

Macroeconomics and Volatility:

Data, Models, and Estimation*

Jesús Fernández-Villaverde[†] Juan F. Rubio-Ramírez[‡]

December 8, 2010

Abstract

One basic feature of aggregate data is the presence of time-varying variance in real and nominal variables. Periods of high volatility are followed by periods of low volatility. For instance, the turbulent 1970s were followed by the much more tranquil times of the great moderation from 1984 to 2007. Modeling these movements in volatility is important to understand the source of aggregate fluctuations, the evolution of the economy, and for policy analysis. In this chapter, we first review the different mechanisms proposed in the literature to generate changes in volatility similar to the ones observed in the data. Second, we document the quantitative importance of time-varying volatility in aggregate time series. Third, we present a prototype business cycle model with time-varying volatility and explain how it can be computed and how it can be taken to the data using likelihood-based methods and non-linear filtering theory. Fourth, we present two “real life” applications. We conclude by summarizing what we know and what we do not know about volatility in macroeconomics and by pointing out some directions for future research.

Keywords: DSGE models, Stochastic volatility, Bayesian methods.

JEL classification numbers: E10, E30, C11.

*We thank Pablo Guerrón, a coauthor in some of the research discussed here, for useful comments, and Béla Személyi for invaluable research assistance. Beyond the usual disclaimer, we must note that any views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Finally, we also thank the NSF for financial support.

[†]University of Pennsylvania, NBER, CEPR, and FEDEA, <jesusfv@econ.upenn.edu>.

[‡]Duke University, Federal Reserve Bank of Atlanta, and FEDEA, <juan.rubio-ramirez@duke.edu>.

1. Introduction

Macroeconomics is concerned with the dynamic effects of shocks. For instance, the real business cycle research program originated with an investigation of the consequences of changes in productivity (Kydland and Prescott, 1982). Later, the new generation of monetary models of the late 1990s and early 2000s was particularly focused on shocks to monetary policy (Christiano, Eichenbaum, and Evans, 2005). In open macroeconomics, considerable attention has been devoted to shocks to the interest rate (Mendoza, 1991) or to the terms of trade (Mendoza, 1995). Similar examples can be cited from dozens of other subfields of macroeconomics, from asset pricing to macro public finance: researchers postulate an exogenous stochastic process and explore the consequences for prices and quantities of innovations to it.

Traditionally, one key feature of these stochastic processes was the assumption of homoscedasticity. More recently, however, economists have started to relax this assumption. In particular, they have started considering shocks to the variance of the innovations of the processes. A first motivation for this new research comes from the realization that time series have a strong time-varying variance component. The most famous of those episodes is the great moderation of aggregate fluctuations in the U.S. between 1984 and 2007, when real aggregate volatility fell by around one third and nominal volatility by more than half. A natural mechanism to generate these changes is to have shocks that also have themselves a time-varying volatility and to trace the effects of changes in volatility on aggregate dynamics.

A second motivation, particularly relevant since the summer of 2007, is that changes to the volatility of shocks can capture the spreading out of distributions of future events, a phenomenon that many observers have emphasized is at the core of the current crisis. For example, an increase in the variance of future paths of fiscal policy (a plausible description of the situation of many European countries) can be incorporated in a parsimonious way by a rise in the variance of the innovations to a fiscal policy rule in an otherwise standard dynamic stochastic general equilibrium (DSGE) model. Similarly, the higher volatility of sovereign debt markets can be included in our models as a higher variance in the innovations to a country-specific spread.

A third, and final motivation, is that, even when the main object of interest is the conditional mean, economists should care about time-varying volatility. As illustrated in two examples by Hamilton (2008), inference about means can be unduly influenced by high variance episodes and standard statistical tests can become misleading. For instance, if we do not control for time-varying variance, a true null hypothesis will be asymptotically rejected with probability one.

Thus, ignoring changes in volatility is simply not an option in many empirical applications even when we do not care about volatility per se.

In this paper, we want to study time-varying volatility with the help of DSGE models, the workhorse of modern macroeconomics and the most common laboratory for policy evaluation. How do we incorporate time-varying volatility in the models? How do we solve models with this time-varying volatility? How do we take them to the data? What are the policy implications of volatility?

To address these questions, the rest of this chapter is organized as follows. First, we review the existing literature. Instead of being exhaustive, we will focus on those papers that have a closer relation with the rest of the chapter. Second, we present data to make the case that time-varying volatilities are an important feature of macroeconomic time series. Then, we present a prototype real business cycle model with time-varying volatility and show how we compute it and take it to the data using a likelihood-based approach. We move them into the summary of two “real life” applications from our own previous work. We conclude by discussing what we know and what we do not know about time-varying volatility and by pointing out directions for future research.

2. Review of the Literature

In one form or another, economists have talked for a long time about time-varying volatility. Just to give an example that mixes theory, data, and policy, David Ricardo, in his defense of free trade on corn in the House of Commons explicitly talked about the volatility of corn prices as an important factor to consider in the design of trade policy (although he dismissed it as an argument for protection).¹ But it was perhaps Haavelmo’s 1944 work that opened the path for the modern understanding of changes in volatility. Haavelmo taught economists to think about observed time series as the realization of a stochastic process. Once this was accomplished, and since nothing in the idea implied that the variance of the stochastic process had to be constant, it was natural to start thinking about processes whose variances changed over time.

Unfortunately, for a long time, most of the procedures that economists used to incorporate time-varying volatility were ad hoc and lacked a sound foundation in probability theory. As late as the mid 1970s, two papers published in the *Journal of Political Economy*, one of the top journals of the profession, when trying to measure the time component in the variance of inflation, resorted to such simple devices as using the absolute value of the first difference of inflation (Khan, 1977)

¹David Ricardo, speech of 9 May 1822. Collected works, volume V, p. 184, Ricardo (2005).

or a moving variance around a moving mean (Klein, 1977). And even these primitive approaches were merely empirical and never made an explicit connection with theoretical models.

A major breakthrough came with Engle's (1982) paper on autoregressive conditional heteroscedasticity, or ARCH. Engle postulated that a fruitful way to study the evolution of variance over time of time series x_t was to model it as an autoregressive process that is hit by the square of the (scaled) innovation on the level of x_t . The beauty of the assumption was that it combined simplicity with its ability to deliver an estimation problem that was straightforward to solve using a scoring iterative maximum likelihood procedure and ordinary least squares. The empirical application in Engle's original paper was the estimation of an ARCH process for British inflation. Engle found that indeed time-varying components were central to understanding the dynamics of inflation.

The profession quickly embraced Engle's contribution. Furthermore, after Bollerslev (1986) expanded the original model and created the Generalized ARCH, or GARCH, researchers joined an arms race to name yet another ARCH that would provide an extra degree of flexibility in modeling the data: Nonlinear GARCH, or NGARCH (Engle and Ng, 1993), Exponential GARCH, or EGARCH (Nelson, 1991), Quadratic GARCH, or QGARCH (Sentana, 1995), or Threshold GARCH, or TGARCH (Zakoïan, 1994) are some of the most popular extensions, but Bollerslev (2010) has recently counted 139 variations.

But it was not in macro where ARCH models came to reign, as one might have guessed from Engle's original application. The true boom was in finance, where the research on volatility took on a life of its own. The reason was simple. Financial institutions are keenly interested in the amount of risk they load onto their books. This risk is a function of the volatility on the return of their assets (in fact, the Basel II regulatory capital requirements depended on the Value-at-Risk of a bank's portfolio and, hence, on the level of variance). Similarly, the price of many assets, such as options depends directly on their volatility. Finally, time-varying volatility is a simple way to generate fat tails in the distribution of asset returns, a salient property of the data. The availability of high frequency data complemented in a perfect way the previously outlined need to describe volatility by providing economists with large samples with which to estimate and test their models.

The situation changed with the publication of the work by Kim and Nelson (1998), McConnell and Pérez-Quirós (2000), and Blanchard and Simon (2001). These influential papers documented that the volatility of U.S. aggregate fluctuations had changed over time. While Kim and Nelson and McConnell and Pérez-Quirós highlighted a change in volatility around 1984, Blanchard and

Simon saw the great moderation as part of a long-run trend toward lower volatility only momentarily interrupted during the 1970s. In a famous review paper, Stock and Watson (2002) named this phenomenon the “great moderation,” a title that became so popular that it even jumped into the popular media (and became rather unfairly attached to economists’ alleged complacency during the real estate boom of the 2000s).

The documentation of the great moderation led to an exploration of its causes and of a need to have models with mechanisms that generated time-varying volatility. McConnell and Pérez-Quirós (2000) had already pointed out the possibility of better inventory control as one possible explanation of the great moderation. Other mechanisms put forward have included financial innovation (Dynan, Elmendorf, and Sichel, 2006) and, in an well-cited study by Clarida, Galí, and Gertler (2000), changes in monetary policy.

A few years later, and in response to the previous work, Sims and Zha (2006) estimated a structural vector autoregression (SVAR) with Markov regime switching both in the autoregressive coefficients and in the variances of the disturbances. They found that the model that best fit the data had changes over time only in the variances of structural disturbances and no variation in the monetary rule or in the private sector of the model. But even when they allowed for policy regime changes, Sims and Zha found that the estimated changes could not account for the evolution of observed volatility. From those results, Sims and Zha concluded that models in which the innovations to the shocks had time-varying volatilities are a key element in the toolbox of applied macroeconomics.²

All of this research has convinced us that 1) time-varying volatility is an important feature of the data and that 2) we need DSGE models that allow us to generate it, quantify its effects, perform welfare analysis, and design optimal policy. First attempts in this direction are Fernández-Villaverde and Rubio-Ramírez (2007) and Justiniano and Primiceri (2008). These papers estimate DSGE economies that incorporate stochastic volatility on the structural shocks and show that such models fit the data considerably better than economies with homoscedastic structural shocks. More recently, Christiano, Motto, and Rostagno (2009) have shown that, in a financial accelerator model, shocks to the volatility of individual firms’ productivity have a significant impact on the business cycle because of their consequences for the level of leverage that firms can take. A related result is found by Arellano, Bai, and Kehoe (2010).

²Sims and Zha’s conclusion is, nevertheless, not incontrovertible. Benati and Surico (2009) illustrate that it is difficult to map between changes in the autoregressive coefficients or in the variance of disturbances in a regime-switching SVARs and equivalent elements in a New Keynesian DSGE model. This would be a key motivation for our application in section 5.

Another strand of the literature starts from the real-option effect of risk. In a situation where investment (in capital, durable goods or any similar item) is subject to frictions such as irreversibilities or non-convex adjustment costs, a change in volatility may have a substantial effect on the investment decision. Think, for example, about a household's decision to buy a new car to substitute its old clunker. If labor market volatility increases, the household may be quite concerned about its own job status in the next few months. By delaying the purchase of a new car, the household loses the differential utility between the services of the old and the new car times the length of the delay. On the other hand, it avoids both the costs of purchasing an expensive item and the risk of facing a liquidity constraint that may force the household to sell the car (with a loss of value) or re-adjust other consumption items. This mechanism is particularly well explored by Bloom (2009) and in Bloom, Jaimovich, and Floetotto (2008).

Guerrón-Quintana (2009) finds that volatility shocks à la Bloom induce depreciations in the real exchange rate in the US, particularly vis-a-vis the Canadian dollar. Fatás (2002) discusses the effects of business cycle volatility on growth. Lee, Ni, and Ratti (1995) show that the conditional volatility of oil prices matter for the effect of oil shocks on the economy. Grier and Perry (2000) and Fountas and Karanasos (2007) relate inflation and output volatility with average output growth, while Elder (2004a and 2004b) links nominal and real volatility.

Of course, the importance of these observations and models is not universally accepted (see Bachmann, Elstner, and Sims, 2010, for a much less sanguine reading of the importance of volatility shocks), but we judge that the preponderance of the evidence is clearly on the side of time-varying volatility. To show this, we start now with a brief summary of some data that will help us to understand better the literature we just discussed.

3. Data

In this section we illustrate the presence of time-varying volatility in two contexts that we will revisit later in the paper: fluctuations in the U.S. economy and fluctuations in the interest rates at which small open emerging economies borrow.

We start with the evolution of aggregate variables in the U.S. In that way we document (once more) the great moderation, which has been the motivating fact of much of the literature on time-varying volatility. In figure 3.1, we plot the absolute deviations of real GDP growth with respect to their mean. In this figure we can see how, since 1984, the absolute deviation rarely crosses 4 percentage points (except in a couple of brief spikes around the 1992 and 2008-2009

recession), while before it did it rather often. Even the great recession of 2008-2009 did not imply a difference in growth rate as big as the two Volcker recessions (although the 2008-2009 recession was longer). Besides, we can also see fat tails in the distribution of deviations.




Figure 3.1: Real GDP Growth, Absolute Deviations from Mean

This change in volatility also appears in nominal variables. Figure 3.2 plots the absolute deviations of the GDP deflator with respect to its mean. Again, we see how the big spikes of the 1970s and early 1980s disappeared after 1984 and they did not come back even briefly in the last recession.




Figure 3.2: GDP Deflator, Absolute Deviations from Mean

In table 3.1, we summarize the graphical information into statistical moments for the sample 1959.Q1 to 2007.Q1 that we will use in section 5 and we add the federal funds rate as a measure of monetary policy. These three variables, inflation, output growth, and the federal funds rate are the most commonly discussed series in monetary models (for example, the “trinity” model so dear to the New Keynesian tradition has only these three variables). We can see in table 3.1 how the standard deviation of inflation falls by 60 percent after 1984.Q1, the standard deviation of output growth by 44 percent, and the standard deviation of the federal funds rate by 39 percent. Again, the evidence of changes in variances over time is rather incontrovertible.

Table 3.1: Changes in Volatility of U.S. Aggregate Variables

	Means			Standard Deviations		
	Inflation	Output Growth	FFR	Inflation	Output Growth	FFR
All sample	3.8170	1.8475	6.0021	2.6181	3.5879	3.3004
Pre 1984.Q1	4.6180	1.9943	6.7179	3.2260	4.3995	3.8665
After 1984.Q1	2.9644	1.6911	5.2401	1.3113	2.4616	2.3560
Post-1984.Q1/pre-1984.Q1	0.6419	0.8480	0.7800	0.4065	0.5595	0.6093

Our second example of time-varying volatility is figure 3.3, where we use the Emerging Markets Bond Index+ (EMBI+) Spread reported by J.P. Morgan at a monthly frequency to plot the country spreads of Argentina, Brazil, Ecuador and Venezuela. This index tracks secondary market prices of actively traded emerging market bonds denominated in U.S. dollars. For comparison purposes, we also plot the real U.S. T-bill rate as a measure of the international risk-free nominal interest rate. We build the real T-bill rate by subtracting expected inflation measured as the average U.S. CPI inflation in the current month and in the eleven preceding months. This is motivated by the observation that U.S. inflation is well approximated by a random walk. The results are nearly identical with more sophisticated methods to back up expected inflation. Both the T-bill rate and the inflation series are obtained from the St. Louis Fed’s FRED database. We use annualized rates in percentage points.

In this figure we can see how the international risk-free real rate is low (with negative interest rates in 2002-2006) and relatively stable over the sample. In comparison, all country spreads are large and volatile, with times of turbulence following much calmer months. The spreads are nearly always larger than the real T-bill rate itself and fluctuate, at least, an order of magnitude more. The most prominent case is Argentina, where the 2001-2002 crisis raised the country spreads to

70 percentage points. In the figure, we also see the problems of Ecuador in 1998-1999 and the turbulence in all four countries during the virulent international turmoil of 1998.




Figure 3.3: Country Spreads and T-Bill Real Rate

Besides the data in these figures, we could present many others, such as those in Bloom (2009). However, we feel we have already made the case for the empirical relevance of time-varying volatility and it seems a better use of our allocated space to jump into the substantive questions by presenting a prototype business cycle model where volatility changes over time.

4. A Prototype Business Cycle Model with Time-Varying Volatility

A simple exercise to illustrate the theoretical, computational, and empirical issues at hand when we deal with DSGE models that incorporate changes in variances is to write down a prototype economy and to introduce in it the minimum modifications required to capture time-varying volatility in a plausible way. The perfect vehicle for such a pedagogical effort is the real business cycle model for two reasons.

First, the stochastic neoclassical growth model is the foundation of modern macroeconomics. Even the more complicated New Keynesian models are built around the core of the neoclassical growth model augmented with nominal and real rigidities. Thus, once we understand how to deal with time-varying volatility in our prototype economy, it will be straightforward to extend it to richer environments. Second, the model is so well known, its working so well understood,

and its computation so thoroughly explored that the role of time-varying volatility in it will be staggeringly transparent.

Once we are done with our basic model, we will move on to analyzing two applications, one in monetary economics and one in international macroeconomics, where changes in volatility play a key role. While these applications are more complicated than our prototype economy, they are explicitly designed to account for a richer set of observations and to demonstrate the usefulness of DSGE models with time-varying volatility in “real life.”

4.1. Environment

To get into the substantive questions as soon as possible, our description of the standard features of our prototype economy will be limited to fixing notation. There is a representative household in the economy, whose preferences over stochastic sequences of consumption, c_t , and work, l_t , are representable by a utility function:

$$U = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, l_t) \quad (1)$$

where $\beta \in (0, 1)$ is the discount factor and \mathbb{E}_0 is the conditional expectation operator. We leave the concrete parameterization of the utility function open since we will consider below the effects of different period utility kernels.

The household’s budget constraint is given by:

$$c_t + i_t + \frac{b_{t+1}}{R_t} = w_t l_t + r_t k_t + b_t$$

where i_t is investment, R_t is the risk-free gross interest rate, b_t is the holding of an uncontingent bond that pays 1 unit of consumption good at time $t+1$, w_t is the wage, l_t is labor, r_t is the rental rate of capital, and k_t is capital. Asset markets are complete and we could have also included in the budget constraint the whole set of Arrow securities. Since we have a representative household, this is not necessary because the net supply of any security must be equal to zero. The uncontingent bond is all we need to derive a pricing kernel for the economy. Capital is accumulated according to the law of motion $k_{t+1} = (1 - \delta)k_t + i_t$ where δ is the depreciation rate.

The final good is produced by a competitive firm with a technology $y_t = e^{z_t} A k_t^\alpha l_t^{1-\alpha}$ where z_t is the productivity level whose evolution we will describe momentarily and A is a constant. Thus, the economy must satisfy the aggregate resource constraint $y_t = c_t + i_t$.

Productivity follows an autoregressive process $z_t = \lambda z_{t-1} + \sigma_t \varepsilon_t$ with $\lambda < 1$ and random innovations $\varepsilon_t \sim \mathcal{N}(0, 1)$. We impose stationarity in the process to save on notation (otherwise we would need to rescale the variables in the model by the level of technology), but besides the notational burden, it would be easy to have a martingale on z_t . Note, and here is where we are introducing time-varying volatility, that the standard deviation of innovations, σ_t , is indexed by the period t . That is, the dispersion of the productivity shocks changes over time: sometimes there are large shocks, sometimes there are smaller shocks. Our specification is extremely simple and we present it only as a default process to start the conversation.

The first question that we need to handle at this point is how to model these changes in volatility. The literature has proposed three alternatives: stochastic volatility, GARCH processes, and Markov regime switching.

The first approach is stochastic volatility, or SV. More concretely, it assumes that σ_t evolves over time as an autoregressive process, for example, with the form:

$$\log \sigma_t = (1 - \rho_\sigma) \log \sigma + \rho_\sigma \log \sigma_{t-1} + \eta u_t, \text{ where } u_t \sim \mathcal{N}(0, 1) \quad (2)$$

The law of motion is expressed in terms of logs to ensure the positivity of σ_t . This is a point that will be important later: by mixing levels (z_t) and logs ($\log \sigma_t$), we create a structure that is inherently non-linear and it twists the distribution of technology. This will have consequences both for the solution and for the estimation of the model.

Our specification (2) is parsimonious and it introduces only two new parameters, ρ_σ , the autoregressive coefficient of the log standard deviation, and η , the standard deviation of the innovations to volatility. At the same time, it is surprisingly powerful in capturing some important features of the data (Shephard, 2008). Another important point is that, with SV, we have two innovations, an innovation to technology, ε_t , and an innovation to the standard deviation of technology, u_t . As we will see below, this will help the researcher to sort out the specific effects of volatility per se.³

The second approach is to specify that the variance of the productivity innovations follows a GARCH process $\sigma_t^2 = \omega + \alpha (\sigma_{t-1} \varepsilon_{t-1})^2 + \beta \sigma_{t-1}^2$, that is, σ_t^2 is a function of its own past and the squared scaled innovation $((\sigma_{t-1} \varepsilon_{t-1})^2)$. As with SV, instead of our simple GARCH, we could

³It is trivial to correlate ε_t and u_t . For example, in the data, times of large volatility such as the 1970s are often also times of low productivity growth. In international macro, times of large spreads are also times of high volatility. This correlation is sometimes called the “leverage effect” of level shocks on volatility shocks because, in asset pricing, one can generate it through the presence of leverage in the firm’s balance sheet.

think about any of the many incarnations of GARCHs mentioned in section 2. Most of what we have to say in the next few lines would be unchanged.

In the GARCH specification there is only one shock driving the dynamics of the level and volatility of technology: ε_t . This means that, when we have a large innovation, we will have a large volatility in the next period. Thus, we cannot separate a volatility shock from a level shock: higher volatilities are triggered only by large level innovations. While this constraint may not be very important when we are dealing with time series from a reduced-form perspective, it is quite restrictive in structural models. In particular, the interconnection of levels and volatilities precludes the use of GARCH models to assess, in a DSGE model, the effects of volatility independently from the effects of level shocks.

Another way to think about it is as follows. In time series analysis, GARCHs are a popular alternative to stochastic volatility because they are much easier to estimate and the loss in empirical fit is minor. In the case of DSGE models, this simplicity advantage disappears because, with either SV or GARCH, we need to solve the model non-linearly. Not only that, but, as we argued before, the presence of two shocks in SV provides the researcher with an extra degree of freedom that can be put to good use.

The third approach to time-varying volatility is Markov regime switching models. For instance, we can postulate that σ_t follows a Markov chain that takes two values, σ^L and σ^H , where L stands for low and H stands for high ($\sigma^L < \sigma^H$), and with transition matrix:

$$\begin{pmatrix} a_1 & 1 - a_1 \\ 1 - a_2 & a_2 \end{pmatrix}$$

where a skillful choice of a_1 and a_2 allows us to introduce a large range of behaviors (for example, $a_1 \gg a_2$ can be read as low volatility being the normal times and high volatility as the rare times). Moreover, there is nothing special about two values of volatility and we could have an arbitrary number of them.

A big difference between this approach and the previous two is the size of the change. We can interpret both SV and GARCH processes as reflecting a continuously changing process that has innovations in every period. In comparison, Markov regime switching models evolve in a more abrupt, discrete way, with sudden jumps interrupted by periods of calm.

In the rest of the paper we will follow the first approach, SV, but we will say a few words about GARCH and Markov regime switching as we move along. As we argued before, we do

not really see any advantage to using a GARCH process instead of SV: it has one less degree of freedom, it prevents us from neatly separating level from volatility shocks, it fits the data worse, and, in the context of DSGE models, it is not any easier to handle. The choice between SV and Markov-regime switching is more subtle. In the real world, the change in the volatility of technology is probably a mix of continuous and discrete events. While there are phenomena affecting technological change that are easier to interpret as a discrete change (for example, the approval of a new patent law), other developments (such as the growth in our understanding of natural laws) are probably better understood as continuous changes. The preference for one or another is an empirical question.

We could even postulate a more encompassing approach that incorporate discrete jumps and continuous changes. The problem with such a model would be that, with the data frequency in macro, we do not have enough observations to tease out these two sources of variation (as we would have, for instance, in finance, where continuous time versions of this process have been taken to the data, see the review in Aït-Sahalia, Hansen, and Scheinkman, 2009). This is disappointing because, as first pointed out by Diebold (1986), ignoring jumps may severely bias the estimates of ρ_σ towards one, creating the misleading impression of non-stationarities and invalidating inference.

One advantage of SV, which we will exploit below and that tips the balance in its favor, is that, since under that specification $\log \sigma_t$ can take any value, we will be able to differentiate the decision rules of the agents in the economy with respect to it, and hence to apply perturbation methods for the computation of the equilibrium dynamics, which are a fast and reliable algorithm.⁴ This is not the case with Markov regime switching models since $\log \sigma_t$ takes only a finite set of values.

However, it is fair to point out that SV has a few problems of its own. A salient one is that, if the real process has a discrete jump, SV will “anticipate” the change by showing changes in volatility before they happen. The reason is that the likelihood (or most other estimating functions) dislikes huge changes in one period and prefers a sequence of smaller u_t over time before and after the actual change to an exceptionally large u_t that captures the jump.⁵

⁴Unfortunately, we do not have proof that the decision rules are differentiable with respect to $\log \sigma_t$. As we will explain later, this is one of the many issues related to volatility that we do not fully understand.

⁵This could also be a virtue. Coming back to our example of a new patent law, we could think about a situation where the volatility of technological change evolves over time as the proposal goes through the legislative process and hence the conditional probability of its approval changes. Whether the anticipation effect is a feature or a bug would depend on the context.

4.2. Equilibrium

The definition of competitive equilibrium of this model is standard and we include it to demonstrate how we are deviating only a minuscule amount from the standard model.

Definition 1. *A competitive equilibrium is a sequence of allocations $\{c_t, l_t, i_t, y_t\}_{t=0}^{\infty}$ and prices $\{w_t, r_t, R_t\}_{t=0}^{\infty}$ such that:*

1. *Given prices $\{w_t, r_t, R_t\}_{t=0}^{\infty}$, the representative household maximizes:*

$$U = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, l_t)$$

$$\text{s.t. } c_t + i_t + \frac{b_{t+1}}{R_t} = w_t l_t + r_t k_t + b_t$$

2. *Given prices $\{w_t, r_t, R_t\}_{t=0}^{\infty}$, the firm minimizes costs given its production function:*

$$y_t = e^{z_t} A k_t^{\alpha} l_t^{1-\alpha} \quad (3)$$

3. *Markets clear:*

$$k_{t+1} = (1 - \delta)k_t + i_t \quad (4)$$

$$y_t = c_t + i_t \quad (5)$$

4. *Productivity follows:*

$$z_t = \lambda z_{t-1} + \sigma_t \varepsilon_t \quad (6)$$

$$\log \sigma_t = (1 - \rho_{\sigma}) \log \sigma + \rho_{\sigma} \log \sigma_{t-1} + \eta u_t \quad (7)$$

The presence of SV does not affect the welfare theorems and this economy is still Pareto optimal. While this is a convenient feature, our analysis of SV will not rely on it. In fact, neither of the economies in the two applications in the sections below will be Pareto-optimal.

4.3. Solution Methods

The solution of models with time-varying volatility presents some challenges. First, the system is, at its very essence, non-linear. If we are employing SV, we are combining a linear process for

the log of technology with a linear process for the log of the standard deviation of technology innovations. Analogously, in the other two specifications we discussed before, GARCH implies a quadratic law of motion and Markov regime switching a discrete support. Second, we have an additional state, $\log \sigma_t$, that agents need to keep track of in order to forecast future volatility.

4.3.1. Value Function Iteration

A first, natural approach is to work with the value function of the social planner problem:

$$V(k_t, z_t, \log \sigma_t) = \max_{c_t, l_t, k_{t+1}} \{u(c_t, l_t) + \beta \mathbb{E}_t V(k_{t+1}, z_{t+1}, \log \sigma_{t+1})\}$$

subject to (3), (4), (5), (6), and (7). This value function can be computed with value function iteration (VFI). The only conceptual difficulty is to ensure that the conditional expectation \mathbb{E}_t is properly evaluated at each point in time.

While VFI is a safe and straightforward procedure, it suffers from two shortcomings. First, it forces us to cast the problem in a recursive form, which may be difficult to do in economies with market imperfections or rigidities. Second, VFI suffers from the “curse of dimensionality” that limits the size of the problems we can handle. The curse of dimensionality is particularly binding when we deal with SV because we double the number of states for each stochastic process that incorporate a time-varying volatility: one state to capture the level of the process and one to keep track of the variance.

4.3.2. Working with the Equilibrium Conditions

A second solution is to work with the equilibrium conditions:

$$\begin{aligned} u_1(c_t, l_t) &= \mathbb{E}_t u_1(c_{t+1}, l_{t+1}) \beta (1 + r_{t+1} - \delta) \\ u_2(c_t, l_t) &= u_1(c_t, l_t) w_t \\ w_t &= (1 - \zeta) e^{z_t} A k_t^\zeta l_t^{1-\zeta} \\ r_t &= \zeta e^{z_t} A k_t^{\zeta-1} l_t^{1-\zeta} \end{aligned}$$

plus (3), (4), (5), (6), and (7). Equilibrium conditions enjoy the advantage that we do not need to rely on any social planner problem or on being able to write the model in terms of a Bellman equation.

The first step is to write the decision rules of the agents as a function of the states, $(k_t, z_{t-1}, \log \sigma_{t-1})$

and the two innovations (ε_t, u_t) . Thus, we have, for the three controls $c_t = c(k_t, z_{t-1}, \log \sigma_{t-1}, \varepsilon_t, u_t)$, $l_t = l(k_t, z_{t-1}, \log \sigma_{t-1}, \varepsilon_t, u_t)$, and $k_{t+1} = k(k_t, z_{t-1}, \log \sigma_{t-1}, \varepsilon_t, u_t)$, and for any other variable x_t defined by the model $x_t = x(k_t, z_{t-1}, \log \sigma_{t-1}, \varepsilon_t, u_t)$. Then, we plug these unknown decision rules into the equilibrium condition and solve the resulting system of functional equations.

This can be accomplished in two ways. The first alternative is to parameterize the unknown functions, for example, as $x_t = \sum_{i=0}^n \theta_i^x \Psi_i^x(k_t, z_{t-1}, \log \sigma_{t-1}, \varepsilon_t, u_t)$, where Ψ_i is a multivariate polynomial built with some combination of univariate polynomials of the 5 state variables (the tensor product of univariate Chebyshev polynomials is a default choice). Next, we plug the parameterized decision rules into the equilibrium conditions and we solve for all the unknown coefficients θ_i^x by making the equilibrium conditions to hold as closely as possible over the state space under some metric (for example, in a collocation, by forcing the equilibrium conditions to be zero at the zeros of the $n + 1$ -th order Chebyshev polynomial).

This approach, called a projection method (because we build a projection of the unknown decision rule into the parameterized approximated decision rule), has the advantage of delivering a high level of accuracy in the whole state space (it is a “global” solution method). As was the case with VFI, the only possible conceptual difficulty is the correct evaluation of the conditional expectation \mathbb{E}_t . On the negative side, we need to solve for a large number of θ_i^x coefficients to achieve a good level of accuracy with a five-dimensional problem, yet another manifestation of the curse of dimensionality.

The second approach to solve for the unknown decision rules in the equilibrium conditions is to build a higher-order perturbation, an approach that has been shown to be both accurate and fast (Aruoba, Fernández-Villaverde, and Rubio-Ramírez, 2006). The main idea is to find a Taylor approximation of the decision rules around the steady state of the model. The first step to doing so is to introduce a new parameter, called the perturbation parameter, Λ , and rewrite the stochastic process (6) and (7) as:

$$z_t = \lambda z_{t-1} + \Lambda \sigma_t \varepsilon_t \tag{8}$$

$$\log \sigma_t = (1 - \rho_\sigma) \log \sigma + \rho_\sigma \log \sigma_{t-1} + \Lambda \eta u_t \tag{9}$$

Then, if we make $\Lambda = 1$, we get back the original formulation of the problem. However, if we set $\Lambda = 0$, we eliminate the sources of uncertainty in the model and the economy will (asymptotically) settle down at the steady state.

The second step is to rewrite all variables in terms of deviations with respect to the steady

state. Thus, we write $\hat{x}_t = x_t - x$ for any arbitrary variable x_t with steady state x , except for $\log \sigma_{t-1}$ where $\hat{\sigma}_{t-1} = \log \sigma_{t-1} - \log \sigma$. Also, define an augmented state vector of the model

$$s_t = \left(\underbrace{\hat{k}_t, z_{t-1}, \hat{\sigma}_{t-1}}_{\mathbb{S}_{t-1}}, \underbrace{\varepsilon_t, u_t}_{\mathbb{W}_t}; \Lambda \right) = (\mathbb{S}_{t-1}, \mathbb{W}_t; \Lambda)$$

where we stack the states in deviations to the mean, \mathbb{S}_{t-1} , and innovations \mathbb{W}_t and we have incorporated the perturbation parameter, Λ , as a pseudo-state (where the “pseudo” is emphasized by the use of a semicolon to separate it from the pure states). Then, the decision rules we are looking for are $\hat{c}_t = c(s_t)$, $\hat{l}_t = l(s_t)$, and $\hat{k}_{t+1} = k(s_t)$.

To approximate them, we will search for the coefficients of the Taylor expansion of these decision rules evaluated at the steady state, $s = \mathbf{0}_{1 \times 5}$. For example, for consumption, we write:

$$\hat{c}_t = c(s_t) = c_{i,ss} s_t^i + \frac{1}{2} c_{ij,ss} s_t^i s_t^j + \frac{1}{6} c_{ijl,ss} s_t^i s_t^j s_t^l + H.O.T.$$

where each term $c_{...,ss}$ is a scalar equal to a derivative of the value function evaluated at the steady state, $c_{i,ss} \equiv c_i(s)$ for $i = 1, \dots, 5$, $c_{ij,ss} \equiv c_{ij}(s)$ for $i, j = 1, \dots, 5$, and $c_{ijl,ss} \equiv c_{ijl}(s)$ for $i, j, l = 1, \dots, 5$, where we follow the tensor notation $c_{i,ss} s_t^i = \sum_{i=1}^5 c_{i,ss} s_{i,t}$, $c_{ij,ss} s_t^i s_t^j = \sum_{i=1}^5 \sum_{j=1}^5 c_{ij,ss} s_{i,t} s_{j,t}$, and $c_{ijl,ss} s_t^i s_t^j s_t^l = \sum_{i=1}^5 \sum_{j=1}^5 \sum_{l=1}^5 c_{ijl,ss} s_{i,t} s_{j,t} s_{l,t}$, that eliminates the symbol $\sum_{i=1}^5$ when no confusion arises, and where we represent all the higher-order terms by $H.O.T.$ (it will become clear momentarily why we were explicit about the first three orders of the solution). We can proceed in analogous ways for all other variables and derive the appropriate formulae.

To find the coefficients $c_{i,ss}$, $c_{ij,ss}$, and $c_{ijl,ss}$, we take derivatives of the equilibrium conditions with respect to each component of s_t and solve for the resulting unknown coefficients that make these derivatives hold. Conveniently, this procedure is recursive; that is, we find the coefficients of each order of the approximation one step at a time. For example, by taking first derivatives of the equilibrium conditions with respect to s_t , we find all the coefficients of the first-order $c_{i,ss}$. Then, we take second derivatives of the equilibrium conditions with respect to s_t , we plug in the coefficients of the first-order $c_{i,ss}$ that we already know and we solve for the coefficients $c_{ij,ss}$, and so on for any arbitrary order. Furthermore, while in the first-order problem we have a quadratic system (with two solutions that satisfy the necessary conditions, one that violates the transversality condition and one that does not), all the higher-order systems are linear and therefore easy to solve.

In addition to all these coefficients, we also need to find a Taylor expansion of the stochastic processes (8) and (9) or in our transformed state variables:

$$z_t = \lambda z_{t-1} + \Lambda \sigma e^{\hat{\sigma}_t} \varepsilon_t \quad (10)$$

$$\hat{\sigma}_t = \rho_\sigma \hat{\sigma}_{t-1} + \Lambda \eta u_t \quad (11)$$

In standard DSGE models solved by linearization, this step is often overlooked because the conventional law of motion for z_t is already linear, but in our case, since we have the term $\sigma e^{\hat{\sigma}_t} \varepsilon_t$, we cannot avoid approximating (10) (equation 11 is already linear in the transformed variables). The reason is that, when we perform a perturbation, all the variables should be perturbed at the same order. This is required by the theorems that ensure that perturbation works (see Jin and Judd, 2002). The unfortunate practice, often seen in the literature, of mixing different orders of approximation, for instance, getting a first-order approximation for consumption and a second-order for the stochastic processes, is wrong.⁶ Beyond its theoretical flaw, mixing orders of approximation is not even particularly accurate and it is simple to show that standard measures as Euler equation errors deteriorate when we follow this practice.

While, theoretically, we could find all the derivatives of the decision rules and the exogenous processes and coefficients by paper and pencil, in practice, we employ some symbolic software to manipulate the equilibrium conditions of the model and take all the relevant derivatives. There are programming languages, such as **Mathematica**, which are particularly suited to these type of manipulations. Also, there is specific software developed in recent years for perturbation such as the **Dynare**, a pre-processor and a collection of **MATLAB** and **GNU Octave** routines that compute up to third-order approximations, or **Dynare++**, a standalone C++ version of **Dynare** that specializes in computing $n - th$ -order approximations.

4.3.3. Structure of the Solution

Our previous discussion gave us an abstract description of how to find the perturbation solution. However, it overlooked the fact that the perturbation solution of the model has a particular pattern that we can exploit. To make this point more generally, we switch in the next few paragraphs to a more abstract notation.

⁶This is also why we solve for consumption, labor, and capital. In principle, given two of these variables, we could find the third one using the resource constraint of the economy. But this would imply that we are solving two variables up to order n and the third one nonlinearly.

The set of equilibrium conditions of a large set of DSGE models, including the real business cycle model with SV in this section, can be written in a compact way as:

$$\mathbb{E}_t f(\mathcal{Y}_{t+1}, \mathcal{Y}_t, \mathcal{S}_{t+1}, \mathcal{S}_t, \mathcal{Z}_{t+1}, \mathcal{Z}_t) = 0 \quad (12)$$

where \mathbb{E}_t is the conditional expectation operator at time t , $\mathcal{Y}_t = (\mathcal{Y}_{1t}, \mathcal{Y}_{2t}, \dots, \mathcal{Y}_{kt})$ is the vector of non-predetermined variables of size k (such as consumption or labor), $\mathcal{S}_t = (\mathcal{S}_{1t}, \mathcal{S}_{2t}, \dots, \mathcal{S}_{nt})$ is the vector of endogenous predetermined variables of size n (such as capital), $\mathcal{Z}_t = (\mathcal{Z}_{1t}, \mathcal{Z}_{2t}, \dots, \mathcal{Z}_{mt})$ is the vector of exogenous predetermined variables of size m , which we refer to as structural shocks (such as productivity), and f is a mapping from $\mathbb{R}^{2 \times k + 2 \times n + 2 \times m}$ into \mathbb{R}^{k+n+m} .

We assume that structural shocks follow an SV process of the form $\mathcal{Z}_{it+1} = \rho_i \mathcal{Z}_{it} + \Lambda \sigma_{it+1} \varepsilon_{it+1}$ where the standard deviation of the innovations evolves as $\log \sigma_{it+1} = \vartheta_i \log \sigma_{it} + \Lambda \eta_i u_{it+1}$ for all $i = \{1, \dots, m\}$ and Λ is still the perturbation parameter. To avoid carrying extra indices, we are assuming that all structural shocks face volatility shocks. By setting the appropriate entries of ϑ_i and η_i to zero, we can easily handle homoscedastic shocks. We are also assuming that the volatility shocks are uncorrelated. This restriction can also be relaxed.

The solution to the system of functional equations defined by (12) can be expressed in terms of two equations, one $\mathcal{S}_{t+1} = h(\mathcal{S}_t, \mathcal{Z}_{t-1}, \Sigma_{t-1}, \mathcal{E}_t, \mathcal{U}_t, \Lambda)$, describing the evolution of predetermined variables, and another, $\mathcal{Y}_t = g(\mathcal{S}_t, \mathcal{Z}_{t-1}, \Sigma_{t-1}, \mathcal{E}_t, \mathcal{U}_t, \Lambda)$, describing the evolution of non-predetermined ones, where $\Sigma_t = (\log \sigma_{1t}, \log \sigma_{2t}, \dots, \log \sigma_{mt})$, $\mathcal{E}_t = (\varepsilon_{1t}, \varepsilon_{2t}, \dots, \varepsilon_{mt})$, and $\mathcal{U}_t = (u_{1t}, u_{2t}, \dots, u_{mt})$. More intuitively, we think of Σ_t as the volatility shocks, \mathcal{E}_t are the innovations to the structural shocks, and \mathcal{U}_t are innovations to volatility shocks.

As we described in the previous subsection, we are seeking a higher-order approximation to the functions $h(\cdot) : \mathbb{R}^{n+(4 \times m)+1} \rightarrow \mathbb{R}^n$ and $g(\cdot) : \mathbb{R}^{n+(4 \times m)+1} \rightarrow \mathbb{R}^k$ around the steady state, $\mathcal{S}_t = \mathcal{S}$ and $\Lambda = 0$. While a general characterization of these functions is difficult, it is surprisingly easy to obtain substantial results regarding the first- and second-order derivatives of the functions $h(\cdot)$ and $g(\cdot)$ evaluated at the steady state.⁷ In particular, we formally show in Fernández-Villaverde, Guerrón-Quintana, and Rubio-Ramírez (2010a) (hereafter, FGR) that the first partial derivative of $h(\cdot)$ and $g(\cdot)$ with respect to any component of \mathcal{U}_t and Σ_{t-1} evaluated at the steady state is zero. In other words, volatility shocks and their innovations do not affect the linear component of the optimal decision rule of the agents for any $i = \{1, \dots, m\}$. The same occurs with the

⁷We conjecture, based on our numerical results, that there exists relatively direct (yet cumbersome to state) extensions of our theorem for higher-order terms.

perturbation parameter Λ . This is not a surprising result since Schmitt-Grohé and Uribe (2004) have stated a similar theorem for the homoscedastic shocks case. The theorem also shows that the second partial derivative of $h(\cdot)$ and $g(\cdot)$ with respect to $u_{i,t}$ and any other variable but $\varepsilon_{i,t}$ is also zero for any $i = \{1, \dots, m\}$.

The interpretation of the theorem is simple. The first part just states that variances or their evolution do not enter in the first-order component of the solution of the model. This is nothing more than certainty equivalence: a first-order approximation is equivalent to a model with quadratic utility functions and where, consequently, agents do not respond to variance. It is only in the second-order component of the solution that we have terms that depend on the variance since those depend on the third derivative of the utility function. In particular, we will have a constant that corrects for risk.

But even in the second-order, time-varying volatilities enter into the solution in a very restricted way: through the interaction term of the innovations to the structural shocks and the innovations to volatility shocks of the same exogenous variable. That is, if we have two different shocks (for instance, one to technology and one to preferences), the only terms different from zero in the second-order perturbation involving volatility would be the term with the innovation to the level of technology times the innovation to the volatility of technology and the term with the innovation to the level of preferences times the innovation to the volatility of preferences.

It is only in the third-order part of the solution (not covered by the theorem) -that is, those terms depending on the fourth derivative of the utility function- that the level of volatility enters without interacting with any other variable. That is why, if we are interested, for instance, in computing the impulse-response function (IRF) of a shock to volatility (as we will be in section 6), we need to compute at least a third-order approximation.

4.4. A Quantitative Example

We now present a quantitative example that clarifies our previous discussion. We start with Greenwood-Hercowitz-Huffman (GHH) preferences $u(c_t, l_t) = \log\left(c_t - \psi \frac{l_t^{1+\zeta}}{1+\zeta}\right)$. We pick these preferences because they do not have a wealth effect. For our illustrative purposes, this is most convenient. Given that the production function is of the form $y_t = e^{z_t} A k_t^\alpha l_t^{1-\alpha}$, an increase in the variance of z_t has a Jensen's inequality effect that induces a change in expected output. GHH kills that effect and avoids distracting elements in the solution. Later, for completeness, we will come back to a CRRA utility function.

The second step is to calibrate the model (below we will discuss how to estimate it, so we can think about this step as just fixing some parameter values for the computation). For our goals here, a conventional calibration will be sufficient. With respect to the preference parameters, we set $\beta = 0.99$ to get an annual interest rate of around 4 percent, we set $\zeta = 0.5$ to get a Frisch elasticity of 2, and $\psi = 3.4641$ to get average labor supply to be 1/3 of available time. With respect to technology, we set $\alpha = 1/3$ to match labor income share in national income, $A = 0.9823$ to normalize $y = 1$, and $\delta = 0.025$ to get a 10 percent annual depreciation. Finally, with respect to the stochastic process parameters, we set $\rho = 0.95$ and $\log \sigma = \log(0.007)$, the standard values for the Solow residual in the U.S. economy, and $\rho_\sigma = 0.95$ and $\eta = 0.1$ as two values that generate changes in volatility similar to the ones observed in the U.S. (since η does not appear in the decision rules up to second-order, its value for our example is less important).

The solution for consumption is then:

$$\begin{aligned}\widehat{c}_t = & 0.055115\widehat{k}_t + 0.576907z_{t-1} + 0.004251\varepsilon_t \\ & -0.000830\widehat{k}_t^2 + 0.036281\widehat{k}_tz_{t-1} + 0.000267\widehat{k}_t\varepsilon_t + 0.315513z_{t-1}^2 + 0.004650z_{t-1}\varepsilon_t \\ & +0.000017\varepsilon_t^2 + 0.004251\varepsilon_t u_t + 0.004038\varepsilon_t\widehat{\sigma}_{t-1} + 0.000013 + H.O.T.,\end{aligned}$$

(where we have already eliminated all terms with zero coefficients), for labor

$$\begin{aligned}\widehat{l}_t = & 0.014040\widehat{k}_t + 0.253333z_{t-1} + 0.001867\varepsilon_t \\ & -0.000444\widehat{k}_t^2 + 0.010671\widehat{k}_tz_{t-1} + 0.000079\widehat{k}_t\varepsilon_t + 0.096267z_{t-1}^2 + 0.001419z_{t-1}\varepsilon_t \\ & +0.000005\varepsilon_t^2 + 0.001867\varepsilon_t u_t + 0.001773\varepsilon_t\widehat{\sigma}_{t-1} + H.O.T.,\end{aligned}$$

and for capital

$$\begin{aligned}\widehat{k}_{t+1} = & 0.983067\widehat{k}_t + 0.563093z_{t-1} + 0.004149\varepsilon_t \\ & -0.0005\widehat{k}_t^2 + 0.035747\widehat{k}_tz_{t-1} + 0.000263\widehat{k}_t\varepsilon_t + 0.3342873z_{t-1}^2 + 0.004926z_{t-1}\varepsilon_t \\ & +0.000018\varepsilon_t^2 + 0.004149\varepsilon_t u_t + 0.003942\varepsilon_t\widehat{\sigma}_{t-1} - 0.000013 + H.O.T.\end{aligned}$$

In this solution, the correction for risk in consumption is reflected by the constant 0.000013, and for capital, by the constant -0.000013 (given the absence of wealth effects in GHH preferences, there is no constant shifting labor). This is because we have two mechanisms that act in different directions. On the one hand, precautionary behavior caused by volatility induces higher saving,

but on the other hand, volatility increases the production risk of capital. In our calibration this second effect predominates. Furthermore, high levels of volatility raise the effects of productivity shocks on consumption, labor, and capital. This is given by the three terms on $\varepsilon_t \widehat{\sigma}_{t-1}$. Finally, shocks to the level and shocks to volatility also reinforce each other (the coefficients on $\varepsilon_t u_t$).

Once we have the solution, there is the question of the quantitative importance of the second- or higher-order terms and, with them, of SV. There are two considerations. First, the size of the effect will depend on the parameters for the SV process. For some countries, a small level of SV may be plausible. For others, larger values are likely. A reasonable prior is that many developed economies would fall into the first group and many emerging economies into the second (our choice of $\eta = 0.1$ gets us closer to developed economies than to emerging ones). Second, the level of accuracy required in a solution is context-dependent. For example, a linear approximation that ignores SV may be good enough to compute some basic business cycle statistics, but it is unlikely to be enough for an accurate evaluation of welfare, and by construction, it is unable to estimate any of the parameters related to SV.

By modifying our utility kernel to the standard log-CRRA form $u(c_t, l_t) = \log c_t - \psi \frac{l_t^{1+\zeta}}{1+\zeta}$. The calibration stays the same except that we readjust $\psi = 4.5425$ to keep $l = 1/3$. The new policy functions for consumption:

$$\begin{aligned} \widehat{c}_t = & 0.043421\widehat{k}_t + 0.199865z_{t-1} + 0.001473\varepsilon_t \\ & -0.000810\widehat{k}_t^2 + 0.005249\widehat{k}_t z_{t-1} + 0.000039\widehat{k}_t \varepsilon_t + 0.053136z_{t-1}^2 + 0.000783z_{t-1}\varepsilon_t \\ & +0.000003\varepsilon_t^2 + 0.001473\varepsilon_t u_t + 0.001399\varepsilon_t \widehat{\sigma}_{t-1} - 0.000003 + H.O.T. \end{aligned}$$

for labor

$$\begin{aligned} \widehat{l}_t = & -0.008735\widehat{k}_t + 0.148498z_{t-1} + 0.001094\varepsilon_t \\ & +0.000449\widehat{k}_t^2 - 0.000676\widehat{k}_t z_{t-1} - 0.000005\widehat{k}_t \varepsilon_t + 0.018944z_{t-1}^2 + 0.000279z_{t-1}\varepsilon_t \\ & +0.000001\varepsilon_t^2 + 0.001094\varepsilon_t u_t + 0.001039\varepsilon_t \widehat{\sigma}_{t-1} + 0.000002 + H.O.T. \end{aligned}$$

and for capital

$$\begin{aligned} \widehat{k}_{t+1} = & 0.949211\widehat{k}_t + 0.730465z_{t-1} + 0.005382\varepsilon_t \\ & -0.000214\widehat{k}_t^2 + 0.017585\widehat{k}_t z_{t-1} + 0.000130\widehat{k}_t \varepsilon_t + 0.351353z_{t-1}^2 + 0.005178z_{t-1}\varepsilon_t \\ & +0.000019\varepsilon_t^2 + 0.005382\varepsilon_t u_t + 0.005113\varepsilon_t \widehat{\sigma}_{t-1} + 0.000006 + H.O.T. \end{aligned}$$

show the consequences of the wealth effect, in particular, the presence of a (small) precautionary behavior for labor, 0.000002, and the switch on the sign of the precautionary behavior for consumption and capital.

Now we can use our solution to form a state space representation, with a transition equation for the states given the innovations:

$$\mathbb{S}_t = f(\mathbb{S}_{t-1}, \mathbb{W}_t; \Psi) \quad (13)$$

that is the law of motion for capital that we just derived and (the second-order expansion of) the laws of motion of the stochastic process for productivity and its volatility, and a measurement equation for observables $\mathbb{Y}_t = g(\mathbb{S}_t, \mathbb{V}_t; \Psi)$ where \mathbb{V}_t is measurement noise (either measurement error or any other shock that affects the observables but not the states). This measurement noise is optional and, in our prototype model, we will not include it (one additional advantage of SV is that, for every stochastic process, we have two innovations, one to the level and one to the volatility) and we can write the simpler version:

$$\mathbb{Y}_t = g(\mathbb{S}_t; \Psi) \quad (14)$$

We index both equations by the vector $\Psi = \{\beta, \psi, \zeta, A, \alpha, \delta, \lambda, \sigma, \rho_\sigma, \eta\}$ of model parameters.

While the transition equation (13) is unique up to an equivalent class, the measurement equation depends on the assumptions about what we observe. For example, in our prototype business cycle model we can assume we observe hours or consumption (or both of them), since the model implies predictions about both variables. The choice should depend on the quality of the observables and on the goal of the empirical exercise. The only constraint is that we must select a number of series less than or equal to dimensionality of $(\mathbb{W}_t, \mathbb{V}_t)$ to avoid stochastic singularities.

4.5. Estimation

The next step in the analysis of our prototype business cycle model is its estimation with observed data. Besides the usual arguments for a rigorous statistical treatment of any model, in this case, a simple calibration exercise suffers from two serious challenges. First, in the presence of higher-order terms, the traditional strategy of selecting parameters by matching moments of the model with steady state values is flawed. When we have non-linearities, the ergodic distribution of the variables is not centered around their steady state, as it would be with a linearization. Instead, it

is translated by the non-linear coefficients. Thus, the only logical stand is to match the moments of the data with the simulated moments of the model, leaving us close to an SMM. Second, and even if we follow an SMM, it is not obvious which moments to select to calibrate the parameters of the SV process. Unfortunately, the experience from many years of methods of moments estimations is that choosing different moments (all of them sensible) may lead to rather different point estimates.

The alternative is to use a likelihood-based approach. The advantages of the likelihood function as the center of inference have been explained in other places (see An and Schorfheide, 2006, Fernández-Villaverde and Rubio-Ramírez, 2004, and Fernández-Villaverde, 2010) and there is not much point in reviewing them here. Suffice it to say that the likelihood is a coherent procedure that respects the likelihood principle and allows us to back up all the parameters of interest, and that has good small and large sample properties. Furthermore, the likelihood function can be easily complemented with presample information in the form of priors, which are particularly useful in macroeconomics, where we have short samples.

The likelihood function $p(\mathbb{Y}^T; \Psi)$ is nothing more than the probability the model assigns to a sequence of observables \mathbb{Y}^T given parameter values Ψ . The challenge with likelihood-based inference is that we need to evaluate that probability. A way to think about how this task can be accomplished for our model is as follows. Given the Markov structure of our state space representation (13)-(14), we factorize the likelihood function as:

$$p(\mathbb{Y}^T; \Psi) = \prod_{t=1}^T p(\mathbb{Y}_t | \mathbb{Y}^{t-1}; \Psi)$$

Then, conditioning on the states \mathbb{S}_t , and the innovation to productivity ε_t , we can write:

$$p(\mathbb{Y}_t | \mathbb{Y}^{t-1}; \Psi) = \int \int p(\mathbb{Y}_t | \mathbb{S}_t, \varepsilon_t; \Psi) p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^{data, t-1}; \Psi) d\mathbb{S}_t d\varepsilon_t \quad (15)$$

except for the first one:

$$p(\mathbb{Y}_1; \Psi) = \int p(\mathbb{Y}_1 | \mathbb{S}_1, \varepsilon_1; \Psi) d\mathbb{S}_1 d\varepsilon_1 \quad (16)$$

If we know \mathbb{S}_t and ε_t , computing $p(\mathbb{Y}_t | \mathbb{S}_t, \varepsilon_t; \Psi)$ is easy: it is just a change of variables implied by the measurement equation. To illustrate this point, imagine that $\mathbb{Y}_t = \widehat{c}_t$,⁸ that is the observable

⁸Note that \widehat{c}_t is equal to the raw data c_t minus the steady state c . Since the evaluation of the likelihood is conditional on some Ψ , we can easily find that steady state and map the raw data c into \widehat{c}_t . In real life, we are likely to have growth in the data, and hence, we will need to solve the model in some (transformed) stationary variable and undo the transformation in the measurement equation.

vector is just consumption which we have solved up to second-order:

$$\begin{aligned}\widehat{c}_t = & a_1\widehat{k}_t + a_2z_{t-1} + a_3\varepsilon_t + \\ & a_4\widehat{k}_t^2 + a_5\widehat{k}_tz_{t-1} + a_6\widehat{k}_t\varepsilon_t + a_7z_{t-1}^2 + a_8z_{t-1}\varepsilon_t + a_9\varepsilon_t^2 + a_{10}\varepsilon_tu_t + a_{11}\varepsilon_t\widehat{\sigma}_{t-1} + a_{12}\end{aligned}$$

where the a_t 's are the coefficients of the perturbation that are complicated non-linear functions of Ψ . Then, given \mathbb{S}_t and ε_t , we find the value of \widetilde{u}_t that accounts for the observation \widehat{c}_t :

$$\widetilde{u}_t = \frac{1}{a_{10}\varepsilon_t} \left[\begin{array}{c} \widehat{c}_t - a_1\widehat{k}_t - a_2z_{t-1} - a_3\varepsilon_t - a_4\widehat{k}_t^2 - a_5\widehat{k}_tz_{t-1} \\ -a_6\widehat{k}_t\varepsilon_t - a_7z_{t-1}^2 - a_8z_{t-1}\varepsilon_t - a_9\varepsilon_t^2 + a_{11}\varepsilon_t\widehat{\sigma}_{t-1} - a_{12} \end{array} \right] \quad (17)$$

By evaluating the p.d.f. of \widetilde{u}_t given Ψ (in our model, just a normal p.d.f.) and applying the change of variables formula, we get $p(\mathbb{Y}_t|\mathbb{S}_t, \varepsilon_t; \Psi)$. This computation of \widetilde{u}_t in (17) takes advantage of the structure of the solution to our model that we characterized before. The result can be generalized to an arbitrary number n of observables and shocks with SV, in which case we would have a linear system of n equations. If we did not know that some coefficients were zero, we would need to solve a quadratic system on u_t , something much harder to do. For example, in the case with n observables, it would be a quadratic system with 2^n solutions, a daunting task.

In the same way, if we know how to draw from $p(\mathbb{S}_1; \Psi)$, we can compute (16) by Monte Carlo. Generating this drawing is usually straightforward, although tedious. As described in Santos and Peralta-Alva (2005), given some parameter values Ψ , we can simulate the model for a sufficiently large path (to wash out the effect of the initial conditions, which we can make equal to the steady state just for simplicity, although other starting points are admissible if convenient) and keep the last realizations as a sample from $p(\mathbb{S}_1; \Psi)$.

Thus, the complication in evaluating (15) is reduced to a) finding the sequence of conditional densities $\{p(\mathbb{S}_t, \varepsilon_t|\mathbb{Y}^{t-1}; \Psi)\}_{t=1}^T$ and b) computing the different integrals. Fortunately, filtering theory aims at providing the user precisely that sequence of conditional densities and ways to compute the required integrals.

Filtering is a recursive procedure that relies on two tools, the Chapman-Kolmogorov equation:

$$p(\mathbb{S}_{t+1}, \varepsilon_{t+1}|\mathbb{Y}^t; \Psi) = \int p(\mathbb{S}_{t+1}, \varepsilon_{t+1}|\mathbb{S}_t, \varepsilon_t; \Psi) p(\mathbb{S}_t, \varepsilon_t|\mathbb{Y}^t; \Psi) d\mathbb{S}_t d\varepsilon_t \quad (18)$$

and Bayes' theorem:

$$p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^t; \Psi) = \frac{p(\mathbb{Y}_t | \mathbb{S}_t, \varepsilon_t; \Psi) p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^{t-1}; \Psi)}{\int p(\mathbb{Y}_t | \mathbb{S}_t, \varepsilon_t; \Psi) p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^{t-1}; \Psi) dS_t d\varepsilon_t} \quad (19)$$

The Chapman-Kolmogorov equation tells that the distribution of states and productivity innovations tomorrow given observations until today, $p(\mathbb{S}_{t+1}, \varepsilon_{t+1} | \mathbb{Y}^t; \Psi)$, is equal to the distribution today, $p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^t; \Psi)$, times the transition probabilities $p(\mathbb{S}_{t+1}, \varepsilon_{t+1} | \mathbb{S}_t, \varepsilon_t; \Psi)$ integrated over all possible events. In other words, the Chapman-Kolmogorov equation just provides the researcher with a forecasting rule for the evolution of states. Given that we have access to the solution of the model, the computation of $p(\mathbb{S}_{t+1}, \varepsilon_{t+1} | \mathbb{S}_t, \varepsilon_t; \Psi)$ is direct given $p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^t; \Psi)$ as an input.

Bayes' theorem updates the distribution of states $p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^t; \Psi)$ when a new observation arrives given its probability $p(\mathbb{Y}_t | \mathbb{S}_t, \varepsilon_t; \Psi)$, which, as we argued above, is also easy to evaluate given our state space representation. Thus, with an input $p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^{t-1}; \Psi)$, the Bayes' theorem gives us $p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^t; \Psi)$. We can see clearly the recursive structure of filtering. Given some initial $p(\mathbb{S}_1, \varepsilon_1; \Psi)$, Bayes' theorem provides us with $p(\mathbb{S}_1, \varepsilon_1 | \mathbb{Y}^1; \Psi)$, which we use as an input for the Chapman-Kolmogorov equation and get $p(\mathbb{S}_2, \varepsilon_2 | \mathbb{Y}^1; \Psi)$, the input for the next application of the Bayes' theorem. By a recursive application of the forecasting and updating steps, we generate the complete sequence $\{p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^{t-1}; \Psi)\}_{t=1}^T$ we are searching for. But while the Chapman-Kolmogorov equation and Bayes' theorem are conceptually straightforward, their practical implementation is cumbersome because they involve the computation of numerous integrals again and again over the sample.

There is, of course, a well-known exception. If the state space representation (13)-(14) were linear and the innovations normally distributed, we could use the Kalman filter to efficiently derive $\{p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^{t-1}; \Psi)\}_{t=1}^T$ and, by taking advantage of the fact that all the appropriate conditional distributions are normal, to solve the required integrals.

Unfortunately, this cannot be done once we have SV since at least one component of (13) is non-linear.⁹ The non-linearity of SV deforms $\{p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^{t-1}; \Psi)\}_{t=1}^T$ in such a way that they do not belong to any known parametric family. Instead, we need to resort to some numerical procedure to compute the relevant integrals. A powerful algorithm for this non-linear filtering is

⁹Even if we kept the linear approximation of the decision rule and cut off its quadratic terms, we would still need to resort to some type of non-linear filtering. We argued before that this mixing of approximation orders (linear for endogenous state variables, non-linear for exogenous ones) violates the theorems that guarantee the convergence of perturbations and it suffers from poor accuracy. Here, we show it does not even save time when estimating the model.

the particle filter, as described, for example, in Fernández-Villaverde and Rubio-Ramírez (2005 and 2007) (see also the technical appendix to Fernández-Villaverde and Rubio-Ramírez, 2007, for alternative algorithms).

The particle filter is a sequential Monte Carlo method that replaces the unknown sequence $\{p(\mathbb{S}_t, \varepsilon_t | \mathbb{Y}^{t-1}; \Psi)\}_{t=1}^T$ with an empirical distribution of N draws $\left\{s_{t|t-1}^i, \varepsilon_{1t}^i\right\}_{i=1}^N$ (where we follow the short-hand notation that a variable $x_{j|m}^i$ is the draw i at time j conditional on the information up to period m) generated by simulation. Then, by an appeal to the Law of Large Numbers, we can substitute the integral in (15) by:

$$p(\mathbb{Y}_t | \mathbb{Y}^{t-1}; \Psi) \simeq \frac{1}{N} \sum_{i=1}^N p(\mathbb{Y}_t | s_{t|t-1}^i, \varepsilon_{1t}^i; \Psi) \quad (20)$$

The key to the success of the particle filter is that the simulation is generated through a procedure known as sequential importance resampling (SIR) with weights:

$$q_t^i = \frac{p(\mathbb{Y}_t | s_{t|t-1}^i, \varepsilon_{1t}^i; \Psi)}{\sum_{i=1}^N p(\mathbb{Y}_t | s_{t|t-1}^i, \varepsilon_{1t}^i; \Psi)} \quad (21)$$

SIR allows us to move from a draw $\left\{s_{t|t-1}^i, \varepsilon_t^i\right\}_{i=1}^N$ to a draw $\left\{s_{t|t}^i, \varepsilon_t^i\right\}_{i=1}^N$ that incorporates information about the observable at period t . The reason is that resampling with weights q_t^i is just equivalent to the application of Bayes' theorem in equation (19): the draw $\left\{s_{t|t-1}^i, \varepsilon_t^i\right\}_{i=1}^N$ is the prior and the weights are the normalized likelihood of \mathbb{Y}_t . SIR guarantees that the Monte Carlo method achieves sufficient accuracy in a reasonable amount of time, something that cannot be achieved without resampling as most draws would wander away from the true unknown state. The forecast step in the Chapman-Kolmogorov equation (18) is extremely simple because we have the law of motion for states given $(s_{t|t-1}^i, \varepsilon_t^i)$, the volatility innovation it implies, and the distribution of the level innovation $p(\varepsilon | \Psi)$. Under weak conditions, the particle filter delivers a consistent estimator of the likelihood function and a central limit theorem applies (Künsch, 2005).

In pseudo-code, this resampling works as follows:

Step 0, Initialization: Set $t \rightsquigarrow 1$. Sample N values $\left\{s_{0|0}^i, \varepsilon_0^i\right\}_{i=1}^N$ from $p(\mathbb{S}_0|\Psi)$ and $p(\varepsilon|\Psi)$.

Step 1, Prediction: Sample N values $\left\{s_{t|t-1}^i, \varepsilon_t^i\right\}_{i=1}^N$ from $p(\mathbb{S}_t, \varepsilon_t|\mathbb{Y}^{t-1}; \gamma)$ using the draw $\left\{s_{t-1|t-1}^i, \varepsilon_{t-1}^i\right\}_{i=1}^N$, the law of motion for states and $p(\varepsilon|\Psi)$.

Step 2, Filtering: Assign to each draw $\left(s_{t|t-1}^i, \varepsilon_t^i\right)$ the weight q_t^i in (21).

Step 3, Sampling: Sample N times with replacement from $\left\{s_{t|t-1}^i, \varepsilon_t^i\right\}_{i=1}^N$ with weights $\{q_t^i\}_{i=1}^N$. Call the new draw $\left\{s_{t|t}^i, \varepsilon_t^i\right\}_{i=1}^N$. If $t < T$, set $t \rightsquigarrow t + 1$ and go to step 2. Otherwise stop.

Once we have evaluated the likelihood function given Ψ , the researcher can either maximize it by searching over the parameter space or we can combine it with a prior $p(\Psi)$ and use a Markov chain Monte Carlo (MCMC) to approximate the posterior:

$$p(\Psi|\mathbb{Y}^T) = \frac{p(\mathbb{Y}^T; \Psi) p(\Psi)}{\int p(\mathbb{Y}^T; \Psi) p(\Psi) d\Psi}$$

An and Schorfheide (2006) is a standard reference for details about how to implement MCMC's. Moreover, the MCMC method (or close relatives such as simulated annealing) can be used for the maximization of the likelihood. One inconvenient consequence of the resampling in the particle filter is that the evaluation of the likelihood is not differentiable with respect to the parameters: a small change in one parameter may imply that we resample a different draw than in the previous pass of the algorithm.¹⁰ Therefore, derivative-based optimization algorithms cannot be applied without further smoothing of the likelihood.

4.6. Implications for Policy

The final step in our discussion is to think about policy implications. The first, and most direct, is that if volatility shocks affect aggregate fluctuations in a significant way, policy makers may need

¹⁰For the maximum likelihood to converge, we need to keep the simulated innovations ε_t and the uniform numbers that enter into the resampling decisions constant as we modified the parameter values. This is required to achieve stochastic equicontinuity. With this property, the pointwise convergence of the likelihood (20) to the exact likelihood is strengthened to uniform convergence and we can swap the argmax and the lim operators (that is, as the number of simulated particles converges to infinity, the MLE also converges). Otherwise, we would suffer numerical instabilities induced by the “chatter” of random numbers. In the Bayesian approach, keeping these random numbers constant is not strictly needed but it improves accuracy.

to consider volatility when implementing fiscal and monetary policy. Imagine, for example, that we extend our model with the need to finance an exogenously given flow of public expenditure and the government only has access to distortionary taxes. This is the same framework as in Chari, Christiano, and Kehoe (1994), except that now technology shocks have SV. A Ramsey optimal policy would prescribe how debt, and fiscal policy in general, needs to respond to volatility shocks. For instance, we conjecture that the presence of SV, by augmenting the risk of having a really bad shock, may imply that governments want to accumulate less public debt on average to leave them enough space to respond to these extreme shocks. Similarly, an optimal interest rate rule followed by the central bank to implement monetary policy could also depend on the level of volatility in addition to the traditional dependence on the levels of inflation and the output gap. In fact, Bekaert, Hoerova, and Lo Duca (2010) have gathered evidence that, in the U.S., the Fed responds to increased stock market volatility by easing monetary policy.

A second policy consideration is that countries subject to volatility shocks require a more sophisticated management of the maturity structure of their debt that takes into account the future paths of the level and volatility of interest rates. This is central in environments with non-contingent public debt, arguably a fair description of reality. Thus, volatility highlights the importance of improving our understanding of the optimal management of government debt in a world with incomplete markets, a field still relatively unexplored.

Now, after our fairly long discussion of the prototype business cycle model, we are ready for our first “real life” application, an exercise in reading the recent monetary history of the U.S. through the lens of DSGE models.

5. Application I: Understanding the Recent Monetary History of the U.S.

As we documented in section 3, around 1984, the U.S. economy entered into a period of low volatility known as the great moderation. Among the many reasons presented in the literature, two have received a considerable amount of attention. One branch of the literature argues that the great moderation was just the consequence of low volatility shocks (for example, Sims and Zha, 2006). Another branch of the literature argues that some other changes in the economy, usually better monetary policy, explain the evolution of aggregate volatility (more famously, Clarida, Galí, and Gertler, 2000, and Lubick and Schorfheide, 2004). The first explanation is pessimistic: we enjoy or suffer periods of low or high volatility, but there is little that policy makers can do about

it. The second one is optimistic: as long as we do not unlearn the lessons of monetary economics, we should expect the great moderation to continue (even after the current turbulence).

Sorting the two different approaches requires that we analyze the question using a model that has both changes in volatility and changes in policy. Moreover, we need an equilibrium model. As shown by Benati and Surico (2009), SVARs may be uninformative for the question at hand since we cannot easily map between changes in variances of the SVAR and changes in variances of the shocks of a DSGE model.

The techniques presented in this paper can help us to fill this gap. In particular, we can build and estimate a medium-scale DSGE model with SV in the structural shocks that drive the economy, parameter drifting in the Taylor rule followed by the monetary authority, and rational expectations of agents regarding these changes. In the next pages, we summarize the material in FGR.

5.1. The Model

We adopt what has become the standard New Keynesian DSGE model, based on Christiano, Eichenbaum, and Evans (2005). Since the model is well known, our description will be brief. In our specification, SV appears in the form of changing standard deviations of the five structural shocks to the model (two shocks to preferences, two shocks to technology, and one shock to monetary policy). Parameter drifting appears in the form of changing values of the parameters in the Taylor policy rule followed by the monetary authority.

In more detail, household j 's preferences are:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t d_t \left\{ \log(c_{jt} - hc_{jt-1}) + v \log\left(\frac{m_{jt}}{p_t}\right) - \varphi_t \psi \frac{l_{jt}^{1+\vartheta}}{1+\vartheta} \right\},$$

which is separable in consumption, c_{jt} , real money balances, m_{jt}/p_t , and hours worked, l_{jt} . In our notation, \mathbb{E}_0 is the conditional expectation operator, β is the discount factor, h controls habit persistence, ϑ is the inverse of the Frisch labor supply elasticity, d_t is a intertemporal preference shock that follows $\log d_t = \rho_d \log d_{t-1} + \sigma_{dt} \varepsilon_{dt}$ where $\varepsilon_{dt} \sim \mathcal{N}(0, 1)$ and φ_t is a labor supply shock that evolves as $\log \varphi_t = \rho_{\varphi} \log \varphi_{t-1} + \sigma_{\varphi t} \varepsilon_{\varphi t}$ where $\varepsilon_{\varphi t} \sim \mathcal{N}(0, 1)$.

As in section 4, the standard deviations, σ_{dt} and $\sigma_{\varphi t}$, of innovations ε_{dt} and $\varepsilon_{\varphi t}$ move according to $\log \sigma_{dt} = (1 - \rho_{\sigma_d}) \log \sigma_d + \rho_{\sigma_d} \log \sigma_{dt-1} + \eta_d u_{dt}$ where $u_{dt} \sim \mathcal{N}(0, 1)$ and $\log \sigma_{\varphi t} = (1 - \rho_{\sigma_{\varphi}}) \log \sigma_{\varphi} + \rho_{\sigma_{\varphi}} \log \sigma_{\varphi t-1} + \eta_{\varphi} u_{\varphi t}$ where $u_{\varphi t} \sim \mathcal{N}(0, 1)$.

All the shocks and innovations are perfectly observed by the agents when they are realized. Agents have, as well, rational expectations about how they evolve over time.

We assume complete financial markets. An amount of state-contingent securities, a_{jt+1} , which pay one unit of consumption in event $\omega_{jt+1,t}$, is traded at time t at unitary price $q_{jt+1,t}$ in terms of the consumption good. In addition, households also hold b_{jt} government bonds that pay a nominal gross interest rate of R_{t-1} . Therefore, the j -th household's budget constraint is given by:

$$\begin{aligned} & c_{jt} + x_{jt} + \frac{m_{jt}}{p_t} + \frac{b_{jt+1}}{p_t} + \int q_{jt+1,t} a_{jt+1} d\omega_{jt+1,t} \\ &= w_{jt} l_{jt} + \left(r_t u_{jt} - \frac{\Phi[u_{jt}]}{\mu_t} \right) k_{jt-1} + \frac{m_{jt-1}}{p_t} + \frac{R_{t-1} b_{jt}}{p_t} + a_{jt} + T_t \end{aligned}$$

where x_t is investment, w_{jt} is the real wage, r_t the real rental price of capital, $u_{jt} > 0$ the rate of use of capital, $\mu_t^{-1} \Phi[u_{jt}]$ is the cost of utilizing capital at rate u_{jt} in terms of the final good, μ_t is an investment-specific technological level, T_t are lump-sum transfers and the profits. We specify $\Phi[\cdot]$ such that it satisfies the conditions that $\Phi[1] = 0$, $\Phi'[\cdot] = 0$, and $\Phi''[\cdot] > 0$. This function carries the normalization that $u = 1$ in the balanced growth path. The capital accumulated by household j at the end of period t is given by:

$$k_{jt} = (1 - \delta) k_{jt-1} + \mu_t (1 - V[x_{jt}/x_{jt-1}]) x_{jt}$$

where δ is the depreciation rate and $V[\cdot]$ is a quadratic adjustment cost function written in deviations with respect to the balanced growth rate of investment, Λ_x . Our third structural shock, the investment-specific technology level μ_t , follows $\log \mu_t = \Lambda_\mu + \log \mu_{t-1} + \sigma_{\mu t} \varepsilon_{\mu t}$, where $\varepsilon_{\mu t} \sim \mathcal{N}(0, 1)$. The standard deviation of the innovation also evolves as $\log \sigma_{\mu t} = \left(1 - \rho_{\sigma_\mu}\right) \log \sigma_\mu + \rho_{\sigma_\mu} \log \sigma_{\mu t-1} + \eta_\mu u_{\mu t}$ where $u_{\mu t} \sim \mathcal{N}(0, 1)$.

The household chooses c_{jt} , b_{jt} , u_{jt} , k_{jt} , and x_{jt} taking prices as given. Labor and wages, l_{jt} and w_{jt} , are chosen in the presence of monopolistic competition and nominal rigidities. Each household j supplies a slightly different type of labor services l_{jt} that are aggregated by a “labor packer” into homogeneous labor l_t^d with the production function:

$$l_t^d = \left(\int_0^1 l_{jt}^{\frac{\eta-1}{\eta}} dj \right)^{\frac{\eta}{\eta-1}}$$

that is rented to intermediate good producers at the wage w_t . The “labor packer” is perfectly

competitive and it takes wages as given. Households follow a Calvo pricing mechanism when they set their wages. Every period a randomly selected fraction $1 - \theta_w$ of households can reoptimize their wages to w_{jt}^* . All other households index their wages given past inflation with an indexation parameter $\chi_w \in [0, 1]$.

There is one final good producer that aggregates a continuum of intermediate goods and it is perfectly competitive and minimizes its costs subject to the production function

$$y_t^d = \left(\int_0^1 y_{it}^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}$$

and taking as given all prices. Each of the intermediate goods is produced by a monopolistic competitor whose technology is given by a production function $y_{it} = A_t k_{it-1}^\alpha (l_{it}^d)^{1-\alpha}$, where k_{it-1} is the capital rented by the firm, l_{it}^d is the amount of the “packed” labor input rented by the firm, and A_t (our fourth structural shock) is neutral productivity that follows $\log A_t = \Lambda_A + \log A_{t-1} + \sigma_{At} \varepsilon_{At}$, where $\varepsilon_{At} \sim \mathcal{N}(0, 1)$. The standard deviation of this innovation evolves following the specification $\log \sigma_{At} = (1 - \rho_{\sigma_A}) \log \sigma_A + \rho_{\sigma_A} \log \sigma_{At-1} + \eta_A u_{At}$ where $u_{At} \sim \mathcal{N}(0, 1)$.

Given the demand function from the final good producer, the intermediate good producers set prices to maximize profits. They also follow a Calvo pricing scheme. In each period, a fraction $1 - \theta_p$ reoptimize their prices to p_t^* . All other firms partially index their prices by past inflation with an indexation parameter χ .

The model is closed by the presence of a monetary authority that sets the nominal interest rates. The monetary authority follows a modified Taylor rule:

$$R_t/R = (R_{t-1}/R)^{\gamma_R} \left((\Pi_t/\Pi)^{\gamma_{\Pi,t}} \left((y_t^d/y_{t-1}^d) / \exp(\Lambda_y) \right)^{\gamma_y} \right)^{1-\gamma_R} \xi_t.$$

The term Π_t/Π , an “inflation gap,” responds to the deviation of inflation from its balanced growth path level Π and the term $(y_t^d/y_{t-1}^d) / \exp(\Lambda_y)$ is a “growth gap” (Λ_y is the growth rate of the economy along its balanced growth path). The term $\log \xi_t = \sigma_{m,t} \varepsilon_{mt}$ is the monetary policy shock. The innovation $\varepsilon_{mt} \sim \mathcal{N}(0, 1)$ to the monetary policy shock has a time-varying standard deviation, $\sigma_{m,t}$, that follows $\log \sigma_{mt} = (1 - \rho_{\sigma_m}) \log \sigma_m + \rho_{\sigma_m} \log \sigma_{mt-1} + \eta_m u_{m,t}$ where $u_{m,t} \sim \mathcal{N}(0, 1)$. In this policy rule, we have a drifting parameter: the response of the monetary authority to the inflation gap, $\gamma_{\Pi,t}$. The parameter drifts over time as $\log \gamma_{\Pi,t} = (1 - \rho_{\gamma_{\Pi}}) \log \gamma_{\Pi} + \rho_{\gamma_{\Pi}} \log \gamma_{\Pi,t-1} + \eta_{\pi} \varepsilon_{\pi t}$ where $\varepsilon_{\pi t} \sim \mathcal{N}(0, 1)$. We assume here that the agents perfectly observe the changes in monetary policy parameters.

5.2. Solution and Estimation

The equilibrium of the model does not have a closed-form solution and we need to resort to a numerical approximation to compute it. For the reasons outlined in section 5, we perform a second-order perturbation around the (rescaled) steady state of the model. The quadratic terms of this approximation allow us to capture, to a large extent, the effects of volatility shocks and parameter drift while keeping computational complexity at a reasonable level.

We estimate our model using five time series for the U.S. economy: 1) the relative price of investment goods with respect to the price of consumption goods, 2) the federal funds rate, 3) real output per capita growth, 4) the consumer price index, and 5) real wages per capita. Our sample covers 1959.Q1 to 2007.Q1, with 192 observations. Then, we follow again section 5 and exploit the structure of the state space representation of the solution of the model to evaluate the likelihood of the model. FGR provide further details.

5.3. The Empirical Findings

We invite the interested reader to check FGR, where all the results are shown in detail and Fernández-Villaverde, Guerrón-Quintana, and Rubio-Ramírez (2010b) where the findings are compared with the historical record. Here, as a summary, we highlight our main findings: 1) there is overwhelming evidence of changes in monetary policy even after controlling for the large amount of stochastic volatility existing in the data; 2) these changes in monetary policy were key for the reduction of average inflation; 3) and the response of monetary policy to inflation under Burns, Miller, and Greenspan was similar, while it was much higher under Volcker.

The first finding can be documented in figure 5.1 with the evolution of the (smoothed) Taylor rule parameter of the response of the monetary authority to inflation that we recover from the data. This figure summarizes how our model understands the recent monetary history of the U.S. The parameter $\gamma_{\Pi t}$ started the sample around its estimated mean, slightly over 1, and it grew more or less steadily during the 1960s until reaching a peak in early 1968. After that year, $\gamma_{\Pi t}$ suffered a fast collapse that pushed it below 1 in 1971, one year after the appointment of Burns as chairman of the Fed in February 1970. The parameter stayed below 1 for all of the 1970s, showing either that monetary policy did not satisfy the Taylor principle or that our postulated monetary policy rule is not a good description of the behavior of the Fed at the time (for example, because the Fed was using real-time data). The arrival of Volcker is quickly picked up by our estimates: $\gamma_{\Pi t}$ increases to over 2 after a few months and stays high during all the years of Volcker's tenure.

Interestingly, our estimate captures well the observation by Goodfriend and King (2007) that monetary policy tightened in the spring of 1980 as inflation and long-run inflation expectations continued to grow. The level of $\gamma_{\Pi t}$ stayed roughly constant at this high during the remainder of Volcker's tenure. But as quickly as $\gamma_{\Pi t}$ rose when Volcker arrived, it went down again when he departed. Greenspan's tenure at the Fed meant that, by 1990, the response of the monetary authority to inflation was again below 1. During all the following years, $\gamma_{\Pi t}$ was low, even below the values that it took during Burns-Miller's time. Moreover, our estimates of $\gamma_{\Pi t}$ are tight, suggesting that posterior uncertainty is not the full explanation behind these movements.




Figure 5.1: Smoothed path for the Taylor rule parameter on inflation \pm 2 standard deviations.

With respect to SV, we plot in figure 5.2 the evolution of the standard deviation of the innovation of the structural shocks, all of them in log-deviations with respect to their estimated means. A first lesson from that figure is that the standard deviation of the intertemporal shock was particularly high in the 1970s and only slowly went down during the 1980s and early 1990s. By the end of the sample, the standard deviation of the intertemporal shock was roughly at the level where it started. This is important to understand the behavior of inflation. A high volatility of intertemporal shocks creates a volatile aggregate demand and, with it, an inflation that is harder to control. Thus, we conclude that a significant component of the volatility of inflation in the 1970s and 1980s was due to the volatility of preferences. In comparison, the standard deviation of

all the other shocks is relatively stable except, perhaps, for the big drop in the standard deviation of the monetary policy shock in the early 1980s and the big changes in the standard deviation of the investment shock during the period of oil price shocks. Hence, the 1970s and the 1980s were more volatile than the 1960s and the 1990s, creating a tougher environment for monetary policy.

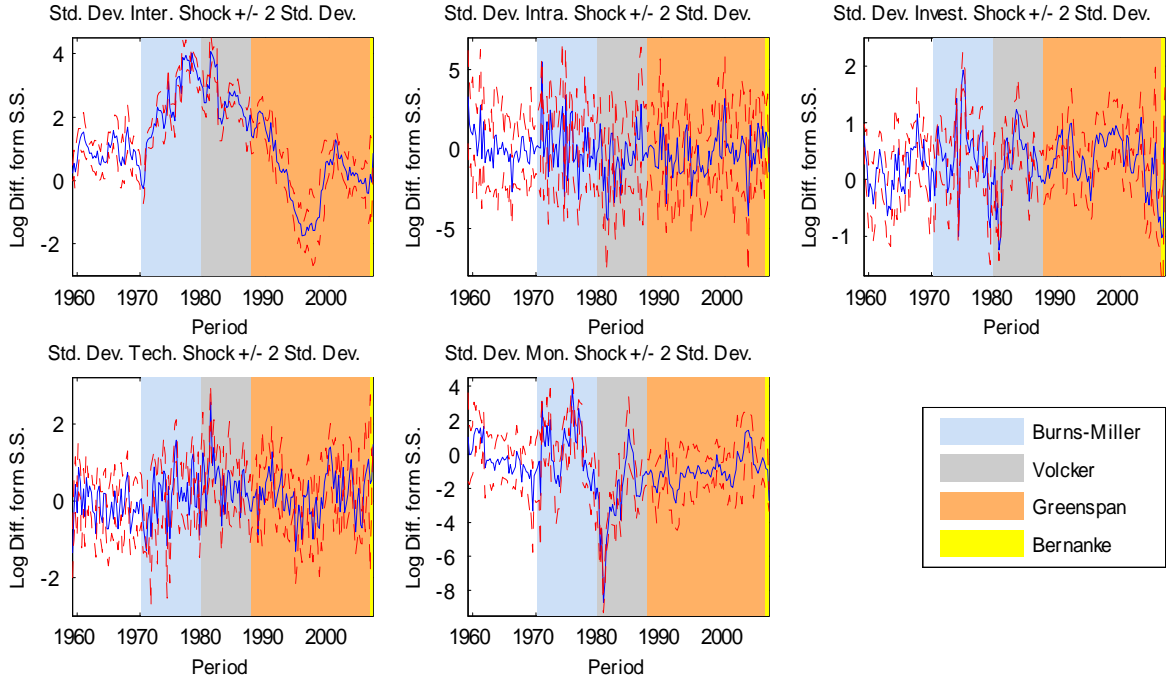


Figure 5.2: Smoothed standard deviation shocks to the intertemporal ($\hat{\sigma}_{dt}$) shock, the intratemporal ($\hat{\sigma}_{\phi t}$) shock, the investment-specific ($\hat{\sigma}_{\mu t}$) shock, the technology ($\hat{\sigma}_{At}$) shock, and the monetary policy ($\hat{\sigma}_{mt}$) shock ± 2 s.d.

One advantage of estimating a structural model is that we can use it to compute counterfactual histories where we remove a source of variation in the data to measure its impact. With one of these counterfactuals, we document our third main finding. We measure that without changes in volatility, the great moderation would have been noticeably smaller. The standard deviation of inflation would have fallen by only 13 percent, the standard deviation of output growth would have fallen by 16 percent, and the standard deviation of the federal funds rate would have fallen by 35 percent, that is, only 33, 20, and 87 percent, respectively, of how much they actually fell.

This application has shown how SV is a fundamental element in our understanding of the recent monetary history of the U.S. and how the methods presented in section 5 can be put to good use in a developed economy. In the next section we show how SV is also important (perhaps even more) for small, open emerging economies.

6. Application II: Small Open Economies

Now we summarize the results in Fernández-Villaverde *et al.* (2009) and show how changes in the volatility of the real interest rate at which emerging economies borrow have a substantial effect on real variables like output, consumption, investment, and hours worked. These effects appear even when the level of the real interest rate itself remains constant.

To prove our case, we use the evidence of time-varying volatility in the real interest rates faced by countries such as Argentina that we briefly showed in figure 3 and that is documented formally in Fernández-Villaverde *et al.* (2009). Then, we feed this time-varying process into an otherwise standard small, open economy business cycle model calibrated to match the data from Argentina. We find that an increase in real interest rate volatility triggers a fall in output, consumption, investment, and hours worked, and a notable change in the current account. Hence, we show that the time-varying volatility of real interest rates might be an important force behind the distinctive size and pattern of business cycle fluctuations of emerging economies.

We do not offer a theory of why real interest rate volatility changes over time. Instead, we model it as an exogenous process. Part of the reason is that an exogenous process focuses our attention on the mechanism through which real interest rate risk shapes the trade-offs of agents in small, open economies. More important, the literature has not developed, even at the prototype level, an equilibrium model to endogenize these volatility shocks. Fortunately, the findings of Uribe and Yue (2006) and Longstaff *et al.* (2007) justify our strategy. The evidence in both papers is strongly supportive of the view that a substantial component of changes in volatility is exogenous to the country. These results should not be a surprise because the aim of the literature on financial contagion is to understand phenomena that distinctively look like exogenous shocks to small open economies (Kaminsky *et al.*, 2003).

6.1. The Model

We postulate a simple small, open economy model with incomplete asset markets. The economy is populated by a representative household with preferences:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-v}}{1-v} - \omega \frac{H_t^{1+\eta}}{1+\eta} \right). \quad (22)$$

Here, \mathbb{E}_0 is the conditional expectations operator, C_t denotes consumption, H_t stands for hours worked, and $\beta \in (0, 1)$ corresponds to the discount factor. The household can invest in two types

of assets: the stock of physical capital, K_t , and an internationally traded bond, D_t . We maintain the convention that positive values of D_t denote debt. Then, the household's budget constraint is given by:

$$\frac{D_{t+1}}{1 + r_t} = D_t - W_t H_t - R_t K_t + C_t + I_t + \frac{\Phi_D}{2} (D_{t+1} - D)^2 \quad (23)$$

where W_t represents the real wage, R_t stands for the real rental rate of capital, I_t is gross domestic investment, $\Phi_D > 0$ is a parameter that controls the costs of holding a net foreign asset position, and D is a parameter that determines debt in the steady state. The cost, assumed to eliminate the unit root otherwise built into the dynamics of the model, is paid to some foreign international institution (for example, an investment bank that handles the issuing of bonds for the household).

We write the real interest rate faced by domestic residents in international markets at time t as $r_t = r + \varepsilon_{tb,t} + \varepsilon_{r,t}$. In this equation, r is the mean of the international risk-free real rate plus the mean of the country-spread. The term $\varepsilon_{tb,t}$ equals the international risk-free real rate subtracted from its mean and $\varepsilon_{r,t}$ equals the country-spread subtracted from its mean. Both $\varepsilon_{tb,t}$ and $\varepsilon_{r,t}$ follow $AR(1)$ processes:

$$\varepsilon_{tb,t} = \rho_{tb} \varepsilon_{tb,t-1} + e^{\sigma_{tb,t}} u_{tb,t}, \text{ where } u_{tb,t} \sim \mathcal{N}(0, 1) \quad (24)$$

$$\varepsilon_{r,t} = \rho_r \varepsilon_{r,t-1} + e^{\sigma_{r,t}} u_{r,t}, \text{ where } u_{r,t} \sim \mathcal{N}(0, 1) \quad (25)$$

The standard deviations $\sigma_{tb,t}$ and $\sigma_{r,t}$ also follow:

$$\sigma_{tb,t} = (1 - \rho_{\sigma_{tb}}) \sigma_{tb} + \rho_{\sigma_{tb}} \sigma_{tb,t-1} + \eta_{tb} u_{\sigma_{tb,t}}, \text{ where } u_{\sigma_{tb,t}} \sim \mathcal{N}(0, 1) \quad (26)$$

$$\sigma_{r,t} = (1 - \rho_{\sigma_r}) \sigma_r + \rho_{\sigma_r} \sigma_{r,t-1} + \eta_r u_{\sigma_{r,t}}, \text{ where } u_{\sigma_{r,t}} \sim \mathcal{N}(0, 1) \quad (27)$$

The parameters σ_{tb} and η_{tb} control the degree of mean volatility and SV in the international risk-free real rate. The same can be said about σ_r and η_r and the mean volatility and SV in the country spread. We call $u_{tb,t}$ and $u_{r,t}$ innovations to the international risk-free real rate and the country-spread, respectively. We call $u_{\sigma_{tb,t}}$ and $u_{\sigma_{r,t}}$ innovations to the volatility of the international risk-free real rate and the country spread, respectively. Sometimes, for simplicity, we call $\sigma_{tb,t}$ and $\sigma_{r,t}$ volatility shocks and $u_{\sigma_{tb,t}}$ and $u_{\sigma_{r,t}}$ innovation to the volatility shocks.

The stock of capital evolves according to $K_{t+1} = (1 - \delta)K_t + \left(1 - \frac{\phi}{2} \left(\frac{I_t}{I_{t-1}} - 1\right)^2\right) I_t$, where δ is the depreciation rate. The parameter $\phi > 0$ controls the size of these adjustment costs. Finally, the representative household is also subject to the typical no-Ponzi-game condition.

Firms rent capital and labor from households to produce output in a competitive environment according to the technology $Y_t = K_t^\alpha (e^{X_t} H_t)^{1-\alpha}$ where $X_t = \rho_x X_{t-1} + e^{\sigma_x} u_{x,t}$ and $u_{x,t} \sim \mathcal{N}(0, 1)$. Firms maximize profits by equating wages and the rental rate of capital to marginal productivities. Thus, we can rewrite equation (23) in terms of net exports NX_t :

$$NX_t = Y_t - C_t - I_t = D_t - \frac{D_{t+1}}{1 + r_t} + \frac{\Phi_D}{2} (D_{t+1} - D)^2$$

6.2. Solving and Calibrating the Model

We solve the model by relying on perturbation methods. We want to measure the effects of a volatility increase (a positive shock to either $u_{\sigma_r,t}$ or $u_{\sigma_{tb},t}$), while keeping the interest rate itself unchanged (fixing $u_{r,t} = 0$ and $u_{tb,t} = 0$). Consequently, we need to obtain a *third* approximation of the policy functions. As we saw in section 4, a first-order approximation to the model would miss all of the dynamics induced by volatility because this approximation is certainty equivalent and a second-order approximation would only capture the volatility effect indirectly via cross product terms of the form $u_{r,t}u_{\sigma_r,t}$ and $u_{tb,t}u_{\sigma_{tb},t}$; that is, up to second-order, volatility does not have an effect as long as the real interest rate does not change. It is only in a third-order approximation that the SV shocks, $u_{\sigma,t}$ and $u_{\sigma_{tb},t}$, enter as independent arguments in the policy functions with a coefficient different from zero. Furthermore, these cubic terms are quantitatively significant.

To calibrate the model, we first estimate the process for the interest rate (24), (25), (26), and (27) using EMBI+ data and a Bayesian approach and we set the parameters for the law of motion of the real interest rate equal to the median of the posterior distributions. Then, we pick the remaining parameters of the model by targeting some moments of the Argentinian economy. Our calibration must target the moments of interest generated by the ergodic distributions and not the moments of the deterministic steady state, since those last ones are not representative of the stochastic dynamics.

6.3. Impulse Response Functions

Now we can analyze the IRFs of shocks to the country spreads and their volatility. In figure 6.1, we plot the IRFs to these shocks (rows) of consumption (first column), investment (second column), output (third column), labor (fourth column), the interest rate (fifth column), and debt (the sixth column). Interest rates are expressed in basis points, while all other variables are expressed as percentage deviations from the mean of their ergodic distributions (computed by simulation).

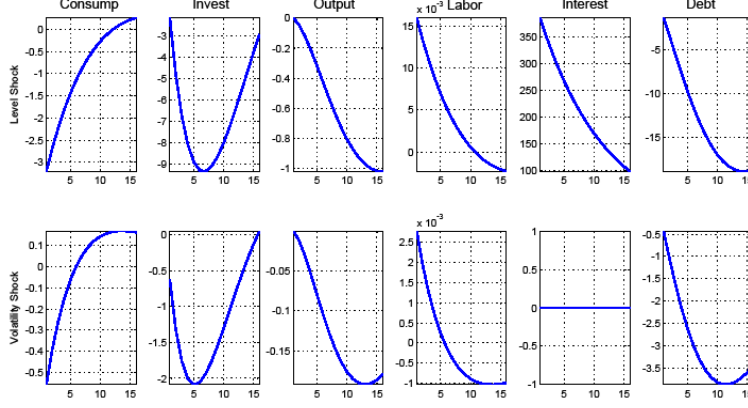


Figure 6.1: IRFs Argentina

The first row of panels plots the IRFs to a one-standard-deviation shock to the Argentinean country spread, $u_{r,t}$. Following an annualized rise of 385 basis points (that corresponds to an increase of nearly 33 basis points at a monthly rate) in Argentina's spread, the country experiences a persistent contraction, with consumption dropping 3.20 percent upon impact and investment falling for seven quarters. Furthermore, the decline in output is highly persistent: after 16 quarters, output is still falling (at that time it is -1.16 percent below its original level). Labor starts by slightly increasing (due to the negative wealth effects) but later falls (by a very small margin given our preferences) due to the reduction in investment and the subsequent decrease of marginal productivity. Debt falls for 14 quarters, with a total reduction of nearly 19 percent of the original value of the liability. The intuition for these movements is well understood. A higher r_t raises the service payment of the debt, reduces consumption, forces a decrease in the level of debt (since now it is more costly to finance it), and lowers investment through a non-arbitrage condition between the returns to physical capital and to foreign assets. This exercise shows that our model delivers the same answers as the standard model when hit by equivalent level shocks and to place in context the size of the IRFs to volatility shocks.

The second row of panels plots the IRFs to a one-standard-deviation shock to the volatility of the Argentinean country spread, $u_{\sigma,t}$. To put a shock of this size in perspective, we estimate that the collapse of LTCM in 1998 meant a positive volatility shock of 1.5 standard deviations and that the 2001 financial troubles amounted to two repeated shocks of roughly 1 standard deviation. First, note that there is no movement on the domestic interest rate faced by Argentina or its expected value. Second, there is a) a contraction in consumption, b) a decrease of investment, c) a slow fall in output, d) labor increases slightly to fall later, and e) debt shrinks upon impact

and keeps declining until it reaches its lowest level, roughly three and a half years after the shock. These IRFs show how increments in risk have real effects on the economy even when the real interest rate remains constant.

The intuition is as follows. Small, open economies rely on foreign debt to smooth consumption and to hedge against idiosyncratic productivity shocks. When the volatility of real interest rates rises, debt becomes riskier as the economy becomes exposed to potentially fast fluctuations in the real interest rate and their associated and unpleasant movements in marginal utility. To reduce this exposure, the economy lowers its outstanding debt by cutting consumption. Moreover, since debt is suddenly a worse hedge for the productivity shocks that drive returns to physical capital, investment falls. A lower investment also reduces output. Interestingly enough, we do not have any of the real-option effects of risk emphasized by the literature, for example, when we have irreversibilities (Bloom, 2009). Introducing those effects would increase the impact of shocks to volatility on investment. Thus, our results are likely to be a lower bound to the implications of time-varying risk.

7. What We Know and What We Do Not Know About Volatility

We arrive now towards the end of our long trip and it seems a fitting conclusion to take stock and enumerate what we know and what we do not know about volatility.

If we try to summarize what we know, we can venture three lessons. First, there is strong evidence that, in many contexts, time series experience time-varying volatility and that an understanding of the behavior of the data requires in consequence an understanding of the behavior of the volatility changes. Second, it is easy to write DSGE models in which volatility changes over time and in which we can measure the impact of these variations in risk. Third, there are a number of contexts where these variations in risk seem sufficiently important from a quantitative perspective as to deserve a more careful consideration.

On the other hand, there are also plenty of issues that we do not understand. First, and foremost, we do not have a good explanation of why aggregate volatility changes over time. In the models that we presented in this chapter, SV was assumed as exogenous. In some more involved models (for instance, where monetary and fiscal policy changes), part of the time-variation in volatility can be endogeneized but, at the same time, it is often the case that the question of why volatility changes is just pushed one step back to some unexplained change in policy. It is fair to note that macroeconomics, in general, lacks a very solid theory of why we have shocks,

either technological, preferences or any other. Much progress has been made just by investigating the consequences of a given exogenous shock without too much attention to its origins. By analogy, much progress may still be made by investigating the consequences of volatility shocks. Second, we do not fully understand many of the theoretical properties of models with SV. Just as an example, we do not have theorems regarding the differentiability of the decision rules with respect to the relevant components of SV beyond some simple cases. Third, there are still many questions regarding the best computational and empirical strategies to take these models to the data, including the best specifications for the structure of the changes of volatility over time. Finally, we know very little about the implications of volatility for optimal policy design.

But, fortunately, we do not see this lack of understanding as a fundamental problem but as a challenge to motivate research for many years to come. We expect to see much work on documenting and measuring the changes in volatility over time, on working out models that generate variations in risk in an endogenous way, and on assessing the implications for policy.

References

- [1] Aït-Sahalia, Y., L.P. Hansen, and J.A. Scheinkman (2009). “Operator Methods for Continuous-Time Markov Processes.” In Y. Aït-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*, North Holland.
- [2] An, S. and F. Schorfheide (2006). “Bayesian Analysis of DSGE Models.” *Econometric Reviews* 26, 113-172.
- [3] Arellano, C., Y. Bai, and P.J. Kehoe (2010). “Financial Markets and Fluctuations in Uncertainty.” *Mimeo*, Federal Reserve Bank of Minneapolis.
- [4] Aruoba, S.B., J. Fernández-Villaverde, and J. Rubio-Ramírez (2006). “Comparing Solution Methods for Dynamic Equilibrium Economies.” *Journal of Economic Dynamics and Control* 30, 2477-2508.
- [5] Bachmann, R., S. Elstner, and E. Sims (2010). “Uncertainty and Economic Activity: Evidence from Business Survey Data.” NBER *Working Paper* 16143.
- [6] Bekaert, G., M. Hoerova, M. and M. Lo Duca (2010). “Risk, Uncertainty and Monetary Policy.” *Mimeo*, Columbia University.
- [7] Benati, L. and P. Surico (2009). “VAR Analysis and the Great Moderation.” *American Economic Review* 99, 1636-1652.
- [8] Blanchard, O.J. and J. Simon (2001). “The Long and Large Decline in U.S. Output Volatility.” *Brookings Papers on Economic Activity* 2001-1, 187-207.

- [9] Bloom, N. (2009). “The Impact of Uncertainty Shocks.” *Econometrica* 77, 623–685.
- [10] Bloom, N., N. Jaimovich, N., and M. Floetotto (2008). “Really Uncertain Business Cycles.” *Mimeo*, Stanford University.
- [11] Bollerslev, T. (1986). “Generalized Autoregressive Conditional Heteroskedasticity.” *Journal of Econometrics* 31, 307-327.
- [12] Bollerslev, T. (2010). “Glossary to ARCH (GARCH)” in *Volatility and Time Series Econometrics Essays in Honor of Robert Engle*. Mark Watson, Tim Bollerslev and Jeffrey Russell (eds.), Oxford University Press.
- [13] Chari, V.V., L.J. Christiano, and P.J. Kehoe (1994), “Optimal Fiscal Policy in a Business Cycle Model.” *Journal of Political Economy* 102, 617-652.
- [14] Christiano, L., M. Eichenbaum, and C.L. Evans (2005). “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy.” *Journal of Political Economy* 113, 1-45.
- [15] Christiano, L., R. Motto, and M. Rostagno (2009). “Financial Factors in Economic Fluctuations.” *Mimeo*, Northwestern University.
- [16] Clarida, R., J. Galí, and M. Gertler (2000). “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory.” *Quarterly Journal of Economics* 115, 147-180.
- [17] Cogley, T. and T.J. Sargent (2005). “Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.” *Review of Economic Dynamics* 8, 262-302.
- [18] Diebold, F.X. (1986). “Modeling the Persistence of Conditional Variances: A Comment.” *Econometric Reviews* 5, 51-56.
- [19] Dynan, K.E., D.W. Elmendorf, and D.E. Sichel (2006). “Can Financial Innovation Help to Explain the Reduced Volatility of Economic Activity?” *Journal of Monetary Economics* 53, 123-150.
- [20] Elder, J. (2004a). “Some Empirical Evidence on the Real Effects of Nominal Volatility.” *Journal of Economics and Finance* 28, 1-13.
- [21] Elder, J. (2004b). “Another Perspective on the Effects of Inflation Uncertainty.” *Journal of Money, Credit and Banking* 36, 911-928.
- [22] Engle, R.F. (1982). “Autoregressive Conditional Heteroskedasticity With Estimates of the Variance of UK Inflation.” *Econometrica* 50, 987-1008.
- [23] Engle, R.F. and V.K. Ng (1993). “Measuring and Testing the Impact of News on Volatility.” *Journal of Finance* 48, 1749-1778.
- [24] Fatás, A. (2002). “The Effects of Business Cycles on Growth.” in Norman Loayza and Raimundo Soto (eds), *Economic Growth: Sources, Trends and Cycles*. Central Bank of Chile.

- [25] Fernández-Villaverde, J. (2010). “The Econometrics of DSGE Models.” *SERIES: Journal of the Spanish Economic Association* 1, 3-49
- [26] Fernández-Villaverde, J. and J. Rubio-Ramírez (2004). “Comparing Dynamic Equilibrium Models to Data: A Bayesian Approach.” *Journal of Econometrics*, 123, 153-187.
- [27] Fernández-Villaverde, J. and J. Rubio-Ramírez (2005). “Estimating Dynamic Equilibrium Economies: Linear Versus Nonlinear Likelihood.” *Journal of Applied Econometrics*, 20, 891-910.
- [28] Fernández-Villaverde, J. and J. Rubio-Ramírez (2007). “Estimating Macroeconomic Models: A Likelihood Approach.” *Review of Economic Studies* 74, 1059-1087.
- [29] Fernández-Villaverde, J., P. Guerrón-Quintana, and J. Rubio-Ramírez (2010a). “Fortune or Virtue: Time-Variant Volatilities Versus Parameter Drifting in U.S. Data.” *NBER Working Paper* 15928.
- [30] Fernández-Villaverde, J., P. Guerrón-Quintana, and J. Rubio-Ramírez (2010b). “Reading the Recent Monetary History of the U.S., 1959-2007” *Federal Reserve Bank of St. Louis Review* 92, 311-338.
- [31] Fernández-Villaverde, J., P.A. Guerrón-Quintana, J. Rubio-Ramírez, and Martín Uribe (2009). “Risk Matters: The Real Effects of Volatility Shocks.” *American Economic Review*, forthcoming.
- [32] Fountas, S. and M. Karanasos (2007). “Inflation, Output Growth, and Nominal and Real Uncertainty: Empirical Evidence for the G7.” *Journal of International Money and Finance* 26, 229-250.
- [33] Grier, K.B. and M.J. Perry (2000). “The Effects of Real and Nominal Uncertainty on Inflation and Output Growth: Some GARCH-M Evidence.” *Journal of Applied Econometrics* 15, 45-58.
- [34] Guerrón-Quintana, P. (2009). “Money Demand Heterogeneity and the Great Moderation External Link.” *Journal of Monetary Economics* 56, 255-266.
- [35] Goodfriend, M. and R. King (2007). “The Incredible Volcker Disinflation.” *Journal of Monetary Economics* 52, 981-1015.
- [36] Hamilton, J. (2008). “Macroeconomics and ARCH.” *Mimeo*, University of California-San Diego.
- [37] Haavelmo, T. (1994). “The Probability Approach in Econometrics.” *Econometrica* 12 (Supplement), iii-vi+1-115.
- [38] Jin, H. and K.L. Judd (2002). “Perturbation Methods for General Dynamic Stochastic Models.” *Mimeo*, Hoover Institution.

- [39] Justiniano A. and G.E. Primiceri (2008). “The Time Varying Volatility of Macroeconomic Fluctuations.” *American Economic Review* 98, 604-641.
- [40] Kaminsky, G.L., C.M. Reinhart and C.A. Vegh (2003). “The Unholy Trinity of Financial Contagion.” *Journal of Economic Perspectives* 17, 51-74.
- [41] Khan, M. S. (1977). “The Variability of Expectations in Hyperinflations.” *Journal of Political Economy* 85, 817-827.
- [42] Kim, C. and C.R. Nelson (1998) “Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle.” *Review of Economics and Statistics* 81, 608-616.
- [43] Klein, B. (1977). “The Demand for Quality-Adjusted Cash Balances: Price Uncertainty in the U.S. Demand for Money Function.” *Journal of Political Economy* 85, 692-715.
- [44] Künsch, H.R. (2005). “Recursive Monte Carlo Filters: Algorithms and Theoretical Analysis.” *Annals of Statistics* 33, 1983-2021.
- [45] Kydland, F., and E. C. Prescott (1982). “Time to Build and Aggregate Fluctuations.” *Econometrica* 50, 1345–1370.
- [46] Lee, K., S. Ni, and R. Ratti (1995). “Oil Shocks and the Macroeconomy: The Role of Price Variability.” *Energy Journal* 16, 39-56.
- [47] Longstaff, F.A., J. Pan, L.H. Pedersen, and K.J. Singleton (2007). “How Sovereign Is Sovereign Credit Risk?” Mimeo.
- [48] Lubick, T. and F. Schorfheide (2004). “Testing for Indeterminacy: An Application to U.S. Monetary Policy.” *American Economic Review* 94, 190-217.
- [49] McConnell, M.M. and G. Pérez-Quirós (2000). “Output Fluctuations in the United States: What Has Changed Since the Early 1980’s?” *American Economic Review* 90, 1464-1476.
- [50] Mendoza, E. (1991). “Real Business Cycles in a Small Open Economy.” *American Economic Review* 81, 797-818.
- [51] Mendoza, E. (1995). “The Terms of Trade, the Real Exchange Rate, and Economic Fluctuations.” *International Economic Review* 36, 101-37.
- [52] Nelson, D.B. (1991). “Conditional Heteroskedasticity in Asset Returns: A New Approach.” *Econometrica* 59, 347-370.
- [53] Ricardo, D. (2005). *The Works and Correspondence of David Ricardo, Vol. 5: Speeches and Evidence*. Edited by Piero Sraffa with the Collaboration of M.H. Dobb. Indianapolis, Liberty Fund.

- [54] Santos, M.S. and A. Peralta-Alva (2005). “Accuracy of Simulations for Stochastic Dynamic Models”. *Econometrica* 73, 1939-1976.
- [55] Schmitt-Grohé, S., and M. Uribe (2004). “Solving Dynamic General Equilibrium Models Using a Second-Order Approximation to the Policy Function.” *Journal of Economic Dynamics and Control* 28, 755-775.
- [56] Sentana, E. (1995). “Quadratic ARCH Models.” *Review of Economic Studies* 62, 639-661.
- [57] Shephard, N. (2008). “Stochastic Volatility.” *The New Palgrave Dictionary of Economics*. Palgrave MacMillan.
- [58] Sims, C.A. and T. Zha (2006). “Were There Regime Switches in U.S. Monetary Policy?” *American Economic Review* 96, 54-81.
- [59] Stock, J.H. and M.W. Watson (2002). “Has the Business Cycle Changed, and Why?” *NBER Macroeconomics Annual* 17, 159-218.
- [60] Uribe M., and V. Yue (2006). “Country Spreads and Emerging Countries: Who Drives Whom?” *Journal of International Economics* 69, 6-36.
- [61] Zakoïan, J.-M. (1994). “Threshold Heteroskedastic Models.” *Journal of Economic Dynamics and Control* 18, 931-955.