Essays in Macroeconometrics

Vassilios S. Bazinas

Balliol College

University of Oxford

A thesis submitted in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

Trinity 2017

To all my teachers.

Statement of Originality

I certify that the intellectual content of this thesis is the product of my own work and that all assistance received in preparing this thesis and relevant sources have been acknowledged. This thesis has not been submitted for any other degree or purposes.

Vassilios S. Bazinas

Acknowledgements

A debt of gratitude is owed to my two supervisors, Professors Bent Nielsen and Dimitrios Tsomocos, with whom two chapters of this dissertation are co-authored. Christopher Bowdler and Kevin Sheedy provided valuable comments in their role as examiners.

Chapter 1 benefitted from useful comments by Guido Ascari, Steve Bond, Andrew Chesher, and Nanny Wermuth as well as from useful discussions with David Cox, David Hendry, and members of the Econometrics Group at the University of Oxford.

Chapter 2 benefitted from helpful comments from Guido Ascari, Charles Goodhart, Leander Heldring, Juan Francisco Martinez, Rahul Nath, Bent Nielsen, Udara Peiris, and seminar participants at the Saïd Business School at the University of Oxford. Support from the Research Assistance Fund at the Saïd Business School is gratefully acknowledged.

Chapter 3 benefitted from helpful comments from Jeremy Chiu, Leander Heldring, Dimitrios Tsomocos, and participants of the Quantitative and Analytical Methods Network seminar series at the International Monetary Fund. I am also grateful to Miguel Segoviano for sharing data and generating useful discussions that helped advance the work.

This thesis was completed through the generous support of Balliol College, the Department of Economics at the University of Oxford, and the David Walton Memorial Fund.

Abstract

This thesis consists of three self-contained chapters.

The first chapter develops a method to explore the causal transmission of a catalyst variable through two endogenous variables of interest. The method is based on the reduced-form system formed from the conditional distribution of two endogenous variables given the catalyst and combines elements from instrumental variable analysis and recursive identification of structural vector autoregressions. Conditions for uniqueness of the causal transmission are provided.

The second chapter identifies and characterises periods of financial fragility over the period 1973–2016 within a simple regime-switching framework. Fragility is driven by high default risk and low profitability of intermediaries, as defined in Tsomocos (2003). Recessions occur primarily when fragility is high, but high fragility can occur without leading to recession. Credit growth is lower under high fragility, but it is not as significant a driver of regime changes when financial fragility is directly taken into account. Shocks have regime-dependent effects that are attributable primarily to switching dynamics as opposed to shock size. Finally, shocks are amplified through financial fragility, suggesting the presence of a debt-deflation channel.

The third chapter assesses the linkages between the real economy and the financial intermediation sector over the period 1999–2015 within a time-varying parameter framework with stochastic volatility. Low volatility and a period of protracted economic growth initiate a feedback loop between the real economy and the financial sector, consistent with the financial instability hypothesis of Minsky (1986). Shocks that adversely affect expectations of future economic growth have the greatest impact when macro-financial linkages are strongest. Augmenting the model with macroprudential capital requirements and monetary policy adjusted for the zero lower bound, we provide evidence that the timing of policy interventions matters. Counter-cyclical

interventions produce a stabilising effect when implemented early but are also capable of exacerbating debt-deflation if employed too late in the financial cycle.

Contents

Introduction

Since Hume's (1758) critique of induction, economists have been wary of explicitly discussing causality within empirical models. The development of regression techniques and the work on structural econometric models in the twentieth century, however, have provided a probabilistic foundation that, at the very least, implies causal language (Hood and Koopmans, 1953; Granger, 1969; Simon, 1977; Sims, 1980; Angrist and Krueger, 2001). Elsewhere in the literature, the advance of graphical modelling techniques in statistics and computer science (Lauritzen, 1996; Cox and Wermuth, 1996; Pearl, 2000; Spirtes et al., 2000) has precipitated a paradigm shift in causal inference that is only gradually seeping into economics (Swanson and Granger, 1997; Demiralp and Hoover, 2003; Moneta, 2008).

The more recent experience of the Global Financial Crisis has precipitated a different kind of shift in applied macroeconomics. It has become painfully clear that linear statistical models are ill-suited for capturing the evolving dynamics and state-dependence of the macro-financial relationship that are predicted by theory (Minsky, 1986; Kiyotaki and Moore, 1997; Bernanke et al., 1999; Tsomocos, 2003). The development of nonlinear methods and the advancement of computational techniques (Hamilton, 1989; Primiceri, 2005) has permitted a recent growth in the literature that examines the effects of financial frictions in the data, but there is no systematic effort to analyse frictions, financial stability, and multiple policy interventions jointly.

This thesis develops the necessary components to investigate the causal transmission of shocks within nonlinear empirical frameworks that can be useful to policymakers concerned with financial stability. The thesis is divided into two parts, consisting of three self-contained chapters.

The first part consists of one chapter and develops (jointly with Bent Nielsen) a theory for testing the causal transmission of shocks through a system of equations. We combine aspects from instrumental variable analysis in applied microeconomics and structural analysis in applied macroeconomics to provide testable conditions in a reduced-form model under which the effect of shocks may induce structure. Thus, rather than imposing structure by assumption—as in the applied macroeconomics literature—or simply identifying effects without uniqueness in direction—as in the applied microeconomics literature—we provide a method for researchers to identify the unique transmission of effects. This theory is developed for general densities in a bivariate case in order to demonstrate concepts clearly.

The second part of this thesis consists of two empirical chapters examining the role of financial fragility within nonlinear models. This part of the thesis draws inspiration from the financial instability hypothesis of Minsky (1986). The first chapter examines (jointly with Dimitrios Tsomocos) the relationship between financial intermediation and the real economy within a Markov-switching model. Historical regimes of financial fragility are accurately identified through high default risk and low profitability of intermediaries, as predicted in general equilibrium models that incorporate endogenous default (Tsomocos, 2003). The second chapter revisits the financial instability hypothesis and delves deeper into explaining transitions between regimes within the context of a time-varying parameter model with stochastic volatility. This chapter provides evidence that attitudes toward risk evolve following periods of protracted growth, thereby increasing the likelihood of amplification mechanisms leading to a state of financial fragility where adverse shocks have relatively larger effects. Within this framework, the time-varying role of macroprudential and monetary policy is explored.

Time-variation in empirical models is necessary for policymakers since decisions to intervene will differ based on the stage of the financial cycle. At the same time, understanding the causal transmission mechanisms through which shocks operate offers a broader perspective on where intervention is possible and what its repercussions will be. Combining these two features will be the subject of future research.

1 Causality in econometrics

In order to introduce the concept of *causal transmission* that is developed in Part I of this thesis, it is instructive to understand the epistemological constraints within which we operate. This section discusses the necessity of an untestable causal assumption required to carry out causal inference, while also distinguishing between association and causation.

1.1 No causes in, no causes out

All discussions of causality in economics begin with David Hume. Hume (1758) characterises the problem of causation as the impossibility of observing natural laws that govern the relationships between objects. Rather, the only notion of causality available to us is that of regularity or "constant conjunction." This leads to the problem of induction; for any object Y observed as following an object X in the past, we believe that this relationship will continue into the future. Hume (1758) writes

> *"When we look about us towards external objects, and consider the operation of causes, we are never able, in a single instance, to discover any power or necessary connexion* [sic.]*; any quality, which binds the effect to the cause, and renders the one an infallible consequence of the other. We only find, that the one does actually, in fact, follow the other."*

Hume discusses singular causal facts as instantiations of an unobservable natural law, and therefore general causal facts cannot be known; they can be inferred, by induction, from singular causal facts. Causal claims are thus reduced to statements about deterministic association and treated with skepticism by Humeans. Nonetheless, Hume appears to describe economics as a causal science. On discussing interest rates and money supply, Hume (1758) writes

> *"...low interest rates and plenty of money are in fact almost inseparable. But still it is of consequence to know the principle whence any phenomenon arises, and to distinguish between a cause and a concomitant effect."*

The impulse to seek causal explanations for economic phenomena is thus tempered by Hume's critique of causal inference by induction and has led to uneasiness about causality in the social sciences.

Humean empiricism has evolved from deterministic to probabilistic regularities (see Suppes, 1970, among others), reducing general causal facts to statements about correlation. Rather than relying on constant conjunction, causes are defined by the statement $\mathsf{P}(Y|X) > \mathsf{P}(Y)$; Y is more likely to occur if X also occurs. Cartwright (1989) argues, however, that while statistical associations can be used as indicators of causation, arbitrary associations are not equivalent to causal claims. A correlation between X and Y can imply any of the following: X causes Y; Y causes X; or Y and X have a common cause W. Rather, causation between Y and X exists if and only if

$$\mathsf{P}(Y|X, W_1, \ldots, W_n) > \mathsf{P}(Y|W_1, \ldots, W_n), \tag{1.1}$$

where the set $\{X, W_1, \ldots, W_n\}$ includes all causes for Y. Antecedent causal information is thus necessary in order to establish whether statistical association can be translated into further general or singular causal facts. This represents the view "no causes in, no causes out"; without some causal facts in hand, no causal inference is possible. The epistemological framework within which we operate is further characterised by distinguishing between association and causation.

1.2 Causality by intervention

Common approaches to causality in economics reflect associational rather than causal influence. For instance, the formulation of causality in Granger (1969) is the time-series analogue of the statement in (1.1), yet it amounts to reducing causality to temporal correlation; cause and effect are linked by the flow of time, rather than by any meaningful economic structure. Causality in the sense of Granger (1969) is thus an associational concept that is useful for predictive purposes but says little about causal structure. How then are we to define an appropriate framework for causal inference?

Pearl (2000) argues that the probabilistic framework on which causality has been developed is inadequate in moving beyond associational concepts. The distinction between association and causation is made based on the joint distribution of variables. Any relationship derived from the joint distribution of the modelled variables describes associational concepts. Causal inference relies on information outside the joint distribution that has an asymmetric influence; outside conditions change and affect the joint distribution of the observed variables, without an opposite effect. Thus, the objective of causal analysis is to infer associations under changing conditions. Various notations are used to distinguish between conditioning by association

$$\mathsf{p}(y|x) \equiv \mathsf{P}(Y = y|X = x), \tag{1.2}$$

and conditioning by intervention

$$\mathsf{p}(y||x) \equiv \mathsf{P}(Y = y|do(X = x)), \tag{1.3}$$

as in Lauritzen (1996) and Pearl (2000). The difference between the two concepts is that (1.2) expresses dependence that occurs naturally between variables over the sample, whereas (1.3) implies that there is an outside intervention that can affect this dependence. We do not use the above notation to indicate conditioning by intervention, since there exist many notations for expressing this within the graphical modelling literature. We imply (1.3) but use standard probabilistic notation, while being clear that we are not operating on a joint distribution but rather on a conditional distribution.

Despite this distinction between association and causation, we must still rely on an untestable causal assumption: that the effect in (1.3) is the result of intervention and not some unknown confounding variable. Since it is practically impossible to include every potential cause of Y, we must rely on some form of antecedent information when selecting our modelled variables and make a causal assumption.

The approach that we follow is a combination of economics knowledge and specification testing. Once the researcher has decided on an appropriate system of variables, he or she can construct a reduced-form model with appropriate lag structure and include dummy variables to satisfy specification tests; see Hendry and Nielsen (2007) for a comprehensive modelling approach. Dummy variables that correspond to known interventions, are treated as natural experiments in the data. We do not discuss the identification of interventions from dummy variables and assume they are known, for instance through the historical narrative approach. The causal assumption we make is that these dummy variables—which we eventually define as *catalysts*—exert asymmetric influence into the modelled system. We are bound by the principle "no causes in, no causes out."

1.3 Causal transmission in reduced-form models

In general, it is difficult to deduce the causal ordering of two observed variables from their joint distribution. However, if we have some causal fact in hand, so that we can assume that a third variable is causal, it may be possible to deduce how the effect of this third variable will transmit between the two variables of interest. By conditioning on a *catalyst*, the joint distribution of a bivariate system can be used to infer a *causal transmission*.

Our method can be thought of as unifying elements of the identification procedures in instrumental variable analysis and structural vector autoregressions. Instrumental variable analysis will, in general, not order the endogenous variables. Rather, we use it to identify a structural relation uniquely. Cholesky decomposition orders endogenous variables, but the ordering is not unique. By carrying out a Cholesky decomposition in the presence of an instrument there is scope for a unique ordering which is interpretable as a causal transmission. In this situation we will refer to the instrument as a catalyst.

The catalyst w may transmit causally through the variables y, z. It is possible that w transmits through z to y or through y to z or, of course, that there is no ordering of the variables. A first set of conditions is needed for establishing that the catalyst w transmits through z to y, say, in a unique fashion. A second set of conditions is needed for showing that w actually affects y. The econometric framework is a reduced form based on the conditional distribution of y, z given w. The theory is formulated for general densities but with special attention to the most common cases, which include the bivariate normal distribution and mixtures of a univariate normal distribution with a logit or probit distribution.

We provide an application to money demand in the UK for illustrative purposes. Today, however, policymakers are consumed by different questions. For instance, how will interventions affect the financial sector? What are the various mechanisms through which this happens? The treatment of such questions requires us to move beyond linearity.

2 Financial stability

The experience and aftermath of the Great Recession underscored the need for nonlinear statistical models that are able to capture the evolving dynamics between the financial and business cycles. There are many definitions for financial stability, however, and many forms of nonlinearity to choose from. In Part II of this thesis, the relationship between the real economy and financial intermediaries is examined in separate empirical frameworks. We suit the particular modelling framework to the question at hand.

Fluctuations in the financial sector have long been recognised as a source of significant perturbations to the business cycle (Bernanke and Gertler, 1989; Kiyotaki and Moore, 1997). Early models took different approaches to account for the impact of financial stress on the macroeconomy, alternatively by introducing financial frictions in DSGE models (Carlstrom and Fuerst, 1997; Bernanke et al., 1999) and by introducing endogenous default in GEI models (Tsomocos, 2003; Dubey et al., 2005). A second generation of models has continued to improve on these early works, with frictions and default in financial intermediation becoming mainstay features (for a recent survey, see Goodhart et al., 2016).

Additionally, there is a growing appreciation in the empirical literature of the time-variation present within the macro-financial relationship and the role of financial shocks (Hubrich and Tetlow, 2015; Prieto et al., 2016; Gambetti and Musso, 2017), as well as the evolving role that monetary policy plays (Galì and Gambetti, 2015). The importance of time-varying shocks to house prices and credit for systemic risk are acknowledged and often treated separately, but these are never studied jointly. This misses an amplification mechanism that is crucial to policymakers. Additionally, no time-varying role is assigned to macroprudential requirements, though there is broad consensus for counter-cyclical measures within constant-parameter empirical studies (for a recent survey, see Galati and Moessner, 2017).

Part II assesses the evidence for the two theorems of the financial instability hypothesis put forth by Minsky (1986): that there exist distinct regimes of financial robustness and fragility; and that over prolonged periods of prosperity, the economy transitions between financial regimes as attitudes toward risk evolve. Drawing on the two aspects of the financial instability hypothesis, different empirical frameworks are employed. The first chapter demonstrates the existence of different financing regimes within a Markov-switching model. The objective in this work is to use economic theory as a reference point for motivating an appropriate definition of financial fragility in terms of default risk and profitability of financial intermediaries. Once it is shown that such a model adequately characterises financial fragility within the historical narrative, we can study how statics and dynamics vary across regimes. This provides a foundation for understanding the evolution between regimes, which is studied in the second chapter by employing a time-varying parameter model. Additionally, this provides a framework within which the effectiveness of macroprudential and monetary policy can be studied in order to mitigate losses following the accumulation of macro-financial linkages.

2.1 Financial fragility in a regime-switching model

We take the view that financial fragility develops endogenously as an equilibrium phenomenon due to the regular functioning of an economy. While at times this may manifest as a financial crisis or as a cyclical component of the data, our aim is to capture all regimes where the state of financial intermediation carries the potential for damaging macroeconomic effects. Ultimately, this becomes an informal test of Hyman Minsky's first hypothesis of financial instability: that the economy undergoes periods of financing stability and instability (Minsky, 1977, 1992). To motivate our reduced-form empirical analysis, we draw on structural economic models that incorporate endogenous default as an equilibrium result in different states (Tsomocos, 2003; Goodhart et al., 2006). We show that financial fragility will arise when the prevailing environment among intermediaries is characterised by high volatility and default risk, while simultaneously lacking sufficient profitability as a buffer to weather adverse conditions. Since equilibrium depends on the realisation of discrete states, we opt for a multivariate regime-switching framework that considers the effects of default, profitability, and credit jointly on economic growth.

Specifically, we ask the following questions. First, can financial fragility be defined in terms of a nonlinear association between the intermediation sector and the real economy? Second, how do the effects of shocks vary when intermediaries are highly fragile? Finally, what is the relative importance of financial fragility as a transmission channel for shocks to economic growth? The answers to these questions should inform the construction of appropriate measures of financial stress, as well as aid policymakers when deciding where and how to intervene.

We employ a Markov-switching vector autoregressive (MS-VAR) model for the United States over the period 1973–2016. To assess real economic effects, we include industrial production and credit extension in our set of macroeconomic variables. The health of the financial intermediation sector is assessed by proxying for default risk and profitability. We proxy for default risk by using the interbank borrowing (TED) spread and for profitability by using the market valuation of the largest US banks. Together, these are indicative of the fragility of the financial sector.

The main results are as follows. First, we find evidence for nonlinearity that is driven by high default risk and low profitability of intermediaries; this is consistent with the definition of fragility in Tsomocos (2003). Second, recessions occur primarily when fragility is high, but high fragility can occur without leading to recession. Third, while credit growth is lower under high fragility, it is not as significant a driver of regime changes when financial fragility is directly taken into account; credit can create financial imbalances over the long term, but is not an adequate measure of the health of financial intermediaries at shorter horizons. This result is reinforced by reproducing the analysis in Schularick and Taylor (2012) for the United States and finding that credit growth alone is not a useful predictor of financial crises. Fourth, shocks have regime-dependent effects that are attributable primarily to switching dynamics; this finding is consistent with Hubrich and Tetlow (2015) and antithetical to Stock and Watson (2012). Fifth, shocks are amplified in the high-fragility regime through default risk and profitability of intermediaries, suggesting the presence of a debt-deflation channel (Fisher, 1933). While we intentionally do not model policy variables and are primarily interested in characterising association rather than causation, this final result lends some weight to unconventional monetary policy that intervenes by providing credit when the private sector is severely constrained (Gertler and Karadi, 2011).

2.2 Assessing time-varying macro-financial linkages

Whereas the first chapter in Part II examines the evidence for different regimes of financial fragility, the second chapter focuses on the transition between regimes. The importance of this lies in the capacity for intervention. Assessing the development of macro-financial linkages allows us to identify turning points in the financial cycle where policy intervention may be useful. It is critically important to allow for time-variation in policy recommendations, however, since policies that are most effectively employed early in the financial cycle may prove disastrous later.

We estimate a statistical model that at once considers the aggregate behaviour of the financial intermediation sector, while allowing for time-varying macro-financial linkages and ask the following questions. First, how does optimism affect the accumulation of leverage in the upswing of the financial cycle? Second, does the importance of adverse shocks to expectations matter more after periods of protracted economic growth? Finally, can interventions in the form of macroprudential capital requirements and monetary policy successfully mitigate the buildup of risk?

To gain answers to these questions, we employ a Bayesian macroeconomic vector autoregressive model for the United States over the period 1999–2015. Our model differs from the standard setup by incorporating a robust measure of systemic risk in the financial intermediation sector that accounts for nonlinear dependence between institutions Goodhart and Segoviano (2009), while also modelling both house prices and credit. Additionally, given our policy-driven objectives, we employ the shadow policy rate of Wu and Xia (2016) to account for monetary policy at the zero lower bound and an aggregated tier 1 capital ratio as a proxy for macroprudential regulation. Thus, our VAR model is comprised of industrial production growth, house price inflation, private credit growth, systemic expected shortfall of financial intermediaries, the aggregate tier 1 capital ratio, and an adjusted Federal Funds rate.

The main results are as follows. First, we find evidence of exceptionally low volatility during the credit buildup that preceded the Great Recession, consistent with the volatility paradox of Brunnermeier and Sannikov (2014). Second, positive shocks to economic growth initiate a feedback loop with credit and housing prices over the period spanning 1999–2006. This occurs because shocks that are interpreted as positive signals of economic performance feed into agents' future expectations, thereby driving over-leveraging. Third, positive signals are found to reduce the repayment rate among financial intermediaries consistently throughout the sample. Fourth, "not unusual" shocks (Minsky, 1992) that reduce the repayment rate have a much more severe effect when macro-financial linkages are high and the financial sector is in a state of fragility. Fifth, we find that policy shocks are most effective when employed early in the financial cycle. Attempting to deflate asset prices and check credit growth either through macroprudential or monetary policy at the peak of the financial cycle can have severe adverse effects.

We differ from most studies of the time-varying macro-financial relationship by studying house prices, credit growth, and systemic risk jointly. This is important in accounting for evolving attitudes toward risk (Minsky, 1986) that are amplified through net worth (Bernanke et al., 1999) and collateralised lending (Kiyotaki and Moore, 1997). Additionally, this helps us evaluate alternative channels of monetary policy through effects on financial stability (Borio and Zhu, 2012; Lin et al., 2015). We are the first to study macroprudential requirements within a time-varying parameter framework.

Our results are consistent with the leverage cycle as described in Kiyotaki and Moore (1997) and Geanakoplos (2010), whereby increasing asset prices lead to higher collateral values, which in turn permits greater credit extension. The effects are highest at the peak of the leverage cycle following protracted periods of economic growth, while lowest following the crash and during recession. This is consistent with the Minsky (1977, 1986, 1992) hypothesis formalised in Bhattacharya et al. (2015). Observing successive periods of growth, agents revise their expectations of a strong economy upward, assigning a higher probability to continued growth. This leads to higher leverage that is employed to finance increasingly riskier projects. Following an adverse shock, the economy reacts much more severely than it would have, had there not been a buildup in leverage. We see this in our results by noting the muted effect of shocks to systemic risk and asset prices following the Great Recession. We argue that this occurs because the financial crisis corrected burgeoning expectations of continued economic growth.

Part I

CHAPTER 1

Causal Transmission in Reduced-Form Models

Chapter Abstract

We propose a method to explore the causal transmission of a catalyst variable through two endogenous variables of interest. The method is based on the reduced-form system formed from the conditional distribution of the two endogenous variables given the catalyst. The method combines elements from instrumental variable analysis and recursive identification of structural vector autoregressions. We give conditions for uniqueness of the causal transmission.

1.1 Introduction

In general, it is difficult to deduce the causal ordering of two observed variables from their joint distribution. However, if we can assume that a third variable is causal, it may be possible to deduce how the effect of this third variable will transmit between the two variables of interest. By conditioning on a *catalyst*, the joint distribution of a bivariate system can be used to infer a *causal transmission*. Our approach allows for different catalysts transmitting through the same two variables in different ways. We formulate this for a general distributional setup.

Philosophers and scientists argue that some background of causal knowledge is required in order to construct new causal facts. The view "no causes in, no causes out" from Cartwright (1989) expresses the concern that we cannot jump from theory to cause without some causal facts in hand. Pearl (2000) similarly underlines the importance of distinguishing between causal and associational concepts, as every causal conclusion relies on a causal assumption that is untested in observational studies. In contrast, Granger (1969) causality is an example of an associational concept seeking to infer correlations from data without a causal assumption. Moreover, Granger causality is concerned with temporal correlations as opposed to the ordering of contemporaneous variables. Causal analysis goes one step further by inferring correlations under changing conditions. In this paper we are particularly interested in such correlations between contemporaneous variables.

Our method can be thought of as unifying instrumental variable analysis and recursive ordering of structural vector autoregressions. Instrumental variable analysis will, in general, not order the endogenous variables. Rather, we use it to identify a structural relation uniquely. Cholesky decomposition orders endogenous variables, but the ordering is not unique. By carrying out a Cholesky decomposition in the presence of an instrument there is scope for a unique ordering which is interpretable as a causal transmission. In this situation we will refer to the instrument as a catalyst.

The catalyst w may transmit causally through the variables y, z. It is possible that w transmits through z to y or through y to z or, of course, that there is no ordering of the variables. We present two sets of testable conditions. A first set of conditions is needed for establishing that the catalyst w transmits through z to y, say, in a unique fashion. A second set of conditions is needed for showing that w actually affects y. The econometric framework is a reduced form based on the conditional distribution of y, z given w. The theory is formulated for general densities but with special attention to the most common cases, which include the bivariate normal distribution and mixtures of a univariate normal distribution with a logit or probit distribution.

The objective of this paper is thus to analyse past events where particular interventions, for instance on the policy rate, that were determined separately from the realisation of endogenous variables can be exploited. We may observe that only a subset of economic variables are affected, while others remain stable. Our interpretation is that there may be some important causal transmission channels at work, which we can learn about through careful econometric analysis. If we are able to find a causal transmission, then we can hope that a future intervention of the same type may transmit in the same manner. We restrict the analysis to the bivariate setup in order to be as clear as possible. Larger systems can have many possible transmission channels, which will be examined in future work.

Our notion of causal transmission bears many similarities to super exogeneity as introduced in Engle et al. (1983). A variable z is super exogenous for y with respect to w under two conditions: first, a structural invariance property, which is by and large similar to our causal transmission; second, weak exogeneity of z for y, which is an estimation property. We do not need the estimation property, which gives a more flexible framework for studying causal transmission. The apparatus that we introduce is also related to graphical model theory that represents independence structures of statistical models as graphs. This, however, is relegated to the concluding discussion, as we define our own framework.

In §1.2 we define and explore causal transmissions. In §1.3 we generalise the idea to situations with multiple catalysts which transmit through the variables of interest in different ways and we offer a structural interpretation that combines Cholesky decomposition and instrumental variable estimation. An empirical illustration using a UK monetary data set follows in §1.4. Proofs are given in an appendix.

1.2 Causal transmission

Causal relations are asymmetric by nature: influence flows one way and cannot be reversed. We analyse a joint conditional probability model for two endogenous variables given a catalyst. We start by exploring ordering within a bivariate normal setup. Subsequently, we give results for unique ordering and non-trivial transmission in a general bivariate distribution setup. From this we define causal transmission.

1.2.1 The uniqueness problem

Suppose we are interested in an economic relationship between two endogenous, or modelled, variables (y, z) given a third variable w. The third variable w is deemed exogenous. Thus, we are interested in the conditional distribution $\mathsf{f}(y, z|w)$. Under normality the distribution of y, z given w is given by

$$\begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} \gamma_{yw} \\ \gamma_{zw} \end{pmatrix} w + \begin{pmatrix} \epsilon_y \\ \epsilon_z \end{pmatrix}, \tag{1.2.1}$$

where the innovations are normally distributed, with positive definite variance

$$\begin{pmatrix} \epsilon_y \\ \epsilon_z \end{pmatrix} \overset{\mathsf{D}}{=} \mathsf{N}\left\{\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{yy} & \sigma_{yz} \\ \sigma_{zy} & \sigma_{zz} \end{pmatrix}\right\}. \tag{1.2.2}$$

There are two different ways of ordering y and z, corresponding to two Cholesky decompositions. First, we can condition y on z to obtain the equations

$$y = \gamma_{yz}z + \gamma_{yw\cdot z}w + \epsilon_{y\cdot z}, \tag{1.2.3}$$

$$z = \gamma_{zw}\ w + \epsilon_z, \tag{1.2.4}$$

with derived parameters $\gamma_{yz} = \sigma_{yz}/\sigma_{zz}$ and $\gamma_{yw\cdot z} = \gamma_{yw} - \gamma_{yz}\gamma_{zw}$, and independent, normal innovations $\epsilon_{y\cdot z}$, ϵ_z with variances $\sigma_{yy\cdot z} = \sigma_{yy} - \sigma^2_{zy}/\sigma_{zz}$, σ_{zz}. Secondly, we can condition z on y to obtain the equations

$$z = \gamma_{zy}y + \gamma_{zw\cdot y}w + \epsilon_{z\cdot y}, \tag{1.2.5}$$

$$y = \gamma_{yw}\ w + \epsilon_y, \tag{1.2.6}$$

where $\gamma_{zy} = \sigma_{zy}/\sigma_{yy}$ and $\gamma_{zw\cdot y} = \gamma_{zw} - \gamma_{zy}\gamma_{yw}$, and the independent, normal innovations $\epsilon_{z\cdot y}$, ϵ_y with variances $\sigma_{zz\cdot y} = \sigma_{zz} - \sigma^2_{yz}/\sigma_{yy}$, σ_{yy}. Without further information the two orderings are equivalent.

An ordering arises from the equations (1.2.3)-(1.2.4) under the restrictions $\gamma_{yw\cdot z} = 0$ and $\gamma_{zw} \neq 0$. The equations (1.2.3)-(1.2.4) then reduce to

$$y = \gamma_{yz}z + \epsilon_{y\cdot z}, \tag{1.2.7}$$

$$z = \gamma_{zw}\ w + \epsilon_z. \tag{1.2.8}$$

This ordering of (y, z) is unique in the sense that it is not possible to have $\gamma_{yw\cdot z} = 0$

and $\gamma_{zw} \neq 0$ so that (1.2.3)-(1.2.4) reduce to (1.2.7)-(1.2.8), while $\gamma_{zw\cdot y} = 0$ in (1.2.5)-(1.2.6). We prove this result for general distributions in §1.2.2.

We note, in passing, that the situation described in equations (1.2.7)-(1.2.8) is different from the usual instrumental variable problem. In that problem, the exclusion restriction $\gamma_{yw\cdot z} = 0$ and the relevance condition $\gamma_{zw} \neq 0$ both hold, but the errors $\epsilon_{y\cdot z}$ and ϵ_z are allowed to be correlated. We will see that the additional independence assumption is what gives the ordering of the variables.

1.2.2 Result for general distributions

For a general joint density of y, z conditional on w, $\mathsf{f}(y, z|w)$, we explore testable restrictions that ensure a unique and non-trivial chain from w through z to y. In §1.2.3 we interpret w as a *catalyst* that initiates a unique *causal transmission* through z to y.

Unique transmission

The natural generalisation of the result for normal distributions is a Markov property. Generally, the joint density of $(y, z|w)$ can be decomposed as

$$\mathsf{f}(y, z|w) = \mathsf{f}(y|z, w)\mathsf{f}(z|w) = \mathsf{f}(z|y, w)\mathsf{f}(y|w). \tag{1.2.9}$$

At this point, there is no natural ordering of the bivariate system. The uniqueness result is inspired by the normal example. It presents a condition under which we can rule out the possibility that both $\mathsf{f}(y|z, w) = \mathsf{f}(y|z)$ and $\mathsf{f}(z|y, w) = \mathsf{f}(z|y)$. In other words, we give a condition that ensures a Markov chain from w to y through z, while excluding a Markov chain from w to z through y.

Theorem 1.2.1 *Suppose the density* $\mathsf{f}(y, z|w)$ *has support on a product space, and it* *is positive on this support. Suppose that, for all* y, z,

$$\mathsf{f}(y|z, w) = \mathsf{f}(y|z). \tag{1.2.10}$$

Then

$$\begin{aligned} \mathsf{f}(z|w) &\neq \mathsf{f}(z) \quad \textit{in a set of } y, z \textit{ with positive probability,} & (1.2.11)\\ \Rightarrow \qquad \mathsf{f}(z|y, w) &\neq \mathsf{f}(z|y) \quad \textit{in a set of } y, z \textit{ with positive probability.} & (1.2.12)\end{aligned}$$

The requirement in Theorem 1.2.1 that the support is a product space is satisfied in a range of common situations, for instance in a normal setup. It allows for the less interesting case where y or z is atomic. If z is atomic, then condition (1.2.11) always fails. If y is atomic, then conclusion (1.2.12) reduces to (1.2.11).

Theorem 1.2.1 gives conditions for a unique Markov structure among the variables. Condition (1.2.10) implies

$$\mathsf{f}(y, w|z) = \mathsf{f}(y|z)\mathsf{f}(w|z). \tag{1.2.13}$$

Theorem 1.2.1 shows that the conditions (1.2.10), (1.2.11) imply (1.2.12), and therefore there is no Markov structure from w through y to z, that is

$$\mathsf{f}(z, w|y) \neq \mathsf{f}(z|y)\mathsf{f}(w|y). \tag{1.2.14}$$

In other words, the conditional model for y, z given w allows for two possible Markov structures, but we can distinguish these through testable assumptions.

The next step is a requirement that the Markov structure is non-trivial.

Definition 1.2.1 *Consider the conditional distribution of* y, z *given* w *with the Markov structure* $\mathsf{f}(y, z|w) = \mathsf{f}(y|z)\mathsf{f}(z|w)$. *If* $\mathsf{f}(y|z) \neq \mathsf{f}(y)$ *and* $\mathsf{f}(z|w) \neq \mathsf{f}(z)$, *we have a* ***non-trivial Markov structure****, which we represent by the undirected graph* w—z—y.

We note that Theorem 1.2.1 implies that the two non-trivial Markov structures w—z—y and w—y—z cannot hold simultaneously. We summarise this as follows.

Theorem 1.2.2 *Suppose the density $\mathsf{f}(y,z|w)$ has support on a product space, and it is positive on this support. Suppose that, for all y, z,*

$$\mathsf{f}(y|z,w) = \mathsf{f}(y|z) \qquad , \tag{1.2.15}$$

and that, for all y, z in a set with positive probability,

$$\mathsf{f}(z|w) \neq \mathsf{f}(z) \qquad \textit{and} \qquad \mathsf{f}(y|z) \neq \mathsf{f}(y). \tag{1.2.16}$$

Then we have a unique and non-trivial Markov structure w—z—y.

The non-trivial Markov structure w—z—y is symmetric in y and w. Thus, in itself it does not give a direction from w to y. A non-trivial Markov structure does not, in general, imply that w and y are dependent, so w may affect z without affecting y. We explore this next and return to the directional issue and causality in §1.2.3.

Non-trivial transmission

Suppose we have a non-trivial Markov structure w—z—y. In a causal context it is of interest if w actually has an effect on y through z. Indeed, the Markov structure w—z—y allows the possibility that y and w are independent. From a causal view point, this is not so exciting. We will seek to characterise when the effect is non-trivial.

Definition 1.2.2 *Consider a non-trivial Markov structure w—z—y. There is a* ***non-trivial transmission*** *between w and y when $\mathsf{f}(y|w) \neq \mathsf{f}(y)$ in a set with positive probability, represented as w—y.*

When exploring the transmission between w and y, it is useful to introduce the notation $\perp\!\!\!\perp$ for independence, so that $y \perp\!\!\!\perp w$ when $\mathsf{f}(y|w) = \mathsf{f}(y)$. Note that the conditional distribution of y given w is a compound distribution of the form

$$\mathsf{f}(y|w) = \int \mathsf{f}(y|z)\mathsf{f}(z|w)dz. \tag{1.2.17}$$

The integral can be interpreted as summation if the dominating measure dz is discrete.

We start with a sufficient condition for a trivial transmission.

Lemma 1.2.1 *Suppose* $\mathsf{f}(y, z|w) = \mathsf{f}(y|z)\mathsf{f}(z|w)$. *Then* $y \perp\!\!\!\perp z$ *or* $z \perp\!\!\!\perp w \Rightarrow y \perp\!\!\!\perp w$.

In general, the sufficient condition is not necessary. From a causal transmission perspective, we are interested in exploring when the sufficient condition is necessary. Indeed, with a non-trivial Markov structure w—z—y, we have the marginal dependence $y \not\!\perp\!\!\!\perp z$ and $z \not\!\perp\!\!\!\perp w$. Our question is when this implies $y \not\!\perp\!\!\!\perp w$. We give some examples.

When z is binary the question relates to collapsibility of contingency tables. Dawid (1980, Theorem 8.3) attributes the following result to Yule.

Lemma 1.2.2 *Suppose w—z—y with binary z. Then $y \not\!\perp\!\!\!\perp w$.*

Moving away from binary z, we find the same result for some common distributions.

Lemma 1.2.3 *Suppose w—z—y with normal $(y, z|w)$ satisfying (1.2.1). Then $y \not\!\perp\!\!\!\perp w$.*

Lemma 1.2.4 *Suppose w—z—y with binary y so that $(y|z)$ is logit,* $\text{logit}\{\mathsf{f}(y = 1|z)\} = \gamma_{yz}z$ *or probit,* $\mathsf{f}(y = 1|z) = \Phi(\gamma_{yz}z)$, *while $(z|w)$ is normal,* $\mathsf{N}(\gamma_{zw}w, \sigma_{zz})$. *Then $y \not\!\perp\!\!\!\perp w$.*

However, the sufficient condition is not necessary in general. An example follows.

Example 1.2.1 *Suppose w—z—y, We construct an example where $y \not\!\perp\!\!\!\perp z$ and $z \not\!\perp\!\!\!\perp w$, yet $y \perp\!\!\!\perp w$. Let w, y be binary, while z takes three values. Describe the conditional distributions* $\mathsf{f}(z|w)$ *and* $\mathsf{f}(y|z, w) = \mathsf{f}(y|z)$ *by the transition matrices*

0	1	2	$z \mid w$
4/8	3/8	1/8	0
4/8	2/8	2/8	1

0	1	$y \mid z$
1/4	3/4	0
2/4	2/4	1
2/4	2/4	2

The conditional distribution $\mathsf{f}(y|w)$*, computed as the product of the transition matrices, satisfies* $\mathsf{f}(y|w) = \mathsf{f}(y)$*, that is*

0	1	$y \mid w$
3/8	5/8	0
3/8	5/8	1

1.2.3 Causal interpretation

Theorem 1.2.2 gave testable conditions ensuring that the conditional distribution $\mathsf{f}(y, z|w)$ reduces to a non-trivial Markov structure w—z—y. This was followed in §1.2.2 by a variety of conditions ensuring a non-trivial transmission between w and y. In the following, we give this a causal interpretation. We will think of the variable w as taking a value that is determined outside the system (y, z). This value then transmits through the system as described by the conditional distribution $\mathsf{f}(y, z|w)$.

Definition 1.2.3 *Consider variables* w, z, y*. Assume that for each realisation of* w*, then* $\mathsf{f}(y, z|w)$ *describes the distribution of outcomes of* y, z*. Let* w *represent an intervention on the system. Then we say that* w *is a* ***catalyst****.*

By an intervention, we mean an external, autonomous change that affects only the specified subset of variables (Pearl, 2000, p. 23). The objective is to separate actions, where variables are assigned values by intervention, and observations, where variables assume values according to a joint distribution. Pearl (2000) assumes, however, that the mechanism that is altered by an intervention is known, as is the nature of the alteration. Directionality within a system of variables is discovered or assumed prior to analysis of interventions by representing a joint distribution with a directed acyclical graph. Contrastingly, we only discover directionality conditional on the presence of an intervention; this corresponds to the notion of a transmission.

Definition 1.2.4 *Consider non-trivial Markov structure w—z—y with non-trivial transmission between w and y and where w is a catalyst. Then we have a* ***causal transmission*** *of the catalyst w to y through z. This is represented by the notation $w \to z \to y$.*

Our definition of a catalyst is related to the notion of a causal effect by Pearl (2000). Pearl's definition is formulated for directed acylical graphs that assume a direction. In contrast, Definitions 1.2.3, 1.2.4 consider the testable and undirected Markov structure w—z—y and merely gives it a causal and directional interpretation. Thus, the important distinction between our exposition and the existing literature is the objective of characterising potential transmission of catalysts using testable assumptions as far as possible. Catalysts will not always be obvious but can potentially be discovered through examination of observational data. Essentially, we are seeking to discover natural experiments and straddle the boundary between observational and experimental frameworks. Definition 1.2.4 has the feature that we are agnostic about the causal relationship between the endogenous variables when a catalyst is not present.

The assumptions required for a causal transmission exclude the situation of a *collider*. A variable z is a collider when $\mathsf{f}(y|z,w) \neq \mathsf{f}(y|z)$ even though $\mathsf{f}(y|w) = \mathsf{f}(y)$. In the graphical modelling literature, this is represented as $w \to z \leftarrow y$. The first condition for a collider contradicts the conditions assumed in Definition 2.1 (w—z—y), while the second condition contradicts the conditions assumed in Definition 2.2 (w—y).

We consider three special cases: a normal model and two types of logit/probit-normal mixtures.

Example 1.2.2 *Suppose* $(y, z|w)$ *has a bivariate normal distribution as in* (1.2.3)-(1.2.4) *or* (1.2.5)-(1.2.6) *with a positive definite covariance matrix. If* $\gamma_{yw\cdot z} = 0$ *while* $\gamma_{zw} \neq 0$*, then Theorem 1.2.1 implies a unique Markov structure. If in addition* $\gamma_{yz} \neq 0$*, then Theorem 1.2.2 and Lemma 1.2.3 imply a non-trivial Markov structure* w—z—y *and a non-trivial transmission* w—y*. When* w *is interpretable as a catalyst, then* $w \to z \to y$*.*

Example 1.2.3 *Suppose* y *is binary and* $(y, z|w)$ *satisfies a logit-normal mixture model or a probit-normal mixture model. That is, the conditional distribution* $(y|z, w)$ *satisfies*

$$\text{logit}\{\mathsf{f}(y = 1|z, w)\} = \gamma_{yz}z + \gamma_{yw\cdot z}w \quad \textit{or} \quad \mathsf{f}(y = 1|z, w) = \Phi(\gamma_{yz}z + \gamma_{yw\cdot z}w),$$

while $(z|w)$ *is* $\mathsf{N}(\gamma_{zw}w, \sigma_{zz})$*. If* $\gamma_{yw\cdot z} = 0$ *while* $\gamma_{zw} \neq 0$*, then Theorem 1.2.1 implies a unique Markov structure. If in addition* $\gamma_{yz} \neq 0$*, then Theorem 1.2.2 and Lemma 1.2.4 imply a non-trivial Markov structure* w—z—y *and a non-trivial transmission* w—y*. When* w *is interpretable as a catalyst, then* $w \to z \to y$*.*

We note that in this situation $\mathsf{f}(y|z, w)$ *is much easier to work with than* $\mathsf{f}(z|y, w)$*. Due to Theorem 1.2.1, we only need to check the first instance to narrow the potential orderings of the system* $(y, z|w)$*.*

Example 1.2.4 *Suppose* z *is binary and* $(y|z, w)$ *is* $\mathsf{N}(\gamma_{yz}z+\gamma_{yw\cdot z}w, \sigma_{yy\cdot z})$*. If* $\gamma_{yw\cdot z} = 0$ *while* $\mathsf{f}(z|w) \neq \mathsf{f}(z)$ *then Theorem 1.2.1 implies a unique Markov structure. If, in addition,* $\gamma_{yz} \neq 0$ *then Theorem 1.2.2 and Lemma 1.2.2 imply a non-trivial Markov structure* w—z—y *and a non-trivial transmission* w—y*. When* w *is interpretable as a catalyst, then* $w \to z \to y$*.*

1.3 Structural considerations

The causal transmission concept unites ideas from Cholesky decompositions within structural vector autoregressions with ideas from instrumental variable estimation. We explore how causal transmission arises as a special case in those two settings. We draw comparisons with the more restrictive concept of super exogeneity before delving into a structural interpretation when multiple catalysts are available.

1.3.1 Cholesky decomposition

Sims (1980) used vector autoregressions to address the haphazard accumulation of restrictions to achieve identification in the large simultaneous equation models of the time. This approach has evolved into the frequently-used structural vector autoregressive (SVAR) approach, where a structural model is identified from the reduced form. In its basic form, this involves a recursive ordering of the variables. We will discuss how Cholesky decomposition relates to causal transmission.

It is well-known that, while useful, recursive orderings are not unique. Causal transmission takes its starting point in recursive orderings but uses a catalyst to establish a unique ordering. If we ignore dynamic features we can explore this using the setup in §1.2.1. The reduced-form system for the variables y, z given w is then given by (1.2.1). Pre-multiplying that system by a square matrix A gives a structural model

$$\begin{pmatrix} a_{1y} & a_{1z} \\ a_{2y} & a_{2z} \end{pmatrix} \begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} b_{1w} \\ b_{2w} \end{pmatrix} w + \begin{pmatrix} e_1 \\ e_2 \end{pmatrix} \tag{1.3.1}$$

where $e = A\epsilon$ has covariance $\Omega_e = A\Sigma_\epsilon A'$ and where Σ_ϵ is the covariance matrix in (1.2.2). A structural model of this general form is not identifiable from the reduced-form model. We therefore consider two Cholesky decompositions where A is triangular and Ω is diagonal. The first possibility is

$$A^z = \begin{pmatrix} 1 & a_{1z} \\ 0 & 1 \end{pmatrix}, \qquad \Omega_e^z = \begin{pmatrix} \omega_{11}^z & 0 \\ 0 & \omega_{22}^z \end{pmatrix}, \tag{1.3.2}$$

which is identifiable from (1.2.3), (1.2.4), when $a_{1z} = -\gamma_{yz} = -\sigma_{yz}/\sigma_{zz}$, while $\omega_{11}^z = \sigma_{yy\cdot z}$ and $\omega_{22}^z = \sigma_{zz}$. The second possibility is

$$A^y = \begin{pmatrix} 1 & 0 \\ a_{2y} & 1 \end{pmatrix}, \qquad \Omega_e^y = \begin{pmatrix} \omega_{11}^y & 0 \\ 0 & \omega_{22}^y \end{pmatrix}, \tag{1.3.3}$$

which is identifiable from (1.2.5), (1.2.6), when $a_{2y} = -\gamma_{zy} = -\sigma_{zy}/\sigma_{yy}$, while $\omega_{11}^y = \sigma_{yy}$ and $\omega_{22}^y = \sigma_{zz\cdot y}$. The Cholesky factors (1.3.2), (1.3.3) are observationally equivalent.

Using the causal transmission analysis we may find, for instance, that $w \to z \to y$ in the reduced-form model. This is consistent with the first Cholesky factor (1.3.2) with the additional restriction that $b_{1w} = 0$, that is

$$\begin{pmatrix} 1 & a_{1z} \\ 0 & 1 \end{pmatrix} \begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ b_{2w} \end{pmatrix} w + \begin{pmatrix} e_1 \\ e_2 \end{pmatrix}, \tag{1.3.4}$$

where the errors e_1 and e_2 are independent. This model is asymmetric. It shows how economic shocks in z can transmit to the structural relation $y + a_{1z}z$. Subtly, the asymmetry is captured by w rather than the errors e_1, e_2, which have a symmetric role. Thus, the interpretation of this structural model is that it shows how, typically, large shocks of the type w move through the economy, which is also subject to, typically, small shocks of the type e_1, e_2. For instance, w may represent the onset of the financial crises or a major government intervention, while the shocks e_1, e_2 represent the minor, daily pulling and pushing forces in the economy. Thus, the structural assumption we need for this analysis is that w is a catalyst. The remaining features of the causal transmission $w \to z \to y$ are testable and discoverable from reduced-form analysis. In §1.3.5 we extend this analysis to a situation with multiple catalysts.

1.3.2 Instrumental variable estimation

The traditional simultaneous equations model has no causal direction. Instead, the focus is to estimate the behavioural equations with the aid of instruments. We discuss this in the context of a simple demand and supply example, with a focus on the demand curve.

Consider the following demand function formulated in terms of the (log) quantity q and the (log) price level p

$$q = a_0 + a_1 p + u_d. \tag{1.3.5}$$

The demand function can be identified with the use of an instrument w that is valid $\mathsf{E}(wu_d) = 0$ and informative $\mathsf{E}(wp) \neq 0$. This corresponds to a supply shock and gives the first-stage equation

$$p = b_0 + b_1 w + u_p, \tag{1.3.6}$$

implying an exclusion restriction in the demand equation. The demand function describes the linear relation between prices and quantities. The variables are jointly determined, so there is no causal direction between them. Thus the equation (1.3.5) can be reversed as

$$p = -\frac{a_0}{a_1} + \frac{1}{a_1} q - \frac{1}{a_1} u_d. \tag{1.3.7}$$

This is reflected when estimating with limited information maximum likelihood. In that case, the product of the estimate for a_1 in equation (1.3.5) and the estimate for $1/a_1$ in equation (1.3.7) is indeed unity. This applies both in the just-identified case where w is univariate and in the over-identified case where w is multivariate.

In general, the structural error u_d in (1.3.5) and the first-stage error u_p in (1.3.6) may be conditionally dependent given w in the instrumental variable problem. However, if the two errors are indeed conditionally independent, then the demand equation (1.3.5) represents the conditional distribution of q given p, w with the property $\mathsf{f}(q|p, w) = \mathsf{f}(q|p)$ so that the first condition of Theorem 1.2.1 is met. The second condition of Theorem 1.2.1 ensures the informativeness of w in the first-stage equation (1.3.6), that is $b_1 \neq 0$. Theorem 1.2.1 then shows that the reversed demand equation (1.3.7) cannot represent the conditional distribution of p given q, w, and there is a unique Markov structure w—p—q: the demand equation does not depend on w, but the conditional distribution must depend on w. In the case of normal errors, w—p—q implies w—q by Lemma 1.2.3. With the additional assumption that w is a catalyst, we arrive at the causal transmission $w \to p \to q$.

We can also start from a reduced-form model and discover Markov structures without imposing structural assumptions. Ignoring intercepts, the starting point is the reduced-form system (1.2.1), that is

$$\begin{pmatrix} q \\ p \end{pmatrix} = \begin{pmatrix} \gamma_{qw} \\ \gamma_{pw} \end{pmatrix} w + \begin{pmatrix} \epsilon_q \\ \epsilon_p \end{pmatrix}, \tag{1.3.8}$$

where $\mathsf{f}(\epsilon_q, \epsilon_p|w) = \mathsf{f}(\epsilon_q, \epsilon_p)$ is normal. Here w is merely an exogenous variable, albeit a candidate for an instrument. The reduced-form system implies an equation that does not depend on the exogenous w

$$q = \frac{\gamma_{qw}}{\gamma_{pw}} p + u \qquad \text{where} \qquad u = \epsilon_q - \frac{\gamma_{qw}}{\gamma_{pw}} \epsilon_p, \tag{1.3.9}$$

when $\gamma_{pw} \neq 0$, so that the instrument is informative. The error term u in equation (1.3.9) has the property that it is independent of the instrument w since $\mathsf{f}(\epsilon_q, \epsilon_p|w) = \mathsf{f}(\epsilon_q, \epsilon_p)$ implies $\mathsf{f}(u|w) = \mathsf{f}(u)$. We note that in this just-identified setup, the ratio of least squares estimators for γ_{qw} and γ_{pw} is the indirect least squares estimator, which is the same as the two-stage least squares estimator or limited information maximum likelihood estimator. If the slope γ_{qw}/γ_{pw} is positive and if we can interpret w as a supply shock, then equation (1.3.9) can be interpreted as a demand equation. By imposing the testable restriction that u and ϵ_p are independent or, equivalently, that $\gamma_{qw}/\gamma_{pw} = \mathsf{Cov}(\epsilon_q, \epsilon_p)/\mathsf{Var}(\epsilon_p)$, the unique Markov structure w—p—q is obtained by Theorem 2.1. If the instrument w can be viewed as a catalyst, it transmits causally through p to the traded quantity q. The likelihood ratio test for the hypothesis of independence of u and ϵ_p can be approximated by the Hausman test for endogeneity.

1.3.3 External instruments

Montiel Olea et al. (2015) identify impulse response functions within structural vector autoregressions by finding instruments for the structural shocks. For comparison, we ignore the lagged dependent variable, which does not play any particular role in the argument, and focus on the first structural shock. The structural model is

$$\begin{pmatrix} y_t \\ z_t \end{pmatrix} = \begin{pmatrix} 1 & H_{21} \\ H_{12} & H_{22} \end{pmatrix} \begin{pmatrix} e_{1,t} \\ e_{2,t} \end{pmatrix}, \qquad (1.3.10)$$

where the structural shocks $e_{1,t}$ and $e_{2,t}$ are assumed to have a diagonal covariance matrix $\Omega_e = diag(\omega_{11}, \omega_{22})$. The parameter H_{12} is of interest, but it is not identifiable from the reduced-form model. The identification strategy sought here is through an external instrument w_t that satisfies $\mathsf{E}(e_{1,t}w_t) = \alpha \neq 0$ for informativeness and $\mathsf{E}(e_{2,t}w_t) = 0$ for validity; diagonality of Ω_e is required for identification of the structural shocks but not for identification of the impulse response function. The coefficient H_{12} can then be found from the covariance of (y_t, z_t) and w_t. It is useful to express the assumptions through the joint density of $\mathsf{f}(e_1, e_2, w)$. If normality is assumed, for instance, the requirements on the instrument and the structural shocks are that

$$\begin{pmatrix} e_t \\ w_t \end{pmatrix} \overset{\mathsf{D}}{=} \mathsf{N}\left\{\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \omega_{11} & 0 & \alpha \\ 0 & \omega_{22} & 0 \\ \alpha & 0 & \sigma_{ww} \end{pmatrix}\right\}. \tag{1.3.11}$$

This approach identifies impulse response functions but does not impose any causal ordering. An ordering would require either $H_{12} = 0$ or $H_{21} = 0$. The instrumental variable assumptions placed on the structural model, however, do not imply either of these situations; they are necessary but not sufficient for causal transmission of w, since no asymmetry is introduced between the endogenous variables.

The recursive ordering achieved through Cholesky decomposition with $H_{21} = 0$ gives direction but no uniqueness as it is observationally equivalent to swapping the roles of the variables y and z. The external instrument approach with $\mathsf{E}(e_{1,t}w_t) = \alpha$ gives uniqueness but no direction; additionally imposing $H_{21} = 0$ would yield a causal transmission $w \to y \to z$, but imposing $H_{12} = 0$ does not due to $\mathsf{E}(e_{2,t}w_t) = 0$. In both identification approaches, the restrictions introduce an asymmetry in the structural model. For example, external instruments impose the asymmetry on the joint distribution of the unobserved structural errors $e_{1,t}, e_{2,t}$ and the instrument w_t. The restrictions are sufficient to just-identify the structural model from the reduced-form model that remains symmetric. It seems necessary to impose the asymmetry on the reduced-form model in order to ensure a unique direction.

1.3.4 Super exogeneity

Broadly speaking, causal transmission is a weaker concept than super exogeneity introduced by Engle et al. (1983): if we have causal transmission and weak exogeneity, we get super exogeneity.

To appreciate the distinction we need to recall the definition of weak exogeneity. In a time series setup write the density $(y_t, z_t|w_t, past)$ as $\mathsf{f}_{\xi,\lambda}(y_t, z_t|w_t, past) = \mathsf{f}_\xi(y_t|z_t, w_t, past)\mathsf{f}_\lambda(z_t|w_t, past)$, with parameters ξ, λ varying in a parameter space Θ. If the parameter space is a product space $\Theta = \Xi \times \Lambda$ so $\xi \in \Xi$ and $\lambda \in \Lambda$, then z is weakly exogenous for ξ. Then the joint likelihood can be optimized by maximising the conditional, partial likelihood for $(y_t|z_t, past)$ and the marginal, partial likelihood for $(z_t|past)$ separately. If, in addition, $\mathsf{f}_\xi(y_t|z_t, w_t, past) = \mathsf{f}_\xi(y_t|z_t, past)$ while $\mathsf{f}_\lambda(z_t|w_t, past) \neq \mathsf{f}_\lambda(z_t, past)$ we get, essentially, super exogeneity of z for y with respect to w. We note that these additional conditions are the conditions (1.2.10), (1.2.11) of the uniqueness Theorem 1.2.1.

Causal transmission is weaker than super exogeneity in the sense that the directions of causal transmission and weak exogeneity need not match. To see this write

$$\mathsf{f}_{\xi,\lambda}(y, z|w) = \mathsf{f}_\xi(y|z, w)\mathsf{f}_\lambda(z|w) = \mathsf{f}_\zeta(z|y, w)\mathsf{f}_\eta(y|w).$$

The model may include further restrictions that are not explicitly stated here, so that ξ, λ vary in a product space $\Xi \times \Lambda$, whereas ζ, η do not vary in a product space. In other words, z is weakly exogenous for y, but y is not weakly exogenous for z. At the same time we may have $w \to y \to z$, so that w transmits causally through y to z. This causal transmission is not compatible with weak exogeneity, so we do not have super exogeneity. Some more subtle differences between the two concepts are that super exogeneity of z for y with respect to w does neither require that $\mathsf{f}_\xi(y|z) \neq \mathsf{f}_\xi(y)$ nor $\mathsf{f}_{\xi,\lambda}(y|w) \neq \mathsf{f}_{\xi,\lambda}(y)$.

1.3.5 Multiple causal transmissions

The concept of causal transmission generalises to multiple catalysts that may flow through the system in different ways. For notational convenience we present this by augmenting the linear, normal system (1.2.1) with two distinct catalysts w_1, w_2, so that

$$\begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} \gamma_{y1} \\ \gamma_{z1} \end{pmatrix} w_1 + \begin{pmatrix} \gamma_{y2} \\ \gamma_{z2} \end{pmatrix} w_2 + \begin{pmatrix} \epsilon_y \\ \epsilon_z \end{pmatrix}, \qquad (1.3.12)$$

where $\mathsf{f}(\epsilon_y, \epsilon_z | w) = \mathsf{f}(\epsilon_y, \epsilon_z)$ is normal as in (1.2.2). The variables w_1, w_2 are observable and may represent two types of shocks to the economy at different points in time.

We now set up the two possibilities for ordering y, z through conditioning. Conditioning y on z gives

$$\begin{aligned} y &= \gamma_{yz} z + \gamma_{y1\cdot z} w_1 + \gamma_{y2\cdot z} w_2 + \epsilon_{y\cdot z}, & (1.3.13) \\ z &= \gamma_{z1} \; w_1 + \gamma_{z2} \; w_2 + \epsilon_z, & (1.3.14) \end{aligned}$$

where $\epsilon_{y\cdot z}$, ϵ_z are independent and $\gamma_{yz} = \sigma_{yz}/\sigma_{zz}$, while conditioning z on y gives

$$\begin{aligned} z &= \gamma_{zy} y + \gamma_{z1\cdot y} w_1 + \gamma_{z2\cdot y} w_2 + \epsilon_{z\cdot y}, & (1.3.15) \\ y &= \gamma_{y1} \; w_1 + \gamma_{y2} \; w_2 + \epsilon_y, & (1.3.16) \end{aligned}$$

where $\epsilon_{z\cdot y}$, ϵ_y are independent and $\gamma_{zy} = \sigma_{zy}/\sigma_{yy}$. Assuming w_1, w_2 are catalysts, we get two causal transmission hypotheses

$$\begin{aligned} \mathsf{H}_1 : &\quad \gamma_{y1\cdot z} = 0 \cap \gamma_{z1} \neq 0 \cap \gamma_{yz} \neq 0 \quad \Rightarrow \quad w_1 \to z \to y, & (1.3.17) \\ \mathsf{H}_2 : &\quad \gamma_{z2\cdot y} = 0 \cap \gamma_{y2} \neq 0 \cap \gamma_{zy} \neq 0 \quad \Rightarrow \quad w_2 \to y \to z. & (1.3.18) \end{aligned}$$

When the hypotheses H_1 and H_2 are both satisfied we get causal transmissions in opposite directions, which we represent by superimposing two directed graphs

$$\mathsf{H}_1 \cap \mathsf{H}_2 \qquad \Rightarrow \qquad w_1 \longrightarrow z \rightleftarrows y \dashleftarrow w_2 \qquad (1.3.19)$$

The joint restrictions imposed by $\mathsf{H}_1 \cap \mathsf{H}_2$ are possibly best expressed in terms of the original system (1.3.12) as

$$\mathsf{H}_1 \cap \mathsf{H}_2 : \quad \gamma_{y1} = \frac{\sigma_{yz}}{\sigma_{zz}} \gamma_{z1}, \quad \gamma_{z2} = \frac{\sigma_{zy}}{\sigma_{yy}} \gamma_{y2}, \quad \gamma_{z1} \neq 0, \quad \gamma_{y2} \neq 0, \quad \sigma_{yz} \neq 0.$$

Written in a vector format, we have the reduced-form model

$$\begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} \sigma_{yz}/\sigma_{zz} \\ 1 \end{pmatrix} \gamma_{z1} w_1 + \begin{pmatrix} 1 \\ \sigma_{zy}/\sigma_{yy} \end{pmatrix} \gamma_{y2} w_2 + \begin{pmatrix} \epsilon_y \\ \epsilon_z \end{pmatrix}, \quad (1.3.20)$$

Following the considerations in §1.3.1, the corresponding structural model is

$$\begin{pmatrix} 1 & -\gamma_{yz} \\ -\gamma_{zy} & 1 \end{pmatrix} \begin{pmatrix} y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ \delta_{21} \end{pmatrix} w_1 + \begin{pmatrix} \delta_{12} \\ 0 \end{pmatrix} w_2 + \begin{pmatrix} e_1 \\ e_2 \end{pmatrix}, \qquad (1.3.21)$$

where $\gamma_{yz} = \sigma_{yz}/\sigma_{zz}$ and $\gamma_{zy} = \sigma_{zy}/\sigma_{yy}$ are multipliers for the catalysts, while $\delta_{21} = (1-\rho^2)\gamma_{z1}$ and $\delta_{12} = (1-\rho^2)\gamma_{y2}$, with $\rho^2 = \sigma_{yz}^2/(\sigma_{yy}\sigma_{zz})$. The innovations of the structural equation (1.3.21) satisfy

$$\begin{pmatrix} e_1 \\ e_2 \end{pmatrix} \overset{\mathsf{D}}{=} \mathsf{N} \left\{ \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{yy\cdot z} & -\sigma_{yz}\left(1-\rho^2\right) \\ -\sigma_{zy}\left(1-\rho^2\right) & \sigma_{zz\cdot y} \end{pmatrix} \right\}, \qquad (1.3.22)$$

with correlation $-\rho = -\sigma_{yz}(1-\rho^2)/(\sigma_{yy\cdot z}\sigma_{zz\cdot y})^{1/2}$. We have identified a structural model with respect to catalysts w_1 and w_2 without imposing any ad hoc restrictions on the causal ordering through the covariance matrix. These catalysts are orthogonal to each other in the structural model, in the sense that w_1 is omitted from the first structural equation and w_2 is omitted from the second structural equation. Structure is, therefore, identified as a linear relationship that remains invariant to large shocks. Rather than imposing structure to identify orthogonal shocks, we use shocks to identify structure. Instead of having a structural model that is ordered for an entire sample, we are only concerned with ordering during periods when large interventions take place. We note that if the parameters γ_{yz}, γ_{zy} of the system (1.3.21) were unrelated to the covariance parameters σ_{yy}, σ_{yz}, σ_{zz} in (1.3.22) we would have a just-identified and undirected, bivariate simultaneous equations model.

Causal transmissions in both directions depending on the type of shock seems compatible with the discussion of shocks in macroeconomics. In many situations, we use indicator variables to represent large external shocks to the economy. When large external shocks arrive in quick succession, it may be difficult to separate the effect of the individual shocks. A pertinent example is the beginning of the financial crisis in 2007-2008 when oil shocks, financial collapse, and large fiscal and monetary policy interventions occurred in quick succession. We envisage that it would be possible to disentangle the effect of these shocks by lining these up, individually, with shocks at other points in time.

1.3.6 Comments

As we have seen, causal transmission is related to traditional approaches such as recursive ordering induced by a Cholesky decomposition and instrumental variable estimation.

As a causal concept, causal transmission is modest in scope: all causal orderings are relative to particular interventions with no attempt to give an overall causal ordering of the variables of interest, y, z. The concept is more modest than the causal inference interpretation of quasi-experiments, where the difference of potential and realised outcomes is estimated using an instrumental variable approach and the causal language from random control trials is applied, see Imbens (2014). Rather than conducting causal inference under an assumption of causal transmission, we are interested in conducting inference about the causal transmission itself. In practice, the consequence is that it becomes clearer that results can only be extrapolated to future interventions insofar as those interventions are comparable with the interventions in the sample.

The type of interventions are somewhat different in nature from the traditional shocks of structural vector autoregressive analysis. Here we focus on large interventions as opposed to the sample variance of the residuals, noting that the residuals simply represent the unmodelled part of the data. The approach reflects a view that it is difficult to disentangle many minor shocks to the economy.

Causal transmission is also restrictive in the sense that we need both a catalyst and a Markov structure. At the same time, it is potentially possible to discover catalysts empirically. An empirical example follows in §1.4.

1.4 Empirical example

We illustrate the causal transmission using the simplified bivariate model of money demand for the UK in Hendry and Nielsen (2007). This has the convenient features of being bivariate, reasonably well-specified, and with two catalysts operating in opposite directions. The data are formed from quarterly observations of log M1 money m, log real total final expenditure x, its log deflator p, and the three-month local authority rate net of the learning-adjusted rate on checkable deposits R_n taken from Hendry and Mizon (1993) over the period from 1963:2 to 1989:2.[1]

To simplify the analysis, we convert the four variables into a bivariate system, modelling the velocity of circulation of money v and the cost of holding money C through

$$v_t = x_t - m_t + p_t, \qquad C_t = \Delta p_t + R_{n,t}.$$

[1]Throughout most of the sample, R_n reflects the three-month local authority rate. The Banking Act of 1984, however, permitted commercial banks to pay interest on checkable deposits, thereby reducing the opportunity cost for money. The interest rate on these deposits is adjusted for learning, in order to reflect the diffusion of knowledge of this legislative change through the economy. For more details, the reader is referred to Hendry and Mizon (1993) and Hendry and Doornik (1994).

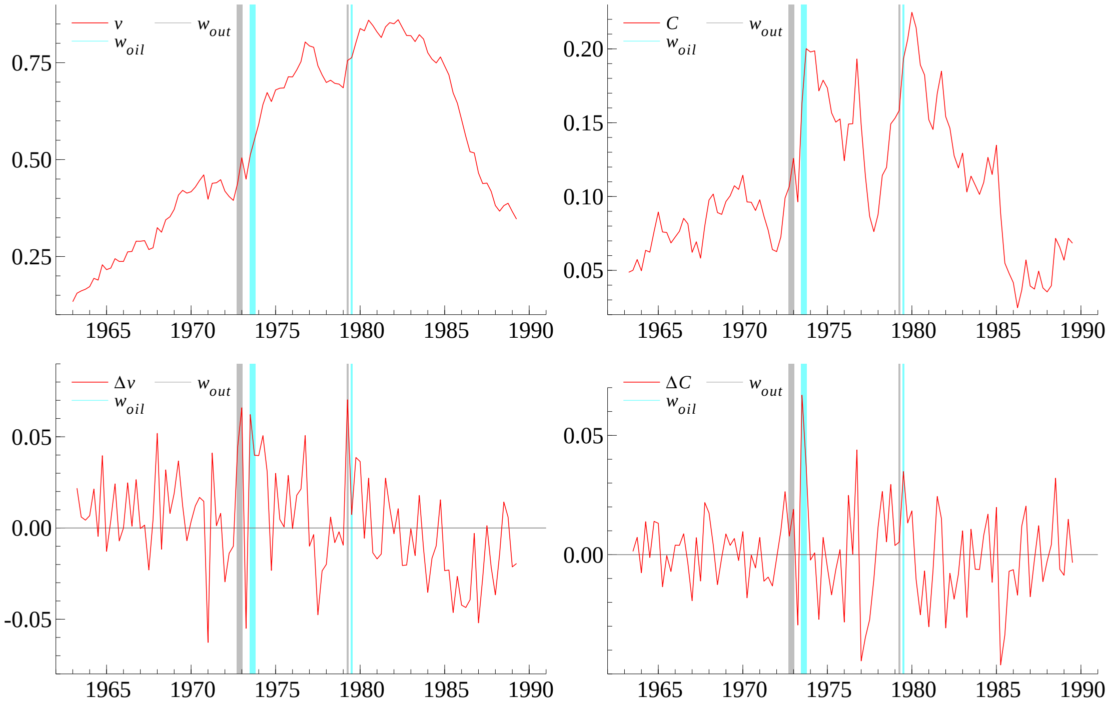

FIGURE 1.4.1: Levels and first-differences of the variables in the system (v, C) plotted with the selected outliers (w_{out}, w_{oil}) from Hendry and Mizon (1993).

We show how the results from the previous sections may be applied in practice to identify multiple causal transmissions. Subsequently, we provide impulse responses for the interventions that are identified. Finally, we address the Lucas critique that asserts that an econometric model may be unstable under changing conditions.

The graphics and subsequent computations were carried out in MATLAB (MATLAB, 2014) and OxMetrics (Hendry and Doornik, 2014).

1.4.1 The unrestricted reduced-form

Figure 1.4.1 shows v_t, C_t in levels and differences. The transformed data series are non-stationary, but their first-order differences have a more stationary appearance. The plots also show two dummy variables $w_{out,t}$, $w_{oil,t}$ representing large fiscal expansions in 1972:4–1973:1 and 1979:2 as well as the oil price shocks in 1973:3–4 and

TABLE 1.4.1: Three models estimated over the period 1963:4 to 1989:2. Standard errors reported in parentheses.

	Joint Model		**Conditional Models**		**Structural Model**	
	Δv_t	ΔC_t	$\Delta v_t \vert \Delta C_t$	$\Delta C_t \vert \Delta v_t$	Δv_t	ΔC_t
Δv_t	–	–	–	0.433	–	0.402
				(0.075)		(0.068)
ΔC_t	–	–	0.606	–	0.579	–
			(0.104)		(0.093)	
Δv_{t-1}	−0.343	−0.048	−0.314	0.100	−0.317	0.092
	(0.095)	(0.081)	(0.082)	(0.074)	(0.091)	(0.076)
ΔC_{t-1}	0.086	0.046	0.058	0.009	0.064	0.003
	(0.117)	(0.099)	(0.100)	(0.085)	(0.108)	(0.090)
v_{t-1}	−0.097	−0.005	−0.094	0.037	−0.095	0.035
	(0.014)	(0.012)	(0.012)	(0.013)	(0.012)	(0.013)
C_{t-1}	0.529	−0.077	0.575	−0.306	0.575	−0.293
	(0.071)	(0.060)	(0.062)	(0.065)	(0.066)	(0.068)
1	−0.004	0.009	−0.009	0.011	−0.009	0.011
	(0.006)	(0.005)	(0.005)	(0.004)	(0.005)	(0.004)
$w_{out,t}$	0.051	0.013	0.044	−0.010	0.039	0
	(0.012)	(0.010)	(0.010)	(0.009)	(0.010)	
$w_{oil,t}$	0.030	0.051	−0.001	0.038	0	0.039
	(0.012)	(0.010)	(0.011)	(0.009)		(0.008)
$\widehat{\sigma}_{vv}^{1/2}$	0.019	–	–	–	–	–
$\widehat{\sigma}_{CC}^{1/2}$	–	0.016	–	–	–	–
$\widehat{\rho}$	0.512		–		0.482	
$\widehat{\sigma}_{vv\cdot C}^{1/2}$	–	–	0.017	–	0.016	–
$\widehat{\sigma}_{CC\cdot v}^{1/2}$	–	–	–	0.014	–	0.014
likelihood	559.31				558.74	

1979:3. They will later be interpreted as catalysts.

Whereas the oil shocks are clearly exogenous to the UK economy, this is less obvious for the fiscal expansions. In fact, what we call fiscal shocks are the expansionary budget of 1972 proposed by Anthony Barber, then Chancellor of the Exchequer, and a significant VAT reduction. While both are likely endogenous to the UK economy, neither shock was set up to influence money demand, and so both can be considered exogenous for the bivariate model of (v, C). Furthermore, while the shocks are different in principle, the effect is the same, so that it becomes possible to extend our conclusions over several types of shocks that are expansionary in nature.

TABLE 1.4.3: Specification tests for unrestricted joint model. p-values reported in brackets.

Test	$\chi^2_{\mathsf{norm}}[2]$	$\mathsf{F}_{\mathsf{ar}(1-5)}[5,91]$	$\mathsf{F}_{\mathsf{arch}(1-4)}[4,95]$	$\mathsf{F}_{\mathsf{het}}[10,92]$	$\max Chow$
Δv_t	2.2 [0.33]	0.4 [0.86]	1.2 [0.32]	0.6 [0.81]	8.4 [0.29]
ΔC_t	1.9 [0.39]	1.9 [0.10]	2.1 [0.08]	1.5 [0.15]	12.3 [0.04]

The dummy variables are taken from Hendry and Mizon (1993). They were originally found through a residual analysis as large outliers. By including dummies for these particular observations the remaining observations appear to match a normal reference distribution, and the model passes standard specification tests including recursive tests. At the same time, these dummies are interpretable as interventions and are in this respect related to the historical narrative approach of Romer and Romer (2010).

The initial specification is a second-order vector autoregressive model including the two dummy variables $w_{out,t}$ and $w_{oil,t}$. The estimated model is the joint model reported in equilibrium-correction form in the first two columns of Table 1.4.1.

Specification tests are reported in Table 1.4.3. The residual specification tests include a cumulant-based test χ^2_{norm} for normality, a test F_{ar} for autoregressive temporal dependence (Godfrey, 1978), a test $\mathsf{F}_{\mathsf{arch}}$ for autoregressive conditional heteroskedasticity (Engle, 1982), a test $\mathsf{F}_{\mathsf{het}}$ for heteroskedasticity (White, 1980), and a test $\max Chow$ based on the maximum of recursive 1-step-ahead Chow (1960) forecast test statistics. We will benefit from this recursive test in §1.4.6. The above references only consider static or stationary models, but the specification tests also apply for non-stationary autoregressions, see Kilian and Demiroglu (2000) for χ^2_{norm}, Nielsen (2006) for F_{ar}, and Nielsen and Whitby (2015) for $\max Chow$. We see that the specification for the velocity equation is very good, while the specification for the cost equation is less good, but tolerable. The two Chow tests take their maximum values in 1971:1 and in 1976:4. These dates correspond to the decimalisation of the Pound and the debt intervention by the International Monetary Fund. Overall, these tests indicate that we cannot reject the model and that the innovations are independent, identically normal.

The dummy variables play a dual role in the subsequent analysis. First, we need the dummy variables to achieve a reasonable specification of the econometric model. Without these the residuals appear too irregular and we cannot perform valid inference. The chosen statistical model is based on the normal distribution and the observations captured by the dummy variables are outliers relative to this reference distribution. Second, the dummy variables help us to distinguish between large and small shocks. The large shocks occur infrequently and are often interpretable as catalysts.

The above specification analysis indicates that the largest shocks after the oil crises and output expansions are the decimalisation of the Pound in 1971:1 and the turmoil around the IMF intervention in 1976:4. In terms of fit, the results in Table 1.4.3 do not suggest that it is necessary to include dummies to represent these events. This could be followed up with a sensitivity analysis for the inference we draw about about the oil shocks and the output expansion. For instance, does it make a difference to include a dummy for the decimalisation? At the same time we could include dummies for the decimalisation and the IMF intervention to explore the transmission of those events. In other words, if we are concerned with a particular macroeconomic intervention we can to some extent search for similar interventions in the past and explore their transmission.

1.4.2 Causal transmission in UK money demand data

We now explore causal transmission. Table 1.4.1 reports the unrestricted reduced-form model is columns 1 and 2. This is a model for v_t, C_t given dummies and the past.

The effect of the oil price shocks can be explored by conditioning v_t on C_t and following Example 1.2.2. The conditional equation for v_t given C_t and the marginal equation for C_t are reported in columns 3 and 2, respectively, of Table 1.4.1. The coefficient for $w_{oil,t}$ is insignificant in the conditional equation, but significant in the marginal equation. Theorem 1.2.1 shows there exists a unique Markov structure, so that v_t and $w_{oil,t}$ are conditionally independent given C_t. Further, the coefficient for ΔC_t is significant in the conditional equation. Theorem 1.2.2 shows the Markov structure is non-trivial so $w_{oil,t}$—C_t—v_t. Lemma 1.2.3 then shows that the transmission between $w_{oil,t}$ and v_t is non-trivial. Correspondingly, the coefficient for $w_{oil,t}$ is significant in the marginal v_t equation. From an economic perspective, it seems reasonable to interpret the oil shocks as catalysts so that $w_{oil,t} \to C_t \to v_t$. The interpretation is that large oil price shocks move prices and in turn the velocity.

To illustrate the uniqueness result we now consider the conditional equation for C_t given v_t in column 4 of Table 1.4.1. Here, $w_{oil,t}$ is significant, so we cannot have a Markov structure from $w_{oil,t}$ through v_t to C_t. This is in line with Theorem 1.2.1.

Turning to the output shock we condition C_t on v_t. The conditional equation for C_t given v_t and the marginal equation for v_t are reported in columns 4 and 1, respectively. We follow Example 1.2.2 again. The output dummy $w_{out,t}$ is significant in the marginal equation and insignificant in the conditional equation. Moreover, velocity, Δv_t, is significant in the conditional equation. Theorems 1.2.1, 1.2.2 imply a non-trivial Markov structure $w_{out,t}$—v_t—C_t. Lemma 1.2.3 shows that the transmission is non-trivial. Interpreting $w_{out,t}$ as a catalyst, we then have $w_{out,t} \to v_t \to C_t$. Economically, large fiscal expansions may impact the velocity of money without having an impact on inflation straight away. The conclusion is, however, less clear than the causal transmission of the oil shocks. Indeed, in line with the discussion in §1.2.2, we check if the fiscal shock $w_{out,t}$ actually has a non-negligible effect on the cost of holding money. The coefficient in the C_t equation has a t-statistic of 1.3, which at best shows marginal significance. Thus, we may very well have $w_{out,t} \to v_t \to C_t$, but evidence for this transmission is weaker than the evidence for the transmission of the oil shocks.

1.4.3 Imposing multiple catalysts

The two causal transmissions $w_{oil,t} \to C_t \to v_t$ and $w_{out,t} \to v_t \to C_t$ can be imposed individually. These are the hypotheses H_1, H_2 of (1.3.17), (1.3.18). Imposing both gives $w_1 \longrightarrow z \rightleftarrows y \dashleftarrow w_2$ as described in §1.3.5. This is a system of seemingly unrelated regressions. When maximising the likelihood, we chose to parametrise it in terms of $\sigma_{yy\cdot z}$, $\sigma_{zz\cdot y}$, ρ and derive standard errors for γ_{yz} and γ_{zy} using the δ-method.

The restricted model is reported in columns 5 and 6 of Table 1.4.1 in the structural form derived from §1.3.5. The likelihood ratio statistic for the two restrictions is $2(559.31-558.74) = 1.14$, which is not significant when compared to a χ^2_2 distribution. The structural estimates largely match those of the conditional models in Table 1.4.1. Writing the model in structural form, it becomes very clear that the dummies $w_{out,t}, w_{oil,t}$ affect distinct linear combinations of the endogenous variables. The first structural equation is interpretable as the monetary quantity relation, showing how money demand reacts to output shocks, while the second structural equation is interpretable as a cost-push relation showing how money demand is driven by price shocks.

1.4.4 Cointegration

The velocity and cost of holding money variables are non-stationary and should possibly be subjected to a cointegration analysis. This is compatible with causal transmission.

Following the maximum likelihood setup of Johansen (1995), the cointegration model with rank one is given by the equilibrium-correction model

$$\begin{pmatrix} \Delta v_t \\ \Delta C_t \end{pmatrix} = \begin{pmatrix} \alpha_v \\ \alpha_C \end{pmatrix} (\beta_v v_{t-1} + \beta_C C_{t-1} + \beta_1) \\ + \begin{pmatrix} \gamma_{vv} & \gamma_{vC} \\ \gamma_{Cv} & \gamma_{CC} \end{pmatrix} \begin{pmatrix} \Delta v_{t-1} \\ \Delta C_{t-1} \end{pmatrix} + \begin{pmatrix} \gamma_{v1} & \gamma_{v2} \\ \gamma_{C1} & \gamma_{C2} \end{pmatrix} \begin{pmatrix} w_{out,t} \\ w_{oil,t} \end{pmatrix} + \begin{pmatrix} \epsilon_{v,t} \\ \epsilon_{C,t} \end{pmatrix} \quad (1.4.1)$$

The model with multiple causal transmissions and cointegration imposed has a likelihood 555.44. For present purposes, we merely consider the likelihood ratio test for the cointegration restriction within the model with multiple causal transmission imposed. The test statistic is $2(558.74 - 555.44) = 6.60$, which should be compared to a 95% critical value of 9.1, see Johansen (1995, Table 15.2).

With a unit cointegration rank, the coefficients to v_{t-1}, C_{t-1} are proportional across the equations. This results in the cointegrating relation $v_{t-1} = 6.239C_{t-1}$, which is interpretable as long-run money demand. The adjustment coefficient in the conditional equation for v_t given C_t is a modest 9.5% per quarter, whereas the adjustment in the marginal equation for C_t is insignificant. We note that in a model without multiple causal transmissions imposed, the constraint $\alpha_C = 0$ would be a hypothesis of weak exogeneity Johansen (1995, §8), but the weak exogeneity is broken when imposing the cross-equation restrictions implied by causal transmission.

1.4.5 Impulse responses

We now carry out an impulse response analysis with respect to the economic shocks represented by $w_{out,t}$ and $w_{oil,t}$. We reconstruct empirical scenarios and compare our results to the data. This offers a distinct advantage over impulse responses created by placing identifying restrictions on the covariance matrix. Figures 1.4.2a, 1.4.2b explore the period around the first oil crisis, where the fiscal expansions in 1972:4–73:1

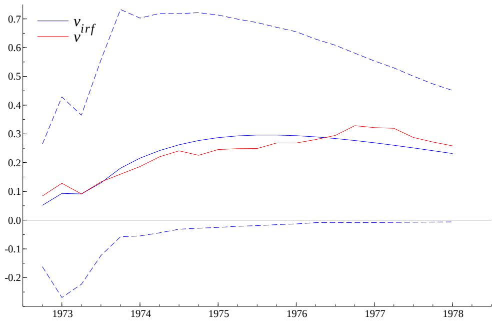

(A) Joint impulse response for v.

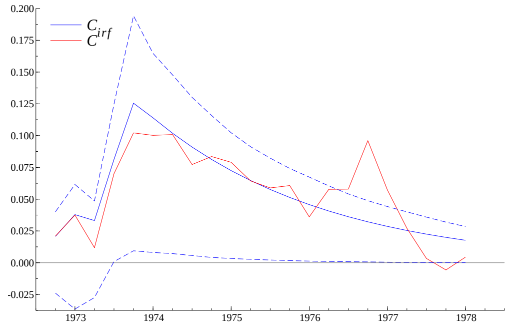

(B) Joint impulse response for C.

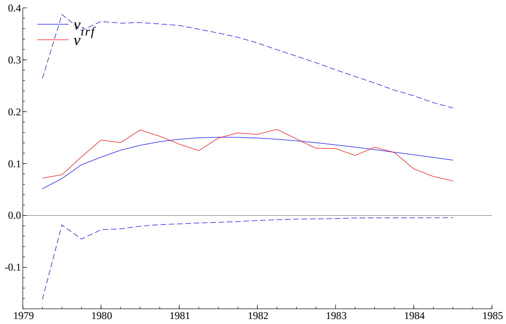

(C) Joint impulse response for v.

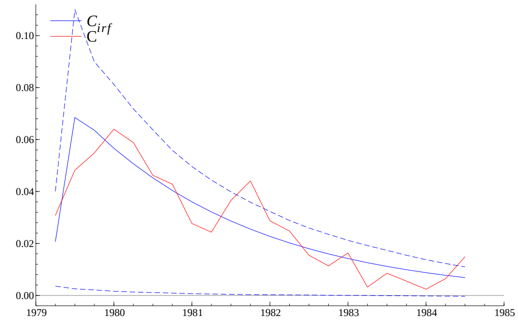

(D) Joint impulse response for C.

FIGURE 1.4.2: Impulse responses matched with the data for the early and late 1970s output and oil shocks. Dashed lines are simulated 90% confidence bands.

are followed by the oil shock in 1973:3–4. Likewise, Figures 1.4.2c, 1.4.2d explore the period around the second oil crisis, where the fiscal expansions in 1979:2 are followed by the oil shock in 1979:3. In both cases, we provide joint impulse responses and compare these to real data over a five-year horizon in Figure 1.4.2. All joint impulses perform remarkably well compared to the scenario under consideration. What is more, the impulse response functions do not decline in performance across each scenario, indicating a temporal stability in causal transmission. This is addressed further in §1.4.6.

1.4.6 Lucas critique

Major shocks like the oil crises and fiscal expansions change the policy environment and, in turn, may influence the behaviour of individual agents. It has long been a concern whether this results in instability for the parameters of an economic model, rendering it useless for analysing the effect of implementing the policy. This is known as the Lucas (1976) critique, although the concern goes back to Frisch and Haavelmo.

Engle et al. (1983) argue that super exogeneity is a sufficient condition for valid policy analysis in the context of a well-specified econometric model that passes recursive specification tests. As remarked in §1.3.4, super exogeneity is essentially a combination of causal transmission and weak exogeneity. We illustrate the use of causal transmission in policy analysis by performing a recursive analysis of the money data.

Previously, $w_{oil,t}$ was constructed as the sum of impulse indicators across the two oil crises. Now, we construct dummies $w_{oil1,t}, w_{oil2,t}$ for 1973:3–4 and 1979:3, respectively, so that $w_{oil,t} = w_{oil1,t} + w_{oil2,t}$. We re-estimate the equations for $(v_t|C_t)$, (C_t) reported in Table 1.4.1 over subsamples 1963:4–1977:2 and 1963:4-1989:3 using the split oil dummy. The results are reported in Table 1.4.5. It is clear that the transmission of the first catalyst $w_{oil1} \to C \to v$ does not differ in a statistically significant way from the transmission of the second catalyst $w_{oil2} \to C \to v$. Deconstructing the catalyst w_{out} provides similar evidence for the stability of the causal transmission of the output shocks. The search for causal transmission in well-specified models therefore seems compatible with the Lucas critique. This does, of course, go hand in hand with the fact that the model in Table 1.4.1 passes recursive specification tests.

1.5 Concluding remarks

Causal transmission has been introduced to capture the idea that large economic shocks may transmit gradually through the macroeconomy.

There are three ingredients to the definition of causal transmission of catalyst w through z to y. First, we need a non-trivial Markov structure w—z—y, that is the

TABLE 1.4.5: Estimation for UK M1 data based on a sub-sample and the full sample.

	1963:4-1977:2		**1963:4-1989:2**	
	$\Delta v_t \vert \Delta C_t$	ΔC_t	$\Delta v_t \vert \Delta C_t$	ΔC_t
ΔC_t	0.542 (0.189)	–	0.605 (0.105)	–
Δv_{t-1}	-0.359 (0.123)	-0.060 (0.094)	-0.311 (0.087)	-0.031 (0.084)
ΔC_{t-1}	-0.004 (0.185)	-0.092 (0.142)	0.057 (0.101)	0.036 (0.099)
v_{t-1}	-0.097 (0.032)	0.007 (0.025)	-0.094 (0.013)	-0.003 (0.012)
C_{t-1}	0.626 (0.145)	-0.146 (0.111)	0.574 (0.063)	-0.084 (0.061)
1	-0.012 (0.008)	0.012 (0.006)	-0.009 (0.005)	0.009 (0.005)
$w_{out,t}$	0.042 (0.015)	0.016 (0.011)	0.044 (0.010)	0.013 (0.010)
$w_{oil1,t}$	0.001 (0.018)	0.058 (0.011)	0.000 (0.014)	0.055 (0.012)
$w_{oil2,t}$	–	–	-0.002 (0.018)	0.043 (0.017)

Markov structure $\mathsf{f}(y, z|w) = \mathsf{f}(y|z)\mathsf{f}(z|w)$ needs to be non-trivial in the sense that y, z are dependent and z, w are dependent. Secondly, we need a non-trivial transmission between w, y, that is w, y are dependent. Thirdly, we need a causal assumption for the catalyst w. When these conditions are satisfied, we write $w \to z \to y$. We have shown how this definition can be extended to the transmission of two unrelated catalysts.

Causal transmission is defined for general densities, and it does not require normality. The first two conditions to the definition of causal transmission are testable using observational data. In standard models, the first condition of a non-trivial Markov structure implies the second condition of a non-trivial transmission. These standard models include normal models and mixtures of normal and logit/probit models.

Causal transmissions also require a catalyst. As in instrumental variable analysis the catalyst can be found as a natural experiment prior to the empirical analysis or it may be discoverable from the empirical analysis of observational data. In this way, Hendry and Santos (2010) give an algorithm for discovering super-exogeneity. This algorithm generalises the robustified least-squares approach used by Hendry and Mizon (1993) in their UK money analysis. A theory for analysing such algorithms is gradually emerging. Indeed, a statistical theory for robustified least-squares is presented in Johansen and Nielsen (2015). The causal transmission relies on an economic interpretation of the catalyst. In this way it relates to the narrative approach of Romer and Romer (2010). Catalysts have to be statistically significant to be discoverable. Likewise, the evidence of causal transmission is stronger if it is found in several situations across time or countries. In these ways, the empirical analysis of causal transmission is consistent with the criteria for causation in Hill (1965).

The ideas presented here are related to graphical models. Graphical model theory seeks to represent dependence structures of statistical models in terms of graphs. There are many quirks to statistical models and distribution theory, which in turn results in a variety of types of graphs; most are directed and some require normality, see Lauritzen (1996) and Wermuth and Sadeghi (2012) for details. Directed graphs and causality are closely related, see Cox and Wermuth (2004). However, the causal transmission idea presented here deviates from that literature in important ways. First, we are concerned with the discovery of transmission, so we postpone the use of directed graphs as far as possible. Second, we establish results for general densities. Third, our causal transmission requires a variety of both (conditional) independence and dependence properties. As a result, we only make limited use of graphs as illustrations and carefully define the properties of those graphs that we use.

The analysis is inspired by Bårdsen et al. (2012). They present 3-year ahead quarterly forecasts from March 2007 generated from their macro-econometric model for Norway. The forecasts track actual unit labour cost, inflation, and unemployment rather well, but fail spectacularly on the short-term interest rate. In 2008, policy-makers in Norway and abroad changed the policy rate dramatically in response to the financial crisis, creating a large shift of the short-term interest rate. It appears that this had the causal impact of offsetting potential big shifts in the labour market in such a way that the macro-econometric model produces good forecasts of unit labour cost, inflation, and unemployment. It is plausible that the effects seen in the forecasts of the Norwegian macro-econometric models of Bårdsen et al. (2012) could be described as a combination of a major financial shock and a subsequent policy reaction calibrated to offset the financial shock in the labour market. In future work, we will seek to generalise the ideas presented here to facilitate such an analysis to guide economic policy.

Appendix

1.A Proofs

The result in Theorem 1.2.1 hinges on the equivalence in the following lemma, of which the left to right implication is related to Lauritzen (1996, Proposition 2.1).

Lemma 1.A.1 *Suppose* $\mathsf{f}(y,z|w)$ *has support on a product space and that it is positive on this support. Then, for all* y,z

$$\left.\begin{array}{rcl} \mathsf{f}(y|z,w) & = & \mathsf{f}(y|z) \\ \mathsf{f}(z|y,w) & = & \mathsf{f}(z|y) \end{array}\right\} \Leftrightarrow \mathsf{f}(y,z|w){=}\mathsf{f}(y,z). \tag{1.A.1}$$

Proof of Lemma 1.A.1. Since the density is positive on a product space, then the marginal densities are also positive.

$\Rightarrow$: By the definition of conditional densities, the first statement on the left hand side of (1.A.1), and the definition of conditional densities

$$\mathsf{f}(y,z|w) = \mathsf{f}(y|z,w)\mathsf{f}(z|w) = \mathsf{f}(y|z)\mathsf{f}(z|w) = \mathsf{f}(y,z)\mathsf{f}(z|w)/\mathsf{f}(z).$$

Swap y,z and use the second statement on the left hand side of (1.A.1) to get

$$\mathsf{f}(y,z|w) = \mathsf{f}(y,z)\mathsf{f}(y|w)/\mathsf{f}(y). \tag{1.A.2}$$

Equating the two expressions we get

$$\mathsf{f}(y|w) = \mathsf{f}(y)\mathsf{f}(z|w)/\mathsf{f}(z).$$

Fixing z this shows that $\mathsf{f}(y|w) = c\mathsf{f}(y)$ for some constant $c = \mathsf{f}(z|w)/\mathsf{f}(z)$, which must be one so that the densities $\mathsf{f}(y|w)$, $\mathsf{f}(y)$ integrate to unity. Insert this in (1.A.2) to get the desired right hand side of (1.A.1).

$\Leftarrow$: We prove the first left right hand side statement. Note that

$$\mathsf{f}(y|z, w) = \mathsf{f}(y, z|w)/\mathsf{f}(z|w).$$

The right hand side of (1.A.1) shows $\mathsf{f}(y, z|w) = \mathsf{f}(y, z)$. Integrate over y to get $\mathsf{f}(z|w) = \mathsf{f}(z)$. Insert these statements above to get

$$\mathsf{f}(y|z, w) = \mathsf{f}(y, z)/\mathsf{f}(z) = \mathsf{f}(y|z),$$

as desired. The other left hand side statement is proved in a similar fashion. ■

Proof of Theorem 1.2.1. Condition (1.2.10) shows $\mathsf{f}(y|z, w) = \mathsf{f}(y|z)$. Thus

$$\mathsf{f}(y, z|w) = \mathsf{f}(y|z, w)\mathsf{f}(z|w) = \mathsf{f}(y|z)\mathsf{f}(z|w). \tag{1.A.3}$$

First, rearrange to get

$$\mathsf{f}(y|z, w) = \mathsf{f}(y, z|w)/\mathsf{f}(z|w) = \mathsf{f}(y|z). \tag{1.A.4}$$

Then, note that Condition (1.2.11) has $\mathsf{f}(z|w) \neq \mathsf{f}(z)$. Insert this in (1.A.3) to get

$$\mathsf{f}(y, z|w) \neq \mathsf{f}(y|z)\mathsf{f}(z) = \mathsf{f}(y, z). \tag{1.A.5}$$

Now apply Lemma 1.A.1. The first statement on the left of (1.A.1) holds through (1.A.4), while the right hand side fails through (1.A.5). Thus, the second statement on the left hand side of (1.A.1) fails as desired. ■

Proof of Theorem 1.2.2. Combine Theorem 1.2.1 and Definition 1.2.1. ■

Proof of Lemma 1.2.1. Consider the compound distribution integral (1.2.17).

If $\mathsf{f}(z|w) = \mathsf{f}(z)$, then $\mathsf{f}(y|w) = \int \mathsf{f}(y|z)\mathsf{f}(z)dz = \int \mathsf{f}(y,z)dz = \mathsf{f}(y)$.

If $\mathsf{f}(y|z) = \mathsf{f}(y)$, then $\mathsf{f}(y|w) = \int \mathsf{f}(y)\mathsf{f}(z|w)dz = \mathsf{f}(y)\int \mathsf{f}(z|w)dz = \mathsf{f}(y)$. ■

Proof of Lemma 1.2.2. Proof by contradiction. When z is binary the compound integral (1.2.17) reduces to $\mathsf{f}(y|w) = \mathsf{f}(y|z=0)p_w + \mathsf{f}(y|z=1)(1-p_w)$. If $y \perp\!\!\!\perp w$ then $D_{w,w^\dagger} = \mathsf{f}(y|w) - \mathsf{f}(y|w^\dagger) = 0$ for all $y, w, w^\dagger$. Combine to get $D_{w,w^\dagger} = \{\mathsf{f}(y|z=0) - \mathsf{f}(y|z=1)\}(p_w - p_{w^\dagger}) = 0$. With the independence structure, y and w vary in a product space. Thus, $p_w = p_{w^\dagger}$ so $z \perp\!\!\!\perp w$ or $\mathsf{f}(y|z=0) = \mathsf{f}(y|z=1)$ so $y \perp\!\!\!\perp z$. ■

Proof of Lemma 1.2.3. Referring to equations (1.2.3), (1.2.4) the Markov assumption implies $0 = \gamma_{yw\cdot z} = \gamma_{yw} - \gamma_{yz}\gamma_{zw}$. Thus, if $\gamma_{yz} \neq 0$ and $\gamma_{zw} \neq 0$ then $\gamma_{yw} \neq 0$. ■

Proof of Lemma 1.2.4. We show that $\mathsf{f}(y=0|w)$ is strictly decreasing in w if and only if $\gamma_{yz} \neq 0$ and $\gamma_{zw} \neq 0$. The partial derivatives of the normal density $\mathsf{f}(z|w)$ and the logit/probit probabilities $\mathsf{f}(y=0|w)$ satisfy, using D as partial derivative symbol,

$$\begin{array}{rrcl} & D_w\mathsf{f}(z|w) &=& (-\gamma_{zw}/\sigma_z^2)D_z\mathsf{f}(z|w), \\ \text{logit:} & D_z\mathsf{f}(y=0|z) &=& -\gamma_{yz}\mathsf{f}(y=0|z)\left\{1-\mathsf{f}(y=0|z)\right\}, \\ \text{probit:} & D_z\mathsf{f}(y=0|z) &=& -\gamma_{yz}\phi(\gamma_{yz}z), \end{array}$$

which are bounded. We can then differentiate the probability $\mathsf{f}(y=0|w)$ and use integration by parts to get

$$\begin{aligned} D_w\mathsf{f}(y=0|w) &= \int_{-\infty}^{\infty} \mathsf{f}(y=0|z)D_w\mathsf{f}(z|w)dz \\ &= \frac{-\gamma_{zw}}{\sigma_z^2}\int_{-\infty}^{\infty} \mathsf{f}(y=0|z)D_z\mathsf{f}(z|w)dz \\ &= \frac{\gamma_{zw}}{\sigma_z^2}\int_{-\infty}^{\infty} \mathsf{f}(z|w)D_z\mathsf{f}(y=0|z)dz, \end{aligned}$$

which is zero if and only if $\gamma_{yz}\gamma_{zw} = 0$. ■

Part II

CHAPTER 2

Financial Fragility in a Regime-Switching Model

Chapter Abstract

We identify and characterise periods of financial fragility over the period 1973–2016 within a simple regime-switching framework. Fragility is driven by high default risk and low profitability of intermediaries, as defined in Tsomocos (2003). Recessions occur primarily when fragility is high, but high fragility can occur without leading to recession. Credit growth is lower under high fragility, but it is not as significant a driver of regime changes when financial fragility is directly taken into account. Shocks have regime-dependent effects that are attributable primarily to switching dynamics as opposed to shock size. Finally, shocks are amplified through financial fragility, suggesting the presence of a debt-deflation channel.

2.1 Introduction

The experience and aftermath of the Global Financial Crisis have convinced economists and policymakers that the relationship between the macroeconomy and fragility in the financial sector cannot be appropriately characterised by a linear statistical model. There are many definitions of fragility, however, and many forms of nonlinearity to choose from. We draw on general equilibrium theory to motivate our empirical framework and find that defining fragility in terms of default risk and profitability in a regime-switching model with production and credit growth uncovers significant nonlinear effects that are consistent with the historical narrative.

Fluctuations in the financial sector have long been recognised as a source of significant perturbations to the business cycle (Bernanke and Gertler, 1989; Kiyotaki and Moore, 1997). Early models took different approaches to account for the impact of financial stress on the macroeconomy, alternatively by introducing financial frictions in DSGE models (Carlstrom and Fuerst, 1997; Bernanke et al., 1999) and by introducing endogenous default in GEI models (Tsomocos, 2003; Dubey et al., 2005). A second generation of models has continued to improve on these early works, with frictions and default in financial intermediation becoming mainstay features (for a recent survey, see Goodhart et al., 2016).

Empirical work on the effect of financial fluctuations on the macroeconomy has been sparse by comparison and largely disconnected from economic theory. This is in part due to the difficulty in measuring financial stress and in part due to the difficulty in uncovering the effects of financial frictions in data (Hubbard, 1998). Instead, focus has been placed in understanding the role that credit cycles play in predicting financial crises (Reinhart and Kaminsky, 1999; Reinhart and Rogoff, 2009; Schularick and Taylor, 2012), primarily by looking at cross-country evidence over time. A rapidly expanding line of work, however, has sought to improve the measurement of financial stress, both on aggregate and among individual intermediaries (for a survey, see Bisias et al., 2012). This has, in turn, facilitated the study of financial stress within statistical models that introduce nonlinearity (Hubrich and Tetlow, 2015; Hartmann et al., 2015). Few models have sought to motivate statistical modelling based on an economic theory, however, so that the choice of nonlinearity and the modelling of financial fragility is somewhat *ad hoc*.

We take the view that financial fragility develops endogenously as an equilibrium phenomenon due to the regular functioning of an economy. While at times this may manifest as a financial crisis or as a cyclical component of the data, our aim is to capture all regimes where the state of financial intermediation carries the potential for damaging macroeconomic effects. Ultimately, this becomes an informal test of Hyman Minsky's first hypothesis of financial instability: that the economy undergoes periods of financing stability and instability (Minsky, 1977, 1992). To motivate our reduced-form empirical analysis, we draw on structural economic models that incorporate endogenous default as an equilibrium result in different states (Tsomocos, 2003; Goodhart et al., 2006). We show that financial fragility will arise when the prevailing environment among intermediaries is characterised by high volatility and default risk, while simultaneously lacking sufficient profitability as a buffer to weather adverse conditions. Since equilibrium depends on the realisation of discrete states, we opt for a multivariate regime-switching framework that considers the effects of default, profitability, and credit jointly on economic growth.

Specifically, we ask the following questions. First, can financial fragility be defined in terms of a nonlinear association between the intermediation sector and the real economy? Second, how do the effects of shocks vary when intermediaries are highly fragile? Finally, what is the relative importance of financial fragility as a transmission channel for shocks to economic growth? The answers to these questions should inform the construction of appropriate measures of financial stress, as well as aid policymakers when deciding where and how to intervene.

We employ a Markov-switching vector autoregressive (MS-VAR) model for the United States over the period 1973–2016. To assess real economic effects, we include industrial production and credit extension in our set of macroeconomic variables. The health of the financial intermediation sector is assessed by proxying for default risk and profitability. We proxy for default risk by using the interbank borrowing (TED) spread and for profitability by using the market valuation of the largest US banks. Together, these are indicative of the fragility of the financial sector.

The main results are as follows. First, we find evidence for nonlinearity that is driven by high default risk and low profitability of intermediaries; this is consistent with the definition of fragility in Tsomocos (2003). Second, recessions occur primarily when fragility is high, but high fragility can occur without leading to recession. Third, while credit growth is lower under high fragility, it is not as significant a driver of regime changes when financial fragility is directly taken into account; credit can create financial imbalances over the long term, but is not an adequate measure of the health of financial intermediaries at shorter horizons. This result is reinforced by reproducing the analysis in Schularick and Taylor (2012) for the United States and finding that credit growth alone is not a useful predictor of financial crises. Fourth, shocks have regime-dependent effects that are attributable primarily to switching dynamics; this finding is consistent with Hubrich and Tetlow (2015) and antithetical to Stock and Watson (2012). Fifth, shocks are amplified in the high-fragility regime through default risk and profitability of intermediaries, suggesting the presence of a debt-deflation channel (Fisher, 1933). While we intentionally do not model policy variables and are primarily interested in characterising association rather than causation, this final result lends some weight to unconventional monetary policy that intervenes by providing credit when the private sector is severely constrained (Gertler and Karadi, 2011).

The paper is organised as follows. The following section briefly examines the relevant empirical literature on financial fragility. In §3 we describe our data and introduce the proxies we will use for default risk and profitability. We present the econometric framework within which we operate in §4 and discuss some empirical issues. §5 presents the empirical results, wherein we identify and characterise the different regimes estimated by our model. Once we reconcile these with the historical narrative, we conduct a structural analysis of shocks in §6 by using a recursive identification assumption. Further information on data, modelling, and our analysis is relegated to the appendix.

2.2 Related literature

Our efforts to model the macro-financial relationship touches upon several strands of literature. The concept of financial fragility that originates at least as early as Minsky (1977) can be formalised in economic models with endogenous default. The episodic nature of fragility necessitates nonlinear modelling that also considers the role of financial cycles and systemic risk. Statistical nonlinear modelling of macro-financial interactions has been particularly active since the Global Financial Crisis, and we briefly survey other approaches to this problem. We then discuss how our approach differs from the existing literature.

Financial fragility. The manifestation of financial stress is alternatively treated as a binary crisis variable resulting from cyclical fluctuations in credit (Reinhart and Rogoff, 2009; Schularick and Taylor, 2012; Jordá et al., 2014), or as a composite index of rates, volatilities, and spreads (Hakkio and Keeton, 2009; Kremer et al., 2012). Both of these interpretations reflect necessary conditions—namely, profitability and default risk— but neither is individually sufficient to characterise financial fragility.

We consider financial fragility as a continuous measure of the health of financial intermediaries and assume the following definition from Tsomocos (2003):

> *"An economic regime is financially fragile when substantial default of a 'number' of households and banks...without necessarily becoming bankrupt occurs and the aggregate profitability of the banking sector decreases significantly."*

To gauge fragility, we must thus examine both default risk and profitability among intermediaries. Default risk encompasses shocks and other events that may render bank balance sheets untenable, while profitability assess to what extent heightened risk can be managed with capital buffers. Regimes of high fragility will occur when the probability of default among intermediaries is relatively high, while simultaneously profitability is relatively low without, however, necessarily implying bankrupcty.

High probability of default—or an increase in any alternative measure of systemic risk—does not necessarily imply high financial fragility. For instance, when paired with high profitability, this can indicate strategically risky behaviour that may contribute to accelerated credit buildup, as described in Minsky (1992). Similarly, low probability of default with low profitability may indicate strategically risk-averse behaviour. This definition of financial fragility is consistent with the general notion of financial instability described in Kindleberger and Aliber (2011), except that we explicitly account for default and bank profitability. Financial fragility thus defined captures at once the inherent buildup of default risk stemming from the optimizing behaviour in the private sector, while also accounting for the capacity of banks to absorb shocks without curbing their normal activities.

The closest test of this definition of fragility within a nonlinear empirical framework occurs in Aspachs et al. (2007), which we discuss subsequently.

Approaches to nonlinearity. There have been several recent efforts to model the macro-financial relationship by incorporating nonlinearity. Among these, Hubrich and Tetlow (2015) and Davig and Hakkio (2010) employ a Markov-switching framework with a financial stress indicator variable in order to examine the effects of financial stress on the US economy over a period beginning 1990. Whereas Davig and Hakkio (2010) examine a rudimentary, bivariate model to conclude that financial stress has larger effects in the distressed regime, the approach in Hubrich and Tetlow (2015) is much more comprehensive. By allowing for breaks in all coefficients as well as for stochastic volatility, Hubrich and Tetlow (2015) conclude that both outsized shocks and agent behaviour contribute to the nonlinear relationship between the real and the financial economy. Nason and Tallman (2015) also employ a Markov-switching framework over a much longer history for the US, using annual data since 1890 to examine the effects of credit shocks but only allow for breaks in the intercept and stochastic volatility. Major economic crises and wars are picked up in the distressed regime, while the authors link economic and financial fluctuations to the identified shocks in credit demand and supply.

Elsewhere in the literature, nonlinearity is modelled through a thresholding approach. Balke (2000) examines the nonlinear effect of credit in the United States over 1960–1997 to show that credit shocks have a larger effect on output when credit conditions are tight. In a more recent study for the United States, Aikman et al. (2016) use a threshold VAR over 1975–2014 to show that financial imbalances in credit, risk appetite, and short-term funding can each lead to an agitated state where shocks negatively impact economic growth. Aspachs et al. (2007), examine the relationship between financial fragility and the real economy using a panel threshold VAR model. For panel data for various countries (excluding the United States) over the period 1990–2004, they find that fluctuations in probability of default and bank profitability have significant, nonlinear effects on welfare. These sorts of models, however, are capable of only one variable displaying a thresholding effect, which limits the ability to characterise the financially stressful state more generally.

Finally, more recent efforts to incorporate nonlinearities in the macro-financial relationship use a time-varying parameter approach. Rather than having discrete changes as in the Markov-switching or thresholding frameworks, a stochastic process for the evolution of the parameters can be assumed and estimated to allow for continuous changes. Most prominently, Prieto et al. (2016) employ this sort of analysis for the US to find that financial shocks bear greater explanatory power during the Great Recession than in normal periods. In particular, they find that housing price shocks were increasingly significant contributors since the early 2000s, while also explaining the slow and weak recovery following the Global Financial Crisis.

Risk measurement. The measurement of default risk is in itself an important question. The literature that has grown out of the financial crisis of 2007-09 has focused on measuring systemic risk; that is, to find a measure that can assess the instability of the financial intermediation sector in aggregate. This has been carried out in a number of ways, both within a macro-financial context, as well as at a granular level examining individual contributions to risk.

A relatively common approach in the macro-financial sphere has been to construct financial stability indicators as (time-varying) weighted averages of interest rates and spreads (Hubrich and Tetlow, 2015; Hartmann et al., 2015; Kremer et al., 2012) in order to pick up any shocks that may affect markets in real time. Alternatively, Gorton and Metrick (2012b) use only the 3-month LIBOR spread over the Overnight Index Swap (OIS) rate to measure risk specifically among financial intermediaries.

Other measures that examine cross-sectional dependence (Huang et al., 2011; Acharya et al., 2012; Adrian and Brunnermeier, 2016) and networks (Billio et al., 2012; Diebold and Yilmaz, 2014) become relevant when we disaggregate the financial system into sectors and individual financial institutions. Recognizing that such linear measures may not pick up the nonlinear dependence during periods of distress, Goodhart and Segoviano (2009) introduce a portfolio-based measure of joint default probability that takes into account stronger co-movement in the distribution tails. Various measures applying this conditional, nonlinear dependence to quantify risk among financial intermediaries are described in Goodhart and Segoviano (2009).

A separate strand of the literature examines credit-based, forward-looking measures of financial stability. Several authors follow the signalling approach of Reinhart and Kaminsky (1999) to introduce nonlinearities in measurement. Alessi and Detken (2011) find that the global M1 and private credit gaps are the best indicators of costly boom-bust cycles. Borio and Drehmann (2009) consider the strength of early-warning indicators, where the property price gap, the equity price gap, and the credit gap exceed certain thresholds that are optimised with respect to a true positive rate, finding that indicators jointly exceeded their thresholds several years before the subprime mortgage crisis.

Credit cycles. The empirical literature elsewhere has taken a long-term approach in order to examine the rapid increase in credit preceding financial crises. There is a consensus in the empirical literature that credit plays a significant role in the buildup of risk over the long run and has the capacity to exacerbate financial shocks. Contributions to the literature vary from historical, cross-country event studies (Bordo and Haubrich, 2010; Hume and Sentance, 2009; Mendoza and Terrones, 2008) that find credit booms to be associated with financial crises, to historical analysis based on new panel datasets (Reinhart and Rogoff, 2009; Schularick and Taylor, 2012). In a similar cross-country approach, Claessens et al. (2012) show that cycles in credit and house prices are complementary, and major recessions coincide with simultaneous downturns in the financial cycle.

In order to capture abrupt changes in the macro-financial relationship, we opt for the Markov-switching framework where parameters and volatility are allowed to switch across regimes. The framework affords us sufficient flexibility to assess both outsized shocks and transmission mechanisms. Since we are interested specifically in capturing default risk within the financial intermediation sector in real time, we use a single proxy for default risk, while also proxying for profitability. This avoids issues related to weighting a basket of rates and spreads, while allowing for a time-varying relationship with the macroeconomic variables in our model. Whereas the ideal measure for default probabilities would rely on CDS-based measures, such as in Goodhart and Segoviano (2009), the lack of data availability prior to 2004 renders this option unattractive since we are interested in exploiting information from as many episodes as possible in the historical narrative. We incorporate credit into our model as a bridge between financial stability and economic output.

Despite evidence for the role of credit in the financial cycle elsewhere in the empirical literature, the experience in the United States is not as clear-cut. For instance, upon closer inspection of the results in Schularick and Taylor (2012), we can find that real credit growth is not sufficient for explaining the occurrence of financial crises in the United States, suggesting that their results are driven by the other countries in the panel; details of this analysis are provided in Appendix 2.D. On the one hand, this may be driven by the definition of crises taken from Reinhart and Rogoff (2009)—financial crises are assumed to have occurred only in 1984 and 2007. We intend to examine the role of credit, however, in a more complete, multivariate setting that also considers the smooth functioning of financial intermediaries.

2.3 Econometric methodology

When deviating from the linear framework, there is typically no obvious form of nonlinearity that we should consider, and there are many ways that we could incorporate nonlinearities into our modelling efforts. Our employment of a Markov-switching setup with changing dynamics and volatility is, however, intuitively appealing both from an empirical and a theoretical perspective. Parameters are allowed to switch discretely between states, rather than drifting, while transition between states follows an exogenous Markov process. From an empirical perspective, this means that we can capture the rapid amplifications that characterise episodes of financial fragility. Discrete switching is also an appropriate assumption in the context of theoretical models that consider financial fragility to be a result of endogenously-determined default in particular nature states. The fact that transition probabilities follow an exogenous Markov process means that we have chosen to concern ourselves with the state-contingent interaction between the real and financial economies, rather than explaining how this change comes about. This is also a widely-held assumption in the setup and calibration of general equilibrium models (Goodhart et al., 2005, 2013).

This section characterises the Markov-switching framework in a generic way, generally following Krolzig (1997) and briefly discusses estimation. One limitation is the difficulty in testing the number of regimes, an issue that stems from parameter identifiability in hidden Markov models. Our motivation is to assess the presence of a financially robust and a financially fragile regime (Minsky, 1992), which is consistent with two regimes; this choice is also supported by the theoretical literature (Tsomocos, 2003; Bhattacharya et al., 2015) and the broad empirical literature (Hubrich and Tetlow, 2015; Davig and Hakkio, 2010; Aspachs et al., 2007; Balke, 2000). The approach to regime-dependent structural analysis is based on Ehrmann et al. (2003).

2.3.1 Markov-switching models

The distinguishing feature of the MS-VAR framework is that the data-generating process of the observable time series for N endogenous variables $\mathbf{y}_t$ depends on the unobservable regime series $s_t \in \{1, 2\}$. In its most general form, an MS-VAR model is described by a mixture of normal distributions

$$\mathbf{y}_t = \begin{cases} \mu_1 + \Gamma_{11}\mathbf{y}_{t-1} + \ldots + \Gamma_{p1}\mathbf{y}_{t-p} + \epsilon_t, & \epsilon_t|s_t \sim N(0, \Sigma_1) \qquad \text{if } s_t = 1 \\ \mu_2 + \Gamma_{12}\mathbf{y}_{t-1} + \ldots + \Gamma_{p2}\mathbf{y}_{t-p} + \epsilon_t, & \epsilon_t|s_t \sim N(0, \Sigma_2) \qquad \text{if } s_t = 2 \end{cases} \tag{2.3.1}$$

where μ_i is a $N \times 1$ vector of switching intercepts, Γ_{pi} are $N \times pN$ matrices of switching autoregressive coefficients up to order p, and Σ_i are $N \times N$ covariance matrices of the $N \times 1$ residual vector ϵ_t for $i = 1, 2$. All parameters are conditional on the regime s_t.

The unobservable regime s_t is assumed to follow a discrete time and discrete state Markov process. This means that the probability of the regime j being realised in the next period conditional on the current regime i

$$\rho_{ij} = Pr(s_{t+1} = j|s_t = i), \quad \sum_{j=1}^{2} \rho_{ij} = 1 \quad \forall i, j \in \{1, 2\} \tag{2.3.2}$$

is exogenous and constant.

2.3.2 Estimation

We estimate all parameters and the Markov process by applying the Expectation-Maximisation algorithm (Hamilton, 1990). To check for convergence to a local maximum, we re-estimate using random starting values to search for a higher log-likelihood. Computations and graphics are carried out in OxMetrics (Hendry and Doornik, 2014) and MATLAB.

2.3.3 Inference

The maximum likelihood estimator for Markov-switching models is generally consistent, asymptotically normal, and efficient, provided parameter identifiability and stability are satisfied (Cappé et al., 2005; Krolzig, 1997). Broadly, this is the requirement that the parameter vectors are distinct for each state and that the Markov chain associated with the unobserved state is irreducible and aperiodic (Leroux, 1992). Teicher (1960) showed that within the family of normal distributions, joint finite mixture models over the mean and variance are identifiable.

The practical implications are the following. If the regularity conditions stated above are satisfied, then tests on joint restrictions within Markov-switching models are asymptotically valid and follow a χ^2 distribution; this result is proved in Cappé et al. (2005). On the other hand, the practical limitations of the identifiability condition are that we cannot test for the number of regimes across models from a likelihood approach (Hamilton, 2016). This occurs because under the null hypothesis, the transition probabilities act as unidentified nuisance parameters and so the information matrix is singular. Thus, the asymptotic distribution under the null will be non-standard. Some solutions to this problem are available for simple models (Hamilton, 2016), but it remains unclear how these apply to more complex MS-VAR models. We avoid this problem by drawing on previous empirical and theoretical results to fix the number of regimes to two.

In the estimated MS-VAR models where we allow both autoregressive parameters and volatility to switch across a fixed number of regimes, identifiability is generally satisfied, since these are joint finite mixtures of normals. Furthermore, stability is assessed by examining the residuals of the model as well as checking whether any of the regimes is a sink or occurs periodically. Our specification tests are thus useful in determining stability and whether statistical inference based on standard asymptotic distributions is valid. The residual tests include the Kolmogorov-Smirnov empirical distribution test for normality (Massey, 1951), tests for autocorrelation (Godfrey, 1978), and a test for autoregressive conditional heterskedasticity (Engle, 1982). Where appropriate, we employ robust standard errors (White, 1980).

2.3.4 Identification of shocks

A separate identification problem arises when moving from a reduced-form model with errors ϵ_t and covariance matrix Σ_i to a structural model with errors e_t, covariance matrix I, and contemporaneous effects A. Since the structural errors are normalised to unity and the reduced-form errors are allowed to switch across regimes, they can be related by $\epsilon_t = A_i e_t$, so that

$$\Sigma_i = E(A_i e_t e_t' A_i') = A_i E(e_t e_t') A_i' = A_i I A_i' = A_i A_i'. \tag{2.3.3}$$

Following Ehrmann et al. (2003), we assume that A_i is lower-triangular so that we can derive impulse responses conditional on the regime. The assumption on A_i implies a recursive ordering that is common in applied macroeconomics, originating in Sims (1980). This permits us to proceed with our central objective of examining the macro-financial relationship across different financing regimes.

In order for this procedure to be valid, we need to ensure that regimes are relatively persistent. This is necessary because we are conditioning on a particular regime over the entire horizon over which the impulse responses are derived. Persistence can be assessed by examining the estimated transition matrix. Additionally, since we are imposing a contemporaneous ordering, we must assess sensitivity of the impulse responses to different orderings. We follow the typical practice, where slow-moving macroeconomic variables are ordered first and fast-moving financial variables are ordered last. A final caveat is that we do not allow for different orderings across the two regimes. Testing for different structure across regimes of financial fragility, however, is within the scope of future work.

2.4 Data

We assess the occurrence of financial fragility in the United States by estimating a regime-switching model over the sample 1973(6)–2016(1) at a monthly frequency. Our primary objective is to understand the association between the real economy and financial intermediaries, and so our model includes the endogenous variables

$$\mathbf{y}_t = [\Delta IP_t, \Delta TC_t, \Delta MV_t, TED_t]',$$

where IP_t denotes log real industrial production, TC_t denotes log real total credit the private non-financial sector, MV_t denotes log real market valuation of US financial intermediaries, and TED_t denotes the TED spread. The ordering of the variables reflects the standard assumption that fast-moving financial variables respond contemporaneously to slow-moving macroeconomic variables.[1]

While the sample period is chosen based on the availability of data, it spans several financial episodes in the United States. These are summarised in Table 2.4.1 and differ substantially from other historical accounts of financial crises in the United States (Reinhart and Rogoff, 2008, 2013; Schularick and Taylor, 2012). We discuss our use of proxies for default risk and profitability of financial intermediaries, while details on data are provided in Appendix 2.A.

[1] With multiple fast-moving financial variables, it is difficult to justify the assumption that one responds to the other after a one-month lag. We examine alternative orderings of the financial variables and find that results do not differ qualitatively.

2.4.1 Timeline of financial events

The length of the time series allows us to examine episodes as early as the 1970s. Dated between 1973–1975, the Bank Capital squeeze resulted in the first major commercial bank failures between 1971 and 1975. Indeed, this period is cited as presenting "the most serious capital problem for US commercial banks in the *whole* postwar period" (Lopez-Salido and Nelson, 2010). The Less-Developed-Country (LDC) Debt crisis originated from the threat that the Mexican government would default on loans to US commercial banks, culminating in the run on and rescue of Continental Illinois in May 1984. This was also the last financial crisis to originate solely in the traditional banking sphere.

Subsequently, major financial events originated in the non-bank sphere. The Savings & Loans crisis—alternatively dated by different sources—originated in the thrift industry in the mid- to late-1980s, whereas the Global Financial crisis originated initially among money market funds (Gorton and Metrick, 2012a) and subsequently spilled over to investment banks in the mid-2000s. Other financial events are characterised by episodes occurring outside the US, such as the Asian financial crisis in 1997 and the subsequent Russian debt default, as well as the European debt crisis in the early 2010s.

Legislation deregulating the bank and non-bank industries allowed the intermediation sector to engage in a wider range of financial activities. Changes to financial regulation are therefore also included as significant financial events to provide historical context.

2.4.2 Proxies

We do not have a perfect measure of the probability of default of financial intermediaries with a sufficiently long history to examine the major financial events discussed

Table 2.4.1: Notable financial events for the United States.

Event	Dates	Source
Bank capital squeeze	1973–1975	Lopez-Salido and Nelson (2010)
DIDMCA	March 1980	
Garn-St. Germain	September 1982	
LDC debt threat	1982–1984	Lopez-Salido and Nelson (2010)
Continental Illinois rescue	May 1984	
Glass-Steagall weakened	December 1986	Sherman (2009)
Stock market crash	October 1987	
	1988–1991	Lopez-Salido and Nelson (2010)
Savings and Loan crisis	1986-1992	Hubrich and Tetlow (2015)
	1984–1991	Reinhart and Rogoff (2013)
FIRREA	August 1989	
Mexican peso crisis	December 1994–1995	
Glass-Steagall weakened	1996	Sherman (2009)
Asian financial crisis	July 1997–1999	
LTCM collapse	May-September 1998	
Russian debt default	August 1998	
GLBA (Glass-Steagall repealed)	November 1999	
Dotcom bubble bursts	March 2000–April 2001	
Argentine debt default	December 2001–2002	
Repo market run	August 2007	Gorton and Metrick (2012a)
Global financial crisis	2008–2009	
Bear Stearns collapse	March 2008	
Lehman bankruptcy	September 2008	
US bank stress tests	February–May 2009	
Dodd-Frank	May 2010	
European debt crisis	May 2010–December 2012	

Notes: Where beginning and end dates of financial events are in contention, we appropriately cite dates from Sherman (2009), Lopez-Salido and Nelson (2010), Gorton and Metrick (2012a), and Hubrich and Tetlow (2015).

in the previous section. Additionally, we do not have an obvious measure for the sector's profitability. We must thus proxy for each of these as indicators of the health of financial intermediaries.

Probability of default. We use the 3-month US interbank lending rate over the 3-month Treasury Bill (TED) spread as a proxy for the probability of default in the financial intermediation sector. The interbank lending rate is the return paid on unsecured interbank loans, while the 3-month Treasury Bill represents a risk-free asset with little credit exposure. Thus, the TED spread at once reflects funding costs for intermediaries on the interbank market, as well as counterparty risk among intermediaries. A similar measure is used by Gorton and Metrick (2012b), who examine the LIBOR spread over the return on an overnight index swap (LIB-OIS). In practice, the TED spread tracks the LIB-OIS spread closely and has the added benefit of offering a much longer history.

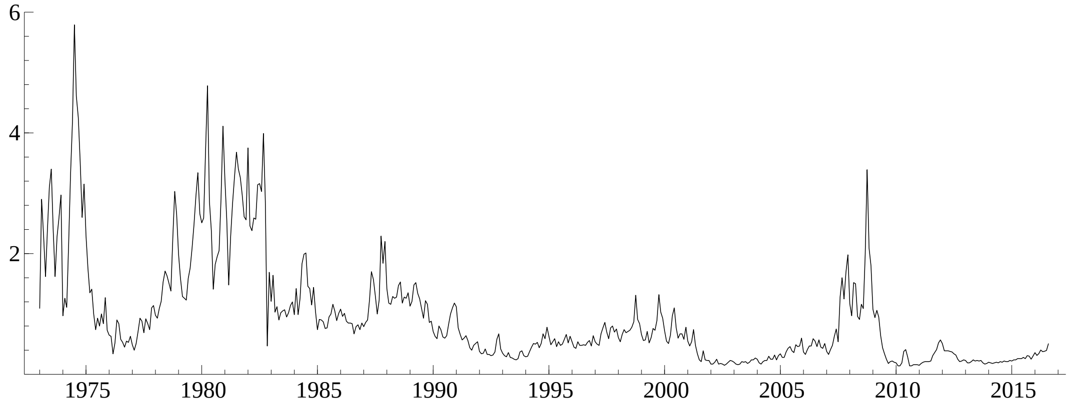

FIGURE 2.4.1: Default risk proxy. The TED spread over the sample 1973–2016. Vertical axis: percentage points. Horizontal axis: years.

The TED spread reflects counterparty risk on unsecured loans between banks. While it increases during most financial events described in Table 2.4.1, the spread peaks noticeably during periods of turmoil among intermediaries. A curiosity of the TED spread is the fact that the mid-1970s and early-1980s peaks exceed the peaks of both the Savings & Loan and the Global Financial crises. This is partially attributable to relatively lower interbank market liquidity prior to financial deregulation in the 1980s that renders the TED spread relatively more responsive than the subsequent portion of the sample. If we re-examine Table 2.4.1, however, there is sufficient evidence to suggest an informative signal despite lower liquidity. Carron and Friedman (1982) cite the problems of Franklin National and other banks as leading to elevated borrowing rates for commercial banks, specifically reflecting the rise in the TED spread. Furthermore, Boyd and Gertler (1994) show that the percent increase of bank failures was highest in the early 1980s, as opposed to the mid- to late-1980s.

Profitability. We use the growth rate of the market valuation of the largest (by assets) US banks and non-banks that increasingly engaged in similar activities as a measure of profitability in the financial intermediation sector. Examining market valuation, rather than simply price as in Aspachs et al. (2007), not only informs on performance but also gives an indication of size. This, in turn, reflects the capability of the sector to withstand shocks through available buffers.

It should be noted that market valuation may be influenced by factors other than current or expected profitability. Nonetheless, alternative proxies for profitability based on balance sheet data are neither available at a desirable frequency, nor are they free from possible manipulation. We thus opt for market-based data.

Financial fragility requires both the TED spread and profitability variables in order to provide reliable indications on the health of the intermediation sector. Since the TED spread tracks the cost of wholesale borrowing on the interbank market, it will increase when other sources of funding—for instance, deposits or equity—become scarce or expensive. Alternatively, the TED spread reflects counterparty risk that may depend on intermediaries' balance sheets. In an economic regime characterised

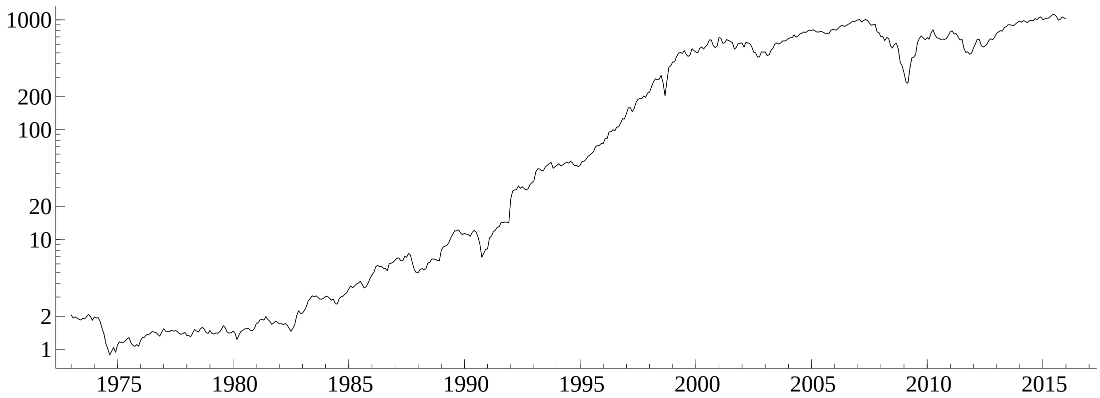

FIGURE 2.4.2: Profitability proxy. Market valuation of the largest US financial intermediaries, including banks and non-banks. Vertical axis: billions US dollars, logarithmic scale. Horizontal axis: years.

by increased volatility, balance sheet positions may become untenable if households and other borrowers optimally default on their loans. This will be passed on as higher default risk among intermediaries who may in turn find it more difficult to repay their (interbank) loans. This effect will be passed on through the interbank market as supply for loanable funds decreases.

In either circumstance, intermediaries' profits can act as a buffer. When funding liquidity is constrained and intermediaries cannot raise funds from alternative sources, profits can be recycled as capital injections to tide over funding requirements. Alternatively, if intermediaries' capital is tied up in unsuccessful loans or risky assets, profitability will indicate to what extent, if at all, an uptick in default risk corresponds to a state of financial fragility. Thus, we must examine default risk and profitability jointly in order to draw inferences on fragility.

2.5 Financial fragility in the United States

In this section, we provide an answer to our first question: financial fragility can be defined in terms of a nonlinear association between the intermediation sector and the real economy. We find evidence for two regimes that are characterised by switching dynamics with respect to default risk and profitability, as well as by switching volatility. The first regime occurs during volatile periods of high default risk and relatively low profitability and so we are able to identify this as the high-fragility regime. This regime captures the most significant financial events in the United States over our sample. Under high fragility, we show that economic growth is nonpositive and credit extension decelerates. Most recessions occur during periods of high fragility, but not all episodes of fragility result in financial crisis and recession.

We assess the occurrence of financial fragility in the United States over the period 1973–2016 by estimating a four-variable model that includes (real) industrial production (IP), (real) total credit to the private non-financial sector (TC), (real) market valuation of US financial intermediaries (MV), and the TED spread (TED); IP, TC, and MV are modelled as monthly growth rates. The data are described in Appendix 2.A.

Several robustness checks are also included in Appendix 2.C. To ensure that these results are not driven entirely by monetary policy (Miranda-Agrippino and Rey, 2015), we estimate a model that includes the Federal Funds rate as a control for monetary policy; policy at the zero lower bound is considered by using the shadow rate from Wu and Xia (2016). Additionally, we include alternative measures for increased default risk among financial intermediaries and re-estimate the initial model. These checks do not affect the results in a substantial way.

2.5.1 Model specification

Existing empirical work that employs nonlinear models using data for the United States overwhelmingly favours two regimes (Hubrich and Tetlow, 2015; Nason and Tallman, 2015; Davig and Hakkio, 2010; Aspachs et al., 2007; Balke, 2000). This is also consistent with our theoretical priors of a good and bad state (Goodhart et al., 2013), and so we opt to simplify our modelling efforts. Three monthly lags are sufficient to capture the serial dependence, given that we permit regime-switching in the intercept, autoregressive coefficients, and the covariance matrix. Details on temporal dependence in the data and residual analysis are provided in Appendices 2.A and 2.B.

We carry out a simple model selection procedure to determine what sort of non-linearities we need to include. Recall the MS-VAR model in (2.3.1), reproduced here for $i = 1, 2$

$$\mathbf{y}_t = \mu_i + \Gamma_{1i}\mathbf{y}_{t-1} + \ldots + \Gamma_{pi}\mathbf{y}_{t-p} + \epsilon_t, \quad \epsilon_t | s_t \sim N(0, \Sigma_i).$$

Column j of each Γ_{pi} matrix represents the vector of autoregressive coefficients of the j^{th} variable at lag p for regime i over all equations. Define $\Gamma_{\bullet i} = \{\Gamma_{1i}, \ldots, \Gamma_{pi}\}$ the set of all $N \times N$ matrices of autoregressive coefficients for regime i. We write $\Gamma^j_{\bullet i}$ to denote the set of j^{th} columns in all matrices Γ_{li}, $l = 1, \ldots, p$. In order to assess which parameters cannot be restricted to equality across the two regimes, we test the following set of hypotheses

$$\mathsf{H}_0^\mu : \mu_1 = \mu_2, \qquad \mathsf{H}_0^\Gamma : \Gamma^j_{\bullet 1} = \Gamma^j_{\bullet 2}, \qquad \mathsf{H}_0^\Sigma : \Sigma_1 = \Sigma_2,$$

where for H_0^Γ we iteratively restrict all autoregressive coefficients in every equation for a subset of the variables in $\mathbf{y}_t$.

Table 2.5.1 provides the comprehensive list of models that we examine; all models allow two regimes and three lags. Initiating a general-to-specific approach, we allow for switching in all autoregressive parameters, the intercept, and the covariance matrix (model M1). Subsequently, we reduce the general model by restricting regime-variation in subsets of variables: models M2–M5 restrict switching in the autoregressive coefficients on one variable; models M6–M11 restrict switching in the autoregressive coefficients on two variables; models M12–M15 restrict switching in the autoregressive coefficients on three variables; models M16–M17 restrict switching in the intercept and covariance matrix; model M18 is the linear case. All models are nested in the general case of model M1. A likelihood ratio test can thus help us test the null hypotheses H_0^{μ}, H_0^{Γ}, and H_0^{Σ}; the resulting χ^2 test statistic is reported in the final column of Table 2.5.1 with p-values in brackets.

Our efforts to reduce the number of switching parameters in model M1 results in the set of models {M2, M3, M4, M6, M8, M11} that are not rejected at the 5% level. We make the following observations. First, there is substantial evidence against linearity, as the reduction M1 $\rightarrow$ M18 is strongly rejected.[2] Second, there is strong evidence to suggest that switching in the size and correlation of shocks is not sufficient to model the data, as the reduction M1 $\rightarrow$ M16 is also rejected strongly. Third, there is strong evidence against constant volatility, since M1 $\rightarrow$ M17 is rejected. From these results, we can draw the preliminary conclusion that some switching in dynamic parameters and volatility is necessary in order to appropriately model the data.

In order to choose among this smaller set of non-nested models, we cannot rely on a likelihood ratio test. We thus compute various information criteria and report these in Table 2.5.3. All criteria point toward the models M6, M8, and M11. In each

[2]A caveat here relates to identifiability as described in §2.3.3, since the number of regimes is being changed in this hypothesis. The test statistic is extremely large, however, and consistent with the progression of tests in Table 2.5.1 that we do not view this result with much skepticism.

TABLE 2.5.1: Model specification from general to specific. Parameters are sequentially restricted from varying across regimes.

Model	T	p	Log-Likelihood	Switching parameters μ_i	$\Gamma^{IP}_{\bullet i}$	$\Gamma^{TC}_{\bullet i}$	$\Gamma^{MV}_{\bullet i}$	$\Gamma^{TED}_{\bullet i}$	Σ_i	LR test χ^2_ν
M1	512	126	-1817.85	✓	✓	✓	✓	✓	✓	
M2	512	114	-1827.24	✓		✓	✓	✓	✓	18.79 [0.09]
M3	512	114	-1828.04	✓	✓		✓	✓	✓	20.39 [0.06]
M4	512	114	-1823.17	✓	✓	✓		✓	✓	10.64 [0.56]
M5	512	114	-1831.38	✓	✓	✓	✓		✓	27.05 [0.01]
M6	512	102	-1833.52	✓			✓	✓	✓	31.35 [0.14]
M7	512	102	-1836.41	✓	✓	✓			✓	37.12 [0.04]
M8	512	102	-1833.95	✓	✓			✓	✓	32.21 [0.12]
M9	512	102	-1841.19	✓	✓		✓		✓	46.69 [0.00]
M10	512	102	-1834.96	✓		✓	✓		✓	34.22 [0.08]
M11	512	102	-1834.31	✓		✓		✓	✓	32.93 [0.11]
M12	512	90	-1846.43	✓	✓				✓	59.07 [0.01]
M13	512	90	-1847.38	✓		✓			✓	57.12 [0.01]
M14	512	90	-1843.37	✓			✓		✓	51.03 [0.04]
M15	512	90	-1844.41	✓				✓	✓	53.12 [0.03]
M16	512	74	-1861.54						✓	87.38 [0.00]
M17	512	116	-2050.34	✓	✓	✓	✓	✓		464.99 [0.00]
M18	512	52	-2261.92							888.14 [0.00]

Notes: Estimation of all models uses monthly data between 1973(6) and 2016(1). A likelihood-ratio test of model reduction is reported with p-values in brackets. This statistic has a χ^2 distribution with ν degrees of freedom equal to the number of restricted parameters.

of these, the intercept, the covariance matrix, and the autoregressive coefficients on the TED spread are allowed to vary across regimes in all equations. Model M6, where the autoregressive coefficients on the growth in market valuation are also allowed to vary, is chosen marginally over the other models.

TABLE 2.5.3: Information criteria for non-rejected model specifications.

Model	SC	HQ	AIC	Switching parameters					
				μ_i	Γ_i^{IP}	Γ_i^{TC}	Γ_i^{MV}	Γ_i^{TED}	Σ_i
M1	8.636	8.002	7.593	✓	✓	✓	✓	✓	✓
M2	8.527	7.953	7.583	✓		✓	✓	✓	✓
M3	8.530	7.956	7.586	✓	✓		✓	✓	✓
M4	8.511	7.937	7.567	✓	✓	✓		✓	✓
M6	**8.405**	**7.892**	**7.561**	✓			✓	✓	✓
M8	8.407	7.893	7.562	✓	✓			✓	✓
M10	8.411	7.897	7.566	✓		✓	✓		✓
M11	8.408	7.895	7.564	✓		✓		✓	✓

Notes: Estimation of all models uses monthly data between 1973(6) and 2016(1). The Schwarz (SC), Hannan-Quinn (HQ), and Aikake (AIC) information criteria are reported; all favor model M6.

The reduction procedure specified here indicates that the nonlinear macro-financial relationship cannot be simply attributed to larger shocks among a single regime of transmission mechanisms. Rather, the evidence presented here suggests that there is a fundamental change in the transmission mechanism of shocks between the financial intermediation sector and the real economy. The fact that we cannot reject switching in the dynamic effects proxying for the health of the financial intermediation sector suggests that we can describe the two regimes in terms of high and low financial fragility. We distinguish between and characterise the high- and low-fragility regimes, while also examining this claim from the perspective of historical events.

2.5.2 The high-fragility regime

In §2.5.1 we provided evidence for outsized shocks across regimes, as well as for a changing transmission mechanism for these shocks. The most significant change in dynamics occurs in the financial variables—profitability and default risk among financial intermediaries—which we may infer act as the drivers of the regime change. Figure 2.5.1 plots both of these series, with the high-fragility regime shaded (Regime 1) and the low-fragility regime unshaded (Regime 2).

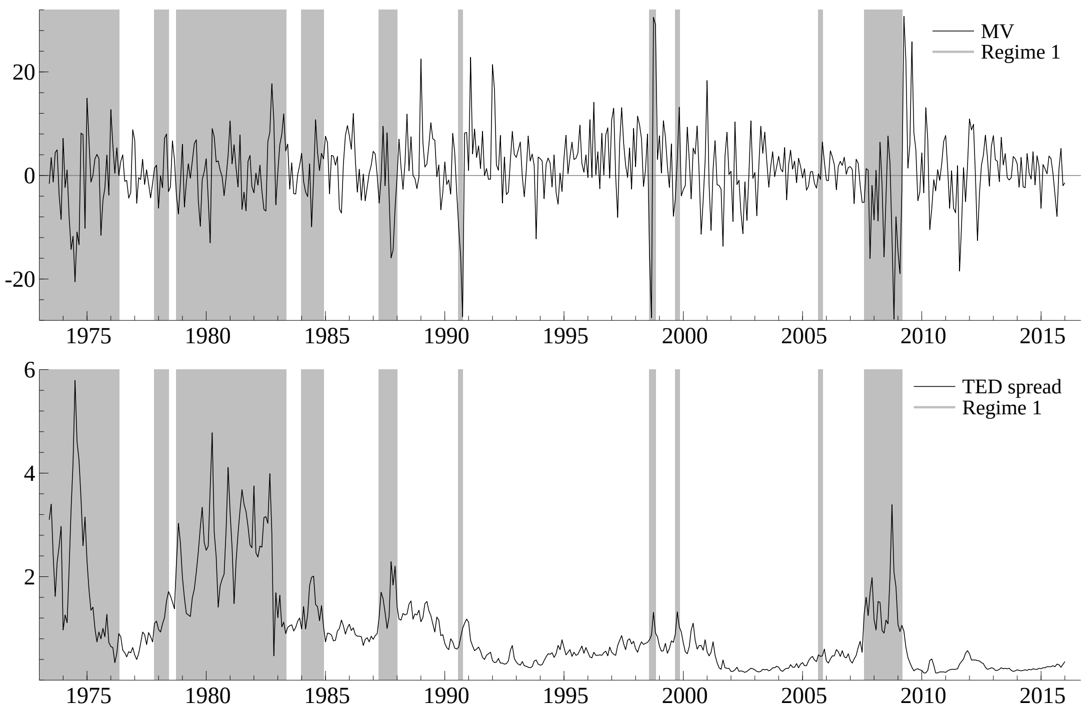

FIGURE 2.5.1: Financial variables by regime. Profitability (growth in MV) and the TED spread are plotted with the high fragility regime. Vertical axis: percentage points. Horizontal axis: years.

We draw attention to three features in Figure 2.5.1. First, the shaded periods corresponding to Regime 1 show greater volatility in both series. Second, most periods of significantly negative profitability and significantly large TED spreads are included in Regime 1. Finally, not all periods of negative/low profitability and heightened TED spreads qualify. What can be inferred from this? Generally, there is a regime switch when both profitability declines and default risk increases. Nonetheless, a sufficiently large change in either of these variables individually will also lead to a regime switch. This is consistent with the financial fragility definition in Tsomocos (2003), and so we characterise Regime 1 as one of "high financial fragility." During this regime, the financial intermediation sector is under greater stress and various financial frictions may impede the sector from performing its regular operations.

It is, however, also important to acknowledge that regime switching also depends on shock size. We do not separately consider regimes for volatility and for dynamic coefficients, but rather only consider regimes where both are restricted to switching simultaneously; greater volatility is implied in the theoretical framework of Goodhart et al. (2005) and made explicit in calibrated models, such as Goodhart et al. (2013). This has the implication that periods which exhibit similarly low profitability and high default risk as those periods classified under Regime 1—for instance, see 2011(6) and 2012(4) in Figure 2.5.1—may not be classified as such. The contrapositive of this statement also holds—see 2005(8) classified under Regime 1. To disentangle volatility and coefficient regimes would lead to an explosion in the number of parameters and would require alternative approaches to estimation. We note that the most obvious periods of financial stress are classified under Regime 1 by inspection of Figure 2.5.1 and leave these questions open for future work.

2.5.3 What characterises financial fragility?

To further understand the properties of the high fragility regime, we plot linear, regime-dependent contemporaneous correlations in Figures 2.5.2 and 2.5.3. The first three panels show the two-dimensional correlation between the financial variables and economic/credit growth. The fourth panel shows a three-dimensional plot of both financial variables along with economic/credit growth. The observations are divided by regimes, with the red dots representing the regime of high fragility, while the blue dots represent the low-fragility regime. The two-dimensional plots are cross-sections of the three-dimensional plot when one of the dimensions is collapsed.

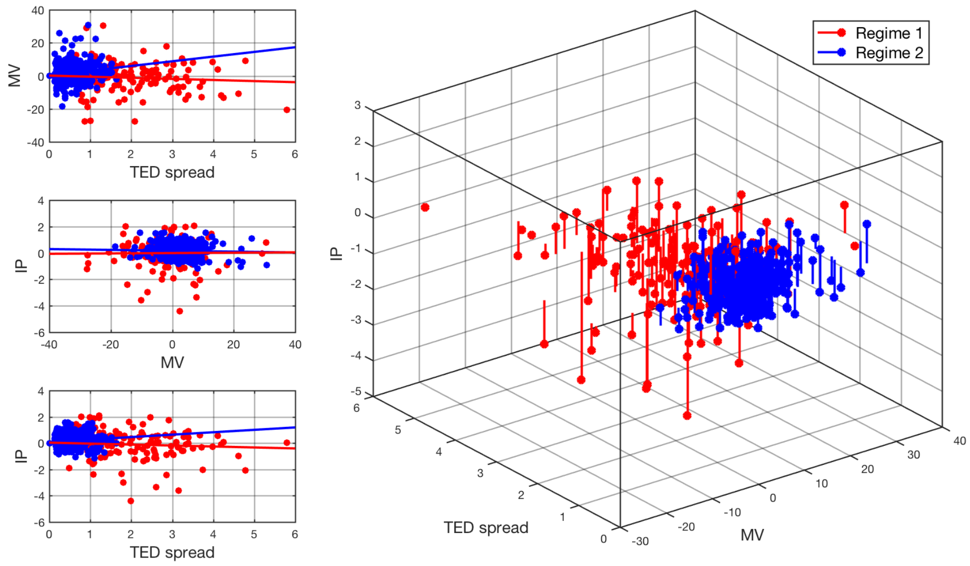

FIGURE 2.5.2: Correlation plots. The contemporaneous correlations between financial and macroeconomic variables plotted by regime. Axes: percentage points.

The first panel in Figure 2.5.2 shows the regime-dependent relationship between the TED spread—proxying for default risk among financial intermediaries—and the growth in market valuation of financial intermediaries—proxying for profitability. Under Regime 1, higher default risk is correlated with lower profitability, while there is significant dispersion; the opposite is true on both counts in Regime 2.

Examining the second and third panels sheds more light on the properties of Regime 1. Profits and growth do not display a significant contemporaneous correlation, but under Regime 1 there is significantly more dispersion in the outcomes. Default risk, on the other hand, has a clear regime-dependent correlation with growth. Under Regime 1, a higher TED spread is associated with lower growth and much more volatility, whereas under Regime 2 the relationship is positive and there is much less volatility. In Regime 2 there is perhaps a tolerance for default risk up to a threshold that is compatible with positive economic growth and profitability. Beyond a certain threshold, however, default risk becomes excessive and occurs with negative economic growth. These effects are also evident in the three-dimensional plot in the fourth panel. The points attributed to Regime 1 are clustered higher on the TED spread axis and lower on the MV axis while occupying the larger portion of the negative range of the IP axis relative to the points attributed to Regime 2; once again, we observer higher volatility under Regime 1. Note that whereas profitability on its own was not significantly correlated with growth in either regime in the second panel, there is a positive correlation once default risk is also taken into consideration due to the partial correlation. This reinforces earlier results in sections §2.5.1-§2.5.2 that suggested regime-switching depends on both profitability and default risk.

Similar conclusions can be drawn by examining the relationship between real credit growth and the financial variables, as depicted in Figure 2.5.3. In the low-fragility regime (Regime 2) in panel 1, increases in production growth are associated positively with credit growth. This association is stronger in the high-fragility regime (Regime 1) with much larger dispersion and asymmetry. Particularly, small negative changes in credit growth are associated with much larger negative changes in production growth. What is more, these negative changes in credit growth have a stronger impact on production than positive changes. Panels 2 and 3 show the association, by regime, between credit growth and the financial variables. In both cases, the correlation between credit growth and the financial variable flips sign between regimes. In panel 2, positive credit growth in the high-fragility regime is associated with higher profitability. In panel 3, once we remove an influential outlier (see the red dotted line), credit growth is negatively associated with default risk in the high-fragility regime, whereas this association is positive in the low-fragility regime, albeit with much less dispersion. At a low level, increases in risk that come about through increased credit growth are tolerable. Beyond a threshold, however, we switch to the high-fragility

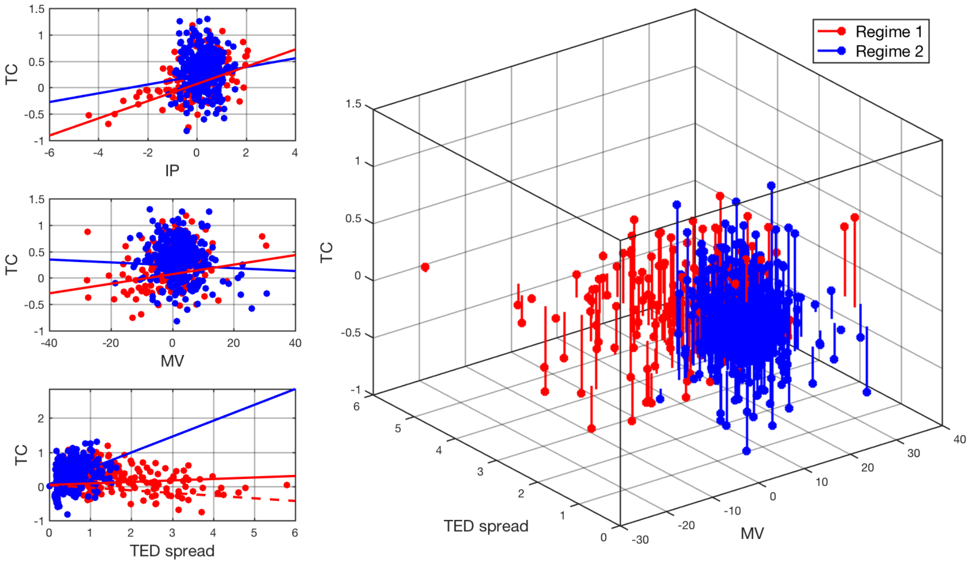

FIGURE 2.5.3: Correlation plots. The contemporaneous correlations between financial and macroeconomic variables plotted by regime. Axes: percentage points.

regime where financial frictions may bind, and thus increased credit acts to assuage default risk rather than increase it.

2.5.4 High fragility in the historical narrative

In this section, we take a closer look at the regime classification relative to the historical narrative. Importantly, we provide an answer to the following question: is high fragility always a precondition for economic recession? Figure 2.5.4 plots the smoothed regime probabilities against all NBER recessions over our sample period.

First, we note the frequent occurrence of highly fragile regimes that far exceeds the number of financial episodes described in Reinhart and Rogoff (2008) and Schularick and Taylor (2012); both include only two financial events in 1984 and 2007 for the United States in the post-war period. Indeed, there is evidence for a highly fragile financial intermediation sector for several of the events described in Table 2.4.1. The first episode of high financial fragility corresponds to the Bank Capital Squeeze begin-

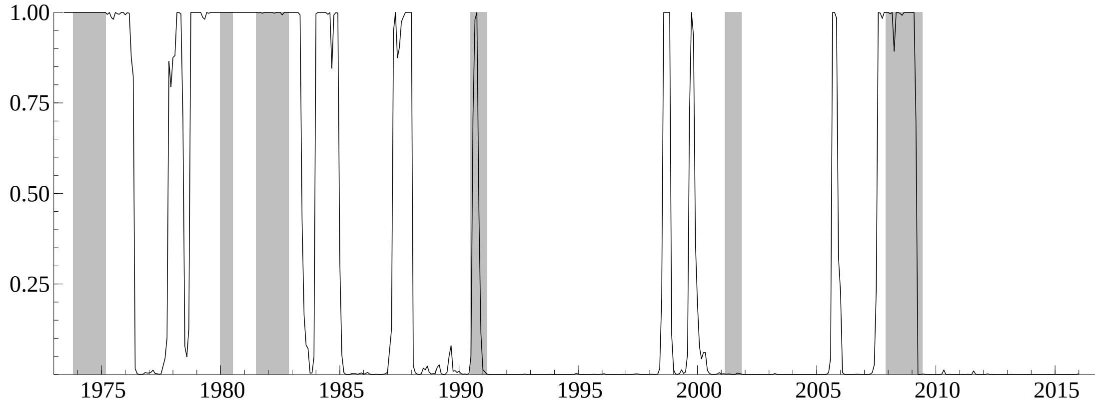

FIGURE 2.5.4: Smoothed regime probabilities. The probability of any period falling under Regime 1 (high fragility) in model M1 is plotted along with recessions in the United States. Vertical axis: probability. Horizontal axis: years. Black line: smoothed regime probability. Grey areas: NBER recessions.

ning in 1973, documented in Lopez-Salido and Nelson (2010) and Boyd and Gertler (1994) among others.

The second episode detected in our model occurs in 1978. It includes the beginnings of the S&L crisis, as well as the LDC debt crisis that culminated in the rescue of Continental Illinois in 1984—the largest bank failure in history, at the time.[3]

The third distinct episode takes place in 1987 just as the S&L crisis was beginning to emerge; while this episode also includes the stock market crash, it is important to note that high fragility was present significantly prior to October. Our dating of the S&L crisis agrees with Hubrich and Tetlow (2015) that place it between 1986–1992. The next detectable episode in 1991 corresponds to the end-date of the S&L crisis and a large credit crunch among banks shrinking their balance sheets to comply

[3]This period is characterised by a series of efforts to deregulate the financial sector. Ceilings on bank deposit rates and reserve requirements meant that it was increasingly difficult for banks to attract deposits and remain profitable. Contrastingly, money market funds did not face interest rate ceilings or reserve requirements and were more able to attract funds. Aggressive monetary policy to combat high inflation also meant that bank profits would be severely squeezed, creating ideal conditions for high fragility. Two pieces of legislation led to the substantial deregulation of banks and thrifts. The Depository Institutions Deregulation and Monetary Control Act (DIDMCA) of 1980 increased deposit insurance and phased out interest rate ceilings, so that banks could pay higher deposit rates. The Garn-St. Germain Act of 1982 deregulated thrifts, allowing them to engage in commercial lending and compete with money market funds.

with regulatory capital buffers (Syron, 1991; Bernanke and Lown, 1991), thereby generating signficant losses.

We find evidence that the Asian financial crisis, as well as the Russian default that precipitated the failure of LTCM, led to periods of high fragility in 1998 and 1999.[4] Finally, we detect the Global Financial crisis from early 2007 through the beginning of 2009.

Figure 2.5.4 shows that most recessions occurred during periods of high fragility, but also that episodes of high fragility do not always result in financial crisis or recession. The weak 2001 recession stands out in particular as one that is not characterised by financial fragility; further results in Appendix 2.C confirm that this result is not spurious. This distinction is worth examining further. Table 2.5.5 breaks down the conditional means and variances over all recessionary periods, as well as periods of high financial fragility. When the distribution moments are sufficiently distinct, it is appropriate to separate the joint distribution into a mixture of distributions conditional on a state variable. This is effectively how we separate the data into two regime-dependent probability distributions of low and high fragility. Across all recessions in our sample for the United States, the TED spread maintained an average of 2%, while month-on-month profitability for financial intermediaries declined 1.5%. This stands in contrast to an average TED spread of 0.8% and monthly profitability of 1.6% outside of recessions. While not all recessions coincide with high financial fragility, a weakened real economy presents threats to the capability of intermediaries to perform their function. Indeed, credit growth grinds to a halt in recessionary periods, compared to 0.4% monthly growth otherwise. Note that while all recessions are—by definition—characterised by negative productive growth, not all recessions are characterised by negative profits and high default risk.

[4] Incremental weakening of Glass-Steagall began in December 1986, when the Federal Reserve permitted banks to derive a fraction of their revenues through investment banking, allowing institutions to hold an increasingly diverse array of assets. This fraction was raised from 5% to 10% in 1996, before Glass-Steagall was effectively repealed in 1999. For an exposition of the history of

Table 2.5.5: A cross-sectional summary for the period 1973–2016.

Dates	IP		TC		MV		TED	
	Mean	StD	Mean	StD	Mean	StD	Mean	StD
1973(11)–1975(3)	-0.79	1.25	-0.13	0.29	-3.46	9.91	2.82	1.36
1980(1)–1980(7)	-0.91	1.07	0.22	0.33	-1.10	7.57	2.89	1.08
1981(7)–1982(11)	-0.49	0.84	0.14	0.25	0.70	6.97	2.88	0.78
1990(7)–1991(3)	-0.44	0.47	0.08	0.28	-1.31	14.85	0.88	0.23
2001(3)–2001(11)	-0.43	0.18	0.55	0.23	-1.17	7.34	0.38	0.18
2007(12)–2009(6)	-1.00	1.20	-0.21	0.24	-2.81	13.70	1.35	0.68
All recessions	-0.70	0.99	-0.01	0.35	-1.47	10.50	1.97	1.29
No recession	0.31	0.54	0.36	0.33	1.58	6.09	0.76	0.63
High fragility	-0.02	1.06	0.22	0.40	-0.84	8.79	1.89	1.03
Low fragility	0.23	0.50	0.34	0.34	1.96	5.90	0.54	0.32

Notes: Rates represent month-on-month average growth rates in percentage points for all differenced variables and monthly averages for variables in levels. The final two rows represent mean values at the intersection of recessionary and financially fragile (stable) periods.

With the exception of the 2001 recession, all other recessions are included under the high-fragility regime; however, other periods that did not result in recession are also included. Consider the conditional means for periods of low financial fragility in the final row of Table 2.5.5. Economic and credit growth are both positive between 0.2-0.3% month-on-month. Real increases in the productive capacity of the economy are more than matched by increases in credit growth, suggesting a possible leveraging effect. The TED spread is positive around 0.5%, but the combination of high profitability at a monthly rate of 2% and lower volatility compared to the no-recession periods—5.9% compared to 6.09% for profitability and 0.32% compared to 0.63% for the TED spread—suggests that the leveraging is a symptom of a well-functioning financial intermediation sector.[5]

deregulation in the US, see Sherman (2009).

[5]This does not necessarily imply that there is not over-leveraging in the low-fragility state, for instance as in Bianchi (2011) or Minsky (1992). It is possible that well-functioning credit intermediation leads to a sustained and self-reinforcing rise in expectations of the future, thereby leading to over-leveraging that magnifies high fragility into a financial crisis (Kindleberger and Aliber, 2011). This is not something we can reliably infer, however, as the scope of the present investigation is to distinguish between regimes of fragility and characterise them appropriately.

Contrast this with the high-fragility regime in the second-last row of Table 2.5.5. This is characterised by non-positive economic growth and profitability, as well as a TED spread of approximately 2% with credit nonetheless increasing on average. Switching between regimes occurs when the TED spread sufficiently nears 2% at the same time as profitability is insufficient— -0.8% month-on-month—to weather shocks to the value of assets on intermediaries' balance sheets, given the increased volatility. Indeed, lower volatility in the financial variables relative to all recessionary periods—8.79% compared to 10.5% for profitability and 1.03% compared to 1.29% for the TED spread—implies marginally higher volatility in the macroeconomic variables. Credit extension continues to grow in the high-fragility regime, albeit at a lower rate than in the low-fragility regime. This deceleration in credit growth suggests that intermediaries' efficiency in conducting their regular activities is impeded as financial frictions and default risk bind.

To shed some light on this, we plot production and real credit growth in Figure 2.5.5 along with the high-fragility regime. In the first panel we see that high financial fragility occurs fairly often with negative economic growth, but this is not always the case. Notable exceptions include 1984, 1987, 1998, and 1999. All of these periods include financial events that contributed to fragility—see §2.5.4 and Table 2.4.1—yet did not lead to a deterioration in economic growth. This reinforces our claim that financial fragility cannot be defined as a binary event that leads to economic crises. Rather, financial fragility must measure the presence of the appropriate conditions among financial intermediaries that *could* lead to an economic crisis.

Additionally, it is not credit growth itself that leads to financial fragility; our results in Appendix 2.D show that the analysis in Schularick and Taylor (2012) does not hold for the case of the United States and credit growth alone is not a significant predictor of financial crises or fragility. Inspecting the second panel of Figure 2.5.5, high-fragility regimes do not systematically occur when credit growth is high. Rather,

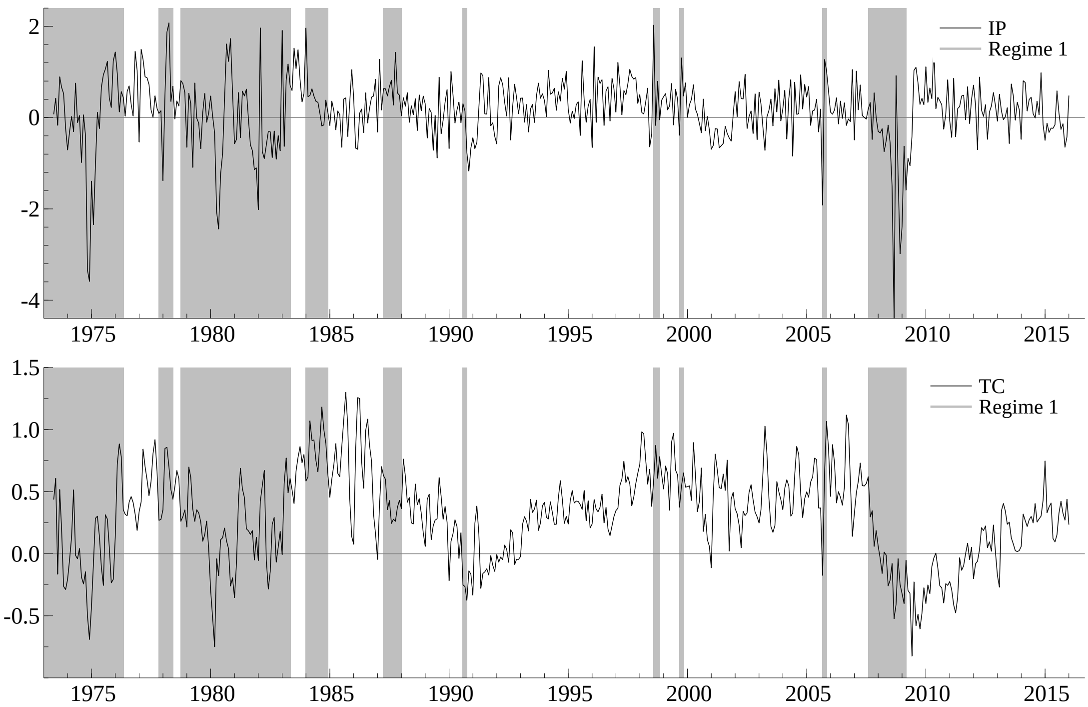

FIGURE 2.5.5: Macroeconomic variables by regime. Industrial production growth and real credit growth plotted with the high fragility regime. Vertical axis: percentage. Horizontal axis: years.

they occur when credit growth decelerates, or shortly thereafter. This does not mean that credit growth is negative (though it could be); it means that credit is still growing, but at a lower rate, as indicated in Table 2.5.5. Finally, note that not all periods of credit deceleration or negative credit growth coincide with high financial fragility—see 1986, 2001, and 2013, for instance. This is consistent with our testing in §2.5.1 where we found that the regime-switching can be restricted to the dynamic coefficients of the financial variables.

These results reflect the statics of different regimes in the data. The subsequent dynamic analysis continues to draw on these static results.

2.6 Shock transmission under financial fragility

How do shocks affect economic growth when intermediaries are highly fragile, and what is the relative importance of the transmission channels? This section provides answers to our second and third questions. The effects of shocks depend heavily on the state of financial intermediaries. Under high fragility, shocks are larger and more persistent. While shock size can overstate the extent to which transmission across states actually varies, our most important results can be traced to changing dynamics. Even though we cannot make direct causal statements, we argue for the relative importance of the financial health of intermediaries as opposed to credit in transmitting shocks.

Our use of a discrete-state MS-VAR model allows the effects of shocks to vary across regimes. Variation can originate in one of three sources: from the contemporaneous impact of the shock; from the dynamic transmission of the shock; and from the size of the shock. The first and third types of variation are permitted through stochastic volatility, while the second source of variation is a result of switching dynamic coefficients.

So far, we have detected nonlinearities in the relationship between the real economy and financial intermediaries and attributed these nonlinearities to fragility. Now, we examine the dynamic properties of this nonlinear behaviour. We compute the effects recursively-identified shocks over a 25-month horizon and condition on each regime separately, employing the approach in Ehrmann et al. (2003) for MS-VAR models described in §2.3.4. Validity of this procedure requires persistence of each regime. This requirement is met, as the transition probabilities for remaining in the current regime are 0.93 for high fragility and 0.97 for low fragility.

2.6.1 Real effects of shocks under financial fragility

The effects of shocks on economic growth depend on whether the financial intermediation sector is in a regime of high or low fragility. In a low-fragility regime, production growth responds mildly and returns to its initial state quickly. Under high fragility, shocks are amplified through the credit and default channels so that effects are larger and more persistent. These dynamics are shown in Figure 2.6.1. Shocks of one standard deviation in each regime are plotted with 90% confidence intervals, along with the impulse response from a linear VAR model for comparison.

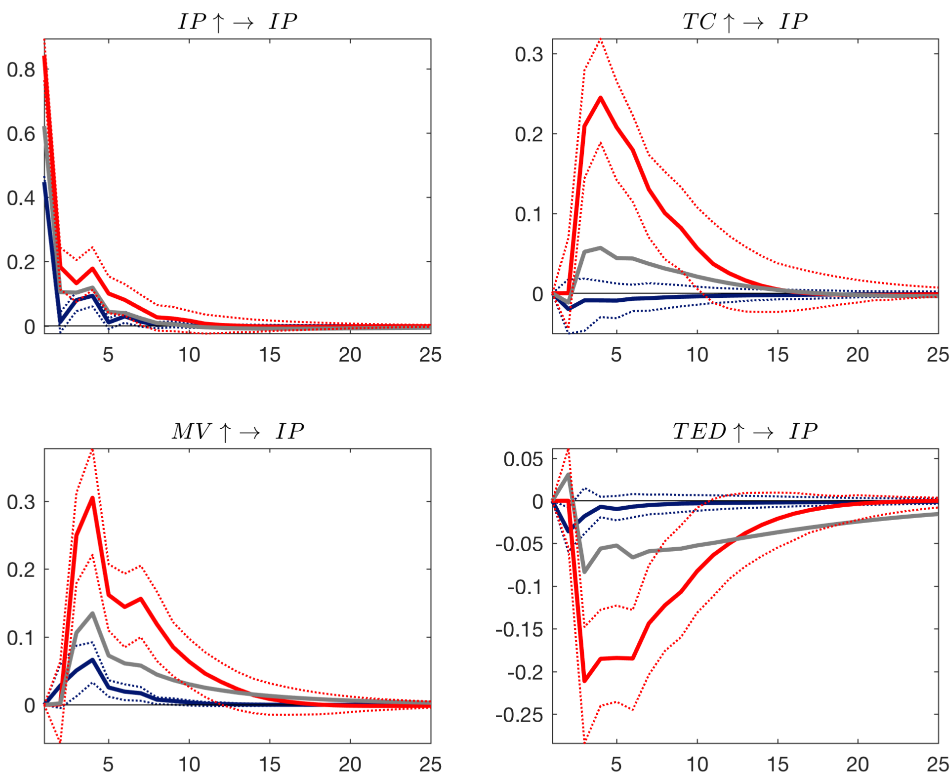

FIGURE 2.6.1: Regime-dependent impulse response functions. IRFs for IP growth to shocks of one standard deviation with 90% confidence intervals. Red: high-fragility regime IRFs. Blue: low-fragility regime IRFs. Grey: linear IRFs. Vertical axis: percentage. Horizontal axis: horizon (months).

All impulse responses are significantly different across regimes, indicating that a nonlinear approach is appropriate; the linear impulse responses are essentially averages of the responses across both regimes. The first panel shows that shocks to production are almost double in size in the high-fragility regime but persistence does not change noticeably; the former observation is an indication that shocks change significantly, while the latter indicates that dynamics do not. Under high fragility, an expansion of credit has significant effects on production growth. Similarly, shocks that mitigate the fragility of intermediaries—higher profitability or lower default risk—will generate a positive effect on production growth that is significantly greater than the effect under a regime of low financial fragility. On first pass, these results suggest that financial frictions bind and there is higher default risk with limited buffers in the high-fragility regime. The difference in responses between the two regimes suggests that the shadow values of credit, profitability, and default risk are significantly higher under high fragility. Marginally increasing credit, profits, or reducing default risk will have larger effects when the constraints bind.

We examine these effects further in Figure 2.6.2. Response variables are stacked by column, whereas shock variables are stacked by row. Plots along the diagonal show the own-response to orthogonalised shocks. The first observation we make is that shock volatility varies significantly for all variables except for credit growth (also see Table 2.5.5). The variation is exceptionally large for the TED spread, where shocks appear almost five times larger in the high-fragility regime. Nonetheless, shocks to credit growth display the greatest persistence.

Next, we note that both financial variables have a significant impact on credit growth and vice versa. High profitability and low default risk will both produce significantly positive effects under high fragility that persist up to 10 months after impact. When total profitability in the financial intermediation sector rises, intermediaries' buffer to adverse shocks increases and this permits balance sheet expansion through increased lending. Similarly, shocks that lower the prospect of default—note

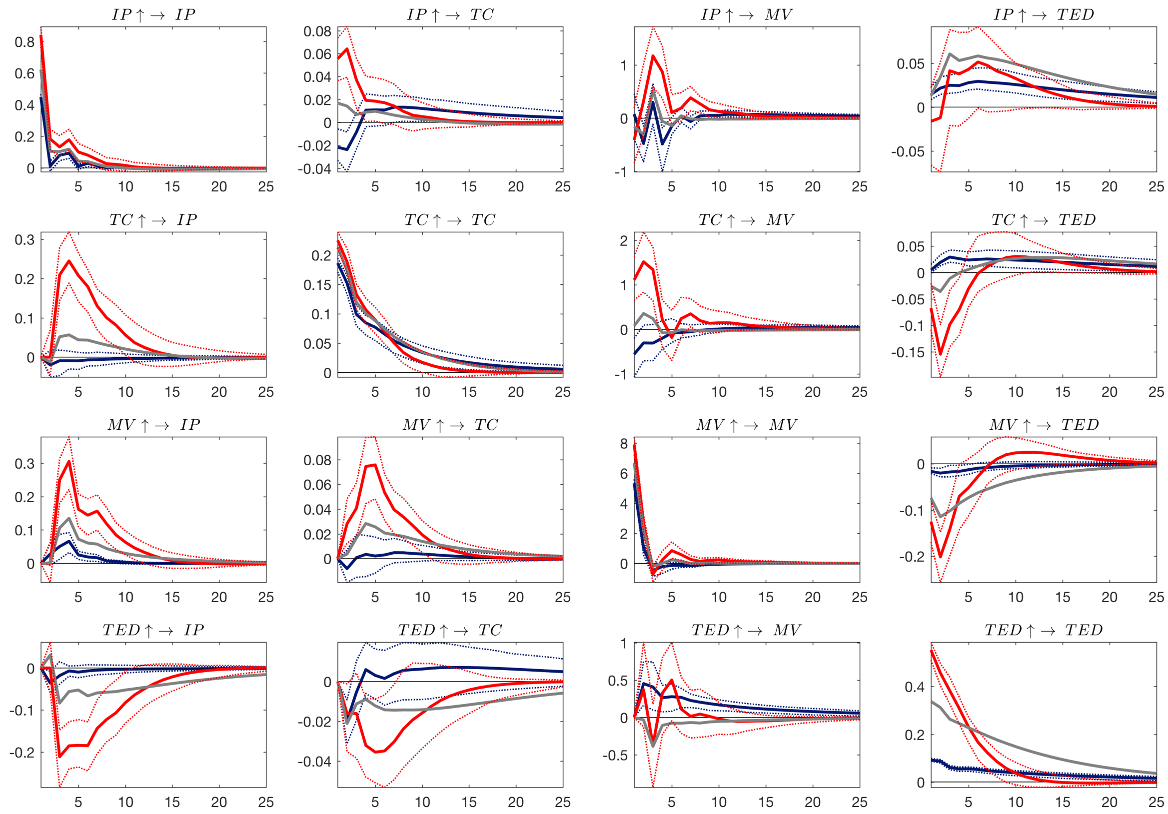

FIGURE 2.6.2: Regime-dependent impulse response functions. IRFs for shocks of one standard deviation with 90% confidence intervals. Red: high-fragility regime IRFs. Blue: low-fragility regime IRFs. Grey: linear IRFs. Vertical axis: percentage. Horizontal axis: horizon (months).

that impulse responses are symmetric in our model—also allow intermediaries to increase lending. The effects are much larger in the high fragility regime, indicating that the shadow price of credit is much higher.

The effects of credit growth are also highly significant, particularly in the high-fragility regime where we expect financial frictions to bind. An increase in credit growth—an acceleration of credit extension for one period—of around 0.2% month-on-month generates a 2% rise in profitability and 0.2% drop in the TED spread under high fragility. Recall from Table 2.5.5 that the monthly growth in credit remained positive yet lower in a high-fragility regime relative to a low-fragility regime. When we expect credit to be relatively scarce, easing credit growth generates a stabilizing effect. The opposite holds in the low-fragility regime, where a similar increase in credit growth generates a relatively small but significant rise in default risk; as credit accumulates, the potential for default rises. Thus, reductions in fragility are reinforced through credit easing when fragility is particularly high.

Finally, we note small but significant interactions between profitability and default risk. In both regimes, an increase in profitability significantly reduces the TED spread; under high fragility, the response in the TED spread to a shock in profitability is almost identical to that of a shock in credit growth. What is more, we see that a slight increase in risk under the low-fragility regime generates an increase in profitability, reinforcing our previous claim that some risk is tolerable. Indeed, during financially stable regimes, shocks associated with a booming economy—an increase in production and credit growth—generate small but persistent rises in the TED spread, suggesting that some default risk is an equilibrium feature.

When the financial sector is fragile, financial frictions generally manifest in the form of credit drying up, weak profitability, and significantly greater volatility. In this environment where credit constraints bind, an increase in credit will have an outsized effect on growth. The same is true of a stronger financial sector: lower default probability and higher profitability will have significantly larger effects in the fragile regime where credit constraints will bind. Indeed, we can argue that a stronger financial sector is a necessary condition for healthy credit expansion.

2.6.2 Changing dynamics or shock size?

In §2.6.1 we made a case for regime-dependent responses to shocks. One feature of our MS-VAR model is that we allow for stochastic volatility, and so the size of shocks will vary across regime. To control for this, we show impulse responses to a unitary shock under either regime in Figure 2.6.3. The linear response is also plotted in order to assess nonlinearity.

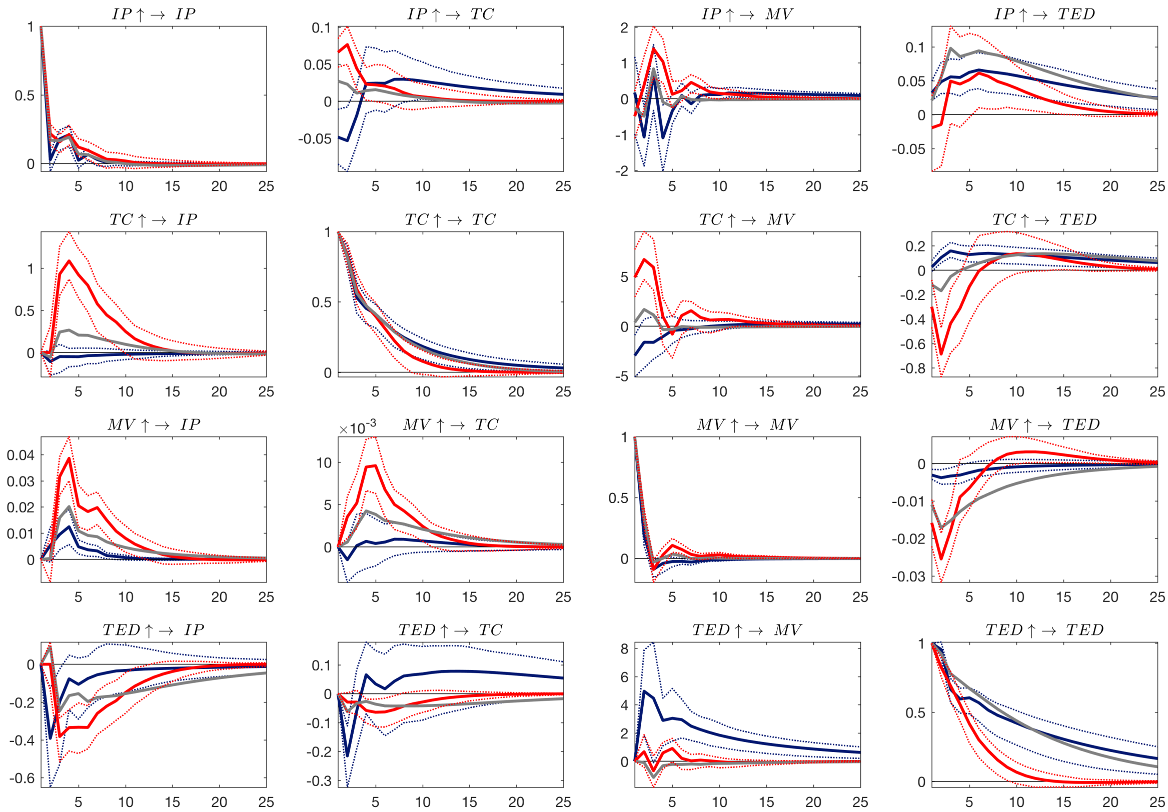

FIGURE 2.6.3: Regime-dependent impulse response functions. IRFs for shocks of one unit (percentage) with 90% confidence intervals. Red: high-fragility regime IRFs. Blue: low-fragility regime IRFs. Grey: linear IRFs. Vertical axis: percentage. Horizontal axis: horizon (months).

The difference among responses is not as stark in Figure 2.6.3 as in Figure 2.6.2, but there remains evidence for regime-dependent dynamics. We make two observations. First, shocks to default risk are substantially less persistent under the high-fragility regime, but there remains evidence for significant persistence relative to the low-fragility regime; this is the case both for industrial production and credit. Second, the comparative responses across regimes for shocks to profitability and credit growth are similar with unitary shocks as they are with shocks of one standard deviation. The relative shock size between regimes varies less with profitability and credit growth than it does for default risk. These observations reinforce the results from Table 2.5.1, indicating that both dynamics and volatility are important in order to capture the size and the transmission of shocks under different regimes of fragility.

2.6.3 Transmission channels

We finally address our final question to show that transmission for shocks to economic growth is much stronger through the financial variables than credit. Under high fragility, shocks to default risk act as an originator of variation, while shocks to default risk and profitability account for a greater share of variation than credit growth. We also find that "turning off" the fragility channel induces a greater reduction in the response of industrial production relative to credit, suggesting that fragility should be the primary concern for policymakers.

In general, it is difficult to make any causal claims about fragility. The scope of this paper has been to assess to the role of high fragility among intermediaries and examine the regime-dependent associations that arise. This is performed in the context of mixing two joint distributions that arise for different latent states of the system. In order to asses causality, we would have to work with a conditional distribution on exogenous perturbations (a theory for this is developed in Bazinas and Nielsen, 2015). We can, however, assess the relative importance of shocks and draw what inferences we can about the transmission channels.

An initial effort to accomplish this is to decompose the mean square forecast error generated by individual shocks. This is shown for all response variables and shocks in Figure 2.6.4. We make three observations. First, own-shocks (plotted along the diagonal) account for the almost the entire variation in industrial production, credit, and profitability under the low-fragility regime and significantly less under the high-fragility regime. Second, this effect is reversed for the TED spread. This indicates that variation in default risk is less dependent on shocks to the other variables than vice versa under high fragility. In this sense, we claim that default risk shocks are the primary source of variation. Thus, even though shocks to industrial production and credit explain a significant portion of the variation in default risk under low fragility, this changes under high fragility. Third, shocks to the health of financial intermediaries—defined jointly as default risk and profitability—account for approximately 30% of the variation in industrial production under high fragility (up from 5% under low fragility), while credit accounts for only 15%. These results are consistent with the notion that fragility behaves as a threshold variable: shocks are only significant beyond a certain threshold and insignificant below this threshold.

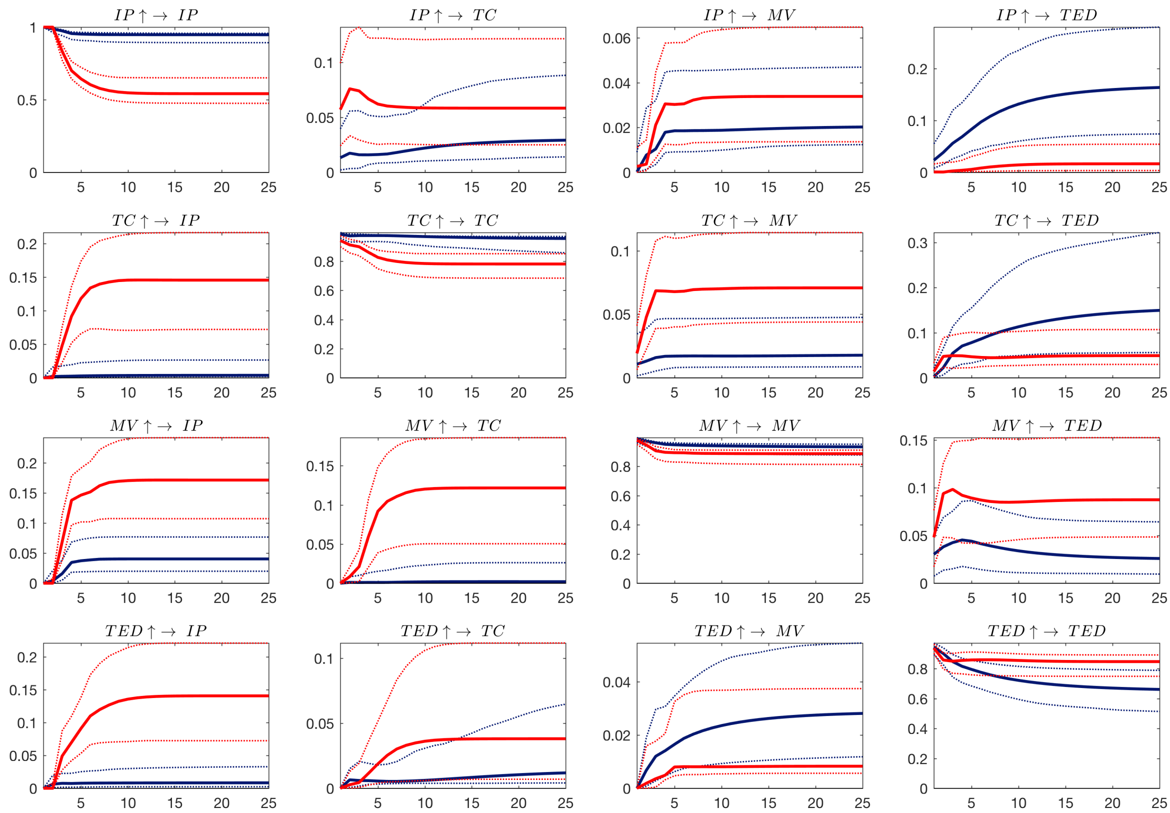

FIGURE 2.6.4: Regime-dependent forecast error variance decompositions. FEVDs for shocks of one standard deviation with 90% confidence intervals. Red: high-fragility regime IRFs. Blue: low-fragility regime IRFs. Vertical axis: percentage. Horizontal axis: horizon (months).

As a secondary effort, we compare shocks to fragility and shocks to credit growth by iteratively shutting off transmission channels. A shock to fragility is defined as a simultaneous standard deviation increase in the TED spread and a standard deviation decrease in profitability. Figure 2.6.5 plots the response of production growth to shocks in fragility and credit growth. In the first panel we also show the response to a shock in fragility when the credit channel is shut off (see dotted lines); the difference in responses is plotted in the second panel. Similarly, the third panel includes the responses to a credit shock when the fragility channel is shut off; the fourth panel plots the difference.

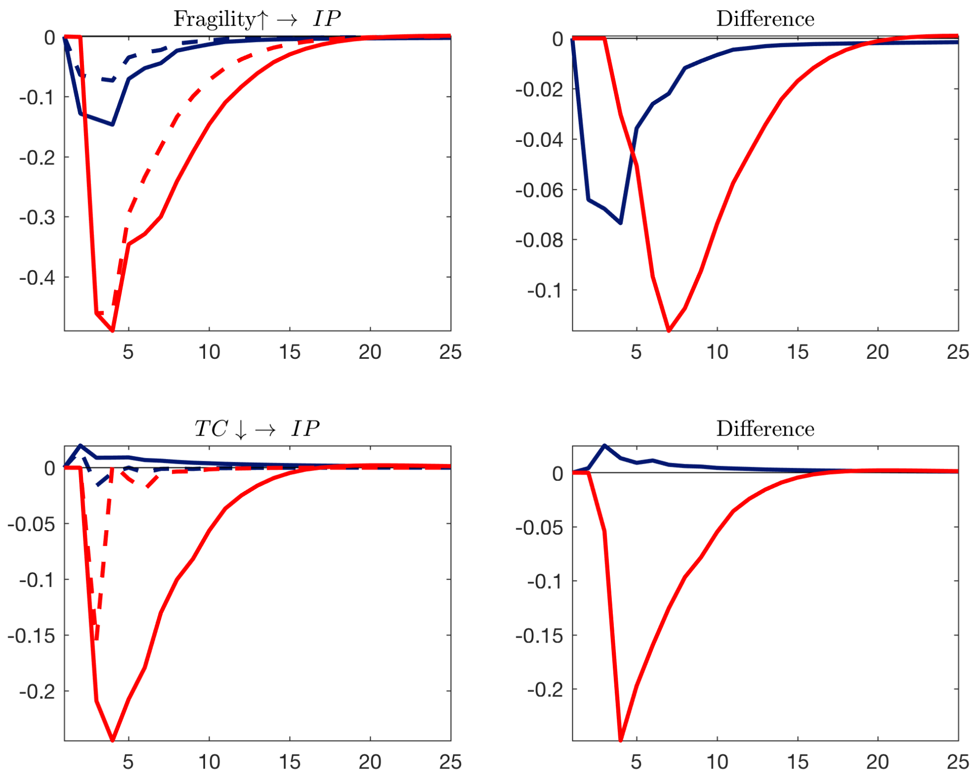

FIGURE 2.6.5: Regime-dependent impulse responses to financial fragility and credit growth. Responses in credit growth and financial fragility are sequentially turned off. Red: high-fragility regime IRFs. Blue: low-fragility regime IRFs. Scale is in percentage points.

In the high-fragility state, a shock to fragility directly affects economic growth, while a shock to credit affects economic growth through its effect on fragility. In the low-fragility state, this result is muted; in fact, shocks to fragility are relatively more likely to transmit to economic growth through the effect on credit.

How do we interpret these findings? In a sense, this is related to dissipating confidence in the high-fragility state. If under a regime of high fragility financial conditions actually worsen, this is likely to trigger an initial wave of defaults since profitability, which acts as a buffer to weather adverse shocks, is already very low. Alternatively, a negative shock to credit will have a similar effect: we are likely to observe relatively more speculative financing in the high-fragility state, so if suddenly credit is restricted, this will trigger a wave of defaults. The reason why this effect is so much larger in the high-fragility state is likely due to a debt-deflation spiral that is initiated after the initial default shock (Fisher, 1933; Minsky, 1992). Higher default on assets will reduce their value, leading to lower profits and therefore also lower default; this mechanism is the same if there is a negative shock to credit that initiates it. Therefore, an otherwise ordinary shock is amplified through the fragility channel, which effectively captures debt-deflationary effects.

From a different perspective, the information in Figure 2.6.5 captures the extent to which policymakers would have to intervene in order to counteract the worsening of financial conditions in the high-fragility state. Stepping in to provide credit in the high-fragility state when the private sector's intermediation activities are hampered is relatively less costly than bailing out intermediaries with weak buffers that are defaulting. Of course, we have not modelled policy variables and so these results are merely preliminary. Nonetheless, our results indicate that there is a role for unconventional monetary policy during times of financial distress (Mishkin, 2009; Gertler and Karadi, 2011). Our results also suggest that buffers accumulated counter-cyclically may mitigate adverse debt-deflationary effects in the high-fragility regimes (Goodhart et al., 2013).

2.7 Concluding remarks

We have demonstrated the existence of robust and fragile financing regimes within a Markov-switching framework that accounts for the nonlinear dynamics between the macroeconomy and financial intermediaries. High fragility is characterised by high default risk and low profitability among intermediaries, while both economic output and credit extension are substantially weaker. Our analysis showed that a linear model is strongly rejected, as it can neither accommodate the changing dynamics from intermediaries to the macroeconomy, nor the stochastic volatility that initiates primarily among intermediaries. We have shown that amplification of shocks during the high-fragility regime depends primarily on default risk and the buffers available to banks; the role of credit as an amplifying mechanism is confined to the low-fragility regime, but an intervention on credit in the high-fragility regime may be effective in reducing fragility.

We have not explicitly modelled policy responses in this paper. Further work is required in order to determine how the financial regulators and the monetary authority should respond and coordinate in order to weigh potential welfare losses during low-fragility episodes relative to potential welfare gains during high-fragility episodes. Preliminary evidence suggests that debt-deflationary effects may be mitigated in the high-fragility state if policymakers provide monetary policy support through credit extension when the private sector is unable to engage in its intermediation activities.

A feature that must be considered in tandem with policy responses, however, is modelling the transition between regimes. Within our Markov-switching framework, we have explicitly assumed that the transition between states follows an exogenous and discrete Markov process. This is an inherent difficulty when modelling nonlinear effects, as there does not exist a theory that allows for testing between various forms of nonlinearity. Further work within a different nonlinear setting would have to be able to partially reproduce the high-fragility regimes, while also describing how the macro-financial relationship evolves between regimes.

Finally, our reduced-form analysis is motivated by an underlying structural model that incorporates default as an endogenous result, but we have not attempted to discuss causality. This is particularly difficult given the nonlinear interactions between the macroeconomy and the financial sector. Our understanding of the macro-financial relationship based on this work can inform our choice for natural experiments that form the foundation of causal analysis.

Appendix

2.A Data

The data used in the main analysis are described in Table 2.B.1. Nominal total credit (TC) and total market capitalisation for the largest US financial intermediaries by asset size[6] (MV) are deflated using the consumer price index (CPI). An extended TED spread is constructed as the difference between the 3-month interbank lending rate (3INTB) and the 3-month Treasury Bill secondary market rate (3TBILL). All series except for the TED spread are log-differenced. The data used for estimation are plotted in Figure 2.A.1.

Table 2.A.1: Data series used in the analysis.

Series	Description	Source
AAA	Moody's seasoned Aaa corporate bond yield, percentage.	Moody's
BAA	Moody's season Baa corporate bond yield, percentage.	Moody's
CPI	Consumer price index, 1982–1984 = 100.	FRED
FFR	Effective Federal Funds rate, percentage.	FRED
IP	Industrial production index, 2012 = 100.	FRED
MV	Total market capitalisation for US banks, millions US dollars.	Thomson Reuters
TC	Total credit to the private non-financial sector adjusted for breaks, billions US dollars.	BIS
TED	TED spread, percentage.	Thomson Reuters
WXSR	Adjusted Federal Funds rate, percentage.	Wu and Xia (2016)
3INTB	3-month interbank lending rate, percentage.	FRED
3TBILL	3-month Treasury Bill secondary market rate, percentage.	FRED

Notes: All macroeconomic data are reported in levels and seasonally adjusted (X12ARIMA).

Some characteristics of the data are presented below. Since we are using data over

[6]We consider the seven largest banks in the United States by assets: JP Morgan Chase & Co., Bank of America Corp., Wells Fargo & Co., Citigroup Inc., Goldman Sachs Group Inc., Morgan Stanley, US Bancorp, and Capital One Financial Corp.

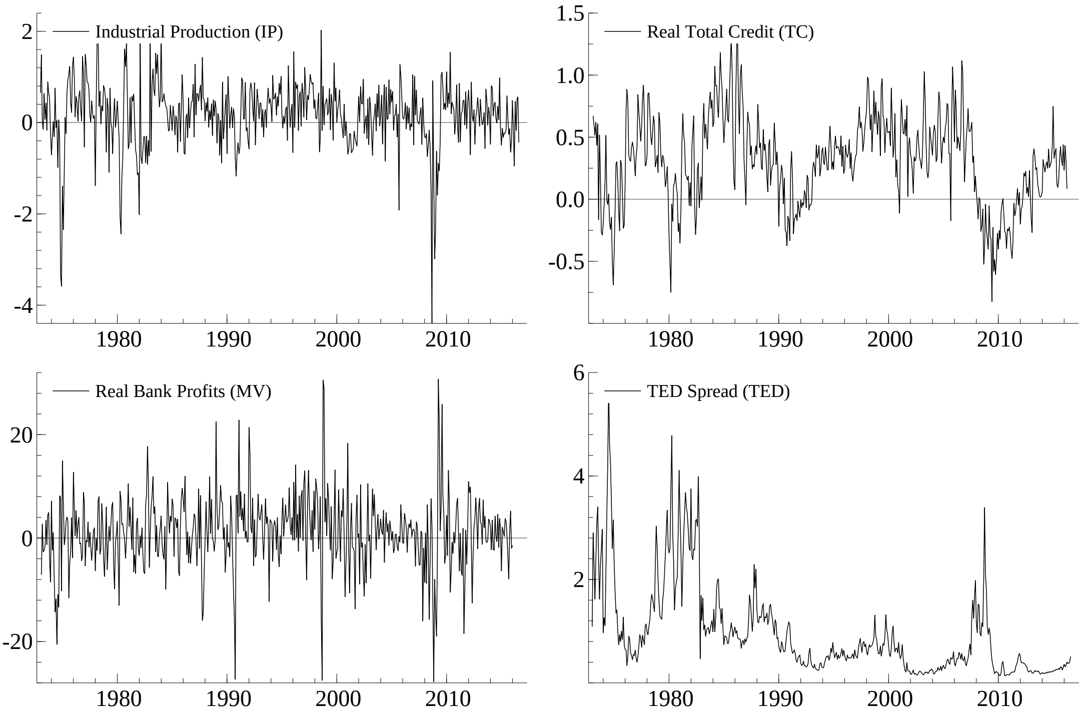

FIGURE 2.A.1: Data. All series are first-differences of log variables, except for the calculated TED spread. Vertical axis: percentage. Horizontal axis: years.

time, we must account for serial dependence in order to avoid spurious correlations. Severe serial dependence will manifest as a unit root, so that any series may vary around a stochastic trend. Figure 2.A.1 suggests that this is not the case, and that the data are most likely stationary. This is confirmed by testing. Thus, we must decide the appropriate number of lags to use in order to capture serial dependence. Figure 2.A.2 shows plots of the autocorrelation and partial autocorrelation function at different lags for each of the modelled time series. The plots suggest that at most 3 monthly lags are sufficient; this is confirmed by testing.

2.B Residual analysis

To ensure that statistical inference is valid, we conduct several residual specification tests. The residual plots are shown in Figure 2.B.1, while the specification tests

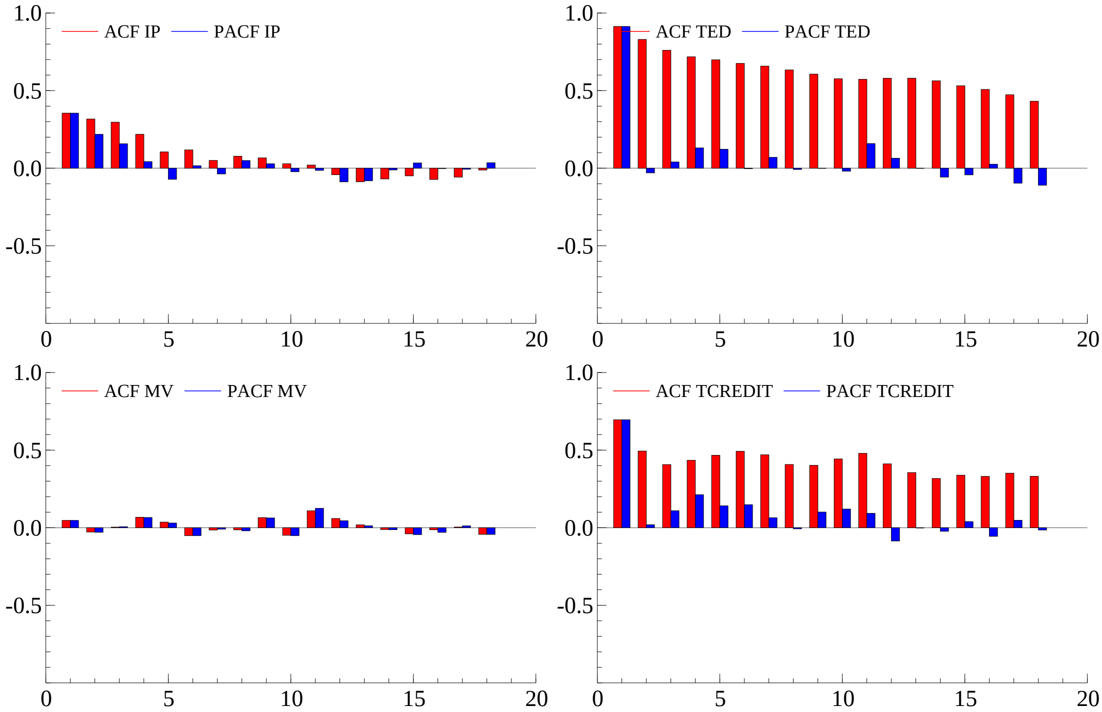

FIGURE 2.A.2: Autocorrelation and partial autocorrelation functions. Vertical axis: correlation. Horizontal axis: lag length. Red bars: autocorrelation at indicated lag. Blue bars: partial autocorrelation at indicated lag.

are provided in Table 2.B.1. Normality is assessed via the Kolmogorov-Smirnov test (Massey, 1951), autocorrelation is assessed using Godfrey (1978), and conditional heteroskedasticity is assed using Engle (1982). A difficulty that we face, however, is that these tests—except for Kolmogorov-Smirnov—rely on linearity. We are using these in a nonlinear setting, and so results must be interpreted carefully; this is why we also examine graphical diagnostics for normality in Figures 2.B.2-2.B.3 and for autocorrelation in Figure 2.B.4.

Table 2.B.1 reports the results for individual and vector tests for the residuals of each equation in our Markov-switching vector autoregression; we note that there is no vector version of the Kolmogorov-Smirnov test. The results indicate that normality is a tenable assumption in each equation. There is some evidence for autocorrelation among the credit residuals, as well as some evidence for conditional heteroskedasticity

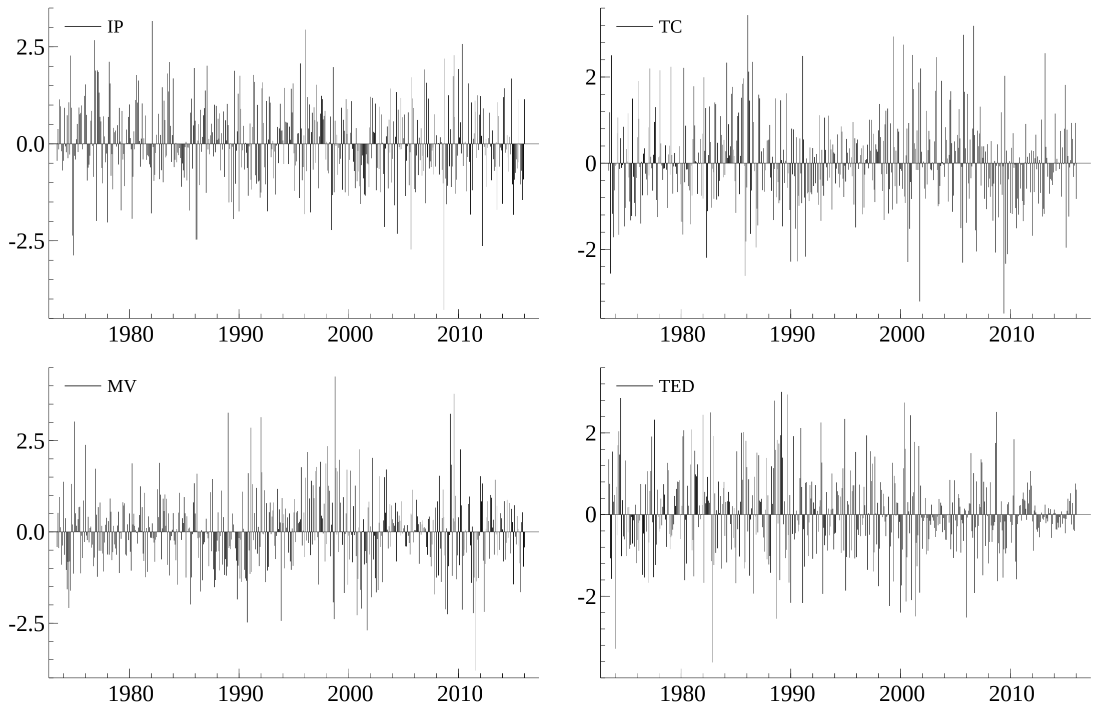

FIGURE 2.B.1: Scaled residuals. The scaled residuals for model M1 are displayed. Vertical axis: scaled residuals. Horizontal axis: years.

in the credit and TED spread residuals. Upon inspection of the credit residuals in the second panel of Figure 2.B.1, we can observe some persistence, as well as a large number of potential outliers; the residual plot for the TED spread in the fourth panel also suggests there may be several outliers. We correct for this by calculating robust standard errors and examine further diagnostic plots.

Figure 2.B.2 shows the residual densities against the standard normal, while Figure 2.B.3 shows the cross-plot of the quantiles for the residual and standard normal densities. In both instances, we find that despite some outliers, the assumption of normality in the residuals cannot be rejected. Furthermore, Figure 2.B.4 shows that the partial autocorrelations are not significant, and that the issues with total credit are likely overstated in our specification tests. Based on these outputs, we can conclude that standard inference will be valid in our main analysis.

Table 2.B.1: Specification tests. p-values are reported in brackets.

Equation	KS	Autocorrelation		ARCH
		$\chi^2_{\mathsf{ar}(1-3)}$	$\chi^2_{\mathsf{ar}(1-12)}$	$\chi^2_{\mathsf{arch}(1-12)}$
IP	0.030 [0.72]	1.1 [0.77]	9.5 [0.66]	24.5 [0.02]
TC	0.059 [0.05]	17.8 [0.00]	74.9 [0.00]	25.7 [0.01]
MV	0.058 [0.06]	0.2 [0.98]	14.2 [0.29]	19.6 [0.08]
TED	0.065 [0.02]	7.4 [0.06]	22.1 [0.04]	49.7 [0.00]
Vector		39.6 [0.80]	267.8 [0.00]	1.1 [0.29]

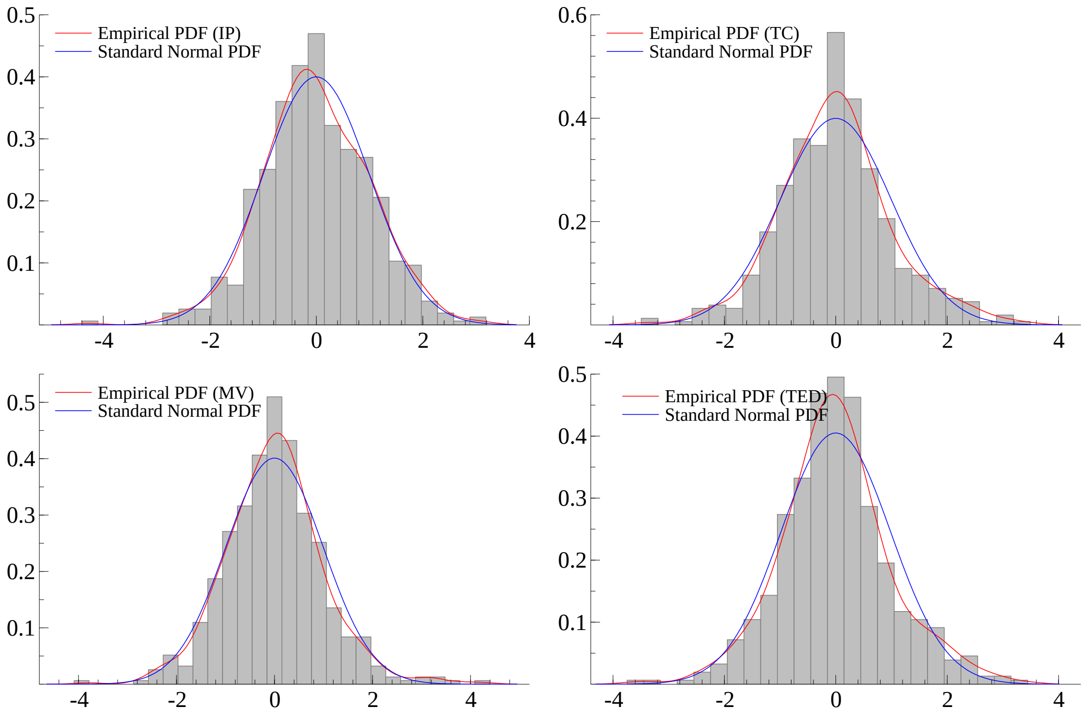

Figure 2.B.2: Residual density plots. The empirical probability density functions of the residuals from model M1 are compared to a standard normal density. Vertical axis: probability density. Horizontal axis: residual bins.

2.C Robustness checks

2.C.1 Monetary policy

Our sample covers the period 1973–2016, over which monetary policy fluctuated significantly. Fluctuating short-term interest rates can impact the ability of interme-

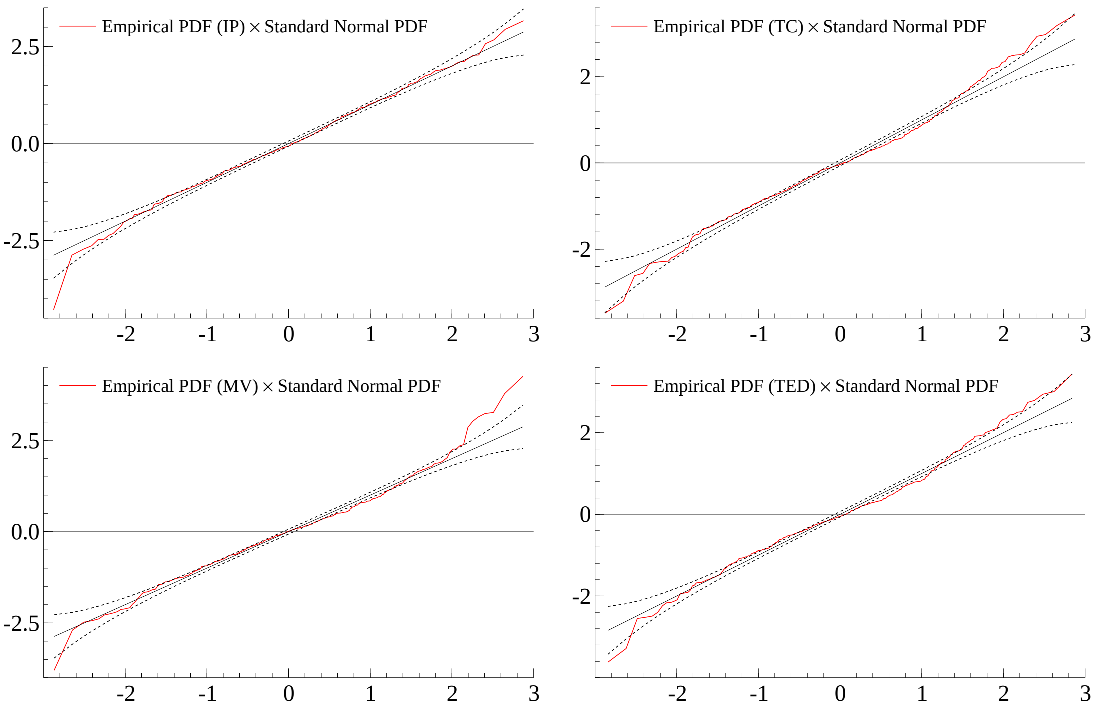

FIGURE 2.B.3: Residual density plots. The quantile-quantile plot for the empirical distribution against the standard normal is displayed along with a 95% confidence interval. Vertical axis: quantiles of residual density. Horizontal axis: quantiles of standard normal density.

diaries to repay their obligations and to continue their credit-extending operations. Additionally, Miranda-Agrippino and Rey (2015) have suggested that US monetary policy drives a global financial cycle with significant cross-border effects on industrial production, credit growth, and asset prices. We conduct a robustness check in order to ensure that monetary policy does not drive regime-switching in our model

We assume that any targets set had an effect on the Federal Funds rate, and so we use this as the measure of policy. In order to account for unusual policy measures at the zero lower bound, we use the shadow rate of Wu and Xia (2016) that uses factor methods to infer where the policy rate would be if there were no lower-bound. Figure 2.C.1 shows the smoothed regime probabilities from model M1 (see Table 2.5.1) for the high-fragility regime and compares these to the similar regime selected in a model

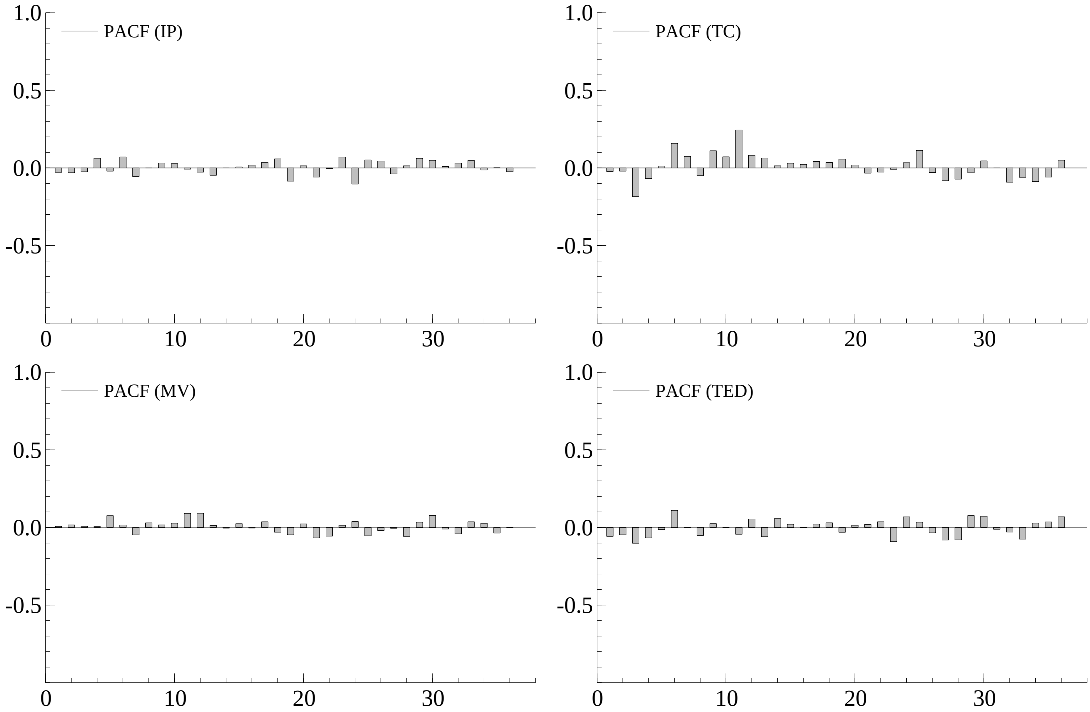

FIGURE 2.B.4: Autocorrelation function plots. The partial autocorrelations for the residuals from model M1 are displayed. Vertical axis: correlation. Horizontal axis: lag length.

that includes monetary policy as a replacement for financial fragility. The objective is to determine whether monetary policy explains the fragility episodes described in this paper.

We make the following observations. First, monetary policy explains only a subset of high-fragility episodes. The main episodes that are captured are the aftermath of the Nixon shock in 1972 and the collapse of Bretton-Woods, the Volcker-era combatting of inflation in the 1980s, and the response to the Global Financial Crisis in 2007-09. Elsewhere, we find outliers in 1984, 1987, and 2005. Regime-selection with monetary policy misses several important episodes of fragility, including the culmination of the LDC debt crisis in 1984, the S&L crisis in the late 1980s, and the collapse of LTCM. Second, except at the beginning of our sample, policy regimes occur several months after the beginning of high-fragility regimes. This may suggest an

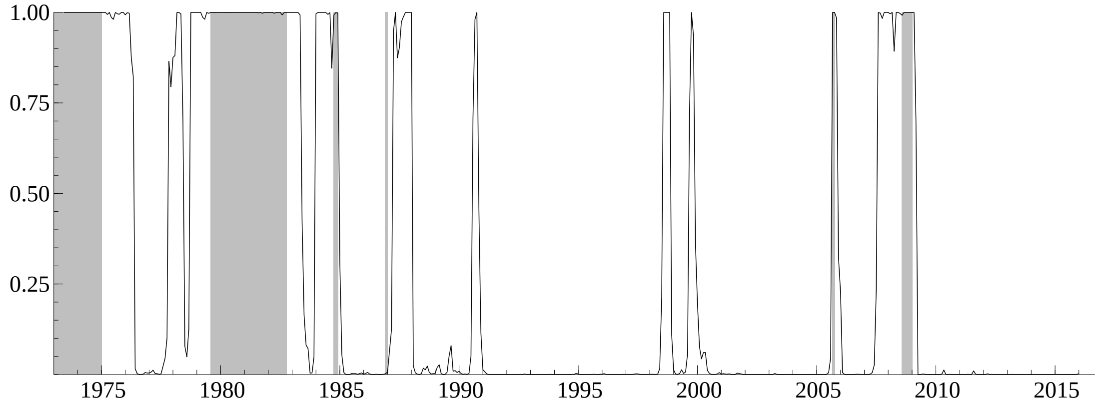

FIGURE 2.C.1: Regime probabilities. Black line: smoothed Regime 1 probabilities in model M1. Grey areas: corresponding Regime 1 classification from an alternate model replacing the financial variables with the adjusted Federal Funds rate from Wu and Xia (2016). Vertical axis: probability. Horizontal axis: years.

endogenous response to underlying factors that also affected the buildup of fragility in the financial sector. Nonetheless, this fragility is detectable several months prior to the response in monetary policy. Further work is required to determine the causal relationship between fragility and monetary policy.

2.C.2 Alternative measures of default risk

The TED spread reflects default risk specifically among financial intermediaries and is often used as a component of financial stress indicators. An alternative measure of risk is the spread in Moody's Aaa and Baa corporate bond yields. If Aaa bonds reflect relatively riskless asset with the generally the same characteristics as Baa bonds, then the spread in yields will reflect the price for higher default risk among the lower-rated bonds. This measure of default risk considers all corporates and not only financial intermediaries, so it will indicate to what extent systemic risk permeates through to the non-financial sector. Figure 2.C.2 shows the smoothed regime probabilities for model M1 (see Table 2.5.1) for the high-fragility regime and compares these to the similar regime selected in a model where the Moody's corporate bond spread proxies for default risk.

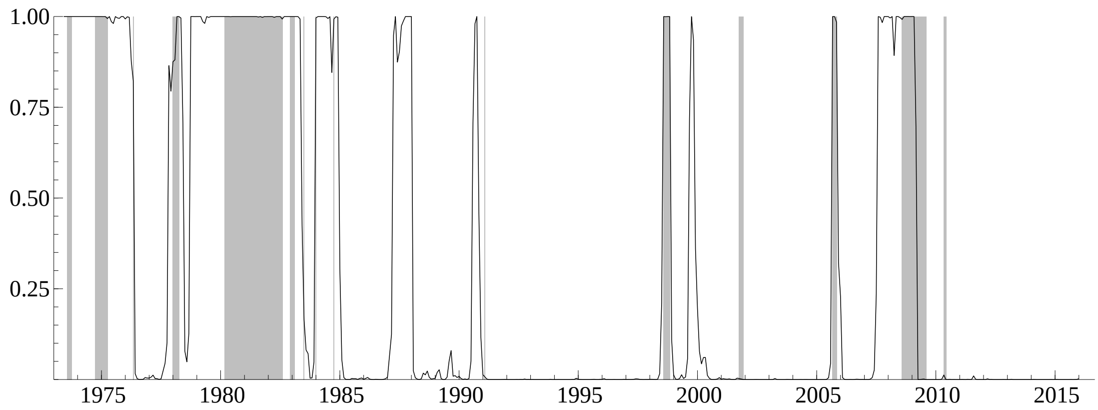

Figure 2.C.2: Regime probabilities. Black line: smoothed Regime 1 probabilities in model M1. Grey areas: corresponding Regime 1 classification from an alternate model replacing the TED spread with the spread on Aaa/Baa corporate bonds from Moody's. Vertical axis: probability. Horizontal axis: years.

The corporate bond spread does not capture fragility in the financial sector to the same extent as the TED spread. It does, however, give an indication of when high fragility among intermediaries may have translated into instability for non-financials. As a rule, this occurs during all major recessions, as well as during the LTCM episode in 1998. Further work might explore this relationship more carefully.

2.C.3 The contribution of default and profitability to fragility

As an additional way of disentangling the contributions of the macroeconomic and the financial variables to the regime selection, we carry out two further experiments. In the first case, we re-estimate the model for the same sample and number of regimes but removing the macroeconomic variables (industrial production growth and real total credit growth). Thus, the model becomes simply a bivariate system for the TED spread and the market valuation growth for intermediaries. The smoothed regime probabilities from this experiment are shown in Figure 2.C.3, along with the corresponding results from the main model M1. Most of the high-fragility regimes

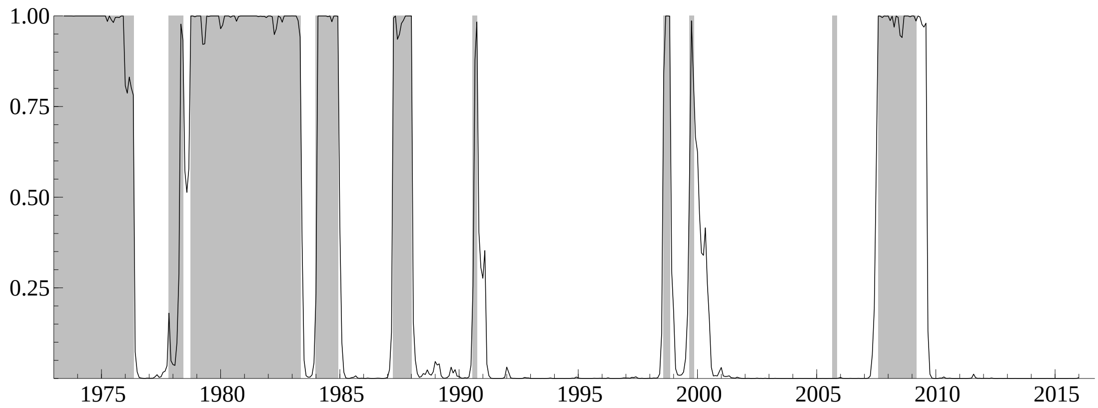

FIGURE 2.C.3: Regime probabilities. Black line: smoothed Regime 1 probabilities from a model comprised only of the TED spread (default risk) and log-differenced MV (profitability). Grey areas: corresponding Regime 1 classification in model M1. Vertical axis: probability. Horizontal axis: years.

depend on the financial variables—default risk and profitability—which we view as a desirable feature for our empirical model. This is reflected in Figure 2.C.3 by the fact that the Regime 1 probabilities from the main model M1 correspond broadly to the Regime 1 probabilities in this simplified model. Two further comments on this are necessary to characterise the outliers. Periods in Regime 1 under the main model M1 that are not selected in this experiment include late 1977 to early 1978 and late 2005.

Additionally, there are some periods in this simplified model that fall under Regime 1 that could potentially be classified as episodes of high fragility but are not included in the main model. These include periods in late 1978, early 1991, early 2000, and early to mid 2009 where the probability of Regime 1 rises in each case. A common feature is that these dates correspond to sharp improvements in the macroeconomic variables either before or after a recession, thereby lowering the probability of high fragility in the main model. When these movements are removed, the information contained in the TED spread and growth in market valuation of intermediaries suggests a highly fragile regime.

We explore this further by reversing the experiment: for the same sample, we simplify the main model by removing the financial variables. The estimated regime probabilities are shown in Figure 2.C.4, along with the corresponding probabilities from the main model M1. The simplified model in this case captures all adverse

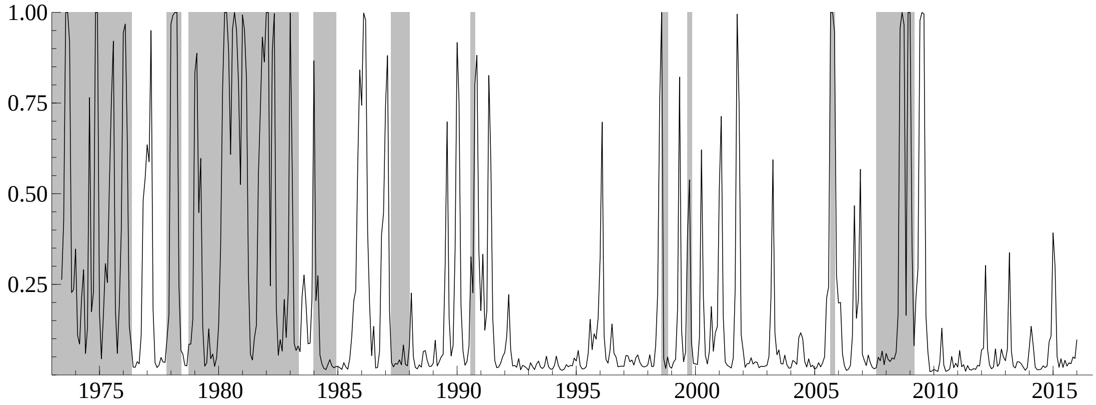

FIGURE 2.C.4: Regime probabilities. Black line: smoothed Regime 1 probabilities from a model comprised only of the log-differenced IP (economic growth) and log-differenced TC (credit growth). Grey areas: corresponding Regime 1 classification in model M1. Vertical axis: probability. Horizontal axis: years.

movements in the macroeconomic variables. Two regimes that are very clearly driven by the macroeconomic variables include late 1977 to early 1978, as well as late 2005; recall that these were not selected in the previous experiment. Both of these periods correspond to large declines in industrial production—in the first case likely as a result of the 1970s oil shocks and in the second case as an unanticipated downturn. It is more difficult to attribute other regimes to movements in macroeconomic or financial variables and require further study in order to attribute causality.

2.D Credit in the United States

Schularick and Taylor (2012) carry out an analysis using annual data on credit, money, and economic activity between 1870-2008 to demonstrate that financial crises—constructed as a dummy variable for 1984 and 2007—are best predicted by lagged credit growth, with other variables adding little explanatory power. Citing a decoupling of broad money to credit after 1945, they split their analysis into pre- and post-war subsamples.

A closer examination shows that their results may not necessarily hold for the US as a stand-alone case. Table 2.D.1 repeats their estimation for the post-war period, using data on credit and prices extended to 2013 from Jordá et al. (2016). The binary dependent variable is alternatively specified as the financial crisis dummy of Reinhart and Rogoff (2009) used in Schularick and Taylor (2012), and as a dummy variable for regimes of high financial fragility, taking a value of 1 when more than 3 months of high fragility occur in a calendar year; the former is labeled as specification 1, while the latter as specification 2. We also use five lags of real credit, as in Schularick and Taylor (2012).

Table 2.D.1: Reproducing the analysis in Schularick and Taylor (2012) for the individual case of the United States. Financial crises defined in Reinhart and Rogoff (2009) and high-fragility episodes identified in this paper are modelled against lagged real credit growth over the period 1951-2013.

Estimation Method Specification	OLS (1)	Logit (1)	Probit (1)	OLS (2)	Logit (2)	Probit (2)
$\log(loans/P)_{t-1}$	0.373 (0.497)	13.030 (19.165)	4.871 (8.241)	0.164 (1.136)	-0.798 (7.960)	-0.254 (4.489)
$\log(loans/P)_{t-2}$	-0.447 (0.546)	-18.148 (24.090)	- 7.622 (10.417)	-0.184 (1.247)	-0.495 (8.824)	-0.493 (4.908)
$\log(loans/P)_{t-3}$	0.240 (0.540)	9.494 (13.438)	5.366 (6.727)	0.474 (1.235)	2.924 (7.499)	1.412 (4.351)
$\log(loans/P)_{t-4}$	-0.582 (0.543)	-21.671 (23.021)	-9.381 (9.980)	-0.616 (1.241)	-2.823 (7.956)	-1.897 (4.555)
$\log(loans/P)_{t-5}$	-0.076 (0.462)	- 0.122 (23.021)	-0.229 (8.009)	-1.490 (1.055)	10.838 (7.634)	-6.306 (4.312)
1	0.056 (0.037)	-3.377^{**} (1.138)	-1.784^{**} (0.487)	0.293^{**} (0.084)	-0.8584 (0.4885)	-0.505 (0.286)
Test statistic	0.551	3.065	2.780	0.801	4.401	4.596
p-value	0.737	0.689	0.734	0.554	0.493	0.467
Observations	62	62	62	62	62	62

Notes: Standard errors reported in parentheses. Test statistic for significance of lagged terms is F for OLS and χ^2 for Logit and Probit. [**] Significant at the 1% level. [*] Significant at the 5% level. Data taken from Jordá et al. (2016).

Since the dependent variable is binary, it can be estimated most appropriately within a generalised linear model (GLM) that includes either a logit or probit link function. We provide both, as well the OLS linear probability model. The test statistic for the OLS model is a joint test of the significance of the coefficients on the covariates and fails to reject. For the logit and probit models, the test statistic is based on the difference between the null and residual deviances. In all of the models, lagged real credit cannot explain financial crises in the US, where financial crises are found to occur only in 1984 and 2007. Additionally, there is no evidence to suggest that real credit growth alone can also explain the regimes of high financial fragility identified in the present paper.

CHAPTER 3

Assessing Time-Varying Macro-Financial Linkages

Chapter Abstract

We assess the linkages between the real economy and the financial intermediation sector over the period 1999–2015 within a time-varying parameter framework with stochastic volatility. Low volatility and a period of protracted economic growth initiate a feedback loop between the real economy and the financial sector, consistent with the financial instability hypothesis of Minsky (1986). Shocks that adversely affect expectations of future economic growth have the greatest impact when macro-financial linkages are strongest. Augmenting the model with macroprudential capital requirements and monetary policy adjusted for the zero lower bound, we provide evidence that the timing of policy interventions matters. Counter-cyclical interventions produce a stabilising effect when implemented early but are also capable of exacerbating debt-deflation if employed too late in the financial cycle.

3.1 Introduction

The experience of the Great Recession in 2007–2009 underscored the need for models that are able to capture the evolving dynamics between the financial and business cycles. Furthermore, policymakers are increasingly responsible for using available instruments not only for stabilising the real economy but the financial sector as well. An understanding of the evolution of macro-financial linkages is critical in knowing when and how to intervene.

There is a growing appreciation of the time-variation present within the macro-financial relationship and the role of financial shocks (Prieto et al., 2016; Gambetti and Musso, 2017), as well as the evolving role that monetary policy plays (Galì and Gambetti, 2015). The importance of time-varying shocks to house prices and credit for systemic risk are acknowledged and often treated separately, but these are never studied jointly. This misses an amplification mechanism that is crucial to policymakers. Additionally, no time-varying role is assigned to macroprudential requirements, though there is broad consensus for counter-cyclical measures within constant-parameter empirical studies (for a recent survey, see Galati and Moessner, 2017).

We estimate a statistical model that at once considers the aggregate behaviour of the financial intermediation sector, while allowing for time-varying macro-financial linkages and ask the following questions. First, how does optimism affect the accumulation of leverage in the upswing of the financial cycle? Second, does the importance of adverse shocks to expectations matter more after periods of protracted economic growth? Finally, can interventions in the form of macroprudential capital requirements and monetary policy successfully mitigate losses following the accumulation of risk?

To gain answers to these questions, we employ a Bayesian macroeconomic vector autoregressive model for the United States over the period 1999–2015. Our model differs from the standard setup by incorporating a robust measure of systemic risk in the financial intermediation sector that accounts for nonlinear dependence between institutions (Goodhart and Segoviano, 2009), while also modelling both house prices and credit. Given our policy-driven objectives, we employ the shadow policy rate of Wu and Xia (2016) to account for monetary policy at the zero lower bound and an aggregated tier 1 capital ratio as a proxy for macroprudential regulation. Thus, our VAR model is comprised of industrial production growth, house price inflation, private credit growth, systemic expected shortfall of financial intermediaries, the aggregate tier 1 capital ratio, and an adjusted Federal Funds rate.

The main results are as follows. First, we find evidence of exceptionally low volatility during the credit buildup that preceded the Great Recession, consistent with the volatility paradox of Brunnermeier and Sannikov (2014). Second, positive shocks to economic growth initiate a feedback loop with credit and housing prices over the period spanning 1999–2006. This occurs because shocks that are interpreted as positive signals of economic performance feed into agents' future expectations, thereby driving over-leveraging. Third, positive signals are found to reduce the repayment rate among financial intermediaries consistently throughout the sample. Fourth, "not unusual" shocks (Minsky, 1992) that reduce the repayment rate have a much more severe effect when macro-financial linkages are high and the financial sector is in a state of fragility. Fifth, we find that policy shocks are most effective when employed early in the financial cycle, which is consistent with the theoretical literature (Goodhart et al., 2013). Attempting to deflate asset prices and check credit growth either through macroprudential or monetary policy at the peak of the financial cycle can have severe adverse effects (Lin et al., 2015).

We differ from most studies of the time-varying macro-financial relationship by studying house prices, credit growth, and systemic risk jointly. This is important in accounting for evolving attitudes toward risk (Minsky, 1986) that are amplified through net worth (Bernanke et al., 1999) and collateralised lending (Kiyotaki and Moore, 1997). This helps us evaluate alternative channels of monetary policy through effects on financial stability (Borio and Zhu, 2012; Lin et al., 2015). We are also the first to study macroprudential requirements within a time-varying parameter framework. Policy identification is achieved through a recursive assumption on the contemporaneous effects of shocks (Sims, 1980). All macro-financial variables respond with a lag to policy shocks, an assumption that is standard in the literature (Christiano et al., 1999; Stock and Watson, 2001; Galì and Gambetti, 2015).

Our results are consistent with the leverage cycle as described in Kiyotaki and Moore (1997) and Geanakoplos (2010), whereby increasing asset prices lead to higher collateral values, which in turn permits greater credit extension. The effects are highest at the peak of the leverage cycle following protracted periods of economic growth, while lowest following the crash and during recession. This is consistent with the Minsky (1977, 1986, 1992) hypothesis formalised in Bhattacharya et al. (2015). Observing successive periods of growth, agents revise their expectations of a strong economy upward, assigning a higher probability to continued growth. This leads to higher leverage that is employed to finance increasingly riskier projects. Following an adverse shock, the economy reacts much more severely than it would have, had there not been a buildup in leverage. We see this in our results by noting the muted effect of shocks to systemic risk and asset prices following the Great Recession. We argue that this occurs because the financial crisis corrected burgeoning expectations of continued economic growth.

The paper is organised as follows. The following section examines the relevant literature, both theoretical and empirical, in assessing the nonlinear relationship between financial and macroeconomic variables. §3 describes Minsky's financial instability hypothesis and presents a formal framework within which protracted expansions can affect attitudes toward risk. §4 describes the econometric methodology in a generic way, while being explicit about the modelling assumptions. The data and the construction of our nonlinear expected shortfall measure are described in §5. §6 presents our results on the transition between financially robust and fragile regimes, while §7 discusses the role of policy interventions. The final section provides some concluding remarks and areas for further research.

3.2 Related literature

Our work combines elements from different areas of applied macroeconomic work and adds to the small but growing empirical literature on macro-financial linkages in a time-varying setup. We also attempt to address potential structural nonlinearities among financial intermediaries by incorporating an appropriate measure of systemic risk that emphasises tail-dependence between institutions. Finally, the effects of macroprudential and monetary policy on financial stability have been studied extensively in recent years, but, to the best of our knowledge, never jointly within a time-varying parameter setup. The relevant literature on time-varying parameter models, the measurement of systemic risk, and the effects of macroprudential and monetary policy on financial stability is briefly surveyed in this section.

Time-varying parameter applications. The aftermath of the Great Recession has spurred the study of the macro-financial relationship within frameworks that allow for time-variation in parameters, especially over the last half-decade. The frameworks can be broadly separated into those with smooth transitions (Primiceri, 2005) and those with discrete transitions (Tong, 1983; Hamilton, 1989). Our work falls into the former category.

The application of time-varying parameter vector autoregressions (TVP-VAR) to the macro-financial relationship is very new, while most of the literature has focused on data from the United States. Prieto et al. (2016) study the effects of shocks in house prices, stock prices, and the corporate credit spread on GDP. They find that shocks to house prices had significant predictive power during the Great Recession, and the subsequent weak recovery is attributed to sluggish credit growth, though credit is not explicitly modelled. The effects of credit supply shocks are examined in Gambetti and Musso (2017), where identification is achieved through sign restrictions on the responses in the macroeconomic variables. They find evidence for time-variation in the effect of credit shocks that are particularly significant during the Great Recession. Elsewhere in the literature, the importance of financial shocks and international spillovers are studied using a time-varying factor-augmented VAR (Abbate et al., 2016) and using a time-varying panel VAR (Ciccarelli et al., 2016).

Models with discrete transition between regimes are also employed to study the importance of financial shocks on macroeconomic variables. This approach is possibly better-suited to capturing dynamics where changes occur rapidly. Recent contributions using data from the United States include Hubrich and Tetlow (2015) who use a financial conditions index embedded within a Markov-switching VAR (MS-VAR) to find evidence for variation in shock volatility and in the transmission mechanism. A longer sample is studied in Nason and Tallman (2015) who also employ an MS-VAR but only allow for breaks in the intercept and stochastic volatility. They find that time-variation in shocks to credit demand and supply are able to explain aggregate fluctuations, while the transmission mechanism remains the same. Balke (2000) studies the effects of credit shocks within a threshold VAR (T-VAR) to find that credit shocks have a larger effect on output when credit conditions are tight. In a more recent study for the United States, Aikman et al. (2016) use a T-VAR over 1975–2014 to show that financial imbalances in credit, risk appetite, and short-term funding can each lead to an agitated state where shocks negatively impact economic growth.

Throughout the literature on time-varying parameter models, the importance of changing shock volatility is a common feature. Additionally, the importance of financial shocks in the form of house prices, credit growth, or proxies for systemic risk/financial conditions is demonstrated but these are never examined jointly. Our approach incorporates all of these features in order to examine the feedback loops inherent in macro-financial linkages.

Policy for financial stability. The study of monetary and macroprudential policy conducted for the purpose of financial stability is relatively new, having grown mostly out of the experiences of the past decade and conducted primarily in response to the Global Financial Crisis. As such, empirical analysis is compounded with the difficulty that there is a dearth of policy interventions in the data, particularly with respect to macroprudential interventions. In many respects, both policy interventions have become intertwined. In what follows, we briefly survey the theoretical and applied literature on macroprudential and monetary policy intervention with the goal of stabilising the financial sector. For a detailed survey of economic models and empirical work in which the role of policy is assessed, see Galati and Moessner (2017).

The need for intervention arises due to the failure of agents, usually financial intermediaries, to internalise the general equilibrium effects of their decisions on asset prices. This results in over-leveraging during normal times that becomes socially expensive during bad times through a credit crunch that initiates a negative feedback loop between debt and asset prices. Various efforts have incorporated these dynamics in economic models, alternatively by introducing endogenous default on loans (Goodhart et al., 2012), endogenous default on collateral (Geanakoplos and Fostel, 2008), time-varying expectations (Bhattacharya et al., 2015), occasionally-binding borrowing constraints (Bianchi and Mendoza, 2010; Benigno et al., 2013), and financial frictions in continuous time (Brunnermeier and Sannikov, 2014).

The broad consensus is to intervene early: by curbing the boom in credit growth during normal times, policymakers can mitigate the fallout from the bust. Goodhart et al. (2013) detail various macroprudential tools that can be employed in order to improve welfare. These include loan-to-value ratios, dynamic provisioning, and capital/margin requirements for intermediaries. Implementation of these tools can address the pro-cyclical development of systemic risk along the time dimension. Nonetheless, we must also be mindful of the development of systemic risk along the structural dimension through interconnectedness of intermediaries. Horváth and Wagner (2013) show that while countercyclical capital requirements can insulate intermediaries from aggregate shocks, this increases their exposure to idiosyncratic shocks and creates an incentive for correlated exposures, thereby increasing systemic risk.

How well have macroprudential tools worked in practice? The effects of macroprudential regulation are typically examined within a cross-country framework in reduced-form panel regressions. Dell'Ariccia et al. (2016) find that macroprudential tools dominate monetary policy, as they are able to build counter-cyclical buffers that limit losses from busts. Cerutti et al. (2017) and Claessens et al. (2013) both find evidence for asymmetric effects during booms and busts. Macroprudential tools are able to manage the upswing of the financial cycle by limiting bank leveraging and asset growth, but they are not as effective during downturns. Along the structural dimension, Gauthier et al. (2012) use Canadian bank data to find that macroprudential measures are able to reduce both individual banks' default probability, as well as the probability of a systemic crisis.

A related question is the importance of monetary policy in addressing financial stability concerns. Mishkin (2009) argues that monetary policy is more potent and less inertial during financial crises, so that aggressive policy can mitigate the effect of feedback loops. In a similar vein, Gertler and Karadi (2011) show that unconventional monetary policy—interpreted as the central bank expanding credit when private intermediaries are unable to do so—can provide significant benefits, particularly at the zero lower bound. On the other hand, monetary policy can also affect the upswing of the financial cycle. Borio and Zhu (2012) discuss a "risk-taking channel" of monetary policy, whereby lower interest rates affect perceived risk and risk tolerance through inflated asset values and higher profits. Expansionary policy in the early upswing of the financial cycle can raise asset prices and profits, thereby lowering (perceived) probabilities of default; this creates a positive feedback, encouraging risk-taking and investment further. Contractionary monetary policy used to prevent over-leveraging later in the upswing can initiate the "debt-deflation channel" of monetary policy described in Lin et al. (2015). A contractionary shock makes it more difficult to re-finance collateralised, long-term loans, so that leverage on these increases. In the following period, the price of collateral falls as there is too much money chasing capital goods, leading to default that initiates the debt-deflation spiral.

Empirical evidence for the risk-taking channel of monetary policy is found in Dell' Ariccia et al. (2017). Using data on internal risk ratings of new bank loans, they find that a decrease in the policy rate leads to high loan risk ratings. Elsewhere, Schularick and Taylor (2012) find that more aggressive monetary policy in the postwar era is likely responsible for preventing debt-deflation spirals on the scale of the Great Depression. Further empirical results are described in the survey articles by Crowe et al. (2011) and De Nicolò et al. (2010) among others.

Attempts to simultaneously address macroprudential and monetary policy considerations within a constant-parameter VAR framework are sparse. These include Kim and Mehrotra (2015) and Glocker and Towbin (2012) who incorporate an aggregated macroprudential policy index and reserve requirements, respectively, along with a monetary policy variable in order to find that both types of interventions tend to limit credit growth. Concerns on endogeneity and reverse causality are addressed through the use of a structural VAR framework for individual countries in both of these studies. In a slightly different vein, Federico et al. (2012) identify exogenous shocks to reserve requirements through the historical narrative approach of Romer and Romer (1989). When identification is carried out in this way, macroprudential policy is found to substitute rather than complement monetary policy.

Macroprudential policy has not been studied within a time-varying parameter framework. Additionally, and to the best of our knowledge, no other study has also accounted for monetary policy at the zero lower bound in such a framework. We include both of these features, since the empirical and theoretical literature suggest that policy interventions will vary over the financial cycle.

Systemic risk measurement. Another question that we address in this paper is the measurement of systemic risk, yet another field spurred on by the experience of the GFC. Systemic risk can result from the pro-cyclicality of default probabilities, for instance, or due to increased interconnectedness among intermediaries; the former is referred to as the "time" dimension, while the latter is referred to as the "structural" dimension. Risk along the time dimension is generally captured by financial stability indicators, such as (time-varying) weighted averages of interest rates and spreads (Hubrich and Tetlow, 2015; Kremer et al., 2012). Other measures examine cross-sectional dependence (Huang et al., 2011; Acharya et al., 2012; Adrian and Brunnermeier, 2016) and networks (Billio et al., 2012; Diebold and Yilmaz, 2014) between individual institutions. Recognising that such linear measures may not pick up the nonlinear dependence during periods of distress, Goodhart and Segoviano (2009) introduce a portfolio-based measure of joint default probability that takes into account stronger co-movement in the distribution tails. Using CDS spreads to determine the marginal default densities, their measures rely on the method of Segoviano (2006) to appropriately account for dependence that varies nonlinearly when institutions reach their default threshold. A survey of these and other measures is available in Bisias et al. (2012).

We employ the method in Goodhart and Segoviano (2009) in order to account for nonlinear structural dependence among financial intermediaries in addition to the pro-cyclicality picked up along the time dimension.

Drawing on different areas of the macroeconomic literature, we develop a time-varying parameter model with stochastic volatility. In addition to this, we control for and examine the effects of macroprudential and monetary policies on financial stability by including a systemic risk measure that accounts for structural and temporal effects.

3.3 The transition between financing regimes

In this section, we motivate the need for time-variation in our empirical framework. The relationship between the macroeconomy and the financial sector is historically characterised by periods of stress (Kindleberger and Aliber, 2011). Minsky (1992) provides a theory for the development of periods where the financial sector is particularly susceptible to amplifying adverse shocks, while a formalisation of this theory is provided in Bhattacharya et al. (2015). We briefly summarise the main ideas in this line of research, in order to demonstrate that a tractable sequence of events can precipitate a transition from a robust to a fragile financing regime.

3.3.1 The financial instability hypothesis

Minsky's theory of financial instability (developed in Minsky, 1977, 1986, 1992) revolves around two hypotheses: first, that there exist robust and fragile financing regimes across which the real economy behaves differently; second, that the transition from a robust to a fragile financing regime takes place over periods of prolonged economic growth. Our chief concern in this paper is the latter hypothesis: how do regimes of financial fragility develop over time?

Financial fragility is described by Minsky (1986) as an attribute of the financial sector, where the regular operations of intermediaries may be disrupted by a "not unusual" event. Fragility arises as an equilibrium phenomenon from the regular profit-optimising behaviour of agents and banks as expectations of economic performance fluctuate. The buildup of optimism following a period of extended economic expansion precipitates an increase in the proportion of short-term debt, as well as an increase in the number firms that engage in speculative—as opposed to hedge—financing.[1] As economic growth and asset prices continue their upward trend during the expansion, attitudes toward risk soften so that speculation rises and firms take on riskier projects. Positive signals create momentum as the value of collateral rises and financial intermediaries transition from hedge to speculative financing. Protracted growth and higher asset prices—for instance in housing—will lead to higher collateral values that borrowers can put up against the value of a loan—for instance a mortgage or a re-financed mortgage—thereby increasing the level of debt circulating in the economy. This will act to further increase asset prices and possibly even boost economic growth despite financial imbalances.

In this state of financial fragility, the economy is especially vulnerable to large losses that can be precipitated by a "not unusual" event. For instance, a deceleration in asset prices or the failure of a large bank can correct the inflated expectations of future economic performance. When expectations suddenly adjust, intermediaries become unwilling to keep extending credit. To the extent that financing activities were reliant on expectations of future economic performance, this may lead to a debt-deflation spiral. When inflated expectations are not realised, actual profits and

[1] *Hedge* financing means that a firm can cover its debt commitments each period without re-financing over the long term. *Speculative* financing means that while a firm may be able to cover the interest portion of its debt commitments in the near term, it will need to re-finance from time to time in order to repay the principal. When firms are unable to cover the interest portion of debt commitments and must re-finance, they are said to engage in *Ponzi* financing.

asset values fall short of firms' commitments and there is default on debt. This reduces aggregate demand, which in turn leads to a decline in asset prices (hence the value of collateral) and thereby an increase in the real value of debt commitments that are defaulted; thus, "each dollar of debt still unpaid becomes a bigger dollar, and if the over-indebtedness with which we started was great enough, liquidation of debts cannot keep up with the fall of prices which it causes," (Fisher, 1933).

Debt-deflation spirals have been rare occurrences in the United States. According to Minsky (1986), this may be attributed to two developments since the Great Depression: a "big government" and a "big bank." The former implies that fiscal policy through the federal budget has the capacity to stimulate aggregate demand. The latter acknowledges the role of the Federal Reserve system as a lender of last resort that can stabilise financial markets during periods of particularly high turbulence. Indeed, both of these policy instruments were applied during the Great Recession of 2007-09. It has been argued elsewhere (Blinder, 2015; Gorton and Metrick, 2012a) that the various measures employed by the Federal Reserve, the Treasury, and Congress were mildly successful at stymying the debt-deflation process that was underway, thus preventing losses on a larger scale than was actually realised.

3.3.2 Time-varying expectations and endogenous default

Minsky's financial instability hypothesis describes how over-leveraging during a protracted expansionary period occurs due to inflated expectations and can lead to greater losses during a financial crisis than in a paradigm where expectations were held in check. In order to formally ground our intuition about this process, we need to examine how expectations of the future can affect the level of default i.e. the expected shortfall in debt repayments.

Consider a partial equilibrium example as in Bhattacharya et al. (2015). Suppose that a continuum of investors can invest in a risk-free and a risky asset to maximise profits by borrowing from a continuum of risk-neutral creditors or by re-investing profits. There are two states of the world $s_t = \{g, b\}$ that evolve according to a Markov process so that $\mathsf{P}(s_{t+1}|s_t, s_{t-1}, \ldots, s_1) = \mathsf{P}(s_{t+1}|s_t)$. All debts are repaid in the good state, but investors can choose to optimally default in the bad state subject to a pecuniary penalty and only repay a fraction $v_t < 1$ on a unit of loans. Equilibrium in the credit market is reached when the lending rate r_t is set, given a repayment rate v_{t+1} and beliefs, so that creditors break even according to

$$\mathsf{E}(v_{t+1})(1 + r_t) = 1. \tag{3.3.1}$$

Information is symmetric, so both investors and creditors share the same beliefs.

Agents' beliefs are constructed by incorporating uncertainty around the probability that the good state is realised, $\mathsf{P}(s_{t+1} = g|s_t) = \theta$, for some unknown θ. For simplicity, assume that $\theta \in \{\theta_1, \theta_2\}$ and that $1 > \theta_1 > \theta_2 > 0$. Agents infer θ at each period after observing s_t and have a subjective belief in period t of a good realisation in period $t + 1$ given by $\pi_{t+1} = \mathsf{E}(\theta|s_t) = \sum_i \mathsf{P}(\theta = \theta_i|s_t)\theta_i$; they assign probability $1 - \pi_{t+1}$ to a bad realisation. These beliefs are updated using Bayes' rule,

$$\mathsf{P}(\theta = \theta_i|s_t) = \frac{\mathsf{P}(s_t|\theta = \theta_i)\mathsf{P}(\theta = \theta_i)}{\mathsf{P}(s_t)}, \tag{3.3.2}$$

for $i = 1, 2$ once the current state s_t is observed. Since agents can observe one of two outcomes in each state, the path at time t with k realisations of the good state g when there is also parameter uncertainty can be described by $\mathsf{P}(s_t) = \sum_i \theta_i^k (1-\theta_i)^{t-k}\mathsf{P}(\theta = \theta_i)$.

This simple setup creates a framework for agent optimism. As the count of good states increases, so does the subjective probability of the good state being realised in the following period. Expectations of the future thus become largely self-fulfilling, and lending terms increasingly cease to reflect fundamentals; the greater the optimism, the greater the expected shortfall in debt repayment. These two results are formalised in the proofs of Propositions 3.3.1 and 3.3.2.

Proposition 3.3.1 *Following a realisation of the good state $s_t = g$, the subjective probability of the good state being realised in the next period increases, so that $\pi_{t+1} > \pi_t$.*

Proof of Proposition 3.3.1 Show that $\mathsf{P}(\theta = \theta_1 | s_{t+1}) > \mathsf{P}(\theta = \theta_1 | s_t)$ and $\mathsf{P}(\theta = \theta_2 | s_{t+1}) < \mathsf{P}(\theta = \theta_2 | s_t)$ when $s_t = g$, (Bhattacharya et al., 2015). $\blacksquare$

Proposition 3.3.2 *Following an increase in the probability of a good realisation, the repayment rate falls.*

Proof of Proposition 3.3.2 Given the subjective beliefs over the outcome of state s_{t+1}, the expected repayment rate satisfies

$$\mathsf{E}(v_{t+1}) = \pi_{t+1} + (1 - \pi_{t+1})\overline{v} = \frac{1}{1 + r_t},$$

where $\overline{v}$ is the mean repayment rate in the bad state. Rewrite this expression as

$$\overline{v} = \frac{1}{1 - \pi_{t+1}} \left(\frac{1}{1 + r_t} - \pi_{t+1} \right),$$

and calculate the derivative with respect to the subjective probability of the good state being realised.

$$\frac{\partial \overline{v}}{\partial \pi_{t+1}} = \left(\frac{1}{1 - \pi_{t+1}} \right)^2 \left(\frac{1}{1 + r_t} - 1 \right) < 0,$$

since $\pi_{t+1} \in (0, 1)$ and $r_t > 0$. $\blacksquare$

In this minimalist setup, we can formally characterise the fundamental aspects of Minsky's financial instability hypothesis: inflated expectations and endogenous fragility. When a sequence of consecutive good states is realised, agents' subjective probabilities of the future are skewed upwards, so that optimism develops. Risk is therefore undervalued, and agents expect default to be milder relative to its realisation. Bhattacharya et al. (2015) show in their model that this skews investment towards the risky asset once expectations improve sufficiently. Since investors are price takers, they do not internalise their impact on default in equilibrium, thereby creating an externality.

3.4 Econometric methodology

Cogley and Sargent (2001) developed the framework for a vector-autoregression with time-varying coefficients. The absence of time-varying volatility, however, created the possibility of excessive time-variation in the autoregressive parameters to compensate (Sims, 2001; Stock, 2001). In response, Cogley and Sargent (2005) incorporated stochastic volatility, while Primiceri (2005) also allowed for variation in the contemporaneous relations between endogenous variables. In this section, we briefly review the framework of Primiceri (2005). We also discuss the identification of impulse responses, as well as some empirical features relating to the calibration of priors.

3.4.1 Time-varying parameter VAR

Due to the evolving nature of economic transmission and the volatility of shocks, we would like to allow greater flexibility in our parameters. In order to allow for variation from all possible sources, we implement a time-varying parameter VAR (TVP-VAR) with multivariate stochastic volatility, as in Primiceri (2005) and Nakajima (2011).

For $\mathbf{y}_t$ a $N \times 1$ vector of endogenous variables, the reduced-form model with time-varying parameters is

$$\mathbf{y}_t = B_{0,t} + \sum_{j=1}^{p} B_{j,t}\mathbf{y}_{t-j} + \boldsymbol{\epsilon}_t, \qquad \boldsymbol{\epsilon}_t \sim N(\mathbf{0}, \Omega_t), \tag{3.4.1}$$

for $t = 1, \ldots, T$, where $B_{0,t}$ is a $N \times 1$ vector of time-varying intercepts, $B_{j,t}$ for $j = 1, \ldots, p$ are $N \times N$ matrices of time-varying coefficients for the lagged terms, and the errors $\boldsymbol{\epsilon}_t$ evolve according to a time-varying covariance matrix Ω_t to allow for multivariate stochastic volatility.

We are interested in orthogonal shocks to understand the dynamics of our model through impulse response analysis. Similarly to the constant parameter model, we achieve this by a triangular factorisation of the error covariance matrix as in Primiceri (2005). The factorisation is given by

$$A_t \Omega_t A_t' = \Sigma_t,$$

where A_t is lower triangular and Σ_t is a diagonal matrix of volatilities

$$A_t = \begin{bmatrix} 1 & 0 & \dots & 0 \\ a_{21,t} & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ a_{N1,t} & \dots & a_{NN-1,t} & 1 \end{bmatrix}, \quad \Sigma_t = \begin{bmatrix} \sigma_{11,t} & 0 & \dots & 0 \\ 0 & \sigma_{22,t} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \sigma_{NN,t} \end{bmatrix}. \tag{3.4.2}$$

The triangular assumption not only identifies the impulse responses but the ordered structure of the model as well.

For notational ease, we can rewrite the model in equation (3.4.1) by stacking coefficients. Let the rows in B_i for $i = 0, \ldots, p$ be stacked into a $(Np+1)N \times 1$ vector $\mathbf{b}$, and let $X_t = \mathbf{I}_N \otimes [1, \mathbf{y}_{t-1}', \ldots, \mathbf{y}_{t-p}']$. Then, using the triangular factorisation of

Ω_t, the model can be written as

$$\mathbf{y}_t = X_t \mathbf{b}_t + A_t^{-1} \Sigma_t^{1/2} \boldsymbol{\varepsilon}_t. \tag{3.4.3}$$

We stack the elements of A into a vector $\mathbf{a}_t = [a_{21,t}, a_{31,t}, a_{32,t}, \dots, a_{NN-1,t}]'$ and let $\mathbf{h}_t = [h_{1,t}, \dots, h_{N,t}]'$ be a vector of the log volatilities in Σ_t, defined $h_{i,t} = \log \sigma_{ii,t}$. The evolution of the parameters follows a driftless random walk, so that

$$\mathbf{b}_t = \mathbf{b}_{t-1} + \mathbf{u}_{b,t},$$

$$\mathbf{a}_t = \mathbf{a}_{t-1} + \mathbf{u}_{a,t},$$

$$\mathbf{h}_t = \mathbf{h}_{t-1} + \mathbf{u}_{h,t},$$

where innovations are jointly normally distributed with covariance matrix

$$\begin{bmatrix} \boldsymbol{\varepsilon}_t \\ \mathbf{u}_{b,t} \\ \mathbf{u}_{a,t} \\ \mathbf{u}_{h,t} \end{bmatrix} \sim N(\mathbf{0}, V), \qquad V = \begin{bmatrix} I_k & O & O & O \\ O & \Sigma_b & O & O \\ O & O & \Sigma_a & O \\ O & O & O & \Sigma_h \end{bmatrix}.$$

Due to the proliferation of parameters that are estimated at every time period, Bayesian sampling methods are used to evaluate the posterior distribution of the parameters by implementing the algorithm of Nakajima (2011). Subsequently, we will occasionally refer to the entire history of parameters by defining

$$\mathbf{b} \equiv \{\mathbf{b}_t\}_{t=p+1}^T, \quad \mathbf{a} \equiv \{\mathbf{a}_t\}_{t=p+1}^T, \quad \mathbf{h} \equiv \{\mathbf{h}_t\}_{t=p+1}^T.$$

A summary of the estimation algorithm is provided in Appendix 3.C.

3.4.2 Impulse response analysis

In the absence of natural experiments in the data, we must identify shocks from observational data. The identification process is important inasmuch as it allows us to conduct experiments. In the present case, the exogenous shocks are identified in a standard way through a lower-triangular assumption on A_t, as in (3.4.2). We can thus calculate impulse responses at a particular horizon $t+h$ to an exogenous shock $\mathbf{s}$ in each of the endogenous variables at time t. This is defined as

$$\mathsf{E}(\mathbf{y}_{t+h}|\overline{B}_t, \mathbf{e}_t = \mathbf{s}) - \mathsf{E}(\mathbf{y}_{t+h}|\overline{B}_t, \mathbf{e}_t = 0) \qquad \text{for} \qquad h \geq 0, \tag{3.4.4}$$

where $\overline{B}_t$ represents the posterior mean of the intercept and autoregressive coefficients at the particular period t in which the impulse responses are calculated.

The triangular structure assumed for A_t ensures that every shock $\mathbf{e}_t = \mathbf{s}$ corresponds to a single linear combination of residuals, $\mathbf{e}_t = A_t \boldsymbol{\epsilon}_t$. The cost of this assumption is a recursive ordering, where higher-ordered variables affect lower-ordered variables contemporaneously, but not the other way around. It is worth noting that this form of identification does not necessarily imply that an exogenous shock originates in a particular variable y_{it}. Since we calculate impulse responses by recursively ordering the residuals in our reduced-form model, rather than strictly identifying a shock to a particular variable y_{it}, we are really identifying shocks to y_{it} that are orthogonal with respect to the remaining variables y_{jt}, $j \neq i$ in our model.

3.4.3 Calibration of priors

The selection of an appropriate prior distribution for the initial state of the model parameters ($\mathbf{b}$, $\mathbf{a}$, $\mathbf{h}$), as well as for the hyperparameters (Σ_b, Σ_a, Σ_h) is important inasmuch as it will influence the posterior distribution of each parameter. While we can (and do) conduct a sensitivity analysis for the calibration of priors, it is appropriate to characterise certain features of the modelling framework described in §3.4.1 that we consider when setting priors.

First, the coefficients are assumed to follow a random-walk process. This means that while the TVP-VAR model is sufficiently flexible to capture both sudden and gradual changes in the parameters, there exists a risk of over-identification. In order to avoid implausible variation in the model parameters, Primiceri (2005) argues for relatively tight priors on Σ_b, Σ_a, and Σ_h.

Second, the appropriate modelling of stochastic volatility is critically important to the estimation of time-varying parameters. Sims (2001) and Stock (2001) argue that the lack of appropriate variation in $\mathbf{a}$ and $\mathbf{h}$ can lead to compensating and exaggerated variation in $\mathbf{b}$. As a result, a tight prior is set for Σ_b, while a more diffuse prior is set for Σ_a and Σ_h. This follows the practice in most of the literature, reflecting the belief that structural shock sizes fluctuate more than autoregressive parameters.

Finally, with the previous considerations in mind, we must decide how to actually set our priors. There are two approaches that we consider in our analysis for setting priors. The approach established in Cogley and Sargent (2001) and Primiceri (2005) sets the mean and variance of the priors for $\mathbf{b}, \mathbf{a}$, and $\mathbf{h}$ equal to the least-squares estimates from a constant-parameter version of (3.4.3) over a training sample preceding the main estimation sample. Priors on the hyperparameters require calibration, with a range of sensible values described in Primiceri (2005). Another approach that is often used is setting reasonably flat priors for the initial state of the model parameters. This uninformative approach incorporates an agnostic view of the model parameters.

3.5 Data

We estimate a model for the US at a quarterly frequency over the period spanning 1999(1)–2015(4). In order to adequately characterise the financial instability hypothesis, while also controlling for fluctuating macroprudential and monetary policy measures, our model includes the endogenous variables

$$\mathbf{y}_t = [\Delta y_t, \Delta c_t, \Delta p_t, \Delta s_t, \Delta r_t, i_t,]',$$

where y_t denotes log real industrial production (IP), c_t denotes log real total credit to the private sector (TPC), p_t denotes the log real Case-Shiller national house price index (CSI), s_t denotes the expected shortfall for an equally-weighted portfolio of selected financial intermediaries,[2] r_t denotes the (log) tier 1 capital to risk-weighted assets ratio for the same intermediaries, and i_t denotes the Federal Funds rate (FFR), adjusted by using the shadow rate of Wu and Xia (2016) at the zero lower bound. The data are plotted in levels and differences in Appendix 3.A.

Our sample period is restricted due to the availability of capital ratio and default data, but it nonetheless spans an eventful period for the US economy. We review notable financial and regulatory events in our sample in §3.5.1, while §3.5.2-3.5.4 describe our approach to measuring systemic risk, macroprudential requirements, and accounting for monetary policy at the zero lower bound.

3.5.1 Timeline of events

Table 3.5.1 provides a reference for the major financial and regulatory events that take place during the sample we consider. The size of the sample is limited by the availability of balance sheet data that can be used as a macroprudential policy target, as well as by the fact that our chosen measure for systemic risk relies on CDS-spreads (Goodhart and Segoviano, 2009). Nonetheless, we are able to capture two reces-

[2]The expected shortfall is alternatively known as the conditional value-at-risk (CVaR); this measures losses in a portfolio's value, given that the losses occur below a specified quantile. The difference in the present approach is that the returns at the specified quantile are obtained from simulations of the joint probability distribution for the portfolio of firms in use, applying the method of Segoviano (2006) and Goodhart and Segoviano (2009). Further details are given in §3.5.2.

sions, at least one complete financial cycle, large foreign/domestic financial shocks, and several expansionary/contractionary shifts in monetary policy. Macroprudential requirements only move in one direction throughout the sample (they are tightened).

Table 3.5.1: Notable financial events for the United States, 1999–2015.

Event	Dates
GLB Act (Glass-Steagall repealed)	November 1999
Dotcom bubble bursts	March 2000–April 2001
US recession	March–November 2001
Argentine debt default	December 2001–2002
Basel II proposed	June 2004
Repo market run	August 2007
US recession	December 2007–June 2009
Global financial crisis	2008–2009
Term Securities Lending Facility (TSLF)	March 2008
Bear Stearns collapse	March 2008
AIG bailout	September 2008
Lehman Brothers bankruptcy	September 2008
Troubled Asset Relief Program (TARP)	October 2008
QE 1	November 2008
US bank stress tests	February–May 2009
Term Asset-backed Securities Loan Facility (TALF)	March 2009
Dodd-Frank	May 2010
Basel III proposed	September 2010
QE 2	November 2010
European debt crisis	May 2010–December 2012
QE 3	December 2012

3.5.2 Measurement of systemic risk

In order to capture the linear and nonlinear dependencies between financial intermediaries, we employ the method of Segoviano (2006) and Goodhart and Segoviano (2009) to compute the joint distribution of asset returns while compensating for increased tail dependence. Thus, when one institution crosses the default threshold, the effect on the asset returns of all other institutions is also accounted for. The intermediaries used for this calculation are chosen by asset size and are shown in Table 3.5.2. Total assets are reported at the beginning and end of the sample, as well as the through-time average. We note that total assets for all intermediaries increase over four times between 1999 and 2015, following the repeal of Glass-Steagall.

TABLE 3.5.2: Financial intermediaries included in the measurement of systemic risk over the sample 1999(1)–2015(4).

Intstitution	Assets (millions USD)		
	1999(1)	2015(4)	Average
JP Morgan Chase & Co.	0.361	2,352	1,514
Bank of America Corp.	0.614	2,144	1,470
Wells Fargo & Co.	0.201	1,788	0.817
Citigroup Inc.	0.691	1,731	1.543
Goldman Sachs Group Inc.	0.231	0.861	0.698
Morgan Stanley	0.322	0.787	0.720
US Bancorp	0.074	0.422	0.245
Capital One Financial Corp.	0.010	0.314	0.139

Figure 3.5.1 plots two variant measures of systemic risk from Goodhart and Segoviano (2009) for our portfolio of financial intermediaries. The first panel shows the joint probability of default, which is defined as the probability that the asset returns of all institutions are in the default region simultaneously. The second panel shows the expected shortfall at the 1% level for a US \$100,000 portfolio—weighted by asset size—among the institutions in Table 3.5.2.

Events that are most clearly picked up by these measures include the Argentine debt default in 2002, the run on the repo market in 2007, followed by the tumultuous events of 2008 that include the collapse of Bear Stearns, the bailout of AIG, and the Lehman Brothers bankruptcy. The various policy interventions between October 2008–March 2009 and the rebounding in housing prices likely contributed to the rapid decline in systemic risk following the GFC. Increases in the joint default probability and the expected shortfall thereafter may be attributed to the intensification of the European debt crisis in 2011. Subsequently and to the present day, systemic risk falls to substantially lower levels that are, nonetheless, higher relative to the pre-GFC era.

It is difficult to say exactly why this relative difference persists. A possible ex-

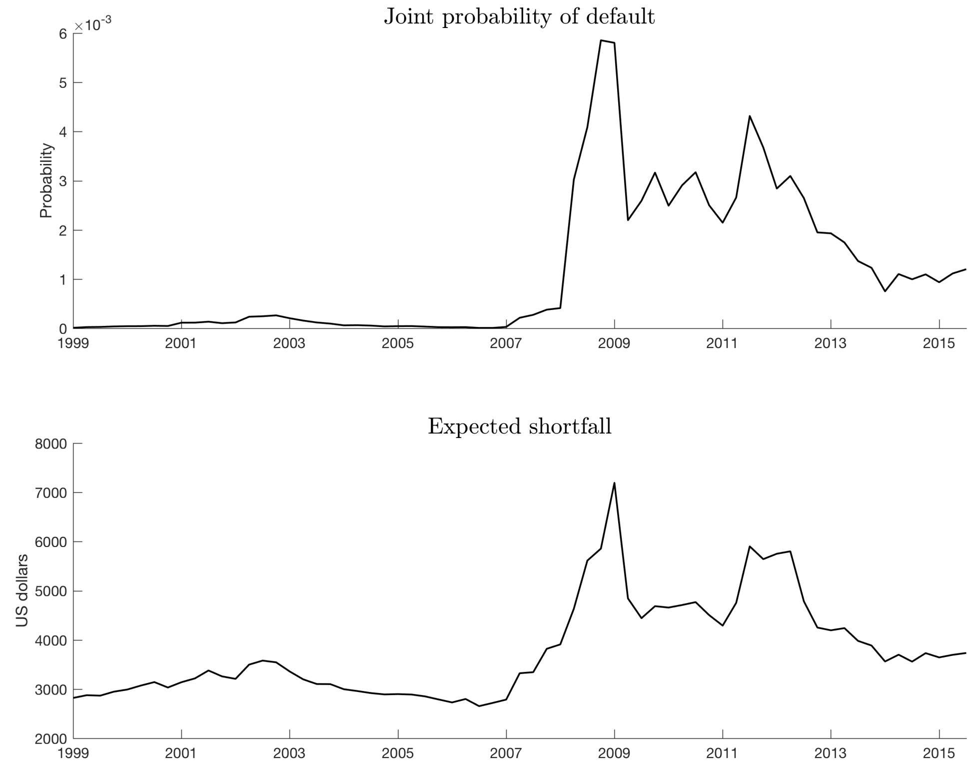

FIGURE 3.5.1: nonlinear systemic risk measures. The joint probability of default and the expected shortfall on a US $100,000 portfolio for financial intermediaries included in the analysis.

planation is that various policy efforts to curb the effect of aggregate shocks leave intermediaries more vulnerable to idiosyncratic shocks, leading to correlated exposures (as argued in Horváth and Wagner, 2013). Additionally, the relatively higher level of systemic risk may reflect in part the concern over asset price bubbles after a protracted period of expansionary monetary policy (Bordo and Landon-Lane, 2014). A third explanation is that there has been a structural break in market expectations after the GFC; this is possible, since the method by Goodhart and Segoviano (2009) relies on CDS-spreads to adjust the posterior multivariate density of asset returns.

3.5.3 Macroprudential capital requirements

The primary objective of macroprudential capital requirements is to improve financial stability on aggregate, where systemic risk arises endogenously through the regular functioning of financial intermediaries. Endogenous systemic risk can occur along the time dimension—due to the pro-cyclicality of the financial sector—and along the structural dimension—due to the interconnectedness of individual financial intermediaries. Thus, the question of how to conduct macroprudential policy and which of the available instruments to use will vary based on which dimension of systemic risk is given more weight.

Macroprudential policy must ultimately force intermediaries to internalise the effects of their borrowing/lending decisions. Within the time dimension, successful policy intervention will curb leveraging during the boom and deleveraging in the bust, thereby also controlling fluctuations in asset prices that are often used as collateral for loans. On the other hand, successful policy must also limit common exposures through the interbank market in order to reduce the likelihood of a domino effect following the failure of one or more financial intermediaries.

An exhaustive list of objective-oriented macroprudential instruments is given in the appendix of Galati and Moessner (2017). From this, various forms of capital requirements fit the bill for addressing systemic issues both along the time and structural dimensions. We use an average of the tier 1 to risk-weighted assets ratio for the financial intermediaries considered in our sample. This variable addresses both dimensions of systemic risk, while it is also an instrument used to implement macroprudential regulatory changes (BIS, 2010). This is plotted in log levels and differences in Figure 3.5.2.

We face several challenges in using the capital ratio as a measure of macroprudential requirements.[3] First, there is a question of endogeneity: to what extent are

[3] In the present analysis, we abstract from potential manipulation of risk weights by financial

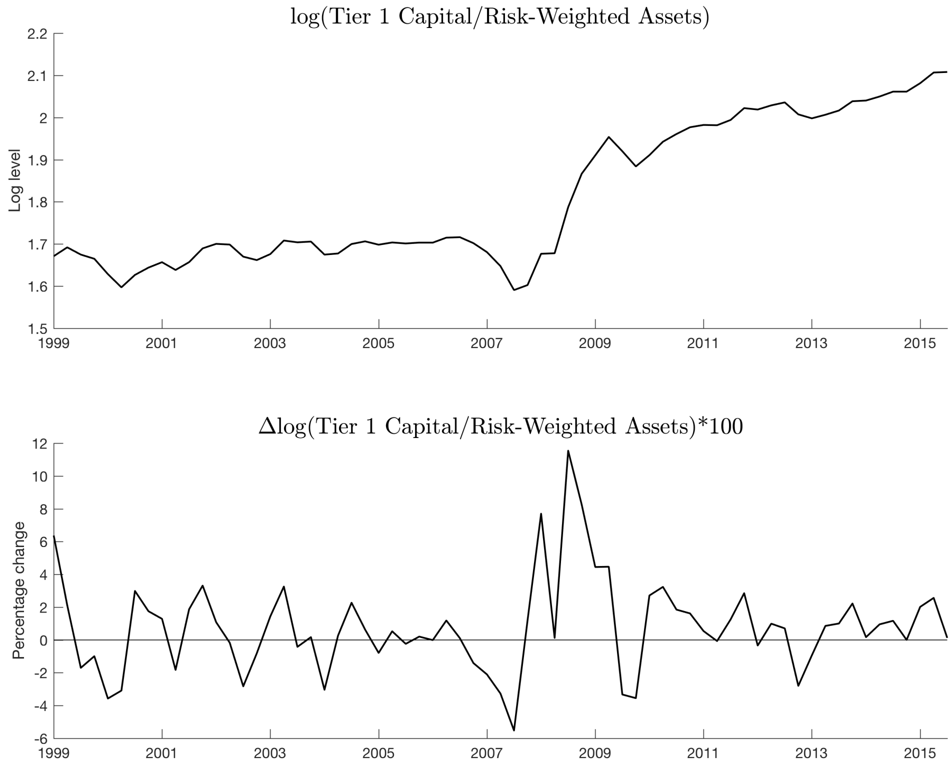

FIGURE 3.5.2: Capital ratio. The ratio of tier 1 capital to risk-weighted assets across all financial intermediaries included in our sample is used to measure changes in macroprudential requirements.

movements in the capital ratio driving or being driven by movements in the other macro-financial variables. For instance, risk-based capital requirements depend on the creditworthiness of borrowers, so that financial intermediaries' assets must be weighted accordingly. An increase in the riskiness of borrowers occurs naturally in a recession, so that assets must decline accordingly in order to meet the minimum capital requirement. Thus, endogeneity between macro-financial variables and capital requirements can even result in pro-cyclicality, thereby further exacerbating a credit crunch during a recession (Catarineu-Rabell et al., 2005; Blinder, 2010). Second, there is a question of disentangling the distinct effects of macroprudential and monetary policy. This is complicated by the fact that monetary policy can systematically

intermediaries.

affect risk-taking (Borio and Zhu, 2012), suggesting potential complementarity and substitutability between both types of intervention. Third, it is not clear to what extent movements in the capital ratio move in anticipation of or in response to policy interventions. For instance, both Basel II and Basel III regulations were published, revised, and implemented over a protracted period of time, allowing intermediaries ample time to adjust to new macroprudential measures even prior to official implementation. In both panels of Figure 3.5.2, we see a rise in the trend of the capital ratio after the financial crisis, suggesting that intermediaries' risk-taking behaviour has been affected. Yet, it is difficult to determine whether this is in response to (anticipated) regulation or due to a correction in future expectations after the GFC.

We address the first two challenges by ordering the policy variables after the macro-financial variables. This is standard practice in the applied macroeconomics literature. It also implies that any shocks to the residuals in the capital ratio equation are orthogonal to contemporaneous shocks in all preceding macro-financial variables; and similarly, that shocks to the residuals in the adjusted Federal Funds rate equation are also orthogonal to contemporaneous shocks in the capital ratio as well. This ordering lets monetary policy respond contemporaneously to all variables, including macroprudential requirements, as in Kim and Mehrotra (2015). The final challenge remains unresolved in this paper and is one that cannot be addressed in a framework where orthogonal ("structural") shocks are identified through a recursive assumption. One avenue to proceed in future work is to examine the impact of shocks to macroprudential requirements identified through the historical narrative approach, as in Federico et al. (2012).

3.5.4 Monetary policy

We study the relationship between monetary policy and a system of macro-financial variables by incorporating the Federal Funds rate into our vector autoregression.

Prominent examples of this approach include Bernanke and Blinder (1992), Christiano et al. (1999), and Stock and Watson (2001). A difficulty that arises in the sample that we consider is the fact that the Federal Funds rate after 2009 is at the zero lower bound and may not incorporate the impact of unconventional monetary policy conducted through asset purchases and forward guidance (see Table 3.5.1).

One solution to this problem proposed by Wu and Xia (2016) is to employ a shadow rate term structure model to infer the stance of monetary policy when the nominal rate is zero. By allowing a negative (shadow) policy rate, it is possible to account for the effects of unconventional monetary policy endogenously. The shadow rate, along with the Federal Funds rate, are plotted in Figure 3.5.3.

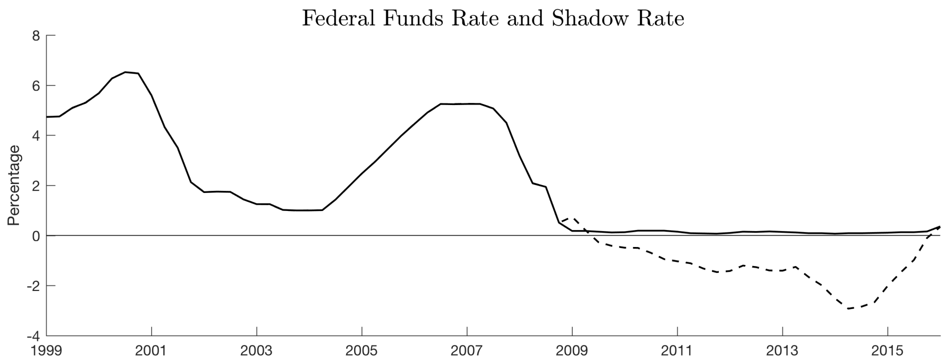

FIGURE 3.5.3: Policy rate. The Federal Funds Rate is adjusted in Wu and Xia (2016) for the zero lower bound.

The same challenges that applied for the identification of macroprudential policy shocks are also relevant here. Problems of endogeneity remain. For instance, the gradual increases in the Federal Funds rate in the boom years leading up to the two recessions in our sample (2001–02, 2007–09) are endogenous responses to the economy overheating. The subsequent dramatic cuts in the policy rate are also endogenous responses to stimulate aggregate demand. Additionally, it is impossible under the present identification scheme to clearly distinguish between macroprudential and monetary policy shocks. The staggering level of intervention (see Table 3.5.1) means that partialling out the contemporaneous effects of shocks to the macro-financial and macroprudential variables is the best we can hope for. Identification of monetary policy shocks through alternative strategies (for instance Romer and Romer, 1989; Kuttner, 2001; Uhlig, 2005) may provide complementary results. Since identification is not a central objective of this paper, however, we leave that for future work.

3.6 Time-varying macro-financial linkages

We examine a quarterly macro-financial model that controls for macroprudential and monetary policy over the period spanning 1999(1)–2015(4). We find evidence for credit-fuelled inflation in housing prices that drives over-leveraging and growth in industrial production. Positive economic signals lead to evolving attitudes toward risk that drive the upward momentum of the financial cycle, up until expectations are corrected. We compare shocks before and after the GFC as an experiment to gauge the importance of upward momentum in expectations. While the largest losses to industrial production occur in the crisis period of 2008, the most persistent losses that arise from strong macro-financial linkages result from shocks that may have occurred at the height of the financial cycle in 2004.

In this section, we discuss our agnostic approach to the selection of prior distributions and estimation of the parameters. We find that there is significant time-variation in volatilities, indicating that the size of structural shocks must be allowed to vary over the financial cycle. Subsequently, we examine the inflationary and deflationary interactions between the macro-financial variables, relating these to the stages of the financial instability hypothesis.

3.6.1 Priors

Bayesian methods are used to evaluate the posterior distribution of the parameters in equation (3.4.3). To this end, a prior assumption on each parameter must be specified. A standard approach in the literature has been to calibrate priors based on the least-squares estimate of a constant-parameter VAR over a training sample (Primiceri, 2005; Cogley and Sargent, 2001). The limited size of our sample, however, prohibitively raises the cost of discarding data in a training sample. Instead, we set reasonably flat, normal priors for the model parameters and tight priors for the hyperparameters with an inverse Wishart distribution

$$\mathbf{b}_0 \sim N(\mathbf{0}, 10 \times I), \tag{3.6.1}$$

$$\mathbf{a}_0 \sim N(\mathbf{0}, 10 \times I), \tag{3.6.2}$$

$$\mathbf{h}_0 \sim N(\mathbf{0}, 10 \times I), \tag{3.6.3}$$

$$\Sigma_b \sim IW(90 \times \kappa_b^2 \times I, 90), \tag{3.6.4}$$

$$\Sigma_a \sim IW(20 \times \kappa_a^2 \times I, 20), \tag{3.6.5}$$

$$\Sigma_h \sim IW(20 \times \kappa_h^2 \times I, 20). \tag{3.6.6}$$

The relatively flat priors for the model parameters imply that we take an agnostic view, so that more relative weight is placed on the available data; in other words, our prior distributional assumption is uninformative. The tight priors for the hyperparameters are intended to limit the degree of in-sample variation. To this end, κ_b, κ_a, and κ_h are important in determining the degree of variation permitted over time. We set $\kappa_b = 0.01$, $\kappa_a = 0.05$, and $\kappa_h = 0.35$, reflecting the importance of volatility (determined by $\mathbf{a}$ and $\mathbf{h}$) in economic models, while remaining broadly consistent with the empirical literature. As a comparison for our choices, Primiceri (2005) sets $\kappa_b = \{0.01, 0.05, 0.1\}$, $\kappa_a = 0.1$, and $\kappa_h = 0.01$, while Prieto et al. (2016) set $\kappa_b = 0.01$, $\kappa_a = 0.1$, and $\kappa_h = 0.1$. Our choices for the location/scaling constants in (3.6.4)-(3.6.6) further reflect the principle outlined in §3.4.3 that a tighter prior is required for Σ_b relative to Σ_a and Σ_h.[4]

[4]We carried out a sensitivity analysis for the calibration of hyperparameters in (3.6.4)-(3.6.6) and found that for values $\kappa_h << 0.35$ volatility was not adequately modelled. Adequacy of estimated

As a sensitivity check, we set priors as in Primiceri (2005) calibrated using the main estimation sample as a training sample as well. While this is not a strictly Bayesian approach, it gives us a way of incorporating informative priors with which to compare our results using uninformative priors. Ultimately, the results produced are qualitatively similar to those using the priors described in (3.6.1)-(3.6.6).[5]

3.6.2 Estimation

We estimate the model based on 150,000 draws, discarding the first 50,000 for convergence of the sampler. Convergence of the Markov chain is assessed formally using the convergence diagnostic (CD) of Geweke (1992) and inefficiency factors (IF) of Chib (2001). Informally, we assess convergence by examining the autocorrelation plots of the MCMC sample. The sampling procedure is described in Appendix 3.C, while results and diagnostics of the estimation algorithm are presented in Appendix 3.D.

3.6.3 Stochastic volatility

Incorporating stochastic volatility into our TVP-VAR plays a dual role. On the one hand, it reduces the possibility of excessive time-variation in the autoregressive parameters that can result from the assumption of constant variance. On the other hand, since our main argument on the buildup of endogenous systemic risk rests on time-varying expectations, it is important to consider both the first and second moments of the data. We find significant variation in volatility over our sample and discuss the economic implications of this.

Figure 3.6.1 shows the posterior mean across all retained draws of the volatilities for each of the endogenous variables. Each of the panels suggests that the size of

volatilities was evaluated by comparing relevant volatilities with the results in Prieto et al. (2016)—for instance, for the policy rate and house price inflation—over a similar sub-sample.

[5] In order to reduce excessive variation in the autoregressive parameters when this method of prior calibration is employed, κ_h must be set in a higher range in order to capture variation in the volatilities, while κ_b must be lowered.

shocks varies significantly over our sample. For all variables, there is relative stability from the beginning of the sample, which increases across the board during the Great Recession before returning to pre-recession levels. The exceptions occur for house prices and expected shortfall; for both, there is a second uptick in volatility around 2011. These results are broadly consistent with those in Prieto et al. (2016) and Abbate et al. (2016), who find large increases in financial shocks for the period of the GFC. Furthermore, our findings lend some support to the Stock and Watson (2012) narrative that the GFC was a result of larger shocks, but we also find evidence for the evolution of macro-financial linkages.

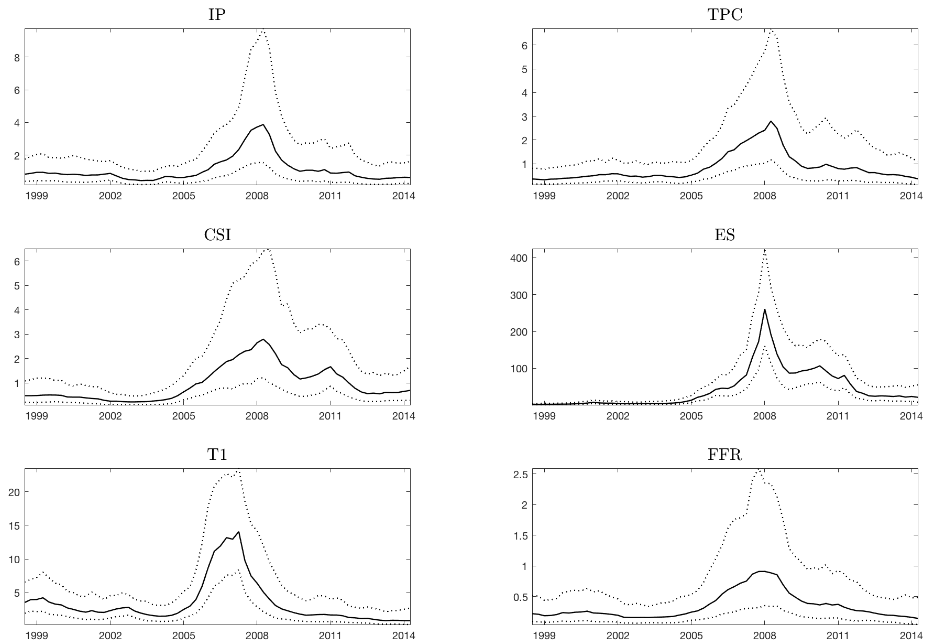

FIGURE 3.6.1: Time-varying volatility. Posterior mean of the variance of the residuals, plotted with bands of 1 standard deviation of the MCMC sample.

These results are also consistent with the "volatility paradox" of Brunnermeier and Sannikov (2014): low exogenous risk contributes to the accumulation of systemic risk. In our empirical framework, "exogenous risk" corresponds to the volatility of the residuals, plotted in Figure 3.6.1. We do not study the direct link between volatility regimes and the accumulation of credit, but we do find that volatility was exceptionally low in the period through 2006, a result that is permissible within the financial instability hypothesis. During the low-volatility environment, agents increase their risk-taking behaviour, eventually exacerbating future economic losses. Direct evidence for this hypothesis is also provided in Danielsson et al. (2016), who find that periods of exceptionally low volatility lead to the accumulation of credit and increase the likelihood of a financial crisis across countries. Subsequently, we examine how credit accumulation within the low-volatility environment takes place.

3.6.4 Inflation

We find evidence for significant time-variation in macro-financial linkages over our sample. Increases in economic growth and house prices are interpreted as positive signals that affect expectations of future economic performance. Positive shocks display greater persistence prior to the GFC, while reducing the repayment rate as well, suggesting that time-varying expectations play a role. These shocks initiate a reinforcing feedback loop between credit and house prices over the period spanning 1999–2006 that all but disappears after 2007 once expectations have corrected. The evidence suggests that house price shocks play the most important role in this narrative, as these are amplified through effects on wealth (Bernanke et al., 1999) and credit (Kiyotaki and Moore, 1997). We abstract from the role of policy in the present discussion and return to this in §3.7.

A central tenet of Minsky's financial instability hypothesis is that protracted periods of economic growth lead to inflated expectations of future economic outcomes. It is generally difficult to measure agents' expectations directly and to assess whether expectations actually translate into macroeconomic shifts. We must thus rely on indirect evidence that expectations are evolving. In §3.3.2, we discussed a framework through which we can interpret our empirical results. Positive signals will increase expectations of future economic performance, and thereby enhance the probability of observing positive signals in the future. Additionally, positive signals will reduce the repayment rate on loans, thereby increasing the expected shortfall among financial intermediaries. We can test these claims against the results. Figures 3.6.2–3.6.4 plot the time-varying impulse response functions to orthogonal shocks in industrial production (IP), real total credit to the private sector (TPC), and the real Case–Shiller national house price index (CSI). We alternatively refer to these as "economic growth," "credit growth," and "house price inflation."

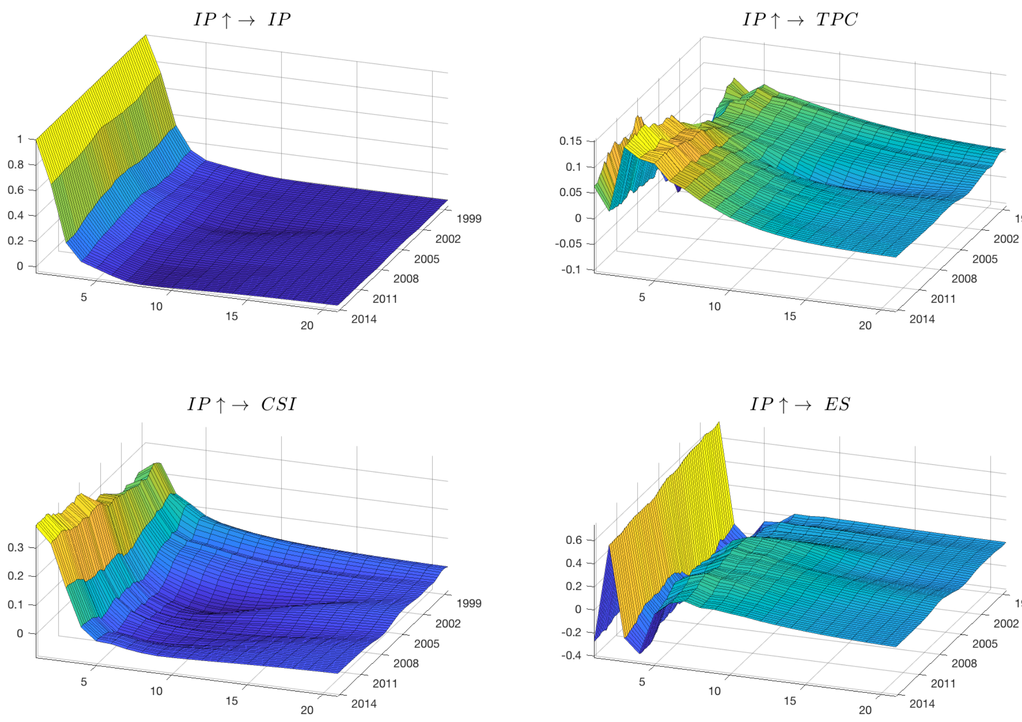

FIGURE 3.6.2: Time-varying impulse response functions. IRFs to a unit increase in the IP growth rate over a horizon of 20 quarters. Vertical axis: percentage change. Front axis left: horizon (quarters). Front axis right: years.

Increases in economic growth do not generate significant time-variation in the own-response depicted in the first panel. Nonetheless, a 1% increase in economic growth generates persistent and positive increases in real credit growth and in real house price inflation in the period up to 2006. While the response in real credit growth rebounds after the GFC, the inflationary response in real house prices is significantly less persistent. The cumulative response in expected shortfall is positive: initially higher growth raises the size of conditional losses, which subsequently acts to reduce the joint probability of default, making large losses less likely.

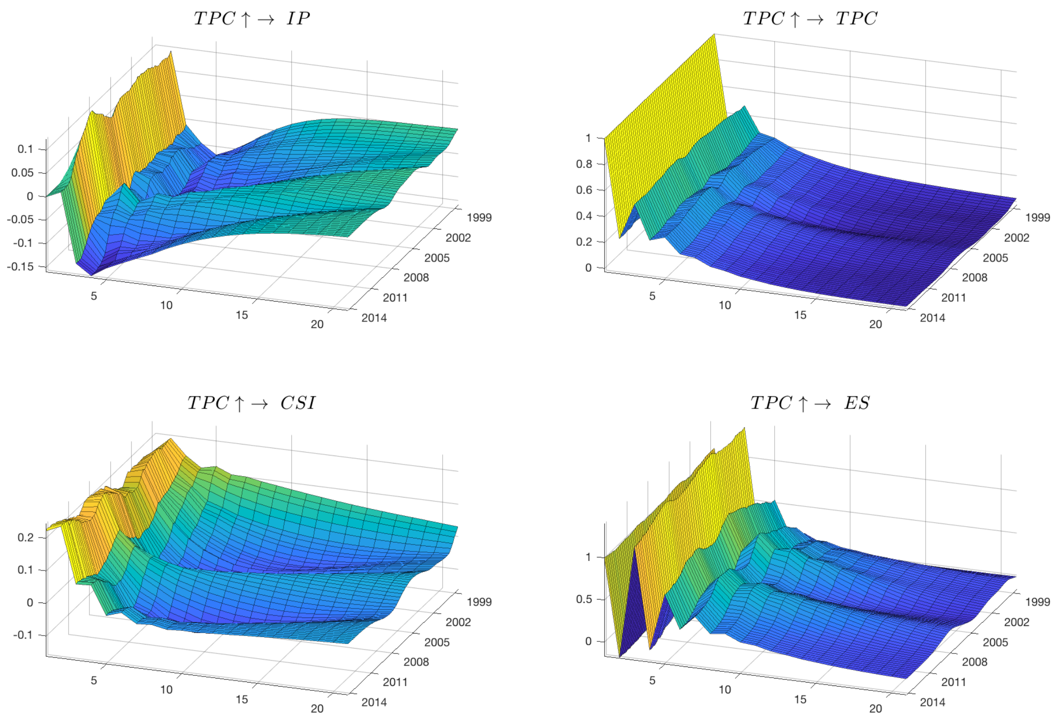

FIGURE 3.6.3: Time-varying impulse response functions. IRFs to a unit increase in the real credit growth rate over a horizon of 20 quarters. Vertical axis: percentage change. Front axis left: horizon (quarters). Front axis right: years.

From Figure 3.6.3, we see that credit initially induces a positive response in economic growth. Greater liquidity in the system can lead to the financing of more projects. Indeed, high credit growth creates upward pressure on house prices that is highly persistent until 2005. Between 2005 and 2008, however, credit cannot sustain house price growth to the same extent. The final panel indicates that credit always increases the expected shortfall of financial intermediaries: as intermediaries

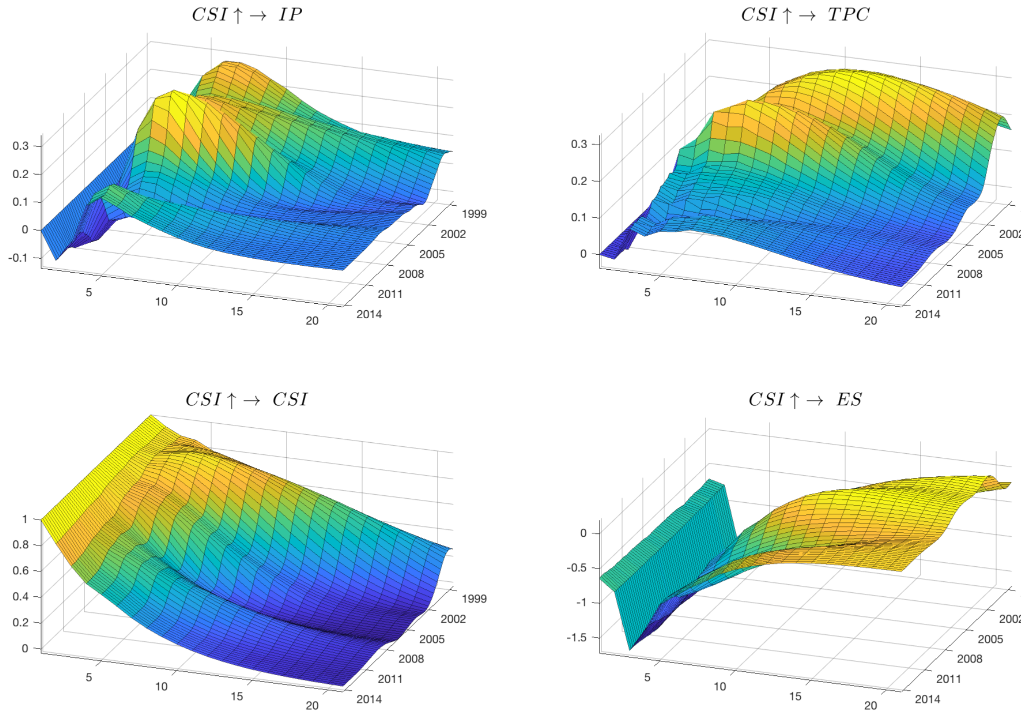

FIGURE 3.6.4: Time-varying impulse response functions. IRFs to a unit increase in real house price inflation over a horizon of 20 quarters. Vertical axis: percentage change. Front axis left: horizon (quarters). Front axis right: years.

increase the amount of loans on their balance sheets, they also increase their exposure to potential losses. This is consistent with the responses in Figure 3.6.2, whereby high economic growth simultaneously promotes credit growth, as well as growth in expected shortfall.

The greatest sources of time-variation among the transmission channels occurs through house price inflation; Figure 3.6.4 shows the impulse response functions to a unit increase. From the first panel, we can observe that house price inflation is a significant driver of growth prior to the 2001 recession, as well as in the buildup to the 2007–09 recession; a percentage increase in house prices was capable of generating up to a 0.3% increase in economic growth. The second panel in Figure 3.6.4 provides evidence that increasing house prices, and thereby increasing collateral values, created a significant expansion in credit. This effect peaks in 2004, which we regard as the peak of the financial cycle leading to the Global Financial Crisis.

Taken cumulatively, these results suggest that positive signals may indeed affect expectations that leads to over-leveraging, even though most of the work is done though shocks to house prices rather than economic growth. We can potentially explain this by examining the mechanisms through which inflated expectations can translate into macroeconomic effects. A higher tolerance for risk can initiate the financial accelerator of Bernanke et al. (1999). For the period over which we analyse data, ever-declining lending standards brought about a transition from hedge to speculative/Ponzi financing and resulted in the subprime mortgage crisis. When households have greater access to credit, this will drive up house prices, thereby inducing an effect on net worth and further increasing house prices. This effect is amplified when there is collateralised lending (Kiyotaki and Moore, 1997). Growing house prices raise the value of collateral for loans, allowing households to increase their borrowings and further exacerbating the inflationary effect on house prices. In fact, Kivedal (2013) finds evidence for a rational house price bubble leading up to the GFC.

In this state of heightened financial fragility, adverse shocks are amplified and can have potentially outsized macroeconomic effects, relative to a regime of financial robustness. We consider further evidence for this in §3.6.5.

3.6.5 Deflation

We consider the effects of "not unusual" shocks that result in a higher expected shortfall. The effect of shocks is much more severe following periods of protracted economic growth, rather than following periods of economic decline. This transmits primarily through a severe contraction in credit and house prices, which subsequently affect economic growth as well, suggesting the presence of a debt-deflation spiral (Fisher, 1933). The evidence is also consistent with predictions in Minsky (1992), Kiyotaki and Moore (1997), and Bernanke et al. (1999): over-leveraging during the upswing of the financial cycle exacerbates losses when the cycle reverses. The impulse response functions are plotted in Figure 3.6.5.

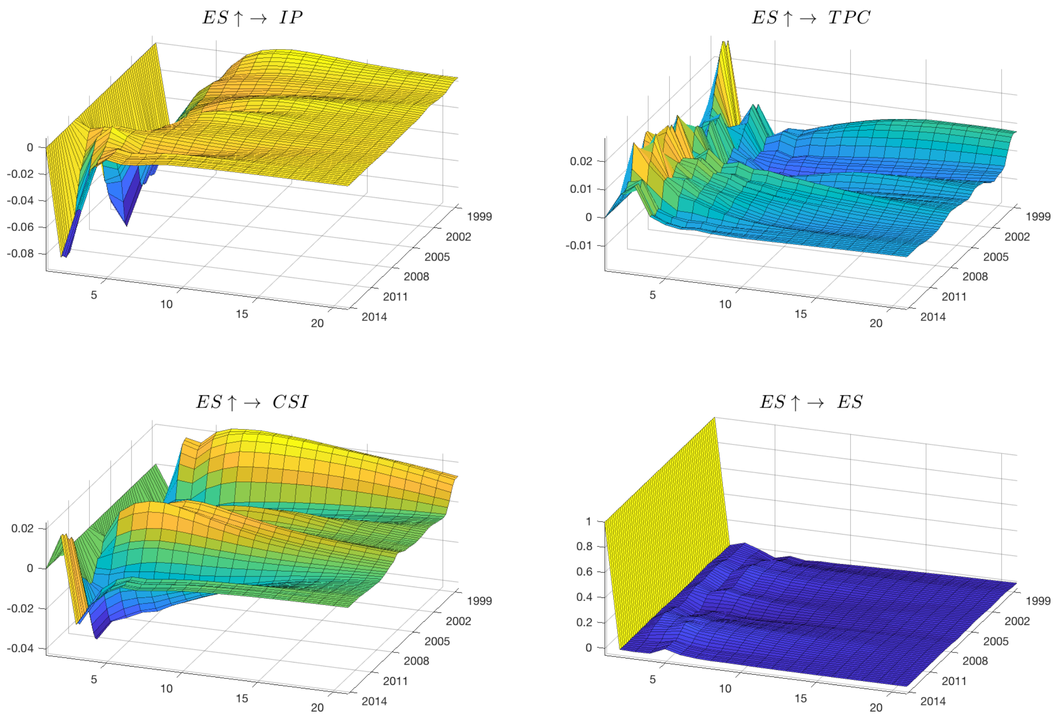

FIGURE 3.6.5: Time-varying impulse response functions. IRFs to a unit increase in expected shortfall growth over a horizon of 20 quarters. Vertical axis: percentage change. Front axis left: horizon (quarters). Front axis right: years.

A rise in expected shortfall among financial intermediaries has the largest effect on economic growth in 2008, as the debt-deflation spiral is unraveled following a protracted period of positive growth: a 1% increase in expected shortfall induces a 0.1% decline in economic growth that persists up to 8 quarters in the period 2003–2008, while the effect and its persistence are halved in 2009. The effect is also relatively less severe in earlier stages of the financial cycle. This can be explained either as a rapid correction in expectations, or alternatively as effective regulatory support.

Nonetheless, the effects of a shock that results in increased expected shortfall has the most persistent effect at the height of the financial cycle in 2004. This is transmitted through a larger contraction in credit and house prices relative to the crisis period in 2008 and can be likely explained through the effect on expectations. Whereas by 2008 sufficient negative signals—through the deceleration in house prices and the repo market run of August 2007, see Table 3.5.1—have occurred and initiated the correction to lower expectations of future economic performance, the upward momentum in expectations is at its peak in 2004. A shock to increase losses given default—either by increasing losses or increasing the joint probability of default—can therefore initiate the debt-deflation spiral precisely at the point where macro-financial linkages are strongest; once again, refer to the first and second panels in Figure 3.6.4.

To some extent, the potential economic losses from the debt-deflation spiral were mitigated following a rapid regulatory response. These included the Troubled Assets Relief Program (TARP) in October 2008, quantitative easing (QE) initiated by the Federal Reserve in November 2008 that saw it accumulate large amounts of mortgage-backed and other securities, as well stricter macroprudential requirements to contain systemic risk brought about through the Dodd-Frank Act in 2010. We examine how some of these policy measures may have affected macro-financial linkages in §3.7.

3.7 Policy interventions

Monetary and macroprudential policy have been very active over the previous decade. In order to obtain some intuition over their effects on the macro-financial relationship, we examine impulse response functions for the ratio of tier 1 capital to risk-weighted assets, as well as the Federal Funds rate, adjusted as in Wu and Xia (2016) to capture unconventional monetary policy measures at the zero lower bound. It is generally very difficult to identify truly exogenous policy shocks. We do not examine the identification of policy in this paper, but rather have assumed a recursive ordering, setting the two policy variables last; this is in keeping with standard practice. In economic terms, this means that both monetary and macroprudential policy respond contemporaneously to shocks in the macro-financial variables in our model. In econometric terms, this means that the identified shocks are the residuals of the policy variables after partialling out the time-varying effects of all other endogenous variables. Thus, the identified policy shocks are orthogonal to shocks in the other variables that arguably capture most of the relevant macro-financial movements in our sample.

We find that the timing of macroprudential measures and monetary policy matters. Early intervention to shore up financial intermediaries' balance sheets has the capacity to trade short-term losses for higher medium and long-term gains. The importance of timing is greater still for monetary policy interventions. We find that monetary policy is primarily useful prior to 2004, the height of the financial cycle. Attempts to deflate the credit-fuelled growth in asset prices via monetary policy past this stage, however, can have unexpected negative effects. This can occur when expectations of economic performance are resilient to monetary policy movements.

We compare impulse responses to a unit increase in the tier 1 capital to risk-weighted assets ratio and the Federal Funds rate adjusted for the zero lower bound for the periods 2000(1), 2004(1), 2008(3), and 2015(3). These periods are chosen in order to clearly illustrate the effects of policy at different stages: early in the financial cycle before the rapid expansion of credit; at the height of the financial cycle when macro-financial linkages are strongest; at the reversion of the financial and business cycles; and when the economy has recovered following policy intervention.

3.7.1 Macroprudential capital requirements

Significant time-variation is evident in the response to macroprudential shocks in Figure 3.7.1. Shocks to the tier 1 capital to risk-weighted assets ratio have varying effects, depending on the financial cycle. An increase in the capital ratio always leads to a fall in economic growth, credit growth, house price inflation, and expected shortfall among intermediaries. Depending on the stage of the financial cycle, however, this can prevent future losses and result in a net positive economic outcome. For instance, a unit increase in the growth rate of the capital ratio taking place in 2000(1) and 2004(1) induces a 0.05% decline in industrial production growth that lasts 4-5 quarters; subsequently, smaller but more persistent gains in growth are made. On the other hand, policy shocks in 2008(3) and 2015(1) have net negative effects.

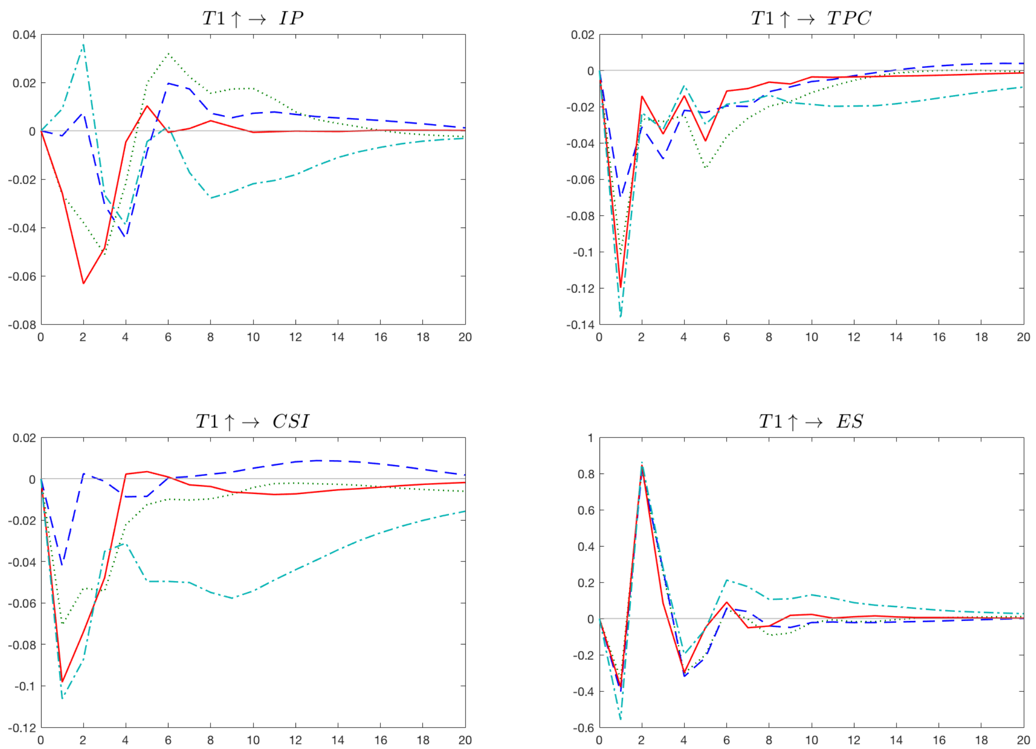

FIGURE 3.7.1: Time-varying impulse response functions. IRFs to a unit increase in the Federal Funds rate, adjusted for the zero lower bound. Vertical axis: percentage change. Horizontal axis: years. Green line: 2000(1). Blue line: 2004(1). Red line: 2008(3). Cyan line: 2015(3).

The timing of the macroprudential policy shocks thus matters. Intervention early in the financial cycle (e.g. in 2000) leads to short-term losses but long-term gains. This occurs due to the ability of tighter macroprudential requirements to rein in credit growth and the resulting house price inflation, thereby curbing agents' expectations of a high-growth economic environment. Eventually, this results in a relatively lower expected shortfall over the impulse horizon on average. Intervention later in the financial cycle (e.g. in 2004) has a relatively smaller upside. Once, however, the momentum in expectations reverses (e.g. in 2008), macroprudential tightening only results in losses. Credit and asset prices have reached unsustainable levels, so that once expectations of future performance adjust, a debt-deflation spiral begins. Intermediaries are offloading assets, thereby driving the value of the assets used as collateral down, further decreasing repayment rates. Increasing capital requirements exacerbates the de-leveraging behaviour and deepens losses. The final impulse response is computed during a period of relative financial stability (e.g. in 2015), when macroprudential and monetary policy have already intervened to bring the economy back to real growth. Increasing capital ratios in this regime undermines these efforts. Without sufficient risk-taking behaviour, asset prices and lending are reduced. While this reduces expected shortfall in the short term, the large negative effect on industrial production likely raises the joint probability of default among intermediaries, thereby increasing expected shortfall in the medium/long term.

3.7.2 Monetary policy

We compare unit shocks to the adjusted Federal Funds rate for the same periods in the sample. We use the shadow rate of Wu and Xia (2016) after 2008 in order to take into account the effects of the significant policy measures taken at the height and in the aftermath of the GFC on the monetary stance. Responses to monetary policy shocks show significant time variation in Figure 3.7.2. A contractionary shock generally leads to a cumulative fall in economic growth, a decline in credit growth, a fall in house price inflation, and an increase in expected shortfall. The size and transmission of monetary policy shocks, however, depends on the timing relative to the financial cycle.

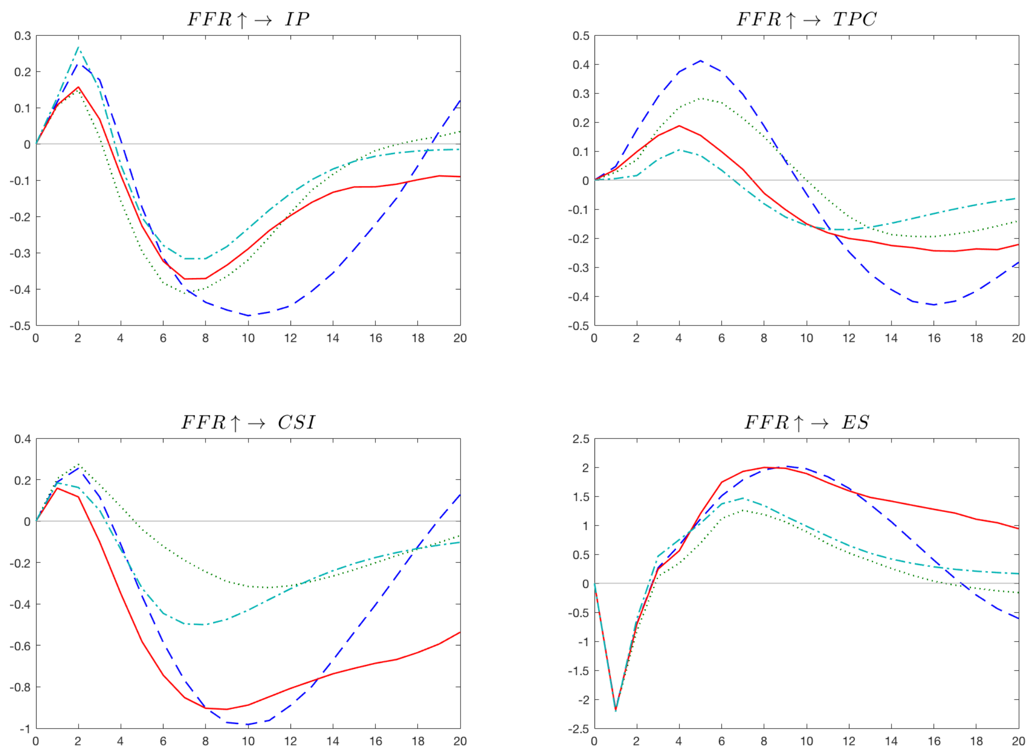

FIGURE 3.7.2: Time-varying impulse response functions. IRFs to a unit increase in the growth rate of tier 1 capital relative to risk-weighted assets for the financial intermediaries in our sample. Vertical axis: percentage change. Horizontal axis: years. Green line: 2000(1). Blue line: 2004(1). Red line: 2008(3). Cyan line: 2015(3).

An appropriately-timed contractionary shock relatively early in the upswing of the financial cycle will front-load costs from lower economic growth while averting the inflation in house prices driven by credit growth, and ultimately resulting in a lower expected shortfall for intermediaries. For instance, a contractionary shock in 2000 could potentially exacerbate the recession following the Dotcom bubble, but this would be offset by future losses by curbing the credit-fuelled growth in housing prices and ensuing recession in 2007-09. On the other hand, an obviously mis-timed policy intervention could wreak havoc. For instance, a contractionary shock at the height of the Global Financial crisis in 2008 produces a significant loss in growth with no expectation of future gains. As expectations over growth in housing prices and future economic growth are correcting, a contractionary shock will render debt repayment more difficult. This will increase default on loans—and thereby expected shortfall among intermediaries—thereby exacerbating the persistent downward spiral between default and housing prices. Surprisingly, the effect on credit growth in 2008 is relatively muted, though this perhaps may be related to other policy interventions (see Table 3.5.1) that are not captured by the shadow policy rate.

In some cases, appropriate timing is more difficult to determine. For instance, intuition would suggest that any shock prior to the correction of expectations in the financial cycle would trade short-term losses for long-term gains. Yet, by examining the effect of such a shock in 2004 at the peak of the financial cycle, we find that future damage to the economy is relatively larger than intervening at other periods. This occurs as credit growth actually increases in response to a contractionary shock, leading to a much larger decline in housing prices relative to a contractionary shock early in the upswing of the financial cycle. The upward momentum in expectations of future economic performance will lead agents to take on more credit despite increased borrowing rates. This exacerbates the correction in prices and credit further as more loans will result in default. The evidence in Figure 3.7.2 suggests that the damage wrought by increasing rates in 2004 is likely greater than increasing rates at the height of the Global Financial Crisis in 2008.

The final case we examine represents a policy intervention while relative stability has been restored. A monetary policy shock in 2015 has similar but relatively milder effects on the economy. The most notable effect is on house prices. The large contraction in house prices relative to other stages of the financial cycle is likely due to a less liquid housing market. This interpretation is supported by the mild effect on credit growth. For instance, the milder contraction in house prices in 2000 was accompanied by an increase in credit that is non-existent in 2015. As attitudes to risk have adjusted over the financial cycle, there is less willingness among financial intermediaries to provide credit buffers. There is less capacity for re-financing, and thus any agents engaged in speculative financing of their operations will suffer disproportionately in the aftermath of a financial correction.

3.7.3 Policy conclusions

We briefly survey our conclusions for policy intervention. These can be separated into efforts for reducing the likelihood or depth of losses from financial instability during the upswing, as well as crisis management tools during the correction.

With regard to macroprudential requirements, early counter-cyclical tightening will curb the extent of credit growth as well as inflation in house prices. The net effect on economic growth is positive, and results from trading short-term losses for long-term gains. During periods of financial instability, however, a loosening of macroprudential requirements can support credit growth when this is most needed and prevent the collapse in house prices. These results reflects the broad consensus in the literature on the use of counter-cyclical macroprudential policy measures (Goodhart et al., 2013; Dell' Ariccia et al., 2017). More work is necessary, however, to determine precisely how macroprudential and monetary policy interact. For instance, the purchase of toxic assets through unconventional monetary policy will remove these from the balance sheets of intermediaries, thereby manufacturing a drop in the capital ratio.

Our findings support calls for expansionary monetary policy during periods of financial instability as a method of mitigating the fallout from a potential debt-deflation spiral (Mishkin, 2009; Gertler and Karadi, 2011). Unconventional intervention that loosens the stance of monetary policy has proven to be somewhat effective. There is also evidence to suggest the presence of a debt-deflation channel of monetary policy

(Lin et al., 2015). A poorly-timed contractionary shock in monetary policy employed to curb over-leveraging relatively late in the upswing of the financial cycle exacerbates losses through an initial rise and contraction in credit, paired with a deep contraction in house prices.

We have attempted to address issues of endogeneity in the identification of policy shocks using a standard approach, while also accounting for unconventional monetary policy. Our efforts to identify the effects of policy interventions are limited, however, by the short sample over which data are available, as well as through our basic identification strategy. Additionally, we have implicitly assumed that positive and negative shocks have effects of the same magnitude, whereas this may not always be the case. Nonetheless, we draw these preliminary conclusions from the current framework in order to aide the direction of future work.

3.8 Concluding remarks

We examine the boom and bust of the most recent financial cycle in the United States and reconcile it empirically with time-varying expectations. The economy is modelled as a macro-financial TVP-VAR augmented with policy variables over the period spanning 1999–2015. The likelihood and size of systemic risk is accounted for with the inclusion of a CDS-based measure that incorporates structural nonlinearities between intermediaries. Policy shocks are identified through a recursive identification assumption, where the appropriate policy variables are ordered last. In order to account for the stance of monetary policy beyond the zero lower bound, we use an adjusted Federal Funds rate derived from a shadow rate term structure model.

The analysis indicates that linkages vary significantly over time. Variation in shock sizes suggests the presence of a volatility paradox that likely contributes to the formation of optimistic expectations during the upswing of the financial cycle.

The effect of transmission channels also changes as optimism rises, initiating riskier investments that translate into an upward spiral in credit and house prices. We find that policy interventions carry different costs when conducted at different stages of the financial cycle, so that any actions must be carefully timed. Early intervention in macroprudential requirements can prevent over-leveraging at the cost of higher expected shortfall and lower economic growth in the short term, while gains are made to compensate for these costs in the long term. Monetary policy is most effective as a crisis management tool. Attempting to disinflate house prices and curb credit growth can have negative repercussions through the debt-deflation channel, thereby increasing the likelihood and severity of a downturn.

This framework can potentially be used to provide policy recommendations that depend on the stage of the financial cycle. In order to improve the reliability of these policy recommendations without running into the Lucas critique, however, further work on causal transmission of shocks is necessary. In this paper, identification was achieved through a recursive ordering of variables in order to orthogonalise contemporaneous effects. This approach necessarily assumes structure in order to identify shocks. An interesting question to answer in future work would be to assess how structure changes in response to identified shocks through alternative approaches.

Appendix

3.A Data

The data used in the main analysis are reported in Table 3.A.1. Nominal total credit is constructed by adding the series for interbank loans (INTB) and total credit to the private non-financial sector (TPNF). Total credit and house prices (CSI) are deflated by the consumer price index (CPI). The data used for estimation are plotted in log levels in Figure 3.A.1 and in log differences in Figure 3.A.2, except for the adjusted Federal Funds rate.

TABLE 3.A.1: Data series used for estimation.

Series	Description	Source
CSI	S&P/ Case-Shiller US national home price index, 2000 = 100.	FRED
CPI	Consumer price index, 1982–1984 = 100.	FRED
ES	Expected shortfall for selected financial intermediaries, thousand US dollars.	IMF
FFR	Effective Federal Funds rate, percentage.	FRED
INTB	Interbank loans for all commercial banks, billions US dollars.	FRED
IP	Industrial production index, 2012 = 100.	FRED
JPOD	Joint probability of default for selected financial intermediaries, percentage.	IMF
TPNF	Total credit to the private non-financial sector adjusted for breaks, billions US dollars.	BIS
WXSR	Adjusted Federal Funds rate, percentage.	Wu and Xia (2016)

Notes: All macroeconomic data are reported in levels and seasonally adjusted (X12ARIMA).

3.B Linear VAR

As a comparison to the TVP-VAR results, we examine a linear model of the same system of variables as in the main model of the paper. Once again, these are

$$\mathbf{y}_t = [\Delta y_t, \Delta c_t, \Delta p_t, \Delta s_t, \Delta r_t, i_t,]',$$

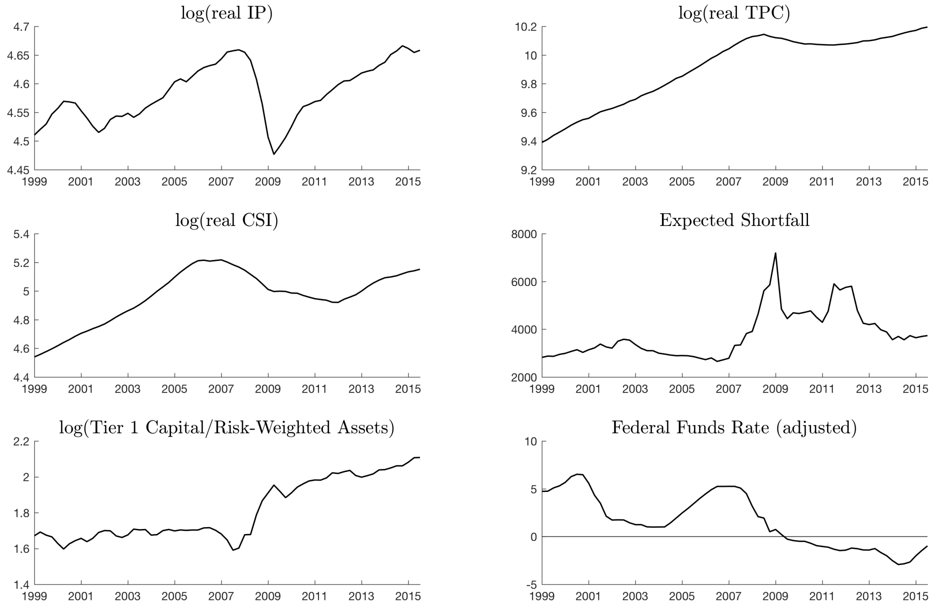

FIGURE 3.A.1: Data. All series are in (log) levels and seasonally adjusted where necessary.

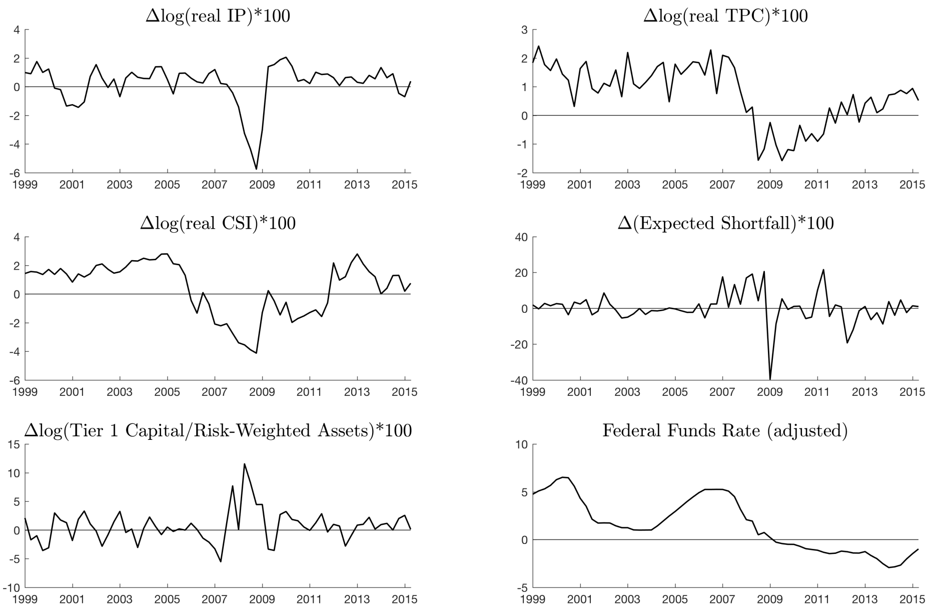

FIGURE 3.A.2: Data. All series are first-differences of log variables, except for the adjusted Federal Funds rate.

where y_t denotes log real industrial production (IP), c_t denotes log real total credit to the private sector (TPC), p_t denotes the log real Case-Shiller national house price index (CSI), s_t denotes the expected shortfall for the portfolio of selected financial intermediaries, r_t denotes the (log) tier 1 capital to risk-weighted assets ratio for the same intermediaries, and i_t denotes the adjusted Federal Funds rate (FFR). The model is estimated with two lags over the same sample period, 1999(1)-2015(4). Impulse responses of one unit (percentage point) to the economic and financial variables are shown in Figure 3.B.1, while the responses from shocks to the policy variables are shown in Figure 3.B.2.

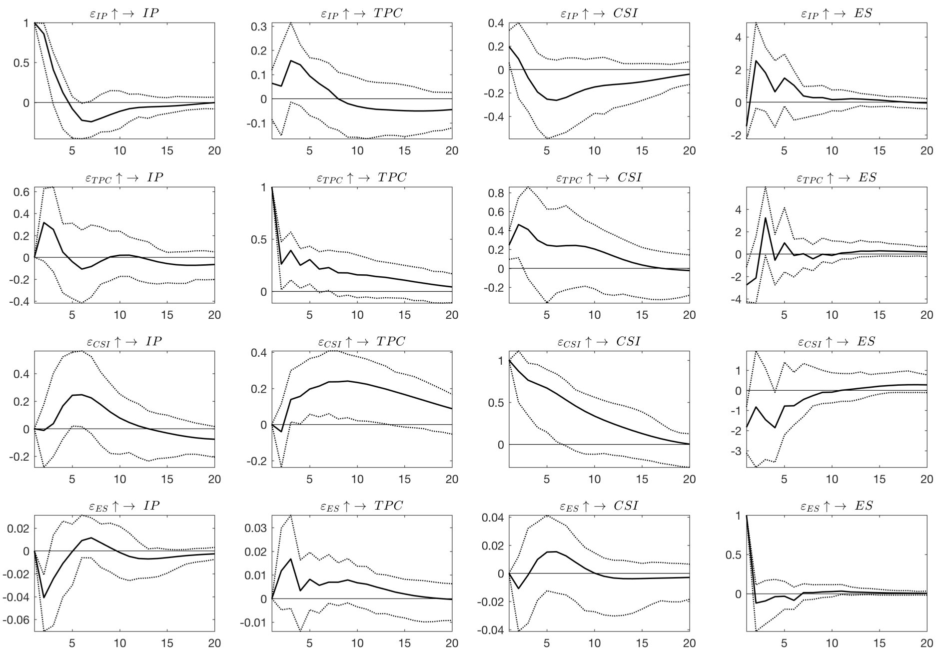

FIGURE 3.B.1: Linear impulse response functions. IRFs to a unit increase in variables along rows, with responses along columns. 95% bootstrapped confidence intervals are provided. Vertical axis: percentage change. Horizontal axis: horizon (quarters).

The linear impulse responses are broadly consistent with the main elements of the financial instability hypothesis. Positive shocks to economic growth (measured via IP) generate an initial positive response in all response variables, while shocks to credit and house prices feed back and reinforce each other, thereby generating greater persistence in the own-responses along the diagonal of Figure 3.B.1. Positive shocks to economic growth increase expected shortfall as positive expectations of future economic performance are reinforced, while positive shocks to house prices reduce expected shortfall as the probability of default among intermediaries declines. Generally, the linear impulse responses in Figure 3.B.1 appear to correspond to the average response in Figures 3.6.2-3.6.5. Notable differences in the effect of economic growth on house prices and the effect of credit on economic growth likely correspond to the absence of time-variation in the autoregressive parameters and stochastic volatility.

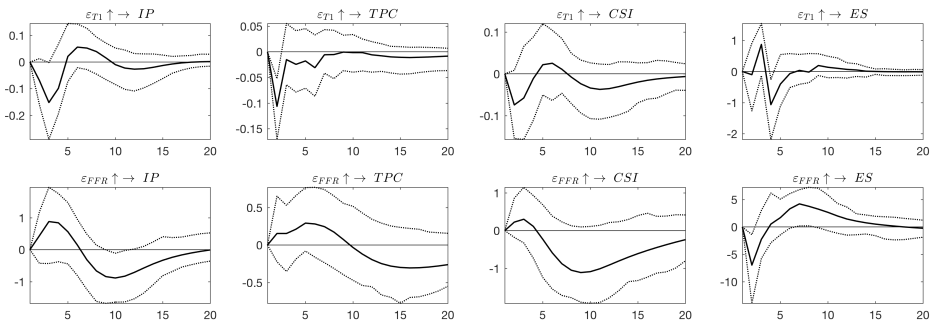

FIGURE 3.B.2: Linear impulse response functions. IRFs to a unit increase in variables along rows, with responses along columns. 95% bootstrapped confidence intervals are provided. Vertical axis: percentage change. Horizontal axis: horizon (quarters).

The linear impulse responses to shocks in the policy variables, plotted in Figure 3.B.2, correspond broadly to the average response in Figures 3.7.1-3.7.2. An important distinction, however, is that a linear model would imply that a policy intervention either impacting the capital ratio or the policy rate would have the same effect at all points of the financial and business cycles. Given the important time-variation in the impulse responses shown in Figures 3.7.1-3.7.2, this assumption of constancy becomes untenable. Policy recommendations must be contingent upon the state of the financial system and the real economy.

3.C Estimation algorithm

The objects to be estimated in the model are the histories of the coefficients $\mathbf{b}$ and $\mathbf{a}$ and volatilities $\mathbf{h}$, as well as the variance-covariance matrices Σ_b, Σ_a, and Σ_h. Let the full sample of data be represented as $\mathbf{y} = \{\mathbf{y}_t\}_{t=1}^T$ and let $\pi(\cdot|\cdot)$ denote a conditional density. The MCMC algorithm follows Nakajima (2011), iterating through the following steps.

1. Initialise $\mathbf{b}$, $\mathbf{a}$, $\mathbf{h}$, Σ_b, Σ_a, and Σ_h.

2. Sample $\mathbf{b}$ from $\pi(\mathbf{b}|\mathbf{a}, \mathbf{h}, \Sigma_b, \mathbf{y})$.

3. Sample Σ_b from $\pi(\Sigma_b|\mathbf{b})$.

4. Sample $\mathbf{a}$ from $\pi(\mathbf{a}|\mathbf{b}, \mathbf{h}, \Sigma_a, \mathbf{y})$.

5. Sample Σ_a from $\pi(\Sigma_a|\mathbf{a})$.

6. Sample $\mathbf{h}$ from $\pi(\mathbf{h}|\mathbf{b}, \mathbf{a}, \Sigma_h, \mathbf{y})$.

7. Sample Σ_h from $\pi(\Sigma_h|\mathbf{h})$.

8. Go to step 2 and repeat.

For further details on the estimation algorithm, the interested reader is referred to Nakajima (2011). Convergence of the algorithm is assessed through the diagnostic checks described in Appendix 3.D.

3.D Diagnostics

This section provides diagnostic outputs to show convergence of the sampling algorithm used to estimate the TVP-VAR. Table 3.D.1 reports the means, standard deviations, and 95% intervals of all estimated parameters, averaged over the entire sample. Additionally, two diagnostics are provided to demonstrate the convergence of the Markov chain. These are the convergence diagnostic (CD) of Geweke (1992) and the inefficiency factor (IF) of Chib (2001).

The convergence diagnostic provides a test for whether the posterior distribution of each parameter has converged. This involves comparing the sample mean and variance for the first T_0 draws and the last T_1 draws from the Markov chain. For each parameter j, this is computed as

$$CD_j = \frac{\overline{x}_{T_0} - \overline{x}_{T_1}}{\sqrt{\hat{\sigma}^2_{T_0}/T_0 - \hat{\sigma}^2_{T_1}/T_1}},$$

where $\overline{x}_i$ and $\hat{\sigma}^2_i$ represent sample means and variances over the first T_0 and last T_1 draws of the MCMC sample. We set both $T_0 = T_1 = 10,000$. The sample variance is calculated using a Parzen window with bandwidth equal to 500. If the posterior distribution has converged, then Geweke (1992) shows that $CD_j \to N(0, 1)$. The columns labeled "CD" in Table 3.D.1 are p-values for the null hypothesis of convergence.

The inefficiency factor is a measure for how well the Markov chain mixes. It approximates the ratio between the variance of the posterior mean from our estimation results relative to the variance of the posterior mean from hypothetical uncorrelated draws. The inefficiency factor for the j^{th} parameter is computed as

$$IF_j = 1 + 2\sum_{l=1}^{L} \rho_j(l),$$

where $l = 1, \ldots, L$ are the lags over which we calculate the sample autocorrelation. We set $L = 500$.

Table 3.D.1 should be read as follows. Parameters in $\mathbf{b}$ are stacked by equation, starting with the intercept, then followed by the lags in increasing order. Parame-

ters in **a** represent the stacked, lower-triangular half of the contemporaneous impact matrix A_t, while parameters in **h** are the stacked log volatilities. Where parameters are prefixed with an "s", these represent the hyperparameters in Σ_b, Σ_a, and Σ_h. For further details, refer to §3.4.1.

For all estimated parameters, we can say at the 5% level that the posterior distribution has converged. Additionally, the diagnostic results indicate efficient sampling from the Markov chain. All inefficiency factors are substantially below the recommended value of 20 (Primiceri, 2005), with most hovering between 1 and 5.

TABLE 3.D.1: Average parameter estimates and diagnostic tests.

Parameter	Mean	Stdev	CD	IF	Parameter	Mean	Stdev	CD	IF
b1	0.2778	0.1276	0.669	1.49	sb1	0.3989	0.018	0.397	8.44
b2	0.6864	0.1159	0.948	1.41	sb2	0.3992	0.0182	0.73	8.9
b3	0.1242	0.1231	0.101	1.32	sb3	0.3993	0.018	0.396	8.89
b4	-0.1006	0.0977	0.921	1.3	sb4	0.3983	0.0178	0.691	8.96
b5	-0.0524	0.0706	0.861	1.12	sb5	0.3941	0.0167	0.048	7.17
b6	-0.0302	0.0814	0.863	1.15	sb6	0.3982	0.0179	0.871	8.51
b7	0.1033	0.1623	0.109	1.48	sb7	0.3977	0.0177	0.896	8.58
b8	-0.2332	0.111	0.859	1.37	sb8	0.3988	0.0181	0.559	8.5
b9	-0.0557	0.1167	0.836	1.73	sb9	0.3992	0.018	0.941	8.1
b10	0.0671	0.0968	0.439	1.28	sb10	0.3981	0.0178	0.786	8.45
b11	-0.0059	0.0709	0.48	1.06	sb11	0.3921	0.0162	0.204	7.15
b12	-0.0349	0.0801	0.509	1.17	sb12	0.398	0.0177	0.572	9.86
b13	-0.1513	0.1611	0.213	1.46	sb13	0.3976	0.0176	0.994	8.63
b14	0.2405	0.114	0.71	1.25	sb14	0.3988	0.0181	0.517	9.4
b15	-0.0734	0.1052	0.518	1.33	sb15	0.3987	0.018	0.223	7.94
b16	0.181	0.1118	0.855	1.3	sb16	0.3986	0.0179	0.81	8.04
b17	-0.0466	0.0993	0.262	1.27	sb17	0.3976	0.0178	0.948	8.97
b18	0.0036	0.0705	0.584	1.13	sb18	0.3918	0.0163	0.837	7.12
b19	-0.094	0.0812	0.388	1.24	sb19	0.3971	0.0175	0.854	8.96
b20	0.0176	0.1318	0.249	1.47	sb20	0.3974	0.0176	0.776	8.47
b21	0.0654	0.1036	0.806	1.31	sb21	0.3985	0.0179	0.455	10.06
b22	0.3272	0.1083	0.82	1.47	sb22	0.3988	0.0179	0.768	8.67
b23	0.2045	0.0953	0.338	1.24	sb23	0.3981	0.0178	0.799	7.58
b24	0.0136	0.0706	0.721	1.11	sb24	0.3915	0.0162	0.671	7.39
b25	0.0068	0.08	0.42	1.17	sb25	0.3976	0.0178	0.919	8.27
b26	0.0089	0.1351	0.473	1.46	sb26	0.3971	0.0174	0.644	7.77
b27	0.1941	0.1359	0.03	1.31	sb27	0.3991	0.0182	0.173	9.24
b28	-0.0311	0.1165	0.662	1.34	sb28	0.3989	0.0181	0.662	8.9
b29	0.0418	0.1191	0.764	1.47	sb29	0.3992	0.0181	0.923	9.45
b30	0.796	0.1182	0.149	1.42	sb30	0.3992	0.0181	0.309	8.22
b31	-0.0135	0.0709	0.636	1.05	sb31	0.3924	0.0165	0.178	8.6
b32	-0.0701	0.0824	0.705	1.24	sb32	0.3974	0.0175	0.457	8.32

Continued on next page

Parameter	Mean	Stdev	CD	IF	Parameter	Mean	Stdev	CD	IF
b33	0.1449	0.1567	0.6	1.57	sb33	0.3981	0.0176	0.932	8.78
b34	-0.1126	0.1115	0.309	1.22	sb34	0.3989	0.018	0.718	9.85
b35	0.1135	0.1221	0.972	1.52	sb35	0.3988	0.018	0.433	8.8
b36	-0.0029	0.1187	0.417	1.34	sb36	0.3987	0.0181	0.093	9.19
b37	0.0058	0.0705	0.847	1.1	sb37	0.3939	0.0167	0.218	7.33
b38	0.0369	0.0801	0.343	1.48	sb38	0.3977	0.0177	0.236	8.38
b39	-0.2614	0.1643	0.721	1.53	sb39	0.3984	0.0177	0.367	7.81
b40	-0.9654	0.8554	0.426	2.54	sb40	0.3991	0.018	0.273	8
b41	1.2259	0.5122	0.104	3.65	sb41	0.3995	0.018	0.604	8.61
b42	0.1489	0.4282	0.647	2.17	sb42	0.3991	0.0182	0.132	8.71
b43	-0.6777	0.5576	0.644	2.53	sb43	0.3993	0.0182	0.634	9.01
b44	0.0732	0.0948	0.652	1.88	sb44	0.399	0.0181	0.915	7.9
b45	-0.4034	0.1576	0.714	1.57	sb45	0.3991	0.0179	0.336	7.61
b46	-2.133	1.0372	0.536	3.86	sb46	0.3991	0.0182	0.303	9.86
b47	0.6755	0.4828	0.758	3.39	sb47	0.3994	0.018	0.381	8.83
b48	0.3152	0.5916	0.294	4.84	sb48	0.3992	0.0181	0.525	8.56
b49	-0.601	0.5335	0.563	2.35	sb49	0.3991	0.0181	0.86	9.09
b50	0.0075	0.0911	0.029	1.6	sb50	0.3994	0.018	0.889	8.06
b51	0.7731	0.1589	0.407	2.54	sb51	0.3994	0.0183	0.723	9.51
b52	2.6584	1.0842	0.458	3.86	sb52	0.3992	0.018	0.732	8.25
b53	1.1916	0.2136	0.984	1.51	sb53	0.3991	0.018	0.668	9.03
b54	-0.666	0.219	0.705	1.99	sb54	0.3992	0.0181	0.871	8.77
b55	-0.3064	0.2327	0.657	1.86	sb55	0.3991	0.0182	0.797	9.19
b56	-0.0562	0.1926	0.52	1.77	sb56	0.399	0.0181	0.868	7.86
b57	0.0948	0.0742	0.449	1.4	sb57	0.4001	0.0185	0.873	10.81
b58	0.1542	0.105	0.623	1.58	sb58	0.3991	0.018	0.794	8.54
b59	-0.5	0.3572	0.96	1.95	sb59	0.3988	0.0181	0.497	8.34
b60	0.0812	0.2054	0.582	2.06	sb60	0.3994	0.0181	0.43	8.5
b61	0.2751	0.2418	0.573	2.29	sb61	0.3993	0.018	0.235	8.49
b62	-0.0178	0.1804	0.504	1.77	sb62	0.3988	0.018	0.998	8.53
b63	0.01	0.0744	0.899	1.46	sb63	0.3961	0.0174	0.595	8.12
b64	-0.2679	0.1031	0.693	1.86	sb64	0.3991	0.0182	0.916	8.98
b65	0.2932	0.3665	0.887	2.23	sb65	0.3989	0.0179	0.303	8.86
b66	-0.0354	0.1022	0.658	1.44	sb66	0.3989	0.018	0.307	7.9
b67	0.1072	0.0996	0.411	1.62	sb67	0.3989	0.0181	0.597	7.96
b68	0.0126	0.1025	0.922	1.73	sb68	0.3985	0.0181	0.956	8.15
b69	-0.1022	0.0948	0.306	1.52	sb69	0.3983	0.0177	0.996	8.67
b70	-0.0003	0.0718	0.346	1.45	sb70	0.3919	0.0161	0.138	6.35
b71	-0.0388	0.0834	0.439	1.86	sb71	0.3968	0.0174	0.36	7.13
b72	1.6078	0.1119	0.786	1.42	sb72	0.397	0.0174	0.896	7.1
b73	-0.0744	0.0973	0.868	1.3	sb73	0.3984	0.0177	0.603	7.95
b74	0.0896	0.1039	0.847	1.54	sb74	0.3989	0.018	0.453	8.57
b75	0.0956	0.0929	0.521	1.59	sb75	0.3981	0.0179	0.174	9.19
b76	0.0016	0.0715	0.638	1.38	sb76	0.3917	0.0163	0.633	7.88
b77	0.0426	0.0826	0.369	1.9	sb77	0.3974	0.0175	0.783	7.2
b78	-0.6972	0.113	0.431	1.88	sb78	0.3969	0.0175	0.997	7.95
a1	0.027	0.3106	0.952	1.53	sa1	0.5039	0.0186	0.987	6.6
a2	-0.075	0.2951	0.243	1.58	sa2	0.5026	0.0179	0.692	6.1
a3	-0.0489	0.3448	0.674	1.54	sa3	0.503	0.018	0.248	6.01
a4	0.0814	0.604	0.125	2.15	sa4	0.5017	0.0178	0.999	6.16

Continued on next page

Parameter	Mean	Stdev	CD	IF	Parameter	Mean	Stdev	CD	IF
a5	-0.1723	0.694	0.709	1.74	sa5	0.5017	0.0178	0.073	7.11
a6	0.2946	0.7164	0.892	2.1	sa6	0.5014	0.0177	0.226	6.25
a7	0.3411	0.42	0.91	1.87	sa7	0.5019	0.0178	0.226	5.04
a8	0.1222	0.467	0.963	1.59	sa8	0.5018	0.0178	0.512	6.17
a9	0.1989	0.3846	0.902	1.9	sa9	0.5017	0.0178	0.958	5.86
a10	-0.0625	0.0669	0.324	1.84	sa10	0.4969	0.0167	0.159	4.97
a11	-0.0021	0.2975	0.972	1.43	sa11	0.5036	0.0182	0.848	6.46
a12	-0.0132	0.3278	0.843	1.33	sa12	0.5034	0.0181	0.394	7.05
a13	-0.0198	0.2801	0.964	1.57	sa13	0.5035	0.0185	0.971	6.83
a14	-0.0192	0.037	0.643	1.97	sa14	0.4986	0.0187	0.956	5.46
a15	-0.0007	0.107	0.915	1.73	sa15	0.5017	0.0183	0.836	5.23
h1	-0.0197	0.3915	0.597	1.74	sh1	0.7838	0.0226	0.54	4.27
h2	-0.3268	0.4166	0.529	1.89	sh2	0.7836	0.0226	0.674	4.01
h3	-0.2744	0.358	0.626	1.91	sh3	0.7837	0.0226	0.766	3.95
h4	2.9882	0.188	0.241	3.33	sh4	0.796	0.024	0.941	5.92
h5	0.9738	0.2468	0.756	3.1	sh5	0.79	0.0235	0.916	4.68
h6	-1.0913	0.4415	0.854	5.38	sh6	0.7832	0.0229	0.721	4.36

Conclusion

1 Final concluding remarks

In this thesis, the main components for future research in the causal transmission of shocks and their effects on financial stability are assembled. The first chapter introduced the idea of causal transmission, in the sense of large, exogenous interventions that induce structure within a system of observed variables. For the bivariate case, three conditions are required in order to show unique causal transmission of the catalyst w through z to y. First, a non-trivial Markov structure is necessary. This means that the factorisation $\mathsf{f}(y, z|w) = \mathsf{f}(y|z)\mathsf{f}(z|w)$ must be non-trivial, in the sense that y depends on z and z depends on w. Second, the transmission of the catalyst itself must be non-trivial, in the sense that y depends on w. Third, we must assume that the effect of w is causal and hence asymmetric.

The second chapter demonstrated the existence of robust and fragile financing regimes and that the dynamics of the macro-financial relationship change over time. High fragility is characterised by high default risk and low profitability among intermediaries, while both economic output and credit extensions are substantially weaker. Amplification of shocks during the high-fragility regime depends primarily on default risk and profitability, suggesting the presence of a debt-deflation channel. Additionally, despite the fact that our model is not equipped to make policy recommendations, there is evidence to suggest that an intervention on credit in the high-fragility regime may be effective in reducing losses that could potentially be incurred.

The third chapter built on the ideas of the second chapter by examining how the economy transitions from a state of financial robustness to one of financial fragility. The economy is modelled as a macro-financial TVP-VAR augmented with policy variables over the period spanning 1999–2015. The analysis suggests that linkages vary significantly over time. Variation in shock sizes suggests the presence of a volatility paradox that likely contributes to the formation of optimistic expectations during the upswing of the financial cycle. The effect of transmission channels also changes as optimism rises, initiating riskier investments that translate into an upward spiral in credit and house prices. We find that policy interventions carry different costs when conducted at different stages of the financial cycle, so that any actions must be carefully timed. Early intervention in macroprudential requirements can prevent over-leveraging at the cost of higher expected shortfall and lower economic growth in the short term, while gains are made to compensate for these costs in the long term. Monetary policy is most effective as a crisis management tool. Attempting to disinflate house prices and curb credit growth can have negative repercussions through the debt-deflation channel, thereby increasing the likelihood and severity of a downturn.

2 Future research agenda

The objective of future research will be to extend and combine the elements of this thesis in order to provide useful policy recommendations. As a first step, the theory in Chapter 1 must be extended to accommodate models of higher dimensionality. Subsequently, it can be applied both in linear and nonlinear models, where recursiveness assumptions can be tested as causal transmissions of shocks, rather than made *a priori* for identification purposes. The combination of these two lines of research will enable us to examine whether and how the causal transmission of shocks evolves over time.

2.1 Causal transmission in high-dimensional reduced-form models

One of the limitations of the work laid out in Chapter 1 is that the theory is proven only for a bivariate case. Finding conditions to guarantee causal transmission with three or more variables is compounded by an additional problem that is related to the grouping of variables. Suppose that a catalyst w intervenes on a system of three observed variables x, y, z defined on a product space with positive density $\mathsf{f}(x, y, z|w)$. In addition to establishing whether the causal transmission $w \to z \to y$ holds, we must now also assign x to a group. For instance, distinct conditions must be established for the cases where, alternatively, $w \to \{z, x\} \to y$, $w \to z \to \{y, x\}$, and $w \to x \to z \to y$. These do not represent an exhaustive list of possibilities but serve to illustrate the issue.

This problem is resolved by the assumption of expert knowledge in the work of Cox and Wermuth (1996) and Wermuth and Sadeghi (2012). In their work, the context is one of biomedical trials where groups of variables can be set on equal footing or otherwise through a sequence of regressions. In economics, treatment and effect are not as clear. Different theories that have explicit structural implications are often in conflict with each other, and so an agnostic approach based on statistical inference is largely necessary. Thus, we need a larger set of conditions to ensure uniqueness of causal transmission, given the grouping problem.

2.2 Applications to macroeconometric models

While nonlinear VAR models permit rich dynamics among the observed variables, all residual-based analysis is constrained by a necessary identifying assumption that usually comes in the form of recursiveness. This approach is limiting in at least two ways. First, the identification of shocks from residuals does not mean that we are calculating structural shocks to the variable of interest. Rather, it means that we are calculating a shock that is contemporaneously uncorrelated with higher-ordered shocks. Second, the recursive identification assumption implies that we are imposing a restrictive causal structure over the entire sample. Despite the flexibility that time-varying parameters afford, ordering is assumed constant.

The application of our method for causal transmission depends crucially on being able to identify shocks in the historical narrative. Whether these correspond to very large shocks that are easily observed in the data or small shocks that require closer attention, exogeneity must be satisfied in order to justify our causal assumption. Once this barrier is eliminated, for instance by applying the method of Romer and Romer (1989) broadly beyond monetary policy identification, we will be able to test for time-varying causal transmission.

References

Abbate, A., Eickmeier, S., Lemke, W., and Marcellino, M. (2016). The Changing International Transmission of Financial Shocks: Evidence from a Classical Time-Varying FAVAR. *Journal of Money, Credit and Banking*, 48(4):573–601.

Acharya, V. V., Pedersen, L. H., Philippon, T., and Richardson, M. P. (2012). Measuring Systemic Risk. CEPR Discussion Papers 8824, C.E.P.R. Discussion Papers.

Adrian, T. and Brunnermeier, M. K. (2016). CoVaR. *American Economic Review*, 106(7):1705–1741.

Aikman, D., Lehnert, A., Liang, J. N., and Modugno, M. (2016). Financial Vulnerabilities, Macroeconomic Dynamics, and Monetary Policy. Finance and Economics Discussion Series 2016-055, Board of Governors of the Federal Reserve System (U.S.).

Alessi, L. and Detken, C. (2011). Quasi real time early warning indicators for costly asset price boom/bust cycles: A role for global liquidity. *European Journal of Political Economy*, 27(3):520–533.

Angrist, J. D. and Krueger, A. B. (2001). Instrumental variables and the search for identification: From supply and demand to natural experiments. *Journal of Economic Perspectives*, 15(4):69–85.

Aspachs, O., Goodhart, C., Tsomocos, D., and Zicchino, L. (2007). Towards a measure of financial fragility. *Annals of Finance*, 3(1):37–74.

Balke, N. S. (2000). Credit and Economic Activity: Credit Regimes and Nonlinear Propagation of Shocks. *The Review of Economics and Statistics*, 82(2):344–349.

Bazinas, V. and Nielsen, B. (2015). Causal transmission in reduced-form models. Economics Papers 2015-W07, Economics Group, Nuffield College, University of Oxford.

Benigno, G., Chen, H., Otrok, C., Rebucci, A., and Young, E. R. (2013). Financial crises and macro-prudential policies. *Journal of International Economics*, 89(2):453–470.

Bernanke, B. and Gertler, M. (1989). Agency Costs, Net Worth, and Business Fluctuations. *American Economic Review*, 79(1):14–31.

Bernanke, B., Gertler, M., and Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1, Part C, chapter 21, pages 1341–1393. Elsevier, 1^{st} edition.

Bernanke, B. and Lown, C. S. (1991). The credit crunch. *Brookings Papers on Economic Activity*, 22(2):205–248.

Bernanke, B. S. and Blinder, A. S. (1992). The Federal Funds Rate and the Channels of Monetary Transmission. *American Economic Review*, 82(4):901–921.

Bhattacharya, S., Goodhart, C. A., Tsomocos, D. P., and Vardoulakis, A. P. (2015). A Reconsideration of Minsky's Financial Instability Hypothesis. *Journal of Money, Credit and Banking*, 47(5):931–973.

Bianchi, J. (2011). Overborrowing and Systemic Externalities in the Business Cycle. *American Economic Review*, 101(7):3400–3426.

Bianchi, J. and Mendoza, E. G. (2010). Overborrowing, Financial Crises and 'Macro-prudential' Taxes. NBER Working Papers 16091, National Bureau of Economic Research, Inc.

Billio, M., Getmansky, M., Lo, A. W., and Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics*, 104(3):535–559.

BIS (2010). Basel III: A global regulatory framework for more resilient banks and banking systems. *Bank for International Settlements*.

Bisias, D., Flood, M., Lo, A. W., and Valavanis, S. (2012). A Survey of Systemic Risk Analytics. *Annual Review of Financial Economics*, 4(1):255–296.

Blinder, A. S. (2010). It's Broke, Let's Fix It: Rethinking Financial Regulation. *International Journal of Central Banking*, 6(34):277–330.

Blinder, A. S. (2015). What Did We Learn from the Financial Crisis, the Great Recession, and the Pathetic Recovery? *The Journal of Economic Education*, 46(2):135–149.

Bordo, M. D. and Haubrich, J. G. (2010). Credit crises, money and contractions: An historical view. *Journal of Monetary Economics*, 57(1):1–18.

Bordo, M. D. and Landon-Lane, J. (2014). Does Expansionary Monetary Policy Cause Asset Price Booms? Some Historical and Empirical Evidence. In Bauducco, S., Christiano, L., and Raddatz, C., editors, *Macroeconomic and Financial Stability: challenges for Monetary Policy*, volume 19 of *Central Banking, Analysis, and Economic Policies Book Series*, chapter 3, pages 61–116. Central Bank of Chile.

Borio, C. and Drehmann, M. (2009). Assessing the risk of banking crises - revisited. *BIS Quarterly Review*.

Borio, C. and Zhu, H. (2012). Capital regulation, risk-taking and monetary policy: A missing link in the transmission mechanism? *Journal of Financial Stability*, 8(4):236–251.

Boyd, J. H. and Gertler, M. (1994). The role of large banks in the recent U.S. banking crisis. *Federal Reserve Bank of Minneapolis Quarterly Review*, 18(1):2–21.

Bårdsen, G., Kolsrud, D., and Nymoen, R. (2012). Forecast robustness in macroeconometric models. Working paper, Norwegian University of Science and Technology.

Brunnermeier, M. K. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.

Cappé, O., Moulines, E., and Ryden, T. (2005). *Inference in Hidden Markov Models*. Springer-Verlag New York.

Carlstrom, C. T. and Fuerst, T. S. (1997). Agency Costs, Net Worth, and Business Fluctuations: A Computable General Equilibrium Analysis. *American Economic Review*, 87(5):893–910.

Carron, A. S. and Friedman, B. M. (1982). Financial Crises: Recent Experience in U.S. and International Markets. *Brookings Papers on Economic Activity*, 13(2):395–422.

Cartwright, N. (1989). *Nature's Capacities and Their Measurement*. Clarendon Press, Oxford.

Catarineu-Rabell, E., Jackson, P., and Tsomocos, D. (2005). Procyclicality and the new Basel Accord - banks? choice of loan rating system. *Economic Theory*, 26(3):537–557.

Cerutti, E., Claessens, S., and Laeven, L. (2017). The use and effectiveness of macroprudential policies: New evidence. *Journal of Financial Stability*, 28(C):203–224.

Chib, S. (2001). Markov chain monte carlo methods: computation and inference. In Heckman, J. and Leamer, E., editors, *Handbook of Econometrics*, volume 5, chapter 57, pages 3569–3649. Elsevier, 1^{st} edition.

Chow, G. C. (1960). Tests of equality between sets of coefficients in two linear regressions. *Econometrica*, 28:591–605.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (1999). Monetary policy shocks: What have we learned and to what end? In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1 of *Handbook of Macroeconomics*, chapter 2, pages 65–148. Elsevier.

Ciccarelli, M., Eva, O., and Teresa, V. M. (2016). Commonalities and cross-country spillovers in macroeconomic-financial linkages. *The B.E. Journal of Macroeconomics*, 16(1):231–275.

Claessens, S., Ghosh, S. R., and Mihet, R. (2013). Macro-prudential policies to mitigate financial system vulnerabilities. *Journal of International Money and Finance*, 39(C):153–185.

Claessens, S., Kose, M. A., and Terrones, M. E. (2012). How do business and financial cycles interact? *Journal of International Economics*, 87(1):178–190.

Cogley, T. and Sargent, T. J. (2001). Evolving Post-World War II U.S. Inflation Dynamics. In *NBER Macroeconomics Annual 2001, Volume 16*, NBER Chapters, pages 331–388. National Bureau of Economic Research, Inc.

Cogley, T. and Sargent, T. J. (2005). Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S. *Review of Economic Dynamics*, 8(2):262–302.

Cox, D. R. and Wermuth, N. (1996). *Multivariate Dependencies: Models, Analysis and Interpretation*. Chapman & Hall, London.

Cox, D. W. and Wermuth, N. (2004). Causality: a statistical view. *International Statistical Review*, 72:285–305.

Crowe, C., Ariccia, G. D., Igan, D., and Rabanal, P. (2011). Policies for Macrofinancial Stability: Options to Deal with Real Estate Booms. Working paper, International Monetary Fund.

Danielsson, J., Valenzuela, M., and Zer, I. (2016). Learning from History : Volatility and Financial Crises. Finance and Economics Discussion Series 2016-093, Board of Governors of the Federal Reserve System (U.S.).

Davig, T. and Hakkio, C. (2010). What is the effect of financial stress on economic activity? *Federal Reserve Bank of Kansas City Economic Review*, 95:35–62.

Dawid, A. P. (1980). Conditional independence for statistical operations. *Annals of Statistics*, 8:598–617.

De Nicolò, G., Ariccia, G. D., Laeven, L., and Valencia, F. (2010). Monetary Policy and Bank Risk Taking. Working paper, International Monetary Fund.

Dell' Ariccia, G., Laeven, L., and Suarez, G. A. (2017). Bank leverage and monetary policy's risk-taking channel: Evidence from the united states. *The Journal of Finance*, 72(2):613–654.

Dell'Ariccia, G., Igan, D., Laeven, L., and Tong, H. (2016). Credit booms and macrofinancial stability. *Economic Policy*, 31(86):299–355.

Demiralp, S. and Hoover, K. D. (2003). Searching for the causal structure of a vector autoregression. *Oxford Bulletin of Economics and Statistics*, 65(s1):745–767.

Diebold, F. X. and Yilmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics*, 182(1):119–134.

Dubey, P., Geanakoplos, J., and Shubik, M. (2005). Default and Punishment in General Equilibrium. *Econometrica*, 73(1):1–37.

Ehrmann, M., Ellison, M., and Valla, N. (2003). Regime-dependent impulse response functions in a Markov-switching vector autoregression model. *Economics Letters*, 78(3):295–299.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50:987–1108.

Engle, R. F., Hendry, D. F., and Richard, J.-F. (1983). Exogeneity. *Econometrica*, 51:277–304.

Federico, P., Vegh, C., and Vuletin, G. (2012). Effects and role of macroprudential policy: evidence from reserve requirements based on a narrative approach. Manuscript, Norges Bank.

Fisher, I. (1933). The Debt-Deflation Theory of Great Depressions. *Econometrica*, 1(4):337–357.

Galati, G. and Moessner, R. (2017). What do we know about the effects of macroprudential policy? *Economica*, pages 1–36.

Galì, J. and Gambetti, L. (2015). The Effects of Monetary Policy on Stock Market Bubbles: Some Evidence. *American Economic Journal: Macroeconomics*, 7(1):233–257.

Gambetti, L. and Musso, A. (2017). Loan supply shocks and the business cycle. *Journal of Applied Econometrics*, 32(4):764–782.

Gauthier, C., Lehar, A., and Souissi, M. (2012). Macroprudential capital requirements and systemic risk. *Journal of Financial Intermediation*, 21(4):594–618.

Geanakoplos, J. (2010). The leverage cycle. In Acemoglu, D., Rogoff, K., and Woodford, M., editors, *NBER Macroeconomics Annual 2009, Volume 24*, pages 1–65. University of Chicago Press.

Geanakoplos, J. and Fostel, A. (2008). Leverage Cycles and the Anxious Economy. *American Economic Review*, 98(4):1211–1244.

Gertler, M. and Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17–34.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In Bernardo, J., Berger, J., Dawid, A., and Smith, J., editors, *Bayesian Statistics*, volume 4, pages 169–188. Oxford University Press.

Glocker, C. and Towbin, P. (2012). The Macroeconomic Effects of Reserve Requirements. Working papers 374, Banque de France.

Godfrey, L. G. (1978). Testing against general autoregressive and moving average error models when the regressors include lagged dependent variables. *Econometrica*, 46:1293–1302.

Goodhart, C., Romanidis, N., Tsomocos, D. P., and Shubik, M. (2016). Macromodelling, Default and Money. Discussion Paper Series 755, London School of Economics.

Goodhart, C. and Segoviano, M. (2009). Banking Stability Measures. FMG Discussion Papers 627, Financial Markets Group.

Goodhart, C., Sunirand, P., and Tsomocos, D. (2006). A model to analyse financial fragility. *Economic Theory*, 27(1):107–142.

Goodhart, C. A., Kashyap, A. K., Tsomocos, D. P., and Vardoulakis, A. P. (2012). Financial Regulation in General Equilibrium. NBER Working Papers 17909, National Bureau of Economic Research, Inc.

Goodhart, C. A., Sunirand, P., and Tsomocos, D. P. (2005). A risk assessment model for banks. *Annals of Finance*, 1(2):197–224.

Goodhart, C. A. E., Kashyap, A. K., Tsomocos, D. P., and Vardoulakis, A. P. (2013). An Integrated Framework for Analyzing Multiple Financial Regulations. *International Journal of Central Banking*, 9(1):109–144.

Gorton, G. and Metrick, A. (2012a). Getting Up to Speed on the Financial Crisis: A One-Weekend-Reader's Guide. *Journal of Economic Literature*, 50(1):128–50.

Gorton, G. and Metrick, A. (2012b). Securitized banking and the run on repo. *Journal of Financial Economics*, 104(3):425–451.

Granger, C. W. J. (1969). Investigating Causal Relations by Econometric Models and Cross-Spectral Methods. *Econometrica*, 37:424–38.

Hakkio, C. S. and Keeton, W. R. (2009). Financial stress: what is it, how can it be measured, and why does it matter? *Economic Review*, (Q II):5–50.

Hamilton, J. (2016). Macroeconomic regimes and regime shifts. *Handbook of Macroeconomics*, 2:163 – 201.

Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57(2):357–84.

Hamilton, J. D. (1990). Analysis of time series subject to changes in regime. *Journal of Econometrics*, 45(1-2):39–70.

Hartmann, P., Hubrich, K., Kremer, M., and Tetlow, R. J. (2015). Melting Down: Systemic Financial Instability and the Macroeconomy. Technical report, European Central Bank.

Hendry, D. and Mizon, G. (1993). Evaluating dynamic econometric models by encompassing the var. In Phillips, P. C., editor, *Models, Methods and Applications of Econometrics*, pages 272–300. Blackwell, Oxford.

Hendry, D. F. and Doornik, J. A. (1994). Modelling linear dynamic econometric systems. *Scottish Journal of Political Economy*, 41(1):1–33.

Hendry, D. F. and Doornik, J. A. (2014). *Empirical Model Discovery and Theory Evaluation: Automatic Selection Methods in Econometrics*. MIT Press, Cambridge, MA.

Hendry, D. F. and Nielsen, B. (2007). *Econometric Modeling: A Likelihood Approach*. Princeton University Press, Princeton-Oxford.

Hendry, D. F. and Santos, C. (2010). An automatic test of super exogeneity. In Bollerslev, T., Russell, J., and Watson, M. W., editors, *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, pages 164–193. Oxford University Press, Oxford.

Hill, A. B. (1965). The environment and disease: Association or causation? *Proceedings of the Royal Society of Medicine*, 58(5):295–300.

Hood, W. C. and Koopmans, T. C. (1953). *Studies in Econometric Methods*, volume 14. New York, NY: John Wiley.

Horváth, B. and Wagner, W. (2013). The disturbing interaction between countercyclical capital requirements and systemic risk. Mimeo, Tilburg University.

Huang, X., Zhou, H., and Zhu, H. (2011). Systemic risk contributions. In for International Settlements, B., editor, *Macroprudential regulation and policy*, volume 60 of *BIS Papers chapters*, pages 36–43. Bank for International Settlements.

Hubbard, R. G. (1998). Capital-Market Imperfections and Investment. *Journal of Economic Literature*, 36(1):193–225.

Hubrich, K. and Tetlow, R. J. (2015). Financial stress and economic dynamics: The transmission of crises. *Journal of Monetary Economics*, 70(C):100–115.

Hume, D. (1758). *Essays and Treatises on Several Subjects*. A. Millar and A. Kincaid & A. Donaldson, Strand and Edinburgh.

Hume, M. and Sentance, A. (2009). The global credit boom: Challenges for macroeconomics and policy. *Journal of International Money and Finance*, 28(8):1426–1461.

Imbens, G. W. (2014). Instrumental variables: an econometrician's perspective. *Statistical Science*, 29:323–358.

Johansen, S. (1995). *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford University Press, Oxford.

Johansen, S. and Nielsen, B. (2015). Outlier detection algorithms for least squares time series. *Scandinavian Journal of Statistics*. To appear.

Jordá, O., Schularick, M., and Taylor, A. M. (2014). The Great Mortgaging: Housing Finance, Crises, and Business Cycles. NBER Working Papers 20501, National Bureau of Economic Research, Inc.

Jordá, O., Schularick, M., and Taylor, A. M. (2016). Macrofinancial History and the New Business Cycle Facts. In *NBER Macroeconomics Annual 2016, Volume 31*, NBER Chapters, pages 213–263. National Bureau of Economic Research, Inc.

Kilian, L. and Demiroglu, U. (2000). Residual based tests for normality in autoregressions: asymptotic theory and simulations. *Journal of Business & Economic Statististics*, 18:40–50.

Kim, S. and Mehrotra, A. (2015). Managing price and financial stability objectives - what can we learn from the Asia-Pacific region? BIS Working Papers 533, Bank for International Settlements.

Kindleberger, C. P. and Aliber, R. Z. (2011). *Manias, panics, and crashes: A History of Financial Crises*. Palgrave Macmillan, 6th edition.

Kivedal, B. K. (2013). Testing for rational bubbles in the us housing market. *Journal of Macroeconomics*, 38(PB):369–381.

Kiyotaki, N. and Moore, J. (1997). Credit Cycles. *Journal of Political Economy*, 105(2):211–48.

Kremer, M., Lo Duca, M., and Holló, D. (2012). CISS - a composite indicator of systemic stress in the financial system. Working Paper Series 1426, European Central Bank.

Krolzig, H.-M. (1997). *Markov-Switching Vector Autoregressions: Modelling, Statistical Inference, and Application to Business Cycle Analysis*. Springer Verlag, Berlin.

Kuttner, K. (2001). Monetary policy surprises and interest rates: Evidence from the fed funds futures market. *Journal of Monetary Economics*, 47(3):523–544.

Lauritzen, S. L. (1996). *Graphical Models*. Oxford University Press, Oxford.

Leroux, B. G. (1992). Consistent estimation of a mixing distribution. *Annals of Statistics*, 20(3):1350–1360.

Lin, L., Tsomocos, D. P., and Vardoulakis, A. P. (2015). Debt deflation effects of monetary policy. *Journal of Financial Stability*, 21(C):81–94.

Lopez-Salido, D. and Nelson, E. (2010). Postwar Financial Crises and Economic Recoveries in the United States. Technical report, Working Paper.

Lucas, R. J. (1976). Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy, Elsevier*, 1:19–46.

Massey, F. (1951). The kolmogorov-smirnov test for goodness of fit. *Journal of the American Statistical Association*, 46(253):68–78.

MATLAB (2014). *Version 8.4.0 (R2014b).* The MathWorks Inc., Natick, MA.

Mendoza, E. G. and Terrones, M. E. (2008). An Anatomy Of Credit Booms: Evidence From Macro Aggregates And Micro Data. NBER Working Papers 14049, National Bureau of Economic Research, Inc.

Minsky, H. P. (1977). A Theory of Systemic Fragility. In Altman, E. I. and Sametz, A. W., editors, *Financial Crises: Institutions and Markets in a Fragile Environment*, volume 1, chapter 6, pages 138–152. John Wiley and Sons.

Minsky, H. P. (1986). *Stabilizing an Unstable Economy.* Yale University Press.

Minsky, H. P. (1992). The Financial Instability Hypothesis. Working Paper 74, The Jerome Levy Economics Institute.

Miranda-Agrippino, S. and Rey, H. (2015). World Asset Markets and the Global Financial Cycle. NBER Working Papers 21722, National Bureau of Economic Research, Inc.

Mishkin, F. S. (2009). Is Monetary Policy Effective during Financial Crises? *American Economic Review*, 99(2):573–577.

Moneta, A. (2008). Graphical causal models and vars: an empirical assessment of the real business cycles hypothesis. *Empirical Economics*, 35(2):275–300.

Montiel Olea, J. L., Stock, J., and Watson, M. W. (2015). Uniform inference in svars with external instruments. mimeo, new york university.

Nakajima, J. (2011). Time-Varying Parameter VAR Model with Stochastic Volatility: An Overview of Methodology and Empirical Applications. IMES Discussion Paper Series 11-E-09, Institute for Monetary and Economic Studies, Bank of Japan.

Nason, J. M. and Tallman, E. W. (2015). Business cycles and financial crises: The roles of credit supply and demand shocks. *Macroeconomic Dynamics*, 19(4):836–882.

Nielsen, B. (2006). Order determination in general vector autoregressions. In H. C. Ho, C. K. I. . T. L. L., editor, *Time series and related topics: In memory of Ching-Zong Wei*, pages 93–112. IMS Lecture Notes and Monograph Series 52.

Nielsen, B. and Whitby, A. (2015). A joint chow test for structural instability. *Econometrics*, 3:156–186.

Pearl, J. (2000). *Causality: Models, Reasoning, and Inference.* Cambridge University Press, Cambridge.

Prieto, E., Eickmeier, S., and Marcellino, M. (2016). Time variation in macro-financial linkages. *Journal of Applied Econometrics*, 31(7):1215–1233.

Primiceri, G. E. (2005). Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies*, 72(3):821–852.

Reinhart, C. M. and Kaminsky, G. L. (1999). The Twin Crises: The Causes of Banking and Balance-of-Payments Problems. *American Economic Review*, 89(3):473–500.

Reinhart, C. M. and Rogoff, K. S. (2008). Is the 2007 US Sub-prime Financial Crisis So Different? An International Historical Comparison. *American Economic Review*, 98(2):339–44.

Reinhart, C. M. and Rogoff, K. S. (2009). *This Time Is Different: Eight Centuries of Financial Folly*. Princeton University Press.

Reinhart, C. M. and Rogoff, K. S. (2013). Banking crises: An equal opportunity menace. *Journal of Banking & Finance*, 37(11):4557–4573.

Romer, C. and Romer, D. (1989). Does monetary policy matter? a new test in the spirit of friedman and schwartz. In *NBER Macroeconomics Annual 1989, Volume 4*, pages 121–184. National Bureau of Economic Research, Inc.

Romer, C. D. and Romer, D. H. (2010). The macroeconomic effects of tax changes: Estimates based on a new measure of fiscal shocks. *American Economic Review*, 100:763–801.

Schularick, M. and Taylor, A. M. (2012). Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008. *American Economic Review*, 102(2):1029–61.

Segoviano, M. (2006). Consistent information multivariate density optimizing methodology. FMG Discussion Papers 557, Financial Markets Group.

Sherman, M. (2009). A Short History of Financial Deregulation in the United States. Technical report, Center for Economic and Policy Research.

Simon, H. A. (1977). *Causal Ordering and Identifiability*, pages 53–80. Springer Netherlands, Dordrecht.

Sims, C. A. (1980). Macroeconomics and Reality. *Econometrica*, 48:1–48.

Sims, C. A. (2001). Comment on Sargent and Cogley's 'Evolving Post-World War II U.S. Inflation Dynamics'. *NBER Macroeconomics Annual*, 16:373–379.

Spirtes, P., Glymour, C. N., and Scheines, R. (2000). *Causation, Prediction, and Search*. MIT Press, Cambridge, MA, 2nd edition.

Stock, J. H. (2001). Discussion of Sargent and Cogley's 'Evolving Post-World War II U.S. Inflation Dynamics'. *NBER Macroeconomics Annual*, 16:379–387.

Stock, J. H. and Watson, M. W. (2001). Vector Autoregressions. *Journal of Economic Perspectives*, 15(4):101–115.

Stock, J. H. and Watson, M. W. (2012). Disentangling the Channels of the 2007-09 Recession. *Brookings Papers on Economic Activity*, 44(1):81–156.

Suppes, P. (1970). *Probabilistic Theory of Causality*. Humanities Press, Atlantic Highlands, NJ.

Swanson, N. R. and Granger, C. W. (1997). Impulse response functions based on a causal approach to residual orthogonalization in vector autoregressions. *Journal of the American Statistical Association*, 92:357–367.

Syron, R. F. (1991). Are we experiencing a credit crunch? *New England Economic Review*, (Jul):3–10.

Teicher, H. (1960). On the mixture of distributions. *Annals of Mathematical Statistics*, 31(1):55–73.

Tong, H. (1983). *Threshold Models in Non-linear Time Series Analysis.* Lecture Notes in Statistics. Springer-Verlag New York.

Tsomocos, D. P. (2003). Equilibrium analysis, banking and financial instability. *Journal of Mathematical Economics*, 39(5-6):619–655.

Uhlig, H. (2005). What are the effects of monetary policy on output? Results from an agnostic identification procedure. *Journal of Monetary Economics*, 52(2):381–419.

Wermuth, N. and Sadeghi, K. (2012). Sequences of regressions and their independences. *TEST*, 21:215–252.

White, H. (1980). A heteroskedastic-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, 48:817–838.

Wu, J. C. and Xia, F. D. (2016). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound. *Journal of Money, Credit and Banking*, 48(2-3):253–291.