yes} &= N\left(\underline{g}(0); \hat{g}_t - \gamma\hat{\sigma}_t^2(\tau-t) + \frac{\sigma_c^2/2}{(T-\tau)(1-\gamma)}, \hat{\sigma}_t^2 - \hat{\sigma}_\tau^2 + \frac{\sigma_c^2}{(T-\tau)^2(1-\gamma)^2}\right) \\ p_t^{no} &= N\left(\underline{g}(0); \hat{g}_t - \gamma[\hat{\sigma}_t^2(T-t) - (T-\tau)\hat{\sigma}_\tau^2] + \frac{\sigma_c^2/2}{(T-\tau)(1-\gamma)}, \hat{\sigma}_t^2 - \hat{\sigma}_\tau^2 + \frac{\sigma_c^2}{(T-\tau)^2(1-\gamma)^2}\right) \end{aligned}$$

Corollary A1 (used in Proposition 6). In the benchmark model for $t \leq \tau$, the volatility of the stochastic discount factor is given by

$$\sigma_{\pi,t} = \gamma\sigma - \frac{1}{\Omega(\hat{g}_t, t)} \frac{\partial \Omega(\hat{g}_t, t)}{\partial \hat{g}_t} \hat{\sigma}_t^2 \sigma^{-1} . \quad (\text{A2})$$

Proposition A2. In the benchmark model for $t \leq \tau$, the stock price for firm i is given by

$$M_t^i = B_t^i \frac{\Phi(\hat{g}_t, t)}{\Omega(\hat{g}_t, t)} , \quad (\text{A3})$$

where

$$\begin{aligned} \Phi(\hat{g}_t, t) &= \bar{p}_t^{yes} K_t^{yes} + (1 - \bar{p}_t^{no}) K_t^{no} \\ K_t^{yes} &= e^{(1-\gamma)\mu(T-t) + (1-\gamma)\hat{g}_t(\tau-t) + \frac{(1-\gamma)^2}{2}((T-\tau)^2\sigma_g^2 + (\tau-t)^2\hat{\sigma}_t^2) - (1-\gamma)\gamma\frac{\sigma_c^2}{2}(T-t)} \\ K_t^{no} &= e^{(1-\gamma)\mu(T-t) + (1-\gamma)\hat{g}_t(T-t) + \frac{(1-\gamma)^2}{2}\hat{\sigma}_t^2(T-t)^2 - (1-\gamma)\gamma\frac{\sigma_c^2}{2}(T-t)} \end{aligned}$$

and

$$\begin{aligned} \bar{p}_t^{yes} &= N\left(\underline{g}(0); \hat{g}_t + (1-\gamma)\hat{\sigma}_t^2(\tau-t) + \frac{\sigma_c^2/2}{(T-\tau)(1-\gamma)}, \hat{\sigma}_t^2 - \hat{\sigma}_\tau^2 + \frac{\sigma_c^2}{(T-\tau)^2(1-\gamma)^2}\right) \\ \bar{p}_t^{no} &= N\left(\underline{g}(0); \hat{g}_t + (1-\gamma)[\hat{\sigma}_t^2(T-t) - (T-\tau)\hat{\sigma}_\tau^2] + \frac{\sigma_c^2/2}{(T-\tau)(1-\gamma)}, \hat{\sigma}_t^2 - \hat{\sigma}_\tau^2 + \frac{\sigma_c^2}{(T-\tau)^2(1-\gamma)^2}\right) \end{aligned}$$

Corollary A2 (used in Proposition 7). In the benchmark model for $t \leq \tau$, the mean and volatility parameters of the return process in Proposition 7 are given by

$$\sigma_{M,t} = \sigma + \left(\frac{\partial \Phi(\hat{g}_t, t) / \partial \hat{g}_t}{\Phi(\hat{g}_t, t)} - \frac{\partial \Omega(\hat{g}_t, t) / \partial \hat{g}_t}{\Omega(\hat{g}_t, t)} \right) \hat{\sigma}_t^2 \sigma^{-1} \quad (\text{A4})$$

$$\mu_M = \sigma_{\pi,t} \sigma_{M,t}, \quad (\text{A5})$$

where $\sigma_{\pi,t}$ and $\sigma_{M,t}$ are given in equations (A2) and (A4), respectively.

Corollary A3 (used in Corollary 9). In the benchmark model for $t < \tau$, the correlation between the returns on any pair of stocks is given by

$$\rho_t = \frac{\sigma_{M,t}^2}{\sigma_1^2 + \sigma_{M,t}^2}. \quad (\text{A6})$$

Proposition A3. Let the timing of the policy change be endogenous as in Section 4. For every $t \in [\tau_i, \tau_{i+1})$, the indirect utility function $V(\hat{g}_t, B_t, t)$ from equation (46) is given by

$$V(\hat{g}_t, B_t, t) = B_t^{1-\gamma} \Phi(\hat{g}_t, t), \quad (\text{A7})$$

where $\Phi(\hat{g}_t, t)$ satisfies the partial differential equation

$$\begin{aligned} 0 = & \frac{\partial \Phi(\hat{g}_t, t)}{\partial t} + \left\{ (1-\gamma)(\mu + \hat{g}_t) - \frac{1}{2} \gamma (1-\gamma) \sigma^2 \right\} \Phi(\hat{g}_t, t) \\ & + \frac{1}{2} \frac{\partial^2 \Phi(\hat{g}_t, t)}{\partial \hat{g}_t^2} (\hat{\sigma}_t^2)^2 \sigma^{-2} + (1-\gamma) \frac{\partial \Phi(\hat{g}_t, t)}{\partial \hat{g}_t} \hat{\sigma}_t^2. \end{aligned} \quad (\text{A8})$$

The boundary conditions at time τ_i are given by

$$\Phi(\hat{g}_{\tau_i}, \tau_i) = E_{\tau_i} \left[\max \left\{ \Phi(\hat{g}_{\tau_i}, \tau_i), \frac{1}{1-\gamma} e^{c+(1-\gamma)\mu(T-\tau_i)+\frac{1}{2}(1-\gamma)^2\sigma_g^2(T-\tau_i)^2-\gamma(1-\gamma)\frac{\sigma^2}{2}(T-\tau_i)} \right\} \right],$$

where the expectation is taken with respect to c just before the policy decision at time τ_i . The final condition at time T is $\Phi(\hat{g}_T, T) = \frac{1}{1-\gamma}$.

Note on Section 5.: Right before time τ , the market value of firm i is given by a weighted average of $M_{\tau+}^{i,yes}$ and $M_{\tau+}^{i,no}$ as in equation (23), except that the weights ω_τ are given by

$$\omega_\tau = \frac{p_\tau}{p_\tau + (1-p_\tau)H_\tau}, \quad (\text{A9})$$

where

$$H_\tau = \frac{E_\tau \left\{ [\alpha_\tau e^{\varepsilon^{no}(\tau,T)} + (1-\alpha_\tau)]^{-\gamma} \right\}}{E_\tau \left\{ [\alpha_\tau e^{\varepsilon^{yes}(\tau,T)} + (1-\alpha_\tau)]^{-\gamma} \right\}}$$

and p_τ is the probability of a policy change as perceived just before time τ (i.e., the probability that the condition (49) holds). This probability is computed numerically.

REFERENCES

- Ait-Sahalia, Yacine, Jochen Andritzky, Andreas Jobst, Sylwia Nowak, and Natalia Tamirisa, 2010, Market response to policy initiatives during the global financial crisis, Working paper, NBER No. 15809.
- Alesina, Alberto, Silvia Ardagna, and Francesco Trebbi, 2006, Who adjusts and when? The political economy of reforms, *IMF Staff Papers* 53, 1–29.
- Alesina, Alberto, and Dani Rodrik, 1994, Distributive politics and economic growth, *Quarterly Journal of Economics* 109, 465–490.
- Anderson, Christopher J., 2007, The end of economic voting? Contingency dilemmas and the limits of democratic accountability, *Annual Review of Political Science* 10, 271–296.
- Belo, Frederico, Vito D. Gala, and Jun Li, 2011, Government spending, political cycles, and the cross-section of stock returns, Working paper, University of Minnesota.
- Bernanke, Ben S., 1983, Irreversibility, uncertainty, and cyclical investment, *Quarterly Journal of Economics* 98, 85–106.
- Bernanke, Ben S. and Kenneth N. Kuttner, 2005, What explains the stock market’s reaction to Federal Reserve policy?, *Journal of Finance* 60, 1221–1257.
- Bloom, Nicholas, 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.
- Boutchkova, Maria, Hitesh Doshi, Art Durnev, and Alexander Molchanov, 2010, Precarious politics and return volatility, Working paper, McGill University.
- Brender, Adi, and Allan Drazen, 2008, How do budget deficits and economic growth affect reelection prospects? Evidence from a large panel of countries, *American Economic Review* 98, 2203–2220.
- Bruno, Michael, and William Easterly, 1996, Inflation’s children: Tales of crises that beget reforms, *American Economic Review P&P* 86, 213–217.
- Callander, Steven, 2008, Searching for good policies, Working paper, Northwestern Univ.
- Coate, Stephen, and Stephen Morris, 1995, On the form of transfers to special interests, *Journal of Political Economy* 103, 1210–1235.
- Croce, Maximiliano M., Howard Kung, Thien T. Nguyen, and Lukas Schmid, 2011, Fiscal policies and asset prices, Working paper, UNC.
- Cutler, David M., 1988, Tax reform and the stock market: An asset price approach, *American Economic Review* 78, 1107–1117.
- Dixit, Avinash, 1989, Entry and exit decisions under uncertainty, *Journal of Political Economy* 97, 620–638.

- Downs, Anthony, 1957, *An Economic Theory of Democracy*, Harper, New York, NY.
- Drazen, Allan, 2000, *Political Economy in Macroeconomics*, Princeton University Press, Princeton, NJ.
- Drazen, Allan, and Elhanan Helpman, 1990, Inflationary consequences of anticipated macroeconomic policies, *Review of Economic Studies* 57, 147–164.
- Drazen, Allan, and William Easterly, 2001, Do crises induce reform? Simple empirical tests of conventional wisdom, *Economics and Politics* 13, 129–157.
- Durnev, Art, 2011, The real effects of political uncertainty, Working paper, McGill Univ.
- Erb, Claude B., Campbell R. Harvey, and Tadas E. Viskanta, 1996, Political risk, economic risk, and financial risk, *Financial Analysts Journal*, December, 29–46.
- Fernandez, Raquel, and Dani Rodrik, 1991, Resistance to reform: Status quo bias in the presence of individual-specific uncertainty, *American Economic Review* 81, 1146–1155.
- Gomes, Francisco J., Laurence J. Kotlikoff, and Luis M. Viceira, 2008, The excess burden of government indecision, Working paper, London Business School.
- Gomes, Francisco J., Alexander Michaelides, and Valery Polkovnichenko, 2009, Fiscal policy in an incomplete markets economy, Working paper, London Business School.
- Grossman, Gene M., and Elhanan Helpman, 1994, Protection for sale, *American Economic Review* 84, 833–850.
- Hassett, Kevin A., and Gilbert E. Metcalf, 1999, Investment with uncertain tax policy: Does random tax policy discourage investment? *The Economic Journal* 109, 372–393.
- Hermes, Niels, and Robert Lensink, 2001, Capital flight and the uncertainty of government policies, *Economics Letters* 71, 377–381.
- Julio, Brandon, and Youngsuk Yook, 2011, Corporate financial policy under political uncertainty: International evidence from national elections, *Journal of Finance*, forthcoming.
- Klingebiel, Daniela, Randy Kroszner, Luc Laeven, and Pieter van Oijen, 2001, “Stock market responses to bank restructuring policies during the East Asian crisis,” World Bank Policy Research Working Paper No. 2571.
- Lewis-Beck, Michael S., and Mary Stegmaier, 2000, Economic determinants of electoral outcomes, *Annual Review of Political Science* 3, 183–219.
- McDonald, Robert, and Daniel Siegel, 1986, The value of waiting to invest, *Quarterly Journal of Economics* 101, 707–728.
- McGrattan, Ellen R., and Edward C. Prescott, 2005, Taxes, regulations, and the value of U.S. and U.K. corporations, *Review of Economic Studies* 72, 767–796.

- Persson, Torsten, and Guido Tabellini, 1994, Is inequality harmful for growth?, *American Economic Review* 84, 600–621.
- Pindyck, Robert S., 1988, Irreversible investment, capacity choice, and the value of the firm, *American Economic Review* 78, 969–985.
- Rigobon, Roberto, and Brian Sack, 2004, The impact of monetary policy on asset prices, *Journal of Monetary Economics* 51, 1553–1575.
- Rodrik, Dani, 1991, Policy uncertainty and private investment in developing countries, *Journal of Development Economics* 36, 229–242.
- Rodrik, Dani, 1996, Understanding economic policy reform, *Journal of Economic Literature* 34, 9–41.
- Rose-Ackerman, Susan, 1999, *Corruption and Government: Causes, Consequences, and Reform*, Cambridge University Press, Cambridge.
- Santa-Clara, Pedro, and Rossen Valkanov, 2003, The presidential puzzle: Political cycles and the stock market, *Journal of Finance* 58, 1841–1872.
- Savor, Pavel, and Mungo I. Wilson, 2010, How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements, Working paper, Wharton.
- Shleifer, Andrei, and Robert W. Vishny, 1993, Corruption, *Quarterly Journal of Economics* 108, 599–617.
- Sialm, Clemens, 2006, Stochastic taxation and asset pricing in dynamic general equilibrium, *Journal of Economic Dynamics and Control* 30, 511–540.
- Sialm, Clemens, 2009, Tax changes and asset pricing, *American Economic Review* 99, 1356–1383.
- Strulovici, Bruno, 2010, Learning while voting: Determinants of collective experimentation, *Econometrica*, forthcoming.
- Thorbecke, Willem, 1997, On stock market returns and monetary policy, *Journal of Finance* 52, 635–654.
- Ulrich, Maxim, 2011, How does the bond market perceive government interventions?, Working paper, Columbia University.
- Yonce, Adam, 2009, Political cycles and corporate investment, Working paper, Berkeley.