

Escaping Volatile Inflation*

Martin Ellison

Tony Yates

University of Warwick and CEPR

Bank of England

June 2005

Abstract

Why has inflation been so stable in developed economies since the early 1990s? In this paper, we answer that the U.S. and other countries may have escaped from a volatile inflation equilibrium. Our argument builds on the story proposed by Tom Sargent in *The Conquest of American Inflation*, where the fall in inflation in the 1980s was attributed to changing government beliefs. To explain the escape in inflation volatility, we unwind one of Sargent's simplifications and allow the government to react to some of the shocks in the economy. In this case, when government beliefs turned against the Phillips curve in the 1980s they not only led to an escape from high inflation, but also stopped government using changes in inflation to offset shocks. Inflation and inflation volatility therefore escaped in tandem. Our analysis also sheds some light on why the escape in inflation occurred at the time it did.

*We are grateful for helpful comments from Kosuke Aoki, Jim Bullard, Andrea Gerali, Francesco Lippi, Tom Sargent, Peter Sinclair, Noah Williams and two anonymous referees. Martin Ellison acknowledges support from the Bank of England and an ESRC Research Fellowship, "Improving Monetary Policy for Macroeconomic Stability in the 21st Century" (RES-000-27-0126). The views in this paper should not be taken to be those of either the Bank of England or the Monetary Policy Committee.

1 Introduction

The current macroeconomic environment of low and stable inflation has prompted both academics and central bankers to address the question of why inflation was so high and volatile in developed economies during the 1970s. Providing a satisfactory answer represents a substantial challenge, since the contrast between the present and the past is stark. Average U.S. inflation came down from 7.1% in the 1970s to 2.9% in the 1990s, with the consensus amongst econometricians being that inflation volatility (as measured by its conditional variance) also fell.¹ Any candidate answer at least has to answer the following questions:

- (i) Why did inflation rise and fall from the 1960s to the 1990s?
- (ii) Why did inflation volatility also rise and fall over the same period?

The publication of *The Conquest of American Inflation* by Tom Sargent in 1999 began a literature that attempts to explain the inflation of the 1970s purely in terms of changing government beliefs. The rise in inflation is attributed to the discovery of the Phillips curve relationship, which tempted the government to increase inflation in an bid to reduce unemployment. High inflation was then sustained by overly pessimistic beliefs about how much unemployment would rise if inflation were to be brought down, which deluded the government into believing its high inflation policy was effective in reducing unemployment. This continued until a rare sequence of shocks led the government to abandon its high inflation policy and adopt a more realistic view of unemployment and policy effectiveness. In the terminology of Cho, Williams and Sargent (2002), the fall in inflation in the 1980s is an example of an escape dynamic.²

The explanation proposed by Sargent (1999) is both elegant and a fine application of simple yet powerful theory. However, it only provides only a partial answer to question (i) of why inflation rose and fell, and is silent on question (ii) concerning the observed changes in inflation

¹Econometricians typically estimate GARCH or stochastic volatility models, and find that the level and conditional volatility of inflation are positively correlated. See Giordani and Söderlind (2003) for a recent innovative study combining this approach with data from the Survey of Professional Forecasters.

²Other recent contributions that attempt to rationalise the past through formal modelling of changing government beliefs include Cogley and Sargent (2004), Orphanides (2003), Primiceri (2005) and Reis (2003).

volatility. If the rise and fall in inflation is to be explained by changing government beliefs, then it is natural to ask why the particular sequence of shocks needed to provoke the fall in inflation occurred at the time it did in the 1980s. The escape dynamic literature pioneered by Cho, Williams and Sargent (2002) identifies the series of shocks needed to change beliefs and trigger the fall, but offers little guidance as to when they are likely to occur. In answer to the question about changes in inflation volatility, Sargent (1999) is silent because his and subsequent work implies that the volatility of inflation should be constant when conditioned on beliefs.

The primary aim of this paper is to reconcile the observed rise and fall in inflation volatility with a simple model of changing government beliefs. The strategy we employ is to unwind one of the simplifications in Sargent (1999), and return to a model in which the government can use stabilisation policy to cushion the effects of some of the shocks that hit the economy. To this end, we assume the presence of frictions in the economy that prevent private agents from optimally adjusting to shocks. This opens the door to stabilisation policy if we further assume that the government is not constrained by frictions and so can react to some of the shocks. Our motivation for generalising the Sargent (1999) story in this way is the long tradition of government using policy to stabilise the economy, as typified by Phelps and Taylor (1977).

Re-introducing a motive for stabilisation policy creates a natural link between the level and volatility of inflation. In our story, the discovery of the Phillips curve not only leads to a higher level of inflation but also creates greater inflation volatility as the government begins to use changes in inflation to offset shocks and stabilise the economy. In the period when inflation is high, the delusion that inflation is effective in reducing unemployment translates into a strong desire for the government to also use stabilisation policy to offset shocks, so the volatility of inflation remains high. When the escape to low inflation eventually occurs, the more realistic view of policy effectiveness adopted by the government forces a reigning in of stabilisation policy, which naturally leads to a concurrent fall in the volatility of inflation.

As a corollary to our analysis, we are also able to shed some light on the question of why the escape to low inflation occurred at the time it did in the 1980s. The insight we offer is that escapes are much more likely to happen in a period when there are relatively few shocks

that can be offset by stabilisation policy. Conversely, an escape is unlikely if there are lots of shocks that the government can offset. This result complements the work of McGough (2005), who shows that favourable shocks to the natural rate of unemployment similarly increase the likelihood of an escape. Taken together, our papers suggest that the escape to low inflation in the 1980s may have been precipitated by a lack of active stabilisation actions and/or a permanent natural rate shock.

To derive our results, we add a new shock to the escape dynamic analysis of Cho, Williams and Sargent (2002) and incorporate price-setting frictions that prevent private agents from adjusting to the shock. The combination of the new shock and frictions creates the desired role for stabilisation policy. To see the effect of this on mean and escape dynamics, we follow Cho, Williams and Sargent (2002) and transform our discrete time model into its continuous time analogue. The transformation allows us to fully describe the expected mean dynamics of the model and, after solving a suitable optimal control problem, gives us a numerical characterisation of the dominant escape path. We support our analytical and numerical results with simulations of the discrete time version of the model.

2 Model

Our model extends that of Cho, Williams and Sargent (2002) by adding a motivation for stabilisation policy. We therefore retain the key features of their model, but follow Phelps and Taylor (1977) and introduce price-setting frictions and a new unemployment shock W_3 that the government can react to. Figure 1 shows the timing of our model. The presence of frictions forces private agents to set prices in advance, based on an expectation \hat{x} of policy formed before the shock is known. The government is not subject to price frictions, so actual policy x can react to the shock and attempt to stabilise the economy.




Figure 1: Timing of the model

In the remainder of this section we sketch out the key features of our model in detail: the actual structure of the economy, the structure of the economy perceived by the government, and the mechanisms by which the government sets policy and updates its beliefs.

2.1 Actual structure of the economy

The structure of the economy is described by the expectations-augmented Phillips curve (2.1), in which unemployment U_n is determined by its natural rate u , the difference between realised inflation π_n and expected inflation \hat{x}_n , and two shocks W_{1n} and W_{3n} . Unexpected inflation affects unemployment due to the presence of price-setting frictions. The parameter θ is assumed to be positive so surprise inflation reduces unemployment. W_{1n} is the unemployment shock in Cho, Williams and Sargent (2002) that occurs after expectations and policy have been set. W_{3n} is our new shock that creates incentives for the government to engage in stabilisation policy. We assume that the government can react to W_{3n} , but price setting frictions prevent private agents adjusting to W_{3n} once their inflation expectations \hat{x}_n have been formed.

$$U_n = u - \theta(\pi_n - \hat{x}_n) + \sigma_1 W_{1n} + \sigma_3 W_{3n} \quad (2.1)$$

The government is assumed to have direct but imperfect control of inflation. Inflation π_n is therefore equal to the level x_n intended by the government, plus a control error W_{2n} . Equation (2.2) defines inflation in our model. We refer to x_n as intended inflation.

$$\pi_n = x_n + \sigma_2 W_{2n} \quad (2.2)$$

2.2 Perceived structure of the economy

Following Cho, Williams and Sargent (2002), we assume that the government does not know the actual structure of the economy. Rather, it has an approximating model which allows for the possibility of a trade-off between unemployment and inflation. We extend the approximating model and permit the government to correctly perceive the effect on employment of our new shock W_{3n} . Although the approximating model is correctly specified in this respect, it remains otherwise misspecified since it assumes a trade-off between unemployment and inflation, when in reality only unexpected inflation matters for unemployment. We write the

perceived structure of the economy as equations (2.3) and (2.4). η_n is an approximation error, capturing any fluctuations in unemployment the government fails to explain by its perceived structure of the economy.

$$U_n = \gamma_{0n} + \gamma_{1n}\pi_n + \sigma_3 W_{3n} + \eta_n \quad (2.3)$$

$$\pi_n = x_n + \sigma_2 W_{2n} \quad (2.4)$$

The government estimates the perceived structure of the economy using standard econometric techniques. Equations (2.5) - (2.8) are recursive formulae for discounted least squares estimation of the coefficients in equation (2.3), with γ_n a 2×1 vector of current parameter estimates $[\gamma_{0n} \ \gamma_{1n}]'$. ε is the discount factor or gain, $g(\gamma_n, \xi_n)$ is the forecast error, and R_n is a 2×2 matrix measuring the precision of the current estimates. The function $M(\gamma_n)$ is introduced to ease notation. Our assumption that the government correctly perceives the effect of the new unemployment shock means that W_{3n} does not appear in the definition of the forecast error (2.7), rendering our recursive estimation scheme identical to that in Cho, Williams and Sargent (2002).

$$\gamma_{n+1} = \gamma_n + \varepsilon R_n^{-1} g(\gamma_n, \xi_n) \quad (2.5)$$

$$R_{n+1} = R_n + \varepsilon (M_n - R_n) \quad (2.6)$$

$$g(\gamma_n, \xi_n) = (u - \theta(\pi_n - \hat{x}_n) + \sigma_1 W_{1n} - \gamma_{0n} - \gamma_{1n}\pi_n) \begin{pmatrix} 1 \\ \pi_n \end{pmatrix} \quad (2.7)$$

$$M(\gamma_n) = \begin{pmatrix} 1 \\ \pi_n \end{pmatrix} \begin{pmatrix} 1 \\ \pi_n \end{pmatrix}' \quad (2.8)$$

In discounting past data, the government implicitly allows for the possibility of structural breaks, even though no such breaks are explicitly present in the model. Discounting with the gain ε gives an exponentially decreasing weight to past data.

2.3 Government policy

The objective of government policy is to minimise the objective function (2.9), where \tilde{E} denotes the expectations operator induced by the perceived structure of the economy (2.3) and (2.4).

The government is penalised quadratically whenever unemployment or inflation deviates from zero. δ is a discount factor.

$$\tilde{E} \sum_{n=0}^{\infty} \delta^n (U_n^2 + \pi_n^2) \quad (2.9)$$

A fully optimal government policy would set intended inflation x_n to minimise the objective function (2.9), subject to the perceived structure of the economy (2.3) - (2.4) and the recursive estimation scheme (2.5) - (2.8). To solve this non-linear dynamic problem requires a high degree of computational sophistication from the government. For this reason, we follow Cho, Williams and Sargent (2002) in invoking anticipated utility as a boundedly rational decision criterion for the government.³ The anticipated utility criterion requires the government to minimise objective function (2.9) subject to perceived structure (2.3) - (2.4), but replaces recursive estimation scheme constraints (2.5) - (2.8) by a simpler assumption that the government believes its current best estimate of the perceived structure is both precise and correct. The anticipated utility policy consequently sets intended inflation x_n to solve the static optimisation problem (2.10).

$$\begin{aligned} \min_{\{x_n\}} \tilde{E} \sum_{n=0}^{\infty} \delta^n (U_n^2 + \pi_n^2) \\ \text{s.t.} \\ U_n = \gamma_{0n} + \gamma_{1n}\pi_n + \sigma_3 W_{3n} + \eta_n \\ \pi_n = x_n + \sigma_2 W_{2n} \\ \gamma_{0n}, \gamma_{1n}, \sigma_3 W_{3n} \text{ given} \end{aligned} \quad (2.10)$$

The solution to optimisation problem (2.10) is the anticipated utility policy (2.11), in which intended inflation x_n leans against the wind created by γ_{0n} and our new shock W_{3n} . The extent to which policy leans against the wind depends on the current estimate γ_{1n} of the trade-off parameter in the perceived structure of the economy.

$$x_n = -\frac{\gamma_{1n}}{1 + \gamma_{1n}^2} (\gamma_{0n} + \sigma_3 W_{3n}) \quad (2.11)$$

³The assumption of anticipated utility maximisation is supported by Kreps (1998), who argues that dynamic problems under uncertainty are fundamentally different to static ones, so require a fundamentally different decision criterion. Also in support is the paper of Cogley, Colacito and Sargent (2005), which suggests that the difference between fully optimal and anticipated utility policies is likely to be small in any case.

Our earlier assumption of price-setting frictions forces private agents to set prices on the expected outcome of policy (2.11). According to our timing protocol, prices are set before our new shock is known, so W_{3n} drops out of expectations. Expected inflation \hat{x}_n is then given by equation (2.12).

$$\hat{x}_n = -\frac{\gamma_{1n}}{1 + \gamma_{1n}^2} \gamma_{0n} \quad (2.12)$$

3 Rising inflation and inflation volatility

The first question to address with our model is the rise in inflation and inflation volatility observed in the 1960s and 1970s. Sargent (1999) attributes the increase in inflation to the evolution in government beliefs after the discovery of the Phillips curve. Our model follows Sargent's line, but demonstrates that rising inflation volatility is a natural consequence of the same changes in government beliefs that caused higher inflation, once we re-introduce a motive for stabilisation policy. We derive the results in this section using stochastic approximation techniques to analyse the *mean dynamics* of the continuous time analogue of our model, thereby tracing out the normal (expected) evolution of government beliefs after the Phillips curve is discovered. In Section 5 we present supporting evidence from simulations of the discrete time model.

3.1 Mean dynamics

The derivation of the mean dynamics of our model begins by re-writing the recursive estimation scheme (2.5) and (2.6) as equations (3.1) and (3.2).

$$\frac{\gamma_{n+1} - \gamma_n}{\varepsilon} = R_n^{-1} g(\gamma_n, \xi_n) \quad (3.1)$$

$$\frac{R_{n+1} - R_n}{\varepsilon} = M_n - R_n \quad (3.2)$$

Equations (3.1) and (3.2) resemble a discrete-time approximation of a continuous time process perturbed by shocks ξ_n . If we take the limit as $\varepsilon \rightarrow 0$, the approximation error tends to zero and a weak law of large numbers ensures that the stochastic element become negligible. In the limit, the mean dynamics of the model can therefore be described by a pair of ordinary

difference equations (3.3) and (3.4), where $\bar{g}(\gamma)$ is the expected value of $g(\gamma, \xi)$ and $\bar{M}(\gamma)$ is the expected value of M . A proof of this is contained in Cho, Williams and Sargent (2002).

$$\dot{\gamma} = R^{-1} \bar{g}(\gamma) \quad (3.3)$$

$$\dot{R} = \bar{M}(\gamma) - R \quad (3.4)$$

Expressions for $\bar{g}(\gamma)$ and $\bar{M}(\gamma)$ can be obtained by taking expectations of the forecast error (2.7) and the precision matrix (2.8), conditional on the true structure of the economy (2.1) - (2.2) and government policy (2.11). The mean dynamics for our model are then given by equations (3.5) and (3.6).

$$\dot{\gamma} = R^{-1} \begin{pmatrix} u - \frac{\gamma_0}{1+\gamma_1^2} \\ (u - \gamma_0 - \gamma_1 \hat{x}) \hat{x} - (\theta + \gamma_1) \left(\sigma_2^2 + \left(\frac{\gamma_1}{1+\gamma_1^2} \right)^2 \sigma_3^2 \right) \end{pmatrix} \quad (3.5)$$

$$\dot{R} = \begin{pmatrix} 1 & \hat{x} \\ \hat{x} & \hat{x}^2 + \sigma_2^2 + \left(\frac{\gamma_1}{1+\gamma_1^2} \right)^2 \sigma_3^2 \end{pmatrix} - R \quad (3.6)$$

The presence of our new shock changes the mean dynamics of the model. In equations (3.5) and (3.6), the new shock W_3 plays a similar role through σ_3 as the inflation control error W_2 does through σ_2 . Intuitively, the shocks have similar effects because they both create unexpected movements in inflation and unemployment. W_2 does so directly, since it is an unpredictable control error. W_3 does so indirectly, since it prompts stabilisation actions that cannot be predicted at the time private agents form their expectations.

3.2 Properties of mean dynamics

The next step in the analysis is to examine how our new shock affects the properties of mean dynamics. For mean dynamics to converge requires the existence of a locally asymptotically stable fixed point of the system of ordinary differential equations (3.5) - (3.6). We begin by checking existence, noting that any fixed point must satisfy the restriction $\dot{\gamma} = \dot{R} = 0$. Applying this restriction to the system of ordinary differential equations, we obtain the unique

fixed point defined by equations (3.7) and (3.8).

$$\bar{\gamma} = \begin{pmatrix} u(1 + \theta^2) \\ -\theta \end{pmatrix} \quad (3.7)$$

$$\bar{R} = \begin{pmatrix} 1 & u\theta \\ u\theta & (u\theta)^2 + \sigma_2^2 + \left(\frac{\theta}{1+\theta^2}\right)^2 \sigma_3^2 \end{pmatrix} \quad (3.8)$$

The level and volatility of unemployment and inflation at the fixed point are given by equations (3.9) - (3.12).

$$\bar{U} = u \quad (3.9)$$

$$\bar{\pi} = u\theta \quad (3.10)$$

$$\overline{\sigma_u^2} = \left(\frac{1}{1 + \theta^2}\right)^2 \sigma_3^2 + \sigma_1^2 \quad (3.11)$$

$$\overline{\sigma_\pi^2} = \left(\frac{\theta}{1 + \theta^2}\right)^2 \sigma_3^2 + \sigma_2^2 \quad (3.12)$$

The fixed point in our model has the same belief pair γ , unemployment level \bar{U} , and inflation level $\bar{\pi}$ as the self-confirming equilibrium (SCE) in Cho, Williams and Sargent (2002). The government believes that higher inflation is effective in reducing unemployment ($\gamma_1 = -\theta$), and is over-pessimistic about how high unemployment would be if inflation were to be brought down ($\gamma_0 = u(1 + \theta^2) < u$). Taken together, these beliefs delude the government into continuing to follow a high inflation policy. This feature of the model is unaffected by our new shock, which changes volatilities but not levels at the self-confirming equilibrium.

Whether the mean dynamics converge to the self-confirming equilibrium depends on the local stability properties of the system of ordinary differential equations (3.5) - (3.6). A sufficient condition for stability is that all the eigenvalues of the Jacobian of the system have negative real parts at the self-confirming equilibrium.⁴ In Appendix A, we show that this reduces to a requirement that the eigenvalues of $D\bar{g}(\gamma)$ have negative real parts. It will be satisfied iff (3.13) holds.

$$\left(\sigma_2^2 + \left(\frac{\theta}{1 + \theta^2}\right)^2 \sigma_3^2\right) > 0 \quad (3.13)$$

⁴Proposition 5.6, Evans and Honkapohja (2001), p. 96.

The system is asymptotically stable in the neighbourhood of the self-confirming equilibrium because condition (3.13) is never violated. The presence of inflation control errors ensures $\sigma_2^2 > 0$ and creates sufficient natural experiments to enable the government to “learn” the structure of the economy, a property identified by El-Gamal and Sundaram (1993) as important for convergence. Furthermore, our new shock increases the slackness in condition (3.13), so once we re-introduce a motive for stabilisation policy the self-confirming equilibrium becomes more stable. Our new shocks complement the natural experiments created by inflation control errors, since they prompt stabilisation actions that create analogous unexpected inflation and unemployment movements.

The speed of convergence to self-confirming equilibrium also depends on the slackness in condition (3.13). Since our new shock increases slackness, it should make mean dynamics converge faster. The easiest way to verify this is through a numerical example. Following Cho, Williams and Sargent (2002), we set $u = 5$, $\theta = 1$ and $\sigma_1 = \sigma_2 = 0.3$, and assume that shocks W_1 and W_2 are Gaussian i.i.d. with mean zero and unit variance. We assume that our new shock W_3 has the same distribution as the other shocks, but set $\sigma_3 = 0$ for a model without the new shock and $\sigma_3 \in (0.3, 0.6, 0.9)$ for models in which the new shock plays a progressively larger role. Figure 2 shows a numerical example of how beliefs about the perceived structure of the economy converge with and without the new shock. Beliefs (γ_0, γ_1) are initialised at $(5, 0)$ to reflect the government’s view at the beginning of the 1960s that

policy was ineffective at reducing unemployment.




Figure 2: Convergence of mean dynamics

The top two lines of Figure 2 trace out the convergence of γ_0 to its self-confirming value $u(1 + \theta^2) = 10$, reflecting increasing government pessimism about the level of unemployment that would prevail if inflation were to be brought down. The bottom two lines show γ_1 converging to $-\theta = -1$ as the government discovers the Phillips curve. As expected, in both cases beliefs converge faster to the self-confirming equilibrium with the new shock in place.

3.3 Inflation and mean dynamics

The behaviour of inflation implied by mean dynamics can be derived from the definition of inflation (2.4) and the equation for government policy (2.11). From the perspective of private agents, expectations are formed before the new shock W_3 is known, so expected inflation and expected inflation volatility are given by equations (3.14) and (3.15) respectively.

$$E(\pi | \gamma_0, \gamma_1) = -\frac{\gamma_0 \gamma_1}{1 + \gamma_1^2} \quad (3.14)$$

$$E(\sigma_\pi | \gamma_0, \gamma_1) = \left[\sigma_2^2 + \left(\frac{\gamma_1}{1 + \gamma_1^2} \right)^2 \sigma_3^2 \right]^{1/2} \quad (3.15)$$

The level of inflation rises as the mean dynamics converge. Two forces are at work here, since both the perceived need for policy and the perceived effectiveness of policy increase along

the convergence path. Higher absolute values of γ_0 and γ_1 create the (deluded) perception that high inflation is needed to reduce unemployment, and that high inflation is effective at reducing unemployment. The behaviour of inflation volatility in turn depends on the precise nature of the shocks in the economy. If there are no W_3 shocks then $\sigma_3 = 0$ and the only source of inflation volatility is the control error W_2 . The model collapses to that of Cho, Williams and Sargent (2002), with the volatility of inflation constant as mean dynamics converge. Once we introduce our new W_3 shock, the volatility of inflation becomes a function of the perceived effectiveness of policy γ_1 . The greater the perceived effectiveness the higher the volatility, since the government becomes increasingly tempted to use changes in inflation to offset the W_3 shocks to unemployment. Rising inflation volatility is then a natural consequence of mean dynamics converging to a point where policy is perceived to be effective.




Figure 3: Convergence of inflation and inflation volatility

The behaviour of inflation and inflation volatility implied by converging mean dynamics is shown in Figure 3, which is plotted for our numerical example. In the top panel, inflation rises from zero to its self-confirming equilibrium value $u\theta = 5$, converging faster in the model with W_3 shocks than without. In the bottom panel, inflation volatility rises if there are W_3 shocks but remains constant otherwise. If volatility rises, it does so slowly since it depends only on the convergence of γ_1 and not γ_0 .

Once we add our new shock to the model, inflation and its volatility rise simultaneously as mean dynamics converge. This is not surprising since the level of inflation and its volatility

have the same root cause in terms of changing government beliefs. We therefore have a partial answer to questions (i) and (ii) posed in the introduction. In our story, inflation and inflation volatility rose from the 1960s onwards because the government became increasingly convinced it could affect unemployment by changing inflation. The level of inflation went up as government tried to bring down unemployment, whereas inflation volatility increased as government attempted to use inflation to offset unemployment shocks and stabilise the economy. In these circumstances, an econometrician conditioning on the beliefs of private agents will identify a relationship between rising inflation and inflation volatility.

4 Falling inflation and inflation volatility

The fall in inflation and inflation volatility in the 1980s cannot be explained by the mean dynamics of the model. Instead, Sargent (1999) accounts for falling inflation by appealing to the model's ability to deviate from its mean dynamics in a significant and predictable way. In his view of the world, inflation fell because a rare sequence of shocks caused the government to abandon its belief that high inflation leads to lower unemployment. We follow this view, but in our story the rare sequence of shocks leads to a fall not just in inflation but also in inflation volatility. Once we re-introduce a motive for stabilisation policy, falling volatility becomes a natural consequence of the same changes in government beliefs that cause inflation to fall. As in the previous section, we obtain our results using the techniques of stochastic approximation to analyse the continuous time analogue of our model. We characterise the *escape dynamics* of the system and derive the way in which the model is most likely to deviate from its mean dynamics. In Section 5 we discuss corresponding simulations of escape dynamics in the discrete time model.

4.1 Escape dynamics

The question posed in escape dynamic analysis is what is the most likely path for beliefs if they deviate significantly (escape) from their mean dynamics. To answer this, we need a way of selecting the most likely path amongst all candidate escape paths. A natural metric is the

likelihood function of the shocks needed to drive beliefs along each escape path. The path that minimises this function is the dominant escape path, representing the path of least resistance for beliefs to escape.

The formal analysis of escape dynamics in economic models is laid out in the pioneering work of Williams (2001), where the dominant escape path is characterised by solving an optimal control problem. The method involves choosing a series of perturbations to mean dynamics that is most likely to cause beliefs to escape from a neighbourhood around the self-confirming equilibrium. Mathematically, the dominant escape path is given by the solution to optimal control problem (4.1).

$$\begin{aligned}
\bar{S} &= \inf_{\dot{v}} \int_0^t \dot{v}(s)' Q(\gamma(s), R(s))^{-1} \dot{v}(s) ds \\
&\quad s.t. \\
\dot{\gamma} &= R^{-1} \bar{g}(\gamma) + \dot{v} \\
\dot{R} &= \bar{M}(\gamma) - R \\
\gamma(0) &= \bar{\gamma}, \quad M(0) = \bar{M}, \quad \gamma(t) \notin G \text{ for some } 0 < t < T
\end{aligned} \tag{4.1}$$

The optimal control problem works by perturbing the mean dynamics of the model (3.3) - (3.4) by a factor \dot{v} and asking which series of perturbations is most likely to cause beliefs to escape. In the objective, $Q(\gamma, R)$ is a weighting function that measures the likelihood of the shocks needed to perturb beliefs by \dot{v} .⁵ We initialise beliefs at their self-confirming values and define a neighbourhood G around the self-confirming equilibrium that beliefs must escape from. The outcome of the optimal control problem is the series of belief perturbations that occur along the dominant escape path.

The re-introduction of stabilisation policy does not change the definition of the optimal control problem (4.1) we need to solve to calculate the dominant escape path in our model. What does change are the functions $Q(\gamma, R)$, $\bar{g}(\gamma)$, $\bar{M}(\gamma)$ and the precision matrix \bar{M} at the self-confirming equilibrium. In our model, beliefs are perturbed by three shocks so Q is a 3×3 matrix-valued function. More details appear in Appendix B.

⁵An analytical expression for $Q(\gamma, R)$ is given in Appendix B.

4.2 Dominant escape path

The dominant escape path solves optimal control problem (4.1). To find the solution we define the Hamiltonian (4.2), where a and λ are co-state vectors for the evolution of γ and R .

$$\mathcal{H} = a \cdot R^{-1} \bar{g}(\gamma) - \frac{1}{2} a' Q(\gamma, R) a + \lambda \cdot (\bar{M}(\gamma) - R) \quad (4.2)$$

The Hamiltonian is convex so first order conditions (4.3) - (4.6) necessarily hold along the dominant escape path.

$$\dot{\gamma} = R^{-1} \bar{g}(\gamma) - Q(\gamma, R) a \quad (4.3)$$

$$\dot{R} = \bar{M}(\gamma) - R \quad (4.4)$$

$$\dot{a} = -a R^{-1} \frac{\partial \bar{g}(\gamma)}{\partial \gamma} + \frac{1}{2} a' \frac{\partial Q(\gamma, R)}{\partial \gamma} a - \lambda \frac{\partial \bar{M}(\gamma)}{\partial \gamma} \quad (4.5)$$

$$\dot{\lambda} = -\mathcal{H}_R \quad (4.6)$$

The first order conditions form a system of ordinary differential equations. They characterise a family of escape paths, with each path being indexed by different initial values of the co-state vectors. The dominant escape path is the member of this family that achieves the escape with the most likely series of belief perturbations. A solution to the optimal control problem can therefore be obtained by searching over all possible initial values of a and λ , applying equations (4.3) - (4.6), and choosing the initial values that imply belief perturbations that are most likely in terms of the $Q(\gamma, R)$ metric.

Re-introducing a motive for stabilisation policy has subtle implications for the nature of the dominant escape path. The functions $Q(\gamma, R)$, $\bar{g}(\gamma)$, $\bar{M}(\gamma)$ all change, so the series of perturbations that cause beliefs to escape is likely to be different. To highlight these differences, we return to our numerical example and calculate the dominant escape path with and without the new shock W_3 that motivates stabilisation policy. We define an escape as happening when beliefs leave a neighbourhood G of 5 Euclidean units around the self-confirming equilibrium. Dominant escape paths with and without the new shock are plotted in Figure 4. The two

middle lines are dominant escape paths with intermediate values of $\sigma_3 = \{0.3, 0.6\}$.⁶




Figure 4: Escape dynamics

The dominant escape path in the model without our new shock is similar to the escape path in Cho, Williams and Sargent (2002).⁷ Beliefs spend a long time near the self-confirming equilibrium before escaping rapidly at $t \approx 6.5$ to new values close to $(\gamma_0, \gamma_1) = (5, 0)$. The mechanism causing beliefs to escape is the same as in Cho, Williams and Sargent (2002). Intuitively, an escape happens when a sequence of shocks makes the government sufficiently doubtful that high inflation is effective at reducing unemployment. At this point the government starts to reduce inflation, an anticipated policy move that does not affect unemployment. The government observes falling inflation having no effect on unemployment, which reinforces the initial doubt that policy may be ineffective. As doubts are reinforced there are further reductions in inflation and a rapid abandonment of the high inflation policy. The sequence of shocks that creates the initial doubt is a series of positively correlated W_1 and W_2 shocks. With positively correlated shocks, inflation varies due to the control error W_2 but unemployment is

⁶We thank Noah Williams for generously providing us with matlab codes to solve the static model in Cho, Williams and Sargent (2002). Following his lead, we simplify the calculations by setting the initial values of the second co-state vector λ to zero. Our extension of the code to allow for the new shock is available from the authors on request.

⁷In replicating the results of Cho, Williams and Sargent (2002), we discovered their escape path is only a local minimum of the optimal control problem (4.1), with a minimised value of the objective $\bar{S} = 4.987 \times 10^{-4}$. The dominant escape path we report is the global minimum with $\bar{S} = 4.458 \times 10^{-4}$.

relatively stable as any movements in unemployment caused by unexpected inflation are offset by the unemployment shock W_1 . The combination of variable inflation and apparently stable unemployment creates sufficient doubt to trigger an escape.

The re-introduction of a motive for stabilisation policy has a clear effect on the dominant escape paths in Figure 4, with the escape occurring later as the new W_3 shock plays an increasingly important role. When $\sigma_3 = 0.9$ the escape is at $t \approx 10.1$, about 55% later than in the model without the new shock. The reason for the delay is that the new shock increases the complexity of the sequence of shocks needed to trigger an escape. In the model with three shocks, we need the unemployment shock W_1 not only to offset the effect of the inflation control errors W_2 but also to offset the change in unemployment brought about by the government's reaction to the new shock W_3 . Only then will unemployment appear stable as inflation varies. The more complex the sequence of shocks needed the less likely it is to occur, so the longer we have to wait for government doubts to surface and the later the escape occurs.

Taken literally, our analysis suggests an economy is more likely to escape in periods when the variance of shocks that can be stabilised is relatively small. This result contributes to a literature that attempts to explain the *timing* of escapes, a question not addressed in Sargent (1999). The most important contributions in this literature to date are by Gerali and Lippi (2002) and McGough (2005), who demonstrate that escapes are more likely in periods when the government is less inflation averse or when there are favourable shocks to the natural rate of unemployment.

4.3 Inflation and escape dynamics

The behaviour of inflation over an escape episode is determined by the definition of inflation (2.4) and the equation for government policy (2.11). As in the mean dynamics analysis of Section 3.3, we take the perspective of private agents and calculate expected inflation and expected inflation volatility using equations (3.14) and (3.15). Figure 5 plots the evolution of inflation and inflation volatility implied by escaping beliefs.

In the top panel, the level of inflation falls rapidly as beliefs escape. At the time of the escape, increasing doubts about whether high inflation can reduce unemployment (i.e., falling

$|\gamma_1|)$ cause the government to abandon its high inflation policy in favour of setting inflation close to zero. The same escape occurs in models with and without the new shock, but with different timing. In the bottom panel, inflation volatility is constant in the model without the new shock but otherwise falls at the time of the escape. For models with the new shock, the escape affects inflation volatility because the doubts that cause government to abandon its high inflation policy also lead it to give up on stabilisation policy. The incentive to use inflation to stabilise shocks disappears and the volatility of inflation escapes to the level implied by inflation control errors.




Figure 5: Escape of inflation and inflation volatility

Our model with the new shock offers a natural explanation for why both inflation and inflation volatility fell in the 1980s, completing our answer to questions (i) and (ii) in the introduction. According to our story, inflation and its volatility fell in an escape episode when a rare sequence of shocks triggered a dramatic change in government beliefs. The escape reversed the increases in inflation and inflation volatility seen in the 1960s. An econometrician conditioning on the beliefs of private agents will identify the same positive correlation between inflation and its volatility in the 1980s as they did in the 1960s and 1970s.

5 Mean and escape dynamics in simulations

The results obtained in previous sections describe the behaviour of the economy as the gain in the learning mechanism tends to zero. In this section, we revisit models in discrete time with positive but small gain parameters. The return to discrete time precludes further analytical results, so we resort to simulations to characterise the models. We report results separately for simulated mean and escape dynamics.

5.1 Mean dynamics

The mean dynamics of the model determine the speed at which beliefs converge to the self-confirming equilibrium. To capture this feature in simulations, we initialise beliefs (γ_0, γ_1) at $(0, 5)$ and simulate the model until beliefs return to a small neighbourhood of their self-confirming values $(-1, 10)$.⁸ Table 1 shows the average time to return to equilibrium in 1000 such simulations. The numbers in parentheses are 90% confidence intervals.

Gain ε	$\sigma_3 = 0$	$\sigma_3 = 0.3$	$\sigma_3 = 0.6$	$\sigma_3 = 0.9$
0.0025	13410 (10501, 24090)	13273 (9776, 24236)	9707 (8447, 11652)	8259 (7403, 9516)
0.0050	5666 (5092, 6555)	5442 (4711, 6433)	5057 (4105, 6916)	4118 (3585, 4996)
0.0075	3694 (3248, 4176)	3509 (3052, 4095)	3183 (2672, 3879)	2777 (2355, 3512)
0.0100	2790 (2426, 3193)	2611 (2256, 3031)	2363 (1984, 2970)	2045 (1716, 2525)

Table 1: Simulated average time for convergence of mean dynamics

The simulation results confirm the prediction of Section 3.2 that re-introducing a motive for stabilisation increases the speed of convergence of mean dynamics. For each value of the gain parameter ε , beliefs converge faster when σ_3 is increased to give more prominence to the new shock in the model. Increasing σ_3 from 0 to $(0.3, 0.6, 0.9)$ reduces the average convergence time by on average (4%, 17%, 29%). Higher values of the gain parameter ε lead to faster convergence

⁸We define a small neighbourhood as beliefs being at most 0.1 Euclidean units away from their self-confirming values. Simulations in which beliefs escape before entering the neighbourhood are ignored.

to self-confirming equilibrium. A higher gain reduces the weight government places on past data, so initial beliefs are “forgotten” quicker.

The behaviour of inflation and inflation volatility implied by mean dynamics was predicted in Section 3.3. Accordingly, there should be a positive correlation between expected inflation and expected inflation volatility once we re-introduce a motive for stabilisation policy. We verify this by noting that the correlation between expected inflation and expected inflation volatility exceeds 0.94 in all simulations where the new shock plays a role. If the new shock is absent then the correlation is zero.

5.2 Escape dynamics

The escape dynamic analysis in Section 4 identifies the most likely path beliefs will take if they deviate significantly from the mean dynamics. To map the dominant escape path results into simulations, we initialise beliefs (γ_0, γ_1) at the self-confirming equilibrium values $(10, -1)$ and simulate the model until beliefs deviate significantly from their initial values.⁹ Table 2 shows average times to first escape in 1000 simulations of different parameterisations of the model. 90% confidence intervals are in parentheses.

Gain ε	$\sigma_3 = 0$	$\sigma_3 = 0.3$	$\sigma_3 = 0.6$	$\sigma_3 = 0.9$
0.0025	1088 (329,2495)	1587 (429,3834)	6852 (937,20016)	228256 (12714,664988)
0.0050	421 (130,981)	570 (168,1317)	1453 (299,4330)	13540 (1017,40156)
0.0075	246 (79,595)	315 (97,735)	773 (187,2020)	3930 (409,11444)
0.0100	170 (54,394)	213 (67,498)	444 (116,1113)	1891 (247,5647)

Table 2: Simulated average time to first escape

Times to first escape in simulations validate the predictions of Section 4.2. The dependence of the dominant escape path on the degree of stabilisation incentives is mirrored in the longer

⁹We define an escape as occurring when beliefs first deviate from their self-confirming values by a Euclidean distance in excess of 5.

times to first escape seen as our new shock becomes more important. A predicted extra delay of 55% in the dominant escape path for the case $\sigma_3 = 0.9$ translates into an average 58% delay in the (logarithmic) time to first escape in simulations.¹⁰ As expected, times to first escape fall as the gain parameter rises and government places less weight on past data.

In Section 4.3, the correlation between expected inflation and expected inflation volatility becomes positive if we re-introduce stabilisation policy along the dominant escape path. This prediction is corroborated by simulation evidence, since the correlation between expected inflation and expected inflation volatility is greater than 0.63 for all simulations with $\sigma_3 > 0$. If $\sigma_3 = 0$ the correlation is zero.

6 Conclusions

At the beginning of this paper, we set ourselves the challenge of explaining why inflation was so high and volatile in the 1970s. Our response is built on the work of Sargent (1999), and shows that changing government beliefs can explain the observed behaviour of inflation and inflation volatility from the 1960s to the 1990s. The story articulated in Sargent's *The Conquest of American Inflation* attributes the rise and fall in inflation to the discovery and subsequent disappearance of the Phillips curve, but is mute on why inflation volatility changes at the same time. To complete the story, we relax one of the simplifications in Sargent (1999) and restore the ability of government to react to some of the shocks hitting the economy. In this case, the discovery of the Phillips curve in the 1960s not only tempts government to raise inflation to reduce unemployment, but also promotes inflation volatility as the government has an incentive to let inflation vary to offset unemployment shocks. Conversely, the disappearance of the Phillips curve in the 1980s leads the government to abandon its high and volatile inflation policy. In our story, the rise and fall in inflation volatility is therefore intrinsically linked to the same changing government beliefs that cause inflation to rise and fall.

Why inflation fell at the time it did in the 1980s is only partially explained by Sargent (1999). According to his story, inflation fell due to a rare sequence of shocks that triggered

¹⁰The translation from continuous time t to discrete time n is given by $t = \log n \approx \log S_0 + \bar{S}/\varepsilon$, see Williams (2001).

government doubts about the existence of a Phillips curve. These doubts were self-confirming in the short run, so the high inflation policy was rapidly abandoned and the Phillips curve disappeared. This application of escape dynamics can explain the fall in inflation, but the question remains of why the rare sequence of shocks that triggered the escape occurred at the time it did. Our paper sheds some light on this. We show that an escape to low inflation is much more likely if there are relatively few shocks that the government can react to. Taken literally, this suggests that the fall in inflation in the 1980s may have been precipitated by a fall in the variance of shocks motivating stabilisation policy. In this respect, we complement the papers by Gerali and Lippi (2002) and McGough (2005) that attempt to explain the timing of escape episodes.

References

- [1] Cho, I-K, Williams, N. and T.J. Sargent, 2002, Escaping Nash Inflation, *Review of Economic Studies*, 69, 1-40.
- [2] Cogley, T., Colacito, R. and T.J. Sargent, 2005, The Benefits of U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas, *mimeo*.
- [3] Cogley, T. and T.J. Sargent, 2004, The Conquest of U.S. Inflation: Learning and Robustness to Model Uncertainty, *mimeo*.
- [4] El-Gamal, M. and R.K. Sundaram, 1993, Bayesian Economists ... Bayesian Agents: An Alternative Approach to Optimal Learning, *Journal of Economic Dynamics and Control*, 17, 355-83.
- [5] Evans, G.W. and S. Honkapohja, 2001, *Learning and Expectations in Macroeconomics*, Princeton University Press.
- [6] Gerali, A. and F. Lippi, 2002, On the ‘Conquest’ of Inflation, *mimeo*.
- [7] Giordani, P. and P. Söderlind, 2003, Inflation Forecast Uncertainty, *European Economic Review*, 47, 1037-59.
- [8] Kreps, D.M., 1998, Anticipated Utility and Dynamic Choice, in *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*, Cambridge University Press.
- [9] McGough, B., 2005, Shocking Escapes, *Economic Journal*, forthcoming.
- [10] Orphanides, A., 2003, The Quest for Prosperity without Inflation, *Journal of Monetary Economics*, 50, 633-63.
- [11] Phelps, E.S. and J.B. Taylor, 1977, Stabilizing Powers of Monetary Policy under Rational Expectations, *Journal of Political Economy*, 85, 163-190.
- [12] Primiceri, G., 2005, Why Inflation Rose and Fell: Policymakers’ Beliefs and US Postwar Stabilization Policy, *NBER Working Paper*, No. 11147.

- [13] Reis, R., 2003, Where Is the Natural Rate? Rational Policy Mistakes and Persistent Deviations of Inflation from Target, *Advances in Macroeconomics, The B. E. Journals in Macroeconomics*, vol. 3 (1), 2003.
- [14] Sargent, T.J., 1999, *The Conquest of American Inflation*, Princeton University Press.
- [15] Williams, N., 2001, Escape Dynamics in Learning Models, *PhD thesis*, University of Chicago.

A Condition for stability of SCE

A sufficient condition for local asymptotic stability is that all the eigenvalues of the Jacobian of equations (3.3) - (3.4) have negative real parts at the SCE. The Jacobian is defined by (A.1).

$$J = \begin{pmatrix} \frac{\partial R^{-1}\bar{g}(\gamma)}{\partial \gamma} & \frac{\partial R^{-1}\bar{g}(\gamma)}{\partial R} \\ \frac{\partial \bar{M}(\gamma)}{\partial \gamma} & -I \end{pmatrix} \quad (\text{A.1})$$

When evaluated at the SCE, the Jacobian is given by (A.2).

$$J|_{SCE} = \begin{pmatrix} \bar{R}^{-1}D\bar{g}(\gamma) & 0 \\ D\bar{M}(\gamma) & -I \end{pmatrix} \quad (\text{A.2})$$

The matrix of zeros off the leading diagonal gives the Jacobian a block recursive structure, implying that its eigenvalues are equal to the eigenvalues of the matrices on the leading diagonal. The eigenvalues of $-I$ trivially have negative real parts, so the stability properties of the SCE depend on the eigenvalues of $\bar{R}^{-1}D\bar{g}(\gamma)$. \bar{R} is a positive definite matrix, in which case stability only requires that the real parts of the eigenvalues of $D\bar{g}(\gamma)$ are negative.¹¹ After simple but tedious calculations, $D\bar{g}(\gamma)$ at the SCE takes the form of equation (A.3).

$$D\bar{g}(\gamma) = \frac{1}{1+\theta^2} \begin{pmatrix} -1 & -2u\theta \\ -u\theta & -2u^2\theta^2 - (1+\theta^2) \left(\sigma_2^2 + \left(\frac{\theta}{1+\theta^2} \right)^2 \sigma_3^2 \right) \end{pmatrix} \quad (\text{A.3})$$

The signs of the real parts of the eigenvalues of (A.3) are the same as those of matrix (A.4), where we use the notation σ' to denote the composite variance term in (A.3).

$$\begin{pmatrix} -1 & -2u\theta \\ -u\theta & -2u^2\theta^2 - \sigma' \end{pmatrix} \quad (\text{A.4})$$

¹¹Since $D\bar{g}(\gamma)$ and \bar{R} are square matrices, the QZ decomposition implies we can write $Q\Lambda Z = D\bar{g}(\gamma)$ and $Q\Omega Z = \bar{R}$. It follows that $\bar{R}^{-1}D\bar{g}(\gamma) = Z^{-1}\Omega^{-1}Q^{-1}Q\Lambda Z = Z^{-1}\Omega^{-1}\Lambda Z = Z^{*-1}\Omega^{*-1}\Lambda^*Z^*$, where Ω^* and Λ^* are diagonal matrices whose elements are the eigenvalues of \bar{R} and $D\bar{g}(\gamma)$. $Z^{*-1}\Omega^{*-1}\Lambda^*Z^*$ is the Jordan canonical form of $\bar{R}^{-1}D\bar{g}(\gamma)$, so the eigenvalues of $\bar{R}^{-1}D\bar{g}(\gamma)$ are equal to the eigenvalues of $D\bar{g}(\gamma)$ scaled by the eigenvalues of \bar{R} . Furthermore, when the matrix \bar{R} is positive definite it has eigenvalues with positive real parts, in which case the sign of the real parts of the eigenvalues of $\bar{R}^{-1}D\bar{g}(\gamma)$ equals the sign of the real parts of the eigenvalues of $D\bar{g}(\gamma)$.

The eigenvalues of matrix (A.4) are given by equations (A.5) and (A.6).

$$\frac{-(1 + 2u^2\theta^2 + \sigma') - \sqrt{(1 + 2u^2\theta^2 + \sigma')^2 - 4\sigma'}}{2} \quad (\text{A.5})$$

$$\frac{-(1 + 2u^2\theta^2 + \sigma') + \sqrt{(1 + 2u^2\theta^2 + \sigma')^2 - 4\sigma'}}{2} \quad (\text{A.6})$$

Simple algebra shows that $(1 + 2u^2\theta^2 + \sigma')^2 - 4\sigma' > 0$, in which case (A.5) always has a negative real part. (A.6) has a negative real part if condition (A.7) is satisfied.

$$(1 + 2u^2\theta^2 + \sigma') > \sqrt{(1 + 2u^2\theta^2 + \sigma')^2 - 4\sigma'} \quad (\text{A.7})$$

A necessary and sufficient condition for (A.7) to be satisfied is $\sigma' > 0$. Returning to the original notation for the composite variance term, the final condition for stability is therefore (A.8).

$$\left(\sigma_2^2 + \left(\frac{\theta}{1 + \theta^2} \right)^2 \sigma_3^2 \right) > 0 \quad (\text{A.8})$$

B Analytical expression for $a'Q(\gamma, R)a$

The cost function $Q(\gamma, R)$ is used to weight belief perturbations along potential escape paths. It is equal to the variance-covariance matrix of belief dynamics $\dot{\gamma}$. As belief dynamics are quadratic forms of Gaussian variables, Q itself is a fourth moment matrix. In static models such as ours, Williams (2001) shows that Q reduces to the logarithm of a moment generating function, meaning the Hamiltonian (4.4) can be derived analytically. We begin by expressing the first part of the Hamiltonian in terms of the corresponding moment generating function (B.1).

$$\frac{1}{2}a'Q(\gamma, R)a = \log E \exp \langle a \cdot R^{-1}(g(\gamma, \xi) - \bar{g}(\gamma)) \rangle \quad (\text{B.1})$$

The right hand side of (B.1) can be calculated explicitly using the definition of the forecast error (2.7), the precision matrix (2.8), the true structure of the economy (2.1) - (2.2) and government policy (2.11). The result is a linear-quadratic expression (B.2) in the three shocks W_1, W_2 and W_3 . The constants $d_1 \dots d_4$ and $n_1 \dots n_4$ are simple functions of the structural parameters $\{u, \theta, \sigma_1, \sigma_2, \sigma_3\}$ and beliefs γ .

$$\log E[e^{d_1 W_1 + d_2 W_2 + d_3 W_1 W_2 + d_4 W_2^2 + n_1 W_3 + n_2 W_1 W_3 + n_3 W_2 W_3 + n_4 W_3^2}] \quad (\text{B.2})$$

The next step is to factorise W_1 out from expression (B.2). The key stage in the factorisation below is the third line, where we exploit the fact that e^{W_1} is log-normally distributed, with expected value half the variance of W_1 .

$$\begin{aligned} & E[e^{d_1 W_1 + d_2 W_2 + d_3 W_1 W_2 + d_4 W_2^2 + n_1 W_3 + n_2 W_1 W_3 + n_3 W_2 W_3 + n_4 W_3^2}] \\ &= E[E(e^{(d_1 + d_3 W_2 + n_2 W_3)W_1} | W_2, W_3) e^{d_2 W_2 + d_4 W_2^2 + n_1 W_3 + n_3 W_2 W_3 + n_4 W_3^2}] \\ &= E[e^{.5(d_1 + d_3 W_2 + n_2 W_3)^2} e^{d_2 W_2 + d_4 W_2^2 + n_1 W_3 + n_3 W_2 W_3 + n_4 W_3^2}] \end{aligned}$$

The outcome of factorisation is an expression (B.3) in only two shocks, W_2 and W_3 . The constants $h_1 \dots h_5$ are simple algebraic manipulations of $d_1 \dots d_4$ and $n_1 \dots n_4$.

$$e^{.5d_1^2} E[e^{h_1 W_2 + h_2 W_3 + h_3 W_2 W_3 + h_4 W_2^2 + h_5 W_3^2}] \quad (\text{B.3})$$

We next factorise W_3 out from the expectation in (B.3). Conditioning on W_2 and collecting

terms in W_3 , we obtain expression (B.4).

$$\begin{aligned}
& E[e^{h_1 W_2 + h_2 W_3 + h_3 W_2 W_3 + h_4 W_2^2 + h_5 W_3^2}] \\
&= E[E(e^{(h_2 + h_3 W_2) W_3 + h_5 W_3^2} | W_2) e^{h_1 W_2 + h_4 W_2^2}]
\end{aligned} \tag{B.4}$$

The conditional expectation in (B.4) can be solved analytically by defining $k_1 = h_2 + h_3 W_2$, $k_2 = h_5 - .5$ and completing the square of $k_1 x + k_2 x^2$. In expression (B.5), we have $a = \sqrt{-2k_2}$, $b = -k_1/a$, $C = -b^2/2$ and the conditional expectation is proportional to $e^{r_1 W_2 + r_2 W_2^2}$ for suitably-defined constants r_1 and r_2 .

$$\begin{aligned}
E(e^{(h_2 + h_3 W_2) W_3 + h_5 W_3^2} | W_2) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} e^{k_1 x + k_2 x^2} dx \\
&= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} e^{-.5(ax+b)^2 - C} dx \\
&= \frac{1}{a} e^{-C} \\
&\propto e^{r_1 W_2 + r_2 W_2^2}
\end{aligned} \tag{B.5}$$

Substituting (B.5) back into (B.4) gives an expression in only the W_2 shock. By completing the square again, we are able to factorise out W_2 and obtain a final analytical expression for $a'Q(\gamma, R)a$ in terms of only structural parameters and beliefs.