Universitat Pompeu Fabra
Departament d'Economia i Empresa

DOCTORAT EN ECONOMIA, FINANCES I EMPRESA

ESSAYS ON ECONOMIC UNCERTAINTY

Tesi Doctoral de:
Matthieu Soupre

Directora de la tesi:
Professor Barbara Rossi

Director del Doctorat en Economia, Finances i Empresa:
Professor José Apesteguia

Setembre 2018

ACKNOWLEDGEMENTS

Those who know me well see this dissertation as a particularly odd achievement. It represents the accomplishment of a long term academic goal, the end – at last – of my studies, and the end of a laborious journey originally supposed be an open sesame to a career as an economist. But it is my farewell letter of sort to the academic world, at least for now, before going on to pursue other radically different endeavors of mine.

My first special thanks go to my advisor, Barbara Rossi. She has supported me and introduced me to exciting research projects throughout my Ph.D. and help me pull through research dry spells every time I needed it. She made me discover my passion for applied macroeconomics and econometrics, a passion that will undoubtedly survive my leaving these fields after I graduate. More importantly, she respected my decision to leave economic research, for which I sincerely thank her.

I am also grateful to the econometrics group at UPF: Majid Al-Sadoon, Christian Brownlees and Geert Mesters. Thank you for your advice, your kindness, and sitting through my less-than-perfect presentations. At least one chapter of this thesis would probably not be there without the help of Tatevik Sekhposyan. She brought me on board of a research paper from day zero, which is an opportunity that is offered to very few lucky doctoral students. I would also like to thank Marta Araque, Laura Agustí and Mariona Novoa for their fantastic help with administrative duties and for being so kind and supportive throughout my time at UPF.

Apart from its professors, Pompeu Fabra gave me the chance to meet wonderful people who all brightened my days as a student: Dani and his warming presence in the office, Óscar and his expertise on thesis titles, Ana and Ana María for bearing my messy office corner, and Korie for joining me in my new journey. I also thank my French friends for their support and their seldom asking the forbidden question for Ph.D. students: “How’s the dissertation progressing?”

Last, I would likely never have been able to bear the mental pressure of these years without my family. My parents, my sisters, brother and brothers-in-law: all of you managed to take off a load of stress off my shoulders in the past five years. And if it were not enough, you are supporting me in my new uncertain projects. And perhaps the one who's had to cope with my struggles the most is my wonderful girlfriend Patricia, whose love, patience and support rescued me from despair more often than I wish to admit. I owe this work to all of you.

Matthieu Soupre
Barcelona
September 2018

Abstract

These essays propose different measures of economic uncertainty and evaluate its impact at the microeconomic and macroeconomic level. The first essay in Chapter 2 proposes a measure of macroeconomic uncertainty that allows to distinguish its various components. Metrics of Knightian uncertainty and risk are proposed, and their respective impact on a number of economic aggregates is evaluated. Chapter 3 extends the classical approach to measuring uncertainty – a mean squared error-based quantity – to entropy methods in econometrics. Several information-theoretic measures of uncertainty are motivated, derived, and estimated on two data sets: the Survey of Professional Forecasters used in Chapter 2, to show that the conclusions hold with this different approach; and the Survey of Economic Expectations, to show how information theoretic measures of uncertainty can help study different situations not afforded by the mean-squared error approach. Chapter 4 studies uncertainty from the point of view of forecasting and propose a measure of forecasting uncertainty to study how business cycles can affect this particular dimension of Knightian uncertainty. Chapter 5 considers the question of the efficacy of fiscal policy in periods of uncertainty, and does so in a way that accounts for the comovements of economic uncertainty with recessions through an conditional adjustment to the classical smooth-transition state dependent models. Chapter 6 concludes.

Resum

Aquesta tesi proposa diferents mesures d'incertesa econòmica i avalua el seu impacte a nivell microeconòmic i macroeconòmic. El primer assaig, al Capítol 2, proposa una mesura de la incertesa macroeconòmica que permet distingir entre les seves múltiples components. Es proposen mètriques d'incertesa i risc de Knight, i se n'avaluen els seus respectius impactes sobre diverses magnituds econòmiques. En el Capítol 3 s'amplia l'enfocament clàssic per la mesura de la incertesa – l'error quadràtic mig – als mètodes d'entropia en econometria. Les diverses mesures d'incertesa que fan servir la teorìa la informació estan motivades, derivades i estimades en dos conjunts de dades: el Survey of Professional Forecasters, que s'utilitza al Capítol 2 per demostrar que les conclusions es mantenen amb aquest nou enfocament; i el Survey of Economic Expectations, que es fa servir per mostrar com aquestes mesures d'informació poden ajudar a estudiar situacions diferents que els metods classics amb error quadràtic mig. El Capítol 4 estudia la incertesa des del punt de vista de la predicció i proposa una mesura d'incertesa de previsió per estudiar com els cicles econòmics poden afectar aquesta dimensió particular de la incertesa knightiana. El Capítol 5 examina la qüestió de l'eficàcia de la política fiscal en períodes d'incertesa, i ho fa de manera que ajusta per als moviments de la incertesa econòmica amb les recessions. També es proposa una nova clase de models depenents de l'estat que inclou condicionalitat. El Capítol 6 conté les conclusions.

Resumen

Esta tesis propone diferentes medidas de incertidumbre económica y evalúa su impacto a nivel microeconómico y macroeconómico. El primer ensayo en el Capítulo 2 propone una medida de la incertidumbre macroeconómica que permite distinguir entre sus diversos componentes. Se proponen métricas de incertidumbre y riesgo de Knight, y se evalúan sus respectivos impactos sobre diversas cantidades económicas. El Capítulo 3 amplía el enfoque clásico para medir la incertidumbre - del error cuadrático medio -, a los métodos de entropía en econometría. Varias medidas de incertidumbre que utilizan la teoría de la información están motivadas, derivadas y estimadas en dos conjuntos de datos: el Survey of Professional Forecasters que se utiliza en el Capítulo 2 para demostrar que las conclusiones se mantienen con este nuevo enfoque y el Survey of Economic Expectations, para mostrar cómo estas medidas de información pueden ayudar a estudiar situaciones diferentes de las que los métodos clásicos con error cuadrático medio permiten. El Capítulo 4 estudia la incertidumbre desde el punto de vista de la predicción y propone una medida de incertidumbre de previsión para estudiar cómo los ciclos económicos pueden afectar a esta dimensión particular de la incertidumbre knightiana. El Capítulo 5 examina la cuestión de la eficacia de la política fiscal en períodos de incertidumbre, y lo hace de una manera que tiene en cuenta los movimientos de la incertidumbre económica con las recesiones. Además, se propone una nueva clase de modelos dependientes del estado que incluye condicionalidad. El Capítulo 6 concluye la tesis.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1

INTRODUCTION

On the other hand, who would doubt that he lives, remembers, understands, wills, thinks, knows, and judges? For even if he doubts, he lives; if he doubts, he remembers why he doubts; if he doubts, he understands that he doubts; if he doubts, he wills to be certain; if he doubts, he thinks; if he doubts, he knows that he does not know; if he doubts, he judges that he ought not to consent rashly. Whoever then doubts about anything else ought never to doubt about all of these; for if they were not, he would be unable to doubt about anything at all.

— Saint Augustine of Hippo, *On The Trinity*

Aristotle's *Metaphysics* opened with the following: "All men by nature desire to know" and for a long time knowledge was confined to pure determinism. Although the concept of uncertainty stands at the crux of epistemology, the science of knowledge, it has been given varying attention throughout the ages. Artistotle's concept of truth was limited by the sensible and measurable, and consequently by the mathematical systems of his time. His epistemology accepted and left room for uncertainty but saw it as belonging to the realm of divine action. Medieval knowledge theory embraced the amount of determinism that escaped the cognitive capacity of men until modern science emerged in the seventeenth century. Perhaps Galileo and Newton's most understated contribution was their putting of Nature and its rules within mathematical law and above simple categories of understanding that they belonged to in the past. By doing so, scientific knowledge on the laws of Nature inherited the strength of mathematical theorems. And it would not be long before the scientific community started having similar demands for other domains of knowledge. Past then, certainty progressively became a convenient crutch to rely on in the quest to fulfill the "desire to know" of the *philosopher*.

Fast forward to modern times, advances in probability and statistics have given a seemingly very precise meaning to the concept of uncertainty. In a desire to understand increasingly complex problems, the second half of the twentieth century saw the development of advanced computational methods applied to a variety of domains followed by an increase in computing power. Together, they allow scientists to turn their efforts to bigger, more complex models consisting of thousands of variables and parameters under the idea that the greater number of parameters, the closer to reality and the smaller the prediction error. But the sheer complexity of these models comes with a greater need for data, a need that is beyond what is available at the moment. Overidentified models' parameters cannot be accurately estimated. Model simplification is frowned upon. In economics and in macroeconomics in particular, the inability to validate intricate theories via observations constitutes an existential crisis for the dismal science. The symptom of this is a profound epistemological uncertainty.

The idea that unequivocal knowledge need not be achievable is the pet peeve of a great many scientists. It haunted Immanuel Kant until he could carry epistemology through the Copernican revolution and qualify the bounds of human knowledge – safe for metaphysics, there is no such thing as certain *a priori* knowledge beyond a reasonable doubt. Economics is not exempt for this caveat. While the discipline has strayed through the unrealistic assumption that agents possess a model of the economy that behaves according to their expectations (a Kantian *a priori* judgment of sort), it has recently come back to the more realistic observation that agents face a good degree of irreducible uncertainty much like real-world scientists do.

This chapter shows how economics has worked with the concept of uncertainty for the past few centuries. Because economics was born at the crossroads of mathematics and philosophy, it is natural to study how both influenced the frameworks retained to study uncertainty in economics. From there, we will see how economics converged towards certain definitions of uncertainty, and how these restrictions affected the conclusions reached on its effects or its measurement. This in turn will justify the approaches retained in this dissertation, which are presented in the last section of this chapter.

A BRIEF HISTORY OF UNCERTAINTY

Uncertainty always was related to probability. The Christian concept of *providence* left the undetectable divine actions appear random to man; any apparent chaos was still a part of the divine purposeful "plan." The seeming of chance is simply a reflection of the complexity of the system in which God acts. A reasonable explanation for putting all chance to an external, supra-human entity, is that lack of a unified theory of uncertain outcomes at the time. The first probabilistic grasp of uncertainty can be attributed to Gerolamo Cardano in his *Liber de ludo aleae* (Book on Games of Chance) in 1563, which was essentially the first version of uniform probability. It was, however, really formalized in the works of Pascal and Fermat in the seventeenth century. A devoted Christian, Pascal did not think that chance and the existence of God conflicted. Through his theory of probability, he "put the certain in the uncertain." Christiaan Huygens in the seventeenth century, Thomas Bayes and Leonhard Euler in the eighteenth century, and Pierre-Simon Laplace in the nineteenth century, all pursued and built on Pascal's earlier breakthrough. In parallel, some progress was made on trying to understand human behavior in front of uncertainty. The first theory of behavior under uncertainty can be attributed to Nicolaus Bernoulli in 1738, explaining the "Saint Petersburg Paradox" proposed by his cousin Daniel, defining the concept of risk aversion for the first time. This theory of expected values would later give birth to that of expected utilities axiomatized by Von Neumann and Morgenstern in 1944 and, albeit very differently, by Savage in 1954. The seventeenth century also saw the emergence of epistemology as it is known today and the first attempts at understanding the limits of human reason. David Hume would assert that the only thing we could have certain knowledge about was the past. The only way to anticipate the future is to evaluate and to interpret our knowledge about the past by assuming the uniformity of the past and future – which is what econometrics would later call "stationarity." Kant would later reject such empiricism and put man at the center of an achievable knowledge that he had so clearly circumscribed. Great progress was made on understanding risks, random events, and knowledge, but little was done, however, on the actual definition of uncertainty. The concept started to be applied to a wider array of sciences, such as physics and biology but was still confined to "what is not known with certainty." And whatever was not (yet) understood under the laws of probability was considered an area that science would later go on to unveil.

Laplace was first to recognize the uncertainty in making predictions but like many he attributed such uncertainty to ignorance.

> We ought then to regard the present state of the universe as the effect of its anterior state and the cause of the one which is to follow. Given for one instant an intelligence which could comprehend all the forces by which nature is animated and the respective situation of the beings who compose it – an intelligence sufficiently vast to submit this data to analysis – it would embrace in the same formula the movements of the greatest bodies of the universe and those of the lightest atom; for it, nothing would be uncertain and the future, as the past, would be present in its eyes.
>
> Laplace, 1814

Laplace speaks of an intelligence that can grasp the movements of all bodies, but does not claim that such intelligence exists and leaves a question mark on the possibility of human omniscience. It is to Frank Knight that we owe the earliest distinction between uncertainty and risk. Knight starts by pursuing Laplace's idea that at least *some* of the observable events unraveling around us do not seem to have easily calculable odds. There exist even more "radical" situations of uncertainty where the possible outcomes are unknown. In an oversimplification of his concepts, risk meant to Knight situations in which one could assign probabilities to outcomes and by uncertainty situations in which one could not. Knight's uncertainty is often thought to be something different than it really is – a homothetic transformation of our ignorance. It is important to note, however, that Knight maintained that (radical) uncertainty was in essence a probabilistic phenomenon:

> It is true, and the fact can hardly be over-emphasized, that a judgment of probability is actually made in such cases.
>
> Knight, 1921

(Knight called these probabilities "estimates".) Although they sharply disagreed on political economy issues, Keynes did borrow a lot from Knight's theory of uncertainty. Keynes view the importance of probabilities more for how decision makers could order them to form a decision, even though the actual numbers may not be possible to figure out. Chapter 12 of the General Theory of Employment, Interest and Money even discusses the implication of radical uncertainty (without giving much credit to Knight). Shackle offered a different take on "non-probabilistic beliefs" by introducing the concept of "potential surprise", echoing Knight's unknown outcomes and odds. Some decades later, Savage will go on to build a theory around subjective probabilities to solve the yet unanswered question of how probabilities come to agents – a stone which Knight had left unturned. Some years later, Savage would build the first alternative to expected utility with an axiomatic characterization of expected subjective utility. Related to the subjectiveness of uncertainty, Ellsberg introduces in 1961 the concept of "ambiguity" to describe a particular type of uncertainty:

> The nature of one's information concerning the relative likelihood of events... a quality depending on the amount, type, reliability and 'unanimity' information, and giving rise to one's degree of 'confidence' in an estimation of relative likelihoods.
>
> Ellsberg, 1961

An important feature that subjective probability theory introduced is that even though agents might not form actual probabilities in their utility maximization, they behave *as if* they did. More and more, models featuring uncertainty became the norm, culminating in macroeconomics with dynamic stochastic general equilibrium (DSGE) models and bounded rationality in microeconomics. All it took was a downturn of the amplitude of the Great Recession to bring it back to the center of the attention of (macro)economists. In a seminal 2009 paper, Bloom analyzed the effects of uncertainty shocks. Back then, uncertainty was defined in a purely probabilistic and aggregate sense, much like macroeconomic volatility.

While I have directed the content of this historical review towards economics, it should be noted that similar concerns emerged in other sciences. Advances in "hard sciences" allowed refinements of the taxonomy of uncertainty – parameter uncertainty, model uncertainty, experimental uncertainty, etc. – all of which are fertile grounds for understanding the uncertainty faced by the *homo œconomicus*. Perhaps the most famous example is Heisenberg's uncertainty principle in quantum mechanics in 1927 which some "heterodox" economists try to include in the discipline. Overall, the second half of the twentieth century was marked by a clear trend in trying to think deeper on the concept of uncertainty, but not until very recently has uncertainty been paid more attention to in economic modeling.

THE GROWING PLACE OF UNCERTAINTY IN MACROECONOMICS

It is often forgotten that Adam Smith's Wealth of Nations already discussed the pernicious effects of "incertitude." Smith pointed that a lack of safety in society could result in money being diverted away from its primary function of facilitating the exchange of present consumption goods and capital, which were the prerequisite to an increase in the wealth of nations. Lowering "incertitude" and increasing security was the main mission of the regalian state. Taxation uncertainty would have the same effect. Jean-Baptiste Say voices a very similar concern:

> The greatest encouragement for circulation is the desire everyone has, especially producers, to lose as little interest as possible on the funds engaged in the exercise of their industry. Circulation slows more due to the obstructions it

> faces than due to an absence of encouragements it might have received. Wars, embargoes, onerous fees to discharge, the danger or difficulty of communication obstruct it. It is also slow in periods of fear and uncertainty, when public order is threatened and all types of enterprise hazardous. It is slow when one expects arbitrary taxation, and is forced to hide his resources. It is slow in periods of speculation when sudden variations caused by wagering on commodities causes some people to hope for a sudden windfall caused by a simple variation in prices. Consequently, merchandise awaits a rise in price and money a fall; and both reflect idle capital, useless to production.
>
> Say, 1803

Contrary to the presentation of early microeconomic theory that is made, uncertainty was not put aside during the Marginal Revolution at the end of the nineteenth century. Jevons considered that uncertainty was the true reason for discounting future utility and not simply the time difference. He even went as far as claiming that future outcomes known with certainty should not be discounted. While he centered his analysis around individual behavior, Jevons claimed that "ignorant" – from not being able to reduce uncertainty – discounting was the root of a sub-optimal savings rate and could explain differences in development and poverty. It is clear, however, that while uncertainty was not forgotten it didn't have nearly the role that it has in modern economic theory. Léon Walras's general equilibrium is, for that matter, completely exempt of such considerations.

It was not until after World War I and macroeconomics was born that economists starting really theorizing on uncertainty. Keynes's analyzes uncertainty in chapters 6 and 26 of the Treatise on Probability in 1921. For Keynes uncertainty is a decreasing function of the weight of evidence rather than a properly defined probability distribution; his vision of uncertainty is sometimes referred to as "distribution intervals." In the General Theory, Keynes actually had a grasp of uncertainty similar to that of Smith:

> By "uncertain" knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty; nor is the prospect of a Victory bond being drawn. Or, again, the expectation of life is only slightly uncertain. Even the weather is only moderately uncertain. The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the position of private wealth-owners in the social system in 1970. About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know. Nevertheless, the necessity for action and for decision compels us as practical men to do our best to overlook this awkward fact and to behave exactly as we should if we had behind us a good Benthamite calculation of a series of prospective advantages and disadvantages, each multiplied by its appropriate probability, waiting to be summed.
>
> Keynes, 1936

Keynes believed that removing all uncertainty was virtually impossible, but additional information can tilt the scale of evidence and help rational decision making. To discuss the impact of uncertainty in the economy, Keynes brought about the concept of "animal spirits": in the face of deep uncertainty, only a manic strong-willed person would put capital at risk. When animal spirits are strong, investment is sufficient to maintain aggregate demand; when they lag, aggregate demand falls, and the economy lapses into depression. It is because animal spirits may not be readily present that uncertainty may make the economy plunge into a recession. It is often thought that "animal spirits" refer to factors that may *hinder* agents from investing, whereas it is precisely the opposite. Animal spirits described the psychological urge to invest in spite of uncertainty; animal spirits for him were neither rational nor irrational. However large the role of uncertainty in Keynesian economics, there is absolutely nothing in the General Theory on how expectations are formed.

The 1950s saw major advancements in mathematical economics and consequently on the modeling of uncertainty. In 1954, Arrow and Debreu proved the existence of competitive equilibria mathematically. Five years later, Debreu extended the framework to uncertain states using what would later be known as "Arrow-Debreu securities." Only a couple of months later John Muth would start what is known as the "Rational Expectations Revolution" in macroeconomics. Not only did Muth make uncertainty a *sine qua non* feature of any respectable macroeconomic model, but it also dictated how agents should perceive and experience uncertainty. Deirdre McCloskey writes about rational expectations:

> Muth's notion was that the professors [of economics], even if correct in their model of man, could do no better in predicting than could the hog farmer or steelmaker or insurance company. The notion is one of intellectual modesty... The common sense is "rationality": therefore Muth called the argument "rational expectations".
>
> McCloskey, 1998

Furthermore, *all* agents form predictions such that they are never surprised by outcomes. Lucas later worked on putting Muth's ideas into application in standardized macroeconomic models. Macroeconomics worked under rational expectations for decades until some started expressing concerns on the veracity of such a framework. Gilboa and Schmeidler in 1989 and several years later Chistopher Sims blew the whistle on the representativeness of such perfect forecasters in economic models. Since then, research on modeling uncertainty in macroeconomics has been very active and the agenda is extremely wide.

Macroeconomics only recently touched the issue of the difficulty of defining uncertainty. Representative agents models miss the mark by making uncertainty a purely exogenous phenomenon that applies to everyone in the same degree. Heterogeneous agents models, too, still struggle to find a credible source of aggregate uncertainty that is not the only the result of agents expectations. This central ambiguity is why Ian Hacking called probability "Janus-faced" in 1984: Probability has a (statistical) connection with the tendency of certain processes to show stable long-run frequencies on repeated trials, and it is also (epistemologically) concerned with how the human agent forms degrees of belief or credence on the basis of knowledge of such frequencies and other things, and hence how he or she decides to act. The same naturally applies to uncertainty. At the same time, the macroeconomics of uncertainty have eschewed the question of defining uncertainty at aggregated and disaggregated levels of the economy. Is macroeconomic uncertainty the uncertainty of a representative agent? Is it the sum of individual uncertainties? Is it measured by how much expectations diverge from one another? All these questions have not been really addressed and in these four essays I try to bring some answers.

A ROADMAP

I summarized and at times bastardized theories that have been developed over the past centuries with the sole purpose of demonstrating that in the study of economic uncertainty, a stance has to be taken. Saint Augustine of Hippo had seen that doubt and uncertainty defined existence. Instead of eschewing those, economics should embrace them and make them the center of the behavior of the agents it studies. This dissertation contributes to this goal by laying out elements of research on the modeling and measure of uncertainty in macroeconomics.

The second chapter represents a first attempt at distinguishing unknown unknowns from predictable risk. The decomposition that is presented relies on the use of density forecasts and their high information content on agents' beliefs. While the first person has been changed for consistency of speech with the rest of this dissertation, all the research presented in the second chapter is joint work with Barbara Rossi and Tatevik Sekhposyan, to which I am immensely grateful for including me in such an exciting project. The third chapter presents a different approach that uses the tools from Claude Shannon's theory of information. Beyond tying the previous chapter's results into perhaps the most influential mathematical theory of uncertainty ever conceived, this chapter generalizes measures of uncertainty to other situations than simple forecast errors. The fourth chapter bounces on the concept of "confidence" and Knightian "impossibility to formulate odds" to offer a measure of forecasting model uncertainty. While a very specific part of uncertainty, forecasting uncertainty has important implications from the point of view of policy makers. The fifth chapter tries to pin down how uncertainty affect economic policy; more specifically if uncertainty warrants government intervention as advised by Keynes himself. The sixth and last chapter concludes.

2

UNDERSTANDING THE SOURCES OF MACROECONOMIC UNCERTAINTY

My biggest concern is concern. The biggest risk we face is uncertainty.

— Patrick T. Harker, 2017

2.1 INTRODUCTION

There are a lot of ways to understand the notion of uncertainty in economics. Beyond the simple intellectual debate, however, it is crucial to understand which type of uncertainty is dealt with because they may very well have different macroeconomic impacts. An increase in predictable risk can be (optimally) insured against. A blurring of the distribution of next period's states cannot. Patrick Harker's paraphrasing of Franklin D. Roosevelt's: "the only thing we have to fear is fear itself" displays a confusing mix of concepts. Uncertainty and risk are very different objects. And whether there is really one to be feared above all others can't be decided upon without a proper measurement of each one's effects.

In this chapter I[1] will adopt the simplifying distinction of Knightian and "non-Knightian" uncertainty, also commonly coined "risk." Risk refers to situations where one can pin down the odds of the unknown with near perfect accuracy,[2] that is, one knows the probability distribution of the stochastic states of nature in the future. Knight, much like Keynes, refers to "uncertainty" as the absence of such knowledge. This could either happen because there is no sensible ways of forming odds on future events, or because the range of possibles is unfathomable. In probability theory terms, either the density or its support is unknown. Another related concept that is often thought to be uncertainty is disagreement, following the logic that if agents disagree on something – e.g., a probability distribution – then they are facing uncertainty.

It should be clear by now that uncertainty is a fundamentally probabilistic concept and that any attempt to measure it should rely on beliefs. Note that while uncertainty appears to fall upon ex-ante predictions, the unraveling of it makes it a phenomenon of ex-post nature all the same.[3] In spite of these well accepted qualifications of uncertainty, the literature has made very sparse use of probability data thus far. Attempts at quantifying uncertainty have been made using either point forecasts (Jurado et al., 2015) or non-probabilistic data (Baker et al., 2015). Most

[1]This chapter is joint work with Barbara Rossi and Tatevik Sekhposyan, with can be found on SSRN as "Understanding The Sources Of Macroeconomic Uncertainty."

[2]Say, the odds of each face of a fair die.

[3]In fact, Knight explains the existence of profit even under near perfect competition by the unraveling of uncertainty. Knight's theory discarded rational expectation equilibria before they even existed.

efforts have been directed at the measurement of economic uncertainty together with quantities that are believed to share a relationship with it, such as disagreement. Yet, none of the measures that were developed depict more than one aspect of economic uncertainty, and it is not clear how they relate to each other to start with. Nor is there a distinction made between ex-ante and ex-post uncertainty.

In this chapter, I propose a decomposition of forecast errors to distinguish between Knightian uncertainty (ambiguity) and risk using survey forecast data from the Survey of Professional Forecasters. The decomposition quantifies overall uncertainty as well as the evolution of the different components of uncertainty over time and investigates their importance for macroeconomic fluctuations. Furthermore, I investigate how the different sources of uncertainty resolve over time as forecasters get closer in time to the event. The behavior and evolution of the various components of the decomposition matches that of a macroeconomic model that features ambiguity and risk, comforting the observations made in the data.

2.2 A NEW SYNTHETIC MEASURE OF UNCERTAINTY

The new uncertainty index measures the distance, on average across forecasters, between the forecast distribution provided by an individual forecaster and the perfect forecast corresponding to the realization, where both are represented by cumulative distribution functions (CDFs).[4] The perfect forecast is denoted by x_{t+h}, which formally is a random variable equal to one when the actual realization y_{t+h} is below some threshold r and it is zero otherwise: $x_{t+h}(r) \equiv 1(y_{t+h} < r)$.[5] Note that $x_{t+h}(r)$ is defined over the support r, $r \in \mathbb{R}$; by varying r, we can focus on different parts of the predictive distribution. Let $P_{s,t+h|t}(r)$ be the probability

[4] As explained later, this measure of uncertainty is similar to a Continuous Rank Probability Score (CRPS).

[5] This notation is consistent with Hersbach (2000).

forecast of the outcome $x_{t+h}(r)$ being equal to one made by forecaster s, $s = 1, ..., N$, i.e. $P_{s,t+h|t}(r) = P(x_{t+h}(r) = 1|\Omega_{s,t}) = E(x_{t+h}(r)|\Omega_{s,t})$, where $\Omega_{s,t}$ is the information set available at time t. We measure the s-th forecaster's uncertainty as the Mean Squared Forecast Error (MSFE) of their probabilistic forecast about a particular outcome, i.e.:[6]

$$\begin{aligned} u_{s,t+h|t}\left(r\right) &= E_Q\left[\left(x_{t+h}\left(r\right) - P_{s,t+h|t}\left(r\right)\right)^2\right] \\ &= \int \left(x_{t+h}\left(r\right) - P_{s,t+h|t}\left(r\right)\right)^2 \mathrm{d}Q_{t+h}, \end{aligned} \tag{2.1}$$

where Q_{t+h} is the true probability distribution. The outcome is compared to the forecaster's probability density forecast for it.

Similarly to Jurado et al.'s measure, Equation (2.1) is an MSFE. In sharp contrast to their forecast error, it is an MSFE applied to a forecast distribution. As such, it measures the unpredictable component associated with each possible value in the domain of the predictive distribution, or in simpler terms, the failure to predict odds with precision. In fact, $u_{s,t+h|t}\left(r\right)$ compares the probability that forecaster s assigns to the different states of nature with the realization, while error-based measures à la Jurado et al. (2015) compare the point forecast with the realization.[7] The overall measure of uncertainty is then defined as the average of the individual uncertainty across forecasters:

$$\begin{aligned} u_{t+h|t}\left(r\right) &= \frac{1}{N}\sum_{s=1}^{N} u_{s,t+h|t}\left(r\right) \\ &= \frac{1}{N}\sum_{s=1}^{N} E_Q\left[\left(x_{t+h}\left(r\right) - P_{s,t+h|t}\left(r\right)\right)^2\right]. \end{aligned} \tag{2.2}$$

The full support of the predictive distribution is explored by letting r vary. The overall measure of uncertainty (which I'll label "Uncertainty")

[6]In the meteorological forecasting literature, this quantity is known as the Brier score and is typically computed on binary forecasts.

[7]In fact, if one associates the value $r \in \mathbb{R}$ with the corresponding quantile of the distribution, the uncertainty index measures an average squared error for that quantile.

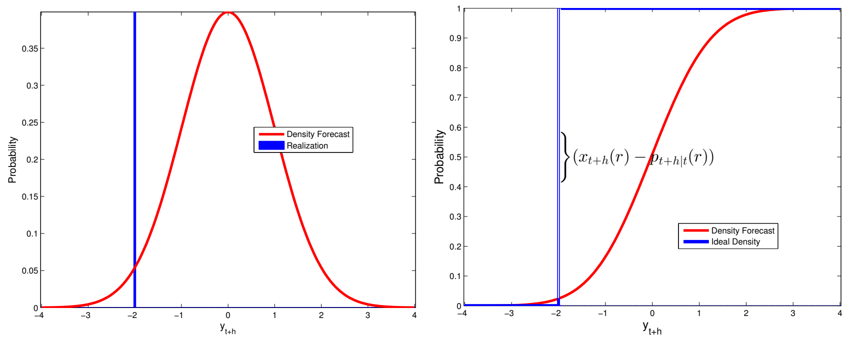

FIGURE 2.1: Brier Score Illustration

integrates the squared forecast errors over the whole domain of the distribution, that is:[8]

$$U_{t+h|t} = \int_{-\infty}^{+\infty} u_{t+h|t}\left(r\right) \mathrm{d}r. \tag{2.3}$$

A graphical interpretation is provided in Figure 2.1. In the figure, the actual realization equals -2, denoted by a vertical bar on the left panel; the predictive density is the Gaussian distribution. The panel on the right shows the CDF of the Normal distribution, as well as that of the perfect forecast, for a particular threshold, $r = -1$. Thus, the perfect forecast assumes ones for values less than -1 (since the realization of -2 is indeed less than -1) and zero otherwise. For any given r, the distance between the CDF of the forecast distribution and the perfect forecast, $\left(x_{t+h}(r) - P_{s,t+h|t}(r)\right)$, is depicted by a solid vertical line. The measure of uncertainty in Equation (2.3) squares this measure and integrates it over the various values of r.

As said in the introduction to this chapter, the existing literature has focused mainly on quantifying and understanding uncertainty as-

[8]Note that Equation (2.3) is the negative of the CRPS, as defined in Gneiting and Raftery (2007). In fact, the CRPS is the integral of Brier scores (Hersbach, 2000, Equation (7)).

sociated with point forecasts, for example by mapping uncertainty to forecasters' prediction errors. An issue that render probabilistic forecast data would be that they are inconsistent with average (point) forecasts, which many practitioners use in practice. Zarnowitz and Lambros (1987) found that individual point forecasts were on average consistent with the weighted mean of their predictive probability distributions, which makes the use of density forecast compelling because of the undoubtedly richer information they contain. The superior informational content of probabilistic forecasts is precisely what allows to quantify Knightian uncertainty and distinguish among various sources of uncertainty. An important difference between this measure of uncertainty and the existing literature is that it uses the probabilistic forecasts provided by the U.S. Survey of Professional Forecasters (SPF) to measure and decompose uncertainty.[9] The focus is on output growth forecasts, which are indicative of business cycle fluctuations and therefore better match what one would understand as "macroeconomic" uncertainty.

Furthermore, a large number of uncertainty measures considered in the literature are ex-post in that they depend *only* on realizations (such as the uncertainty measures recently proposed by Jurado et al., 2015; Rossi and Sekhposyan, 2015 and Rossi and Sekhposyan, 2016; and Scotti, 2013. Such ex-post measures are arguably difficult to square with the notion of economic agents' forward looking decision making, and as stressed in the introduction to this dissertation, they tell only half of the story. Zarnowitz and Lambros (1987) define uncertainty as

[9]The analysis can be done with any predictive density. We choose to use predictive densities from the SPF since they are produced by professional forecasters monitoring a wider range of indicators rather than a specific parametric model. Furthermore, the SPF is known for its superior forecasting performance from a point forecasting point of view, as shown in Giannone et al. (2008) and McCracken et al. (2015), among others.

the "difuseness" of a forecaster's predictive density as follows:

$$\int_{-\infty}^{+\infty} E\left[\left(x_{t+h}(r) - P_{s,t+h|t}(t)\right)^2 |\Omega_{s,t}\right] \mathrm{d}r = \tag{2.4}$$

$$\int_{-\infty}^{+\infty} P_{s,t+h|t}(t)(1 - P_{s,t+h|t}(t))\, \mathrm{d}r; \tag{2.5}$$

and they emphasize how uncertainty differs from disagreement. They do not, however, consider Knightian uncertainty. As I will show below, this new framework is able to distinguish between ex-post measures of uncertainty (for instance, realized risk or bias) and ex-ante risk (also termed ambiguity). The capacity of this new measure of uncertainty to distinguish between Knightian uncertainty (which is essentially ex-ante) and risk (ex-post).

Forecast densities and their probabilistic dimension are the cornerstone of that feature.[10] This will become clearer with the decomposition that I will now expose. One of the goals of this chapter is to link existing measures of uncertainty based on aggregate data with uncertainty measures based on disagreement among forecasters. To do so, define an aggregate probability density $\left\{P_{t+h|t}\left(r\right), r \in \mathbb{R}\right\}$, which is related to the individual ones $\left\{\left(P_{s,t+h|t}\left(r\right)\right), 1 \leqslant s \leqslant N, r \in \mathbb{R}\right\}$ by:

$$P_{t+h|t}\left(r\right) = \frac{1}{N}\sum_{s=1}^{N} P_{s,t+h|t}\left(r\right). \tag{2.6}$$

The corresponding uncertainty measure for the aggregate predictive density is:

$$u_{t+h}^{A}\left(r\right) \equiv E_Q\left[\left(x_{t+h}\left(r\right) - P_{t+h|t}\left(r\right)\right)^2\right].$$

[10]Knightian uncertainty is defined as the agents' inability to correctly characterize probability distributions or their disagreement on them. Clearly, it is impossible to quantify uncertainty associated with the agents inability to characterize all possible states of nature or situations where they have no opinions on the probability distributions associated with known states of the nature. Thus, one can think of this Knightian uncertainty measure as a lower bound on the actual Knightian uncertainty present in the economy.

I show in the Appendix to this chapter that the overall uncertainty measure can be broker down as follows:

$$\begin{aligned} u_{t+h|t}(r) &= E_Q\left[\left(x_{t+h}(r) - P_{t+h|t}(r)\right)^2\right]\dots \\ &\quad + E_Q\left[\frac{1}{N}\sum_{s=1}^{N}\left(P_{t+h|t}(r) - P_{s,t+h|t}(r)\right)^2\right] \\ &= u^A_{t+h|t}(r) + d_{t+h|t}(r), \end{aligned} \tag{2.7}$$

where:

$$d_{t+h|t}(r) \equiv \frac{1}{N}\sum_{s=1}^{N} E_Q\left[\left(P_{t+h|t}(r) - P_{s,t+h|t}(r)\right)^2\right] \tag{2.8}$$

measures the disagreement between individual forecast densities and the aggregate forecast density, similar to the disagreement defined in Patton and Timmermann (2010) for point forecasts. Lahiri and Sheng(2010, eq. 18) discuss a similar decomposition for point forecasts.

Note that the decomposition in Equation (2.7) holds for a particular threshold r, thus it accounts for a forecast error associated with the binary outcome $1\left(y_{t+h} < r\right)$. The overall measure of uncertainty accounts for uncertainty at all possible values of r by considering the integral of the decomposition in Equation (2.7) over r. "Uncertainty" breaks down into "Aggregate Uncertainty" and "Disagreement":[11]

$$\begin{aligned} U_{t+h|t} &\equiv \int_{-\infty}^{+\infty} u_{t+h|t}(r)\,\mathrm{d}r = \int_{-\infty}^{+\infty} u^A_{t+h|t}(r)\,\mathrm{d}r + \int_{-\infty}^{+\infty} d_{t+h|t}(r)\,\mathrm{d}r \\ &\equiv \underbrace{U^A_{t+h|t}}_{\text{"Aggregate Uncertainty"}} + \underbrace{D_{t+h|t}}_{\text{"Disagreement"}} \end{aligned} \tag{2.9}$$

[11] A reason why the aggregate probability distribution, measured with a simple average of the individual probability distributions, is a good measure of aggregate uncertainty is that, as in the context of point forecasts, combinations constructed by simple averages result in more accurately calibrated densities. Furthermore, the average of probability distributions is a measure widely used in a variety of central banks and policy institutions and a (surprisingly) well performing forecast.

This decomposition represents a first step towards the separation between ex-ante and ex-post uncertainty. While the aggregate uncertainty term is by construction ex-post, disagreement is purely ex-ante and mirrors the inability of forecasters to pin down the "correct" probabilities of future states of the world.

A finer distinction between ex-post and ex-ante uncertainty can be obtained by breaking down the aggregate uncertainty term. As shown in the Appendix, the aggregate uncertainty, $U^A_{t+h|t}(r)$ decomposes into components that measure mean bias, dispersion of probability forecasts, realized risk and a covariance term between the forecast and the ideal distribution as follows:

$$\begin{aligned} u^A_{t+h}(r) = &\left(\left[E_Q\left(P_{t+h|t}(r)\right) - E_Q\left(x_{t+h}(r)\right)\right]^2\right)\dots \\ &+ V_Q(P_{t+h|t}(r))\dots \\ &+ V_Q\left(x_{t+h}(r)\right)\dots \\ &- 2\mathrm{Cov}_Q(x_{t+h}(r), P_{t+h|t}(r)), \end{aligned} \tag{2.10}$$

where $V_Q(.)$ denotes the variance taken with respect to the probability measure Q_{t+h}. Since the covariance term turns out to be rather small empirically, we summarize aggregate uncertainty with the following additive decomposition:

$$U^A_{t+h|t} \approx \underbrace{B_{t+h|t}}_{\text{"Mean-Bias"}} + \underbrace{V_{t+h|t}}_{\text{"Dispersion"}} + \underbrace{\mathrm{Vol}_{t+h|t}}_{\text{"(Realized) Risk"}} \tag{2.11}$$

where:

- $B_{t+h|t} \equiv \int_{-\infty}^{\infty} E_Q\left[\left(P_{t+h|t}(r)\right) - E_Q\left(x_{t+h}(r)\right)^2\right] \mathrm{d}r$ is the mean squared bias of the forecast distribution;

- $V_{t+h|t} \equiv \int_{-\infty}^{\infty} V_Q(P_{t+h|t}(r))\, \mathrm{d}r$ is the uncertainty about the ex-ante subjective probabilities in the aggregate distributional forecast

– $\mathrm{Vol}_{t+h|t} \equiv \int_{-\infty}^{\infty} V_Q\left(x_{t+h}\left(r\right)\right) \mathrm{d}r$ is the realized variance of the binary outcome, $x_{t+h}\left(r\right) \equiv 1\left(y_{t+h} < r\right)$, and thus stands for the inherent risk in the data.

The three component decomposition in Equation (2.11) has an interesting interpretation. The realized volatility component $\mathrm{Vol}_{t+h|t}$ is a measure of the underlying uncertainty in the data, and thus a measure of realized risk. On the other end, the bias component $B_{t+h|t}$ measures how far the predictive density is from the perfect prediction on average, and the dispersion, $V_{t+h|t}$, gives an estimate of the variability in the predictive density. As it will be shown later $V_{t+h|t}$ is empirically small, so it can be ignored. Knightian uncertainty is proxied as the sum of bias, dispersion and disagreement, as all these terms represent a different "incapacity" to perfectly estimate odds. The realized variance or realized volatility, instead, is a measure of risk. In short, we have the following "Knightian uncertainty/(Realized) Risk" decomposition:

$$U_{t+h|t} \approx \underbrace{\mathrm{Vol}_{t+h|t}}_{\text{“(Realized) Risk”}} + \underbrace{B_{t+h|t} + D_{t+h|t}}_{\text{“Knightian Uncertainty”}} .$$

2.3 EX-ANTE V. EX-POST UNCERTAINTY

It is important to note that the proposed measure of uncertainty, $U_{t+h|t}$, as well as aggregate uncertainty $U^A_{t+h|t}$, are constructed using ex-post realizations of the data. Thus, it is interesting to refine our measure by distinguishing between an ex-ante component (that does not include the realizations) and an ex-post component (which does). Also, one might wonder how the expected mean and the variance embedded in the forecast distribution affect our measure of uncertainty. Let the aggregate predictive distribution for the forecast of y_{t+h} made at time t be Normal with mean $\mu_{t+h|t}$ and variance $\sigma^2_{t+h|t}$ and the data be i.i.d. We have the following "Ex-ante/Ex-post" decomposition:

$$
\begin{aligned}
U^{A}_{t+h|t} = E_Q &\left[2\sigma_{t+h|t}\phi\left(\frac{y_{t+h} - \mu_{t+h|t}}{\sigma_{t+h|t}}\right)\right] \ldots \\
&+ E_Q\left[\left(y_{t+h} - \mu_{t+h|t}\right)\left(2\Phi\left(\frac{y_{t+h} - \mu_{t+h|t}}{\sigma_{t+h|t}}\right) - 1\right)\right] \ldots \\
&- \frac{\sigma_{t+h|t}}{\sqrt{\pi}}; \qquad (2.12)
\end{aligned}
$$

where $\phi\,(.)$ and $\Phi\,(.)$ denote the PDF and the CDF of the Normal distribution, respectively. The first two terms (in square brackets) are *ex-post* in that they depend on the realization of the forecast; the last term is purely *ex-ante*. The proof is provided in the Appendix and follows Gneiting and Raftery (2007).[12]

The rightmost component, $\sigma_{t+h|t}/\sqrt{\pi}$, is the only component that is not affected by the realization, hence its "ex-ante"denomination. In fact, as the proof suggests, this is the component that arises from the average distance of random draws from a given predictive distribution. Moreover, it is a function of the standard deviation of the forecaster's density forecasts, and a common measure used in the uncertainty literature as a measure of ex-ante uncertainty. Note that the ex-ante measure of uncertainty is simply $\sigma_{t+h|t}/\sqrt{\pi}$, which, under normality, is a monotone function of the width of the predictive distribution. Thus, the ex-ante measure is linked to the inter-quantile range measure proposed by Zarnowitz and Lambros (1987), among others.[13] The ex-ante component might be viewed as a measure of ex-ante risk. Note that, from Equations (2.11) and (2.12), we have that *Ex-post* $\approx B_{t+h|t} + V_{t+h|t} + \mathrm{Vol}_{t+h|t} +$ Ex-Ante. Thus, the ex-post measure of aggregate uncertainty combines components of Knightian uncertainty, $B_{t+h|t} + V_{t+h|t}$, realized risk (measured by the volatility in the economy, $\mathrm{Vol}_{t+h|t}$) and ex-ante risk (measured by the variance of the predictive

[12] Note that even if $U^{A}_{t+h|t}$ is the difference of two components, it is always positive; thus, the ex-post component is always bigger than the ex-ante one.

[13] For a Gaussian distribution, the inter-quantile range is 1.34σ.

densities of the forecasters, *Ex-Ante*). Note the difference between $V_{t+h|t}$ and *Ex-ante*: the first measures the variability of the probability distribution, while the second measures the width of the distribution at a particular point in time. Thus, if the aggregate density forecast does not changed over time, $V_{t+h|t}$ would be zero. However, *Ex-Ante* will not be zero as long as the forecasters provide a distributional forecast.

There is a major difference between the two decompositions in that the "Ex-ante"– "Ex-post" decomposition is written in terms of the moments of the original predictive distribution, while the "Knightian Uncertainty – (Realized) Risk" decomposition is in terms of distributional quantile outcomes summarized by $x_{t+h}(r)$. As such, the latter decomposition could be applied to general situations (general forms of distribution and non-iid data), while the former one relies unequivocally on the assumption of Gaussianity and independence in the underlying predictive distribution. D'Amico and Orphanides (2014) and Giordani and Söderlind (2003) provide empirical support in favor of Gaussianity for the Survey of Professional Forecasters, and the iid assumption would be satisfied for correctly calibrated density forecasts.

A general remark that applies to all proposed decompositions is that the resulting components are a priori not orthogonal to each other. This is in line with the rest of the empirical literature which typically finds that a variety of uncertainty measures, constructed from different sources and measuring different aspects of uncertainty, are correlated with each other.

2.4 THE DATA

We use density forecasts from the Survey of Professional Forecasters (SPF) to calculate our uncertainty measures. The Federal Reserve Bank of Philadelphia provides the aggregate (mean probability distribution) forecasts, as well as the underlying disaggregate density forecasts of a panel of professional forecasters.[14] I use the real GNP/GDP growth density forecasts to extract measures of macroeconomic uncertainty, as real GNP/GDP fluctuations are indicative of the state of the business cycle, and are therefore reflective of macroeconomic uncertainty (Stock and Watson, 1998).

SPF forecasters are asked to assign a probability value (over predefined intervals) to inflation and output growth for the current and the following (one-year-ahead) calendar years. The growth rate is defined as the rate of change in the average GDP from one year to another. The forecasters update the assigned probabilities for the current-year and the one-year-ahead forecasts on a quarterly basis. Thus, by construction, SPF forecasters provide four quarterly forecasts of the same target variable each year; this type of forecasts are typically referred to in the literature as "fixed-event "or "moving-horizon" forecasts. Being fixed-event forecasts, their horizon changes over the quarter. Dovern et al. (2012) propose a method to transform SPF fixed-event forecasts into fixed-horizon forecasts by constructing a weighted average of current-year and next-year forecasts. More specifically, for each quarter the survey contains a pair of "fixed-event" density forecasts for the current-year, denoted by $\widehat{f}^{FE}_{t+k|t}$, and for the following year, labeled $\widehat{f}^{FE}_{t+k+4|t}$. The four-quarter-ahead (fixed-horizon) forecast at time t $\widehat{f}^{FH}_{t+4|t}$ is calculated as the average of the two fixed event forecasts using weights that are proportional to their share of the overlap with the forecast horizon. Let k denote the number of quarters from time t until the end of the year. In quarter one, $k = 4$, while in quarter four, $k = 1$. For instance, in the third quarter of the year, the four-step-ahead fixed-horizon forecast overlaps with the current year forecasts and next year forecasts 50% of the time. A natural estimate for 4-quarter growth averages the two-fixed event forecasts with weights equal to 2/4 and 2/4. In general, for

[14]The composition of the forecasters can change over time.

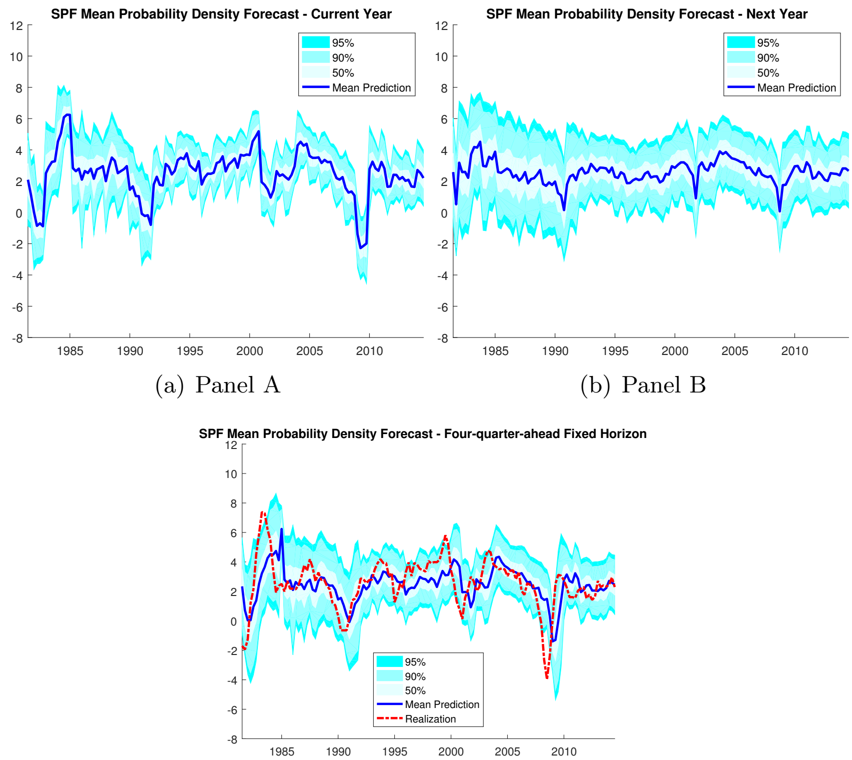

(a) Panel A (b) Panel B

(c) Panel C

Figure 2.2: Brier Score Illustration

$k = 1, 2, 3, 4$:

$$\widehat{f}^{FH}_{t+4|t} = \frac{k}{4}\widehat{f}^{FE}_{t+k|t} + \frac{4-k}{4}\widehat{f}^{FE}_{t+k+4|t}. \tag{2.13}$$

Panels A and B in Figure 2.2 show the evolution of the current and next year densities over time. The figures plot the mean as well as several quantiles of the distribution, together with the realization. Panel C, on the other hand shows the fixed horizon forecast, Equation (2.13). The fixed-horizon forecast is by construction less smooth than the fixed-event forecasts. However, both share the same feature that ex-ante uncertainty was higher earlier in the sample, in the sense that both density forecasts have a wider distribution prior to the mid-1980s relative to the later part of the sample; this suggest that forecasters noticed the Great Moderation starting mid-1980s. There appears to be no dramatic shift in the densities forecast after the Great Recession. Some descriptive statistics on the SPF distributions is provided in the Appendix to this chapter.

The analysis of SPF probability distributions is complicated since the SPF questionnaire has changed over time in various dimensions: there have been changes in the definition of the variables, the intervals over which probabilities have been assigned, as well as the time horizon for which forecasts have been made. To mitigate the impact of these problematic issues, I truncate the data set and consider only the period 1981:III-2014:II.[15]

As noted, our uncertainty measure depends on realizations. The realized values of output growth are from the real-time data set for macroeconomists, also available from the Federal Reserve Bank of Philadelphia. The four-quarter-ahead growth rates of output and prices are calculated from the first release of the realization. For instance, in order to get the 4-quarter-ahead realization at the start of our sample, 1981:III, the growth rate between 1982:III and 1981:III is computed using the 1982:IV vintage of the data.

2.5 CLASSIFYING UNCERTAINTY OVER TIME

Figure 3, Panel A, shows the evolution of the estimated measure of uncertainty and its components, aggregate uncertainty and disagreement, over time. The figure highlights two interesting facts: disagreement is, in magnitude, only a small portion of the overall measure of uncertainty;[16]

[15]See instead Ferrara and Guérin (2015) for a high-frequency analysis of uncertainty shocks.

[16]The magnitudes of $U_{h+h|t}$ and $U^A_{h+h|t}$ are reported on the y-axis on the left while that of disagreement is reported on the y-axis on the right. The magnitude of disagreement is small. This could be due to the fact that, unlike the existing measures

in addition, it is trending down until the financial crisis of 2007. This is in sharp contrast with the overall measure of uncertainty, as well as aggregate uncertainty, which have clear spikes in the early 1980s, early 2000s and the financial crisis. Using disagreement as a measure of uncertainty may result in underestimating both the overall level of uncertainty in the economy as well as its fluctuations over time, as currently the level of disagreement is similar to what it was in the mid-1990s and lower than its value in the late 1980s. In addition, most would agree that the early 2007-2008 were probably the most uncertain times in the latest decades, and while disagreement increases during that period, it peaks only much later – after the end of the recession, in 2009. Thus, disagreement (i.e., the component of Knightian uncertainty due to disagreement among forecasters) may not be a timely measure of macroeconomic uncertainty. Note that this result is not an artifact of constructing disagreement measures based on density forecasts: Sill (2014, Figure 1) shows a similar delay. In particular, Sill (2014) plots the dispersion of the mean one-year-ahead real GDP growth rate forecasts measured by the inter-quantile range: the first peak in the disagreement does not appear until the middle of the recession.

Panel B in Figure 2.3 depicts the decomposition of aggregate uncertainty into Knightian uncertainty and realized risk. The figure suggests that realized risk (measured by $\mathrm{Vol}_{t+h|t}$) was an important component of uncertainty throughout the last three decades, as was Knightian uncertainty, measured by the mean bias component. Some differences between the two are important to note, however. The realized risk component was high during the latest financial crisis, and sharply decreased as soon as the recession was over; Knightian uncertainty (measured by the mean bias component, $B_{t+h|t}$) remained persistently high even after the end of

of disagreement on point forecasts, we measure disagreement in probabilities, not in the mean forecast. Another possible explanation is that professional forecasters all use similar models or have comparable information sets, making their forecasts agree for the most part.

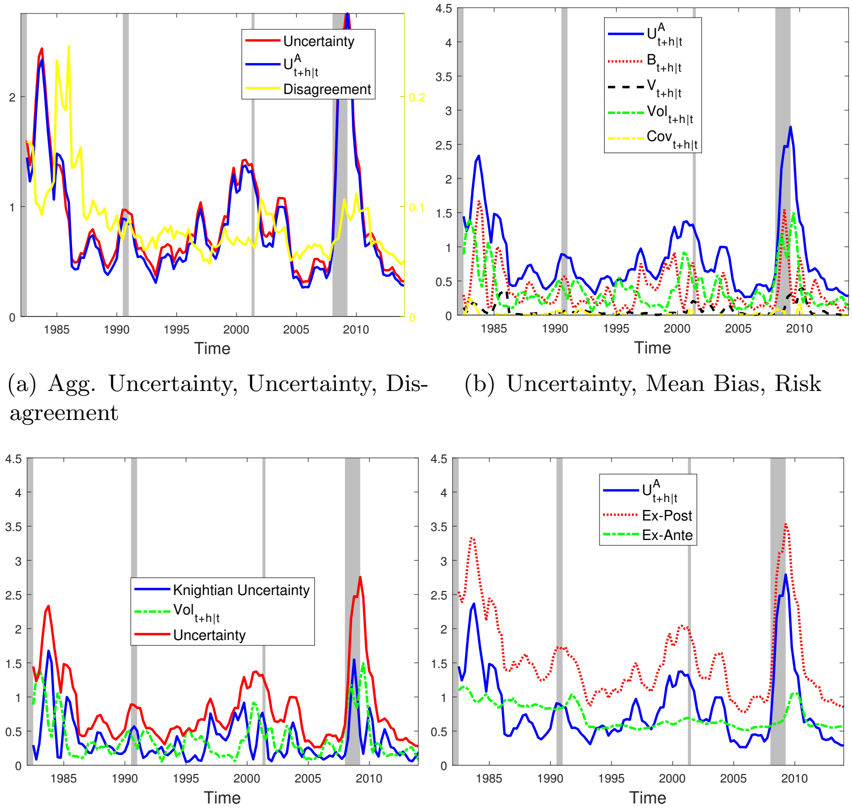

(a) Agg. Uncertainty, Uncertainty, Disagreement

(b) Uncertainty, Mean Bias, Risk

(c) Knightian, Risk

(d) Ex-Ante, Ex-Post

FIGURE 2.3: Decomposing Uncertainty

the crisis. Overall uncertainty remained persistently high after the end of the latest recession mostly because of forecasters' errors as opposed to risk being high. The role of dispersion in probability forecasts ($V_{t+h|t}$) as well as the co-movement between prediction and realization ($\text{Cov}_{t+h|t}$) are negligible for the cyclical dynamics of aggregate uncertainty.

Turning to the ex-ante and ex-post components, depicted in Panel C of Figure 3 together with the aggregate uncertainty measure ($U^A_{t+h|t}$), it is interesting to note that ex-ante uncertainty is quite constant in the 1980s and up to 2007. This provides support to the idea that movements in uncertainty during that period cannot be attributed to changes in ex-ante uncertainty. Ex-ante uncertainty does increase during the latest recession, but only towards its end, and spikes much later than the peak of the recession. This suggests that measures of volatility in the forecasters' predictive distributions are, themselves, not timely measures of uncertainty, and reinforces the idea that risk (realized volatility) and forecast uncertainty are different objects.

Finally, it is also of interest to investigate how the various components of uncertainty evolve as the forecasters get closer in time to the realization date, that is, as the forecast horizon becomes shorter. We separately consider forecasts for $h = 1, 2, ..., 7, 8$ and compare them with the fixed-event realization. Both uncertainty as well as aggregate uncertainty decrease as the forecast horizon increases (Panel A in Figure 2.4, top left and right graphs). It may seem counter-intuitive that uncertainty decreases at longer horizons. One would think that the longer the forecast horizon, the harder it is to forecast and the higher the uncertainty. This surprising finding, however, can be better understood by looking at the types and sources of uncertainty. Clearly, disagreement decreases as forecasters get closer to the realization: in fact, disagreement decreases on average as the horizon decreases (cf. bottom graph in Figure 2.4, Panel A). This finding is reassuring, as it is reminiscent of what Patton and Timmermann (2010) discovered for point forecasts, and our results show that similar results hold for disagreement calculated on density forecasts. The mean bias also decreases as the horizon decreases (Panel B in Figure 2.4). On the other hand, the dispersion of the density forecasts increases, thus increasing the aggregate uncertainty. The realized variance and covariance are constant over the horizons, and the latter hovers around zero.

The most striking patterns are displayed by ex-ante and ex-post uncertainty, depicted in Figure 2.4, Panel C. Clearly, ex-ante uncertainty decreases monotonically as the forecast horizon decreases; that is, forecasters' predictive densities become more spread out when the forecast horizon increases, thus reflecting more uncertainty in the economy when looking at events that are further in the future. However, there is no clear pattern in ex-post uncertainty. This means that, even though the forecasters' predictive densities become tighter as the realization gets closer in time, the uncertainty in the actual realizations does not diminish, as the size of the forecast errors does not diminish with the horizon. The closer forecasters find themselves to the actual realization, the more "clear-cut" their forecast, and the more potential error if one considers the *whole* density forecast.

Comparing the evolution of the ex-ante uncertainty in Panel C and the dispersion of the aggregate predictive density, $V_{t+h|t}$ in Panel B, it can be noted that, although forecasters, on average, become less confident about the future as the forecast horizon increases, their views about uncertainty does not seem to be updated often for forecasts that are further in the future, thus resulting in the low variability of the predictive distribution over time. Moreover, as the distribution becomes more spread out with the forecast horizon, it has a higher chance of including the realization, thus resulting in a decline in the aggregate and overall uncertainty.

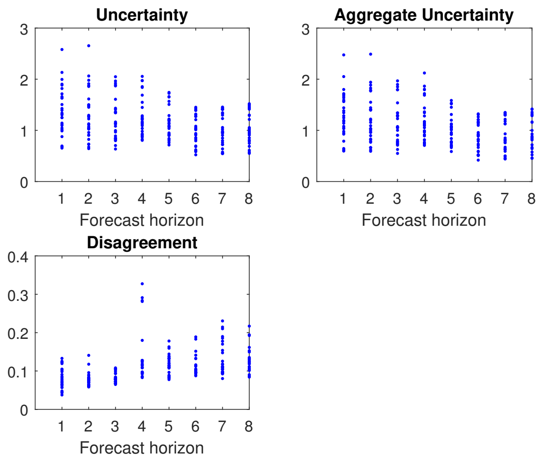

(a) Agg. Uncertainty, Uncertainty, Disagreement

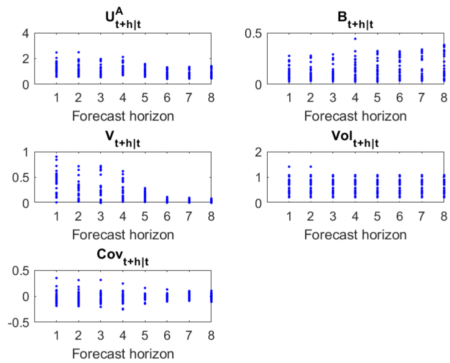

(b) Knightian Uncertainty and Risk

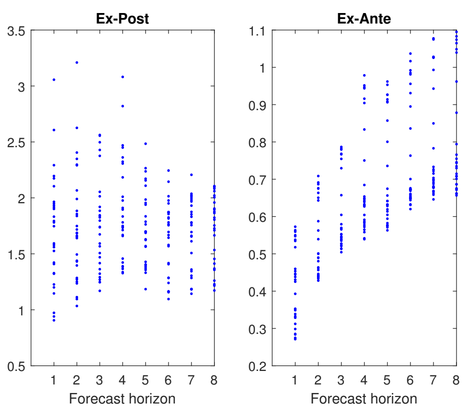

(c) Knightian Uncertainty and Risk

Figure 2.4: Uncertainty At Different Horizons

2.6 THE IMPACT OF KNIGHTIAN UNCERTAINTY AND RISK

Let us now shift the analysis to of the effects of uncertainty "shocks." The decomposition allows to distinguish between disagreement and aggregate uncertainty on the one hand,[17] and between measures of realized volatility, ex-ante uncertainty and bias on the other. These various components have all been used in the literature as measures of uncertainty, but only the decomposition presented in this chapter allows distinctions to be drawn among them and understand their relationship.

How do the various components relate to existing measures of uncertainty? The top panel in Figure 2.5 plots Jurado et al.'s (2015) uncertainty measure together with Baker et al.'s (2015) index.[18] Both indices are standardized for comparison. The figure shows that the former is overall smaller than the latter until 1995, then it becomes overall larger, and in particular spikes up earlier than the latter during the latest financial crisis of 2007-2008. The lower panel plots the decomposition of aggregate uncertainty index into ex-ante and ex-post components. The ex-post component is lower than the ex-ante component up to mid-1992, then it becomes systematically more prevalent, and spikes up around 2007-2008, behaving similarly to how Jurado et al.'s (2015) behaves relative to Baker et al.'s (2015). Thus, it seems that the Baker et al. (2015) uncertainty measure is driven more by ex-ante uncertainty, while the Jurado et al. (2015) uncertainty measure is clearly affected by ex-post uncertainty, namely uncertainty due to misspecification in the predictions.

To estimate the effects of the uncertainty and its components on the economy, we estimate a Vector Autoregression (VAR) that includes

[17]similar to that of Lahiri and Sheng (2010), who consider the relationship between aggregate uncertainty and disagreement over the business cycle, yet measure it in terms of uncertainty and disagreement about the mean of the distribution, as opposed to the whole distribution.

[18]Using Jurado, Ludvigson and Ng's (2015) one-year-ahead uncertainty index.

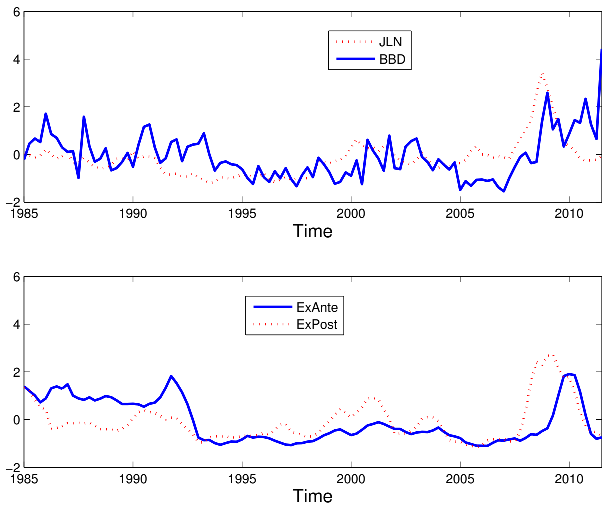

FIGURE 2.5: Comparison of Uncertainty Measures

Notes: The figure compares the Jurado, Ludvigson and Ng (2015) and Baker, Bloom and Davis (2016) uncertainty indices (top panel) with the ex-ante and ex-post components of our uncertainty measure in the bottom panel.

(the log of) real GDP, (the log of) employment, the Federal Funds rate, (the log of) stock prices and the specific uncertainty indices one at a time. Identification is achieved via a classic Cholesky procedure, which follows the order in which the variables are listed. The VAR specification is the same as in Baker et al. (2015), although ours is at the quarterly frequency, and accordingly we use GDP instead of real industrial production. The variables are ordered as in Jurado et al.'s (2015) benchmark specification, i.e. from slow to fast moving. For completeness, the robustness of the results in a larger VAR are exposed in the Appendix. To better interpret and compare the magnitude of the effects of the uncertainty indices, the uncertainty indices are demeaned and standardized.

Panel A in Figure 2.6 shows the effects of our uncertainty index on the economy. Clearly, an increase in uncertainty has recessionary effects: both GDP and employment decrease, as well as the interest rate and the S&P 500. Panels B and C describe the effects of each of the components in the decomposition. Panel B shows the effects of a shock to aggregate uncertainty, which is in line with that of uncertainty since aggregate uncertainty is the main determinant of the total. Panel C focuses on disagreement; it also decreases employment although by a smaller magnitude; at the same time, it has no significant effects on the remaining variables.

Figure 2.7 shows the effects of uncertainty measured by mean bias, realized volatility and the dispersion in the probability forecasts. The mean bias and realized volatility appear to have recessionary effects (Panels A and C). Dispersion in the density forecasts (Panel B) drives down employment, while it increases stock prices and output. It is important to note that, in magnitude, the mean bias and realized volatility have similar macroeconomic impact, though these effects are statistically significant for the first but not for the second.

The effects of ex-ante and ex-post uncertainty on other macroeconomic variables are depicted in Figure 2.8. They both lead to decreases in employment, interest rates and stock prices of similar magnitude; an increase in ex-ante uncertainty, however, has a small negative impact effect on GDP, while the medium run effect is positive and small, and the longer run effect is again negative; the effects of ex-post uncertainty on GDP are, instead, negative and large.

Figure 2.9 compares the results with those in the existing literature; the latter are also obtained by estimating VARs that include (the log of) real GDP, (the log of) employment, the Federal Funds rate, (the log of) stock prices, and the alternative uncertainty index, which is demeaned and standardized as well. The alternative uncertainty indices include: Bloom (2009), labeled “VXO”; Baker et al.’s (2015) policy uncertainty index, labeled “BBD”; Jurado et aL. (2015), labeled “JLN”; and the Scotti (2013) macroeconomic surprise-based uncertainty index.

Panel A in Figure 2.9 shows that the VXO and BBD indices have

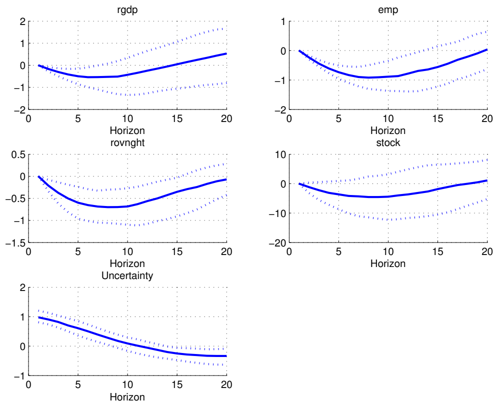

(a) Uncertainty

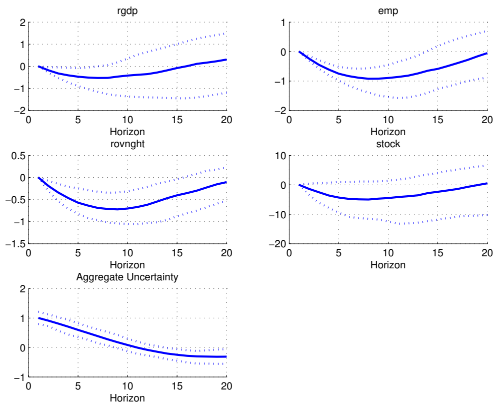

(b) Agg. Uncertainty

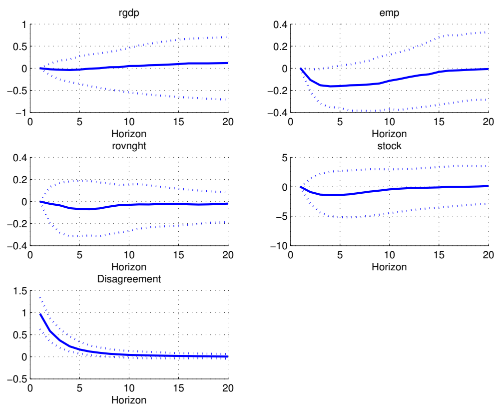

(c) Disagreement

Figure 2.6: Effects Of Uncertainty On The Economy

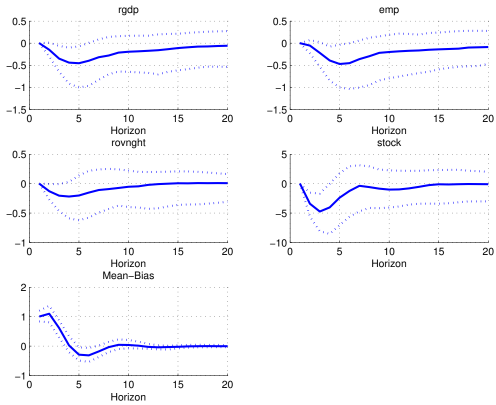

(a) Mean-Bias

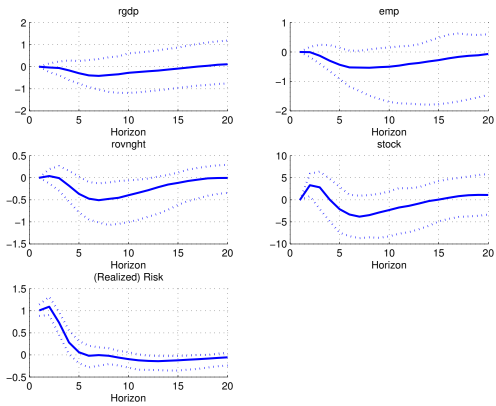

(b) Risk

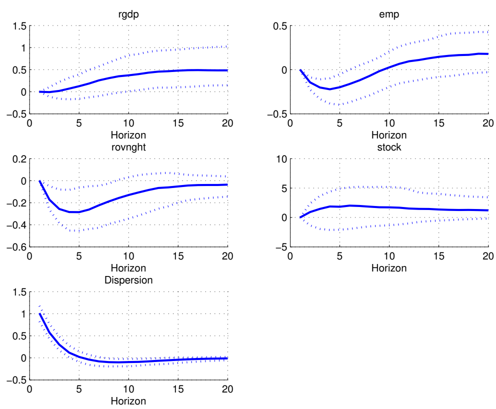

(c) Dispersion

FIGURE 2.7: Effects Of Uncertainty On The Economy, Continued

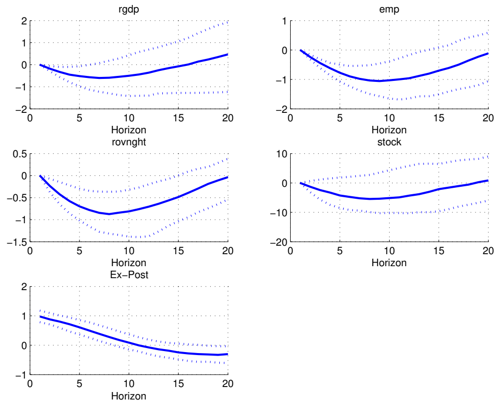

(a) Ex-Post

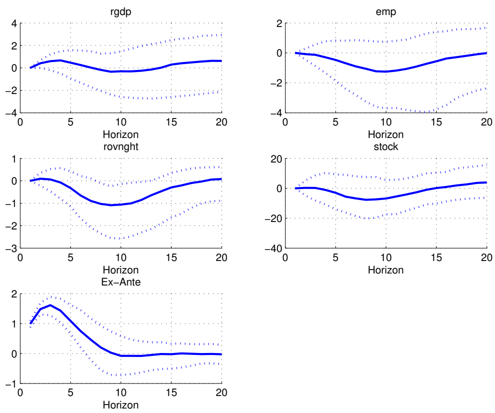

(b) Ex-Ante

Figure 2.8: Effects Of Uncertainty On The Economy, Continued

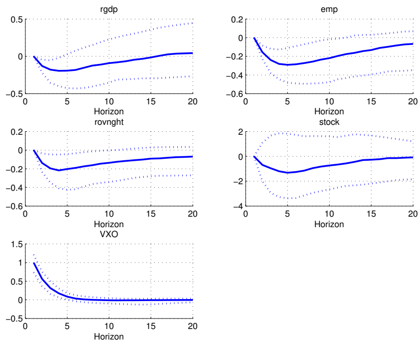

(a) VXO

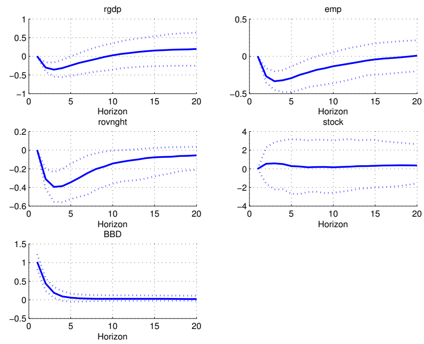

(b) BBD

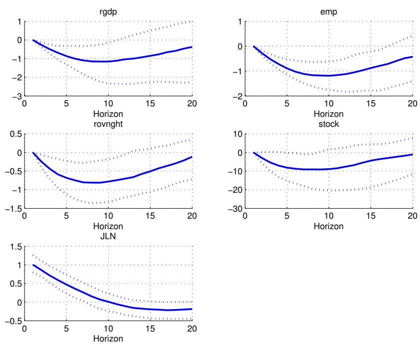

(c) JLN

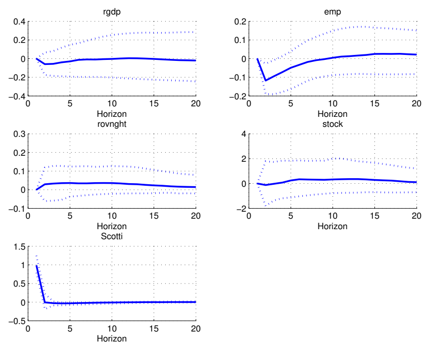

(d) Scotti

Figure 2.9: Effects Of Uncertainty On The Economy, Continued

similar effects on the economy, while an increase in uncertainty measured by the Jurado, Ludvigson and Ng's (2015) index are qualitatively similar but much larger in magnitude, and, thus, are similar to the effects uncovered ex-post uncertainty. The effects of Scotti's index are again recessionary for GDP, employment and stock markets, and lead to an increase in the interest rate. The effects of this index are small and overall insignificant. The effects of our realized volatility measure are more similar to those of the VXO.

In sum, the decomposition of uncertainty puts several components into light that are related to uncertainty measures used in the literature. This analysis helps understand why the various measures of uncertainty differ from each other, which phenomenon they actually measure, and which one is more appropriate to use depending on the goals of the researcher. This represents the second contribution of this chapter to the understanding of the effects of economic uncertainty.

2.7 Knightian uncertainty and risk in an economic model

Last, I now discuss the interpretation of the components of uncertainty in the presence of time-varying macroeconomic risk and ambiguity using a stylized model featuring these aspects. This allows the mapping of the various components in the decompositions to the sources of uncertainty that the model controls for, therefore giving support to the interpretation of the different components of uncertainty that were estimated.

The model of ambiguity follows that of Ilut and Schneider (2014). The model is built as follows. GDP growth, Z_{t+1}, evolves according to an autoregressive model with a time varying mean, μ_t^*:

$$Z_{t+1} = \rho_z Z_t + \mu_t^* + u_{t+1}, \tag{2.14}$$

where u_{t+1} is i.i.d. $\mathcal{N}(0, \sigma_u^2)$ and μ_t^* is deterministic such that its empirical sequence converges to an iid stochastic process $\mathcal{N}(0, \sigma_\mu^2)$, where $\sigma_\mu^2 = \sigma_z^2 - \sigma_u^2$. Consequently, the observed values of $z_t \equiv Z_{t+1} - \rho_z Z_t$ look like realizations from an i.i.d. process with mean zero and variance σ_z^2. For all practical purposes, we treat μ_t^* as a realization from a stochastic process $\mathcal{N}(0, \sigma_\mu^2)$. Moreover, μ_t^* and u_t are assumed to be independent. Thus, GDP growth is driven by two sources of uncertainty in the economy: the first source is the unpredictable shocks, u_{t+1}; the second, μ_t^*, is a proxy for ambiguity, as we discuss below.

Agents in this model know that the data generating process for GDP growth is autoregressive with persistence ρ_z, and are subject to two sources of uncertainty. They do not observe, however, either μ_t^* or u_{t+1}, even though they know the probability distribution of u_{t+1}. They gather intangible information about μ_t^*, which sometimes makes them relatively confident that the correct forecast of future GDP growth is $\rho_z Z_t$, and sometimes less confident, i.e. the signal is can be more or less ambiguous. One could think of a situation where the agents acquire poor quality information or conflicting news from newspapers or professional forecasters. The ambiguity is modeled by letting agents form their beliefs about GDP growth dynamics based on the following law of motion:

$$Z_{i,t+1} = \rho_z Z_{i,t} + \mu_{i,t} + u_{t+1}, \ i = 1, 2, ..., N \tag{2.15}$$

where $\mu_{it}^* \in [-a_{i,t}, -a_{i,t} + 2|a_{i,t}|]$, N is the total number of agents (equal to 100), and u_{t+1} is i.i.d. $\mathcal{N}(0, \sigma_u^2)$. The bounds on μ_{it}^* formalize the idea that sometimes agents are more ambiguous regarding the second source of disturbance to output growth: those situations are associated, in the model, with a larger value of $a_{i,t}$, which implies a larger set of beliefs and more ambiguity perceived at time t by agent i. Thus, $a_{i,t}$ is akin to the ambiguous component of uncertainty, i.e. Knightian uncertainty.[19]

Furthermore, agents receive signals about $\mu_{i,t}^*$ from the process:

$$a_{i,t+1} - \bar{a}_i = \rho_{a,i}(a_{i,t} - \bar{a}_i) + \sigma_{a,i}\varepsilon_{t+1}^a, \tag{2.16}$$

where ε_{t+1}^a is iid $\mathcal{N}(0, 1)$. One can view ε_{t+1}^a as a signal that the agent gets about the ambiguity component, whose volatility depends on $\sigma_{a,i}$. In some periods the signal results in a higher $a_{i,t}$; in such cases, there is more ambiguity and the set of beliefs is larger. In other periods, depending on the information received, the set can be smaller, in which case the agents are less ambiguous about the stochastic disturbances in the data generating process. To ensure that the average ambiguity is

[19]This notion of Knightian uncertainty is similar to that of Ilut and Schneider (2014). It should be noted, however, that they assume that the total factor productivity shocks are ambiguous, while here the same is done for output growth.

less than the total uncertainty about the process of Z_{t+1}, ambiguity is restricted by the constraints that $\bar{a}_i = n_i \sigma_z$ and $\sigma_{a,i} = \sigma_n \sigma_z$ for $n \in (0,1)$ for every agent i, where n_i and σ_n are parameters (one can think of $\bar{a}_i$ as the unconditional mean and $\sigma^2_{a,i}$ as the unconditional variance of the shock to perceived ambiguity). In particular, $n_i \overset{\text{iid}}{\sim} \mathcal{N}(n, \sigma^2_{n,I}$, where $\sigma^2_{n,I}$ controls the cross sectional variability of N.[20]

Finally, when faced with ambiguity, modeled with Equation (2.16), the agents choose:

$$\mu^{**}_{i,t} = \min([-a_{i,t}, -a_{i,t} + 2|a_{i,t}|]). \tag{2.17}$$

The effective perceived law of motion for agent i becomes:

$$Z_{i,t+1} = \rho_z Z_{i,t} + \mu^{**}_{i,t} + u_{t+1}. \tag{2.18}$$

Note that when $a_{i,t}$ is bigger, ambiguity is higher, the set of beliefs is bigger, and the wider interval implies a lower worst case mean that the agents choose.

The model is a simplification of Ilut and Schneider (2014). To be sure, they model ambiguity and risk about the technology process. However, under the assumption of fixed inputs, this would directly translate into a similar output growth dynamics. Thus, for simplicity, only the dynamics of output growth are modeled and the parameters of the output growth process, ρ_z and σ_z, are calibrated using an AR(1) model estimated on the quarterly growth rate for the U.S. GDP. On the other hand, the ambiguity parameters, i.e. $\rho_{a,}$ n and σ_n, are borrowed from their posterior mode estimates with the caveat that their estimates apply to the ambiguity in total factor productivity rather than output growth. Table 1 summarizes the baseline parameter values. Since μ_t and u_{t+1} can not be identified separately, the values for their respective variances are assigned arbitrarily. $\sigma_\mu = 0.5$ is set, while σ_u is assigned a value to

[20] Alternatively, one could model the level of ambiguity to be uniformly distributed across the forecasters. This would attenuate ambiguity.

ρ_z	0.625	Estimated from an AR(1) model fitted to GDP growth
ρ_a	0.887	Ilut and Schneider's (2014) mode
n	0.995	Ilut and Schneider's (2014) mode
σ_u	0.780	Estimated from an AR(1) model fitted to GDP growth (σ_z)
σ_μ	0.500	Arbitrary, as the parameter is not separately identified
σ_n	0.134	Ilut and Schneider (2014) mode

Table 2.1: Baseline Parameter Values

match the total conditional volatility in the output growth observed in the data, σ_z.

Four scenarios are considered. In the first three scenarios, there is no cross-sectional heterogeneity in ambiguity, i.e. $\sigma^2_{n,I} = 0$ and $n_i = n$ for every agent; in the fourth scenario we consider heterogeneity by letting $n_i \neq n$.

Scenario 1: Ambiguity. Only the level of ambiguity increases in the model, i.e. the level of n. More specifically n shifts from 0.2 to 0.8. While the data is generated by equation (2.14), the agents forecast output growth using the law of motion in Equation 2.18. In this context the set of possible values that μ_t can take changes: as n increases, both the conditional and unconditional means of a_{t+1} increase – see Equation (2.14), and the signals the agents get about the additional source of uncertainty, denoted by the set $[-a_t, -a_t + 2|a_t|]$), become noisier.

Scenario 2: Risk and ambiguity. The level of risk by increases with the value of σ_u going from 0.3 to 1. In this experiment the model is still described by Equation (2.14), the perceived law of motion is described by Equation (2.18), while learning under ambiguity occurs under Equation (2.16). In this case, increasing the level of uncertainty increases both the objective and perceived level of uncertainty. However, given that $n_i = n$, $\bar{a} = n\sigma_z$ and $\sigma_a = \sigma_n\sigma_z$ for $n \in (0, 1)$, where $\sigma_z^2 = \sigma_\mu^2 + \sigma_u^2$, both the level of ambiguity ($\bar{a}$) and the uncertainty about ambiguity (σ_a) increase.

An increase in σ_u increases both risk and Knightian uncertainty in the model.

Scenario 3: Risk but no ambiguity. Similar to Scenario 2, the level of risk increases, but agents are now forecasting based on the true model: $\mu_t^{**} = \mu_t^*$, effectively turning down ambiguity. In other words, the true model is still the one described by Equation (2.14), while the model used for forecasting is not determined by Equation (2.18), but instead by equation (2.14) itself. The design in this scenario intends to explore how the ex-post and overall uncertainty evolve when there is no ambiguity.

Scenario 4: Disagreement. The variance of ambiguity across agents in the model, i.e. $\sigma_{n,I}$, now increases, implying that agents are not all equally ambiguous about the signal. $\sigma_{n,I}$ goes from 0.5 to 1 and $\rho_{a,i} \sim \mathcal{N}(\rho_a, 0.01)$ is left heterogeneous across agents. Agents differ both in the volatility of the signal they receive and in its persistence. Note that, in this case, agents disagree on the level of ambiguity, although the aggregate level of ambiguity in the data is unchanged; that is, on average, $\bar{a}$ equals $n\sigma_z$, which does not change as $\sigma_{n,I}$ increases.

The model is simulated for 254 periods for each of these scenarios (using an additional 100 periods as a burn-in sample). The components of the proposed decompositions are then computed from the simulated data and and plotted over time.

Panel A in Figure 2.10 depicts the results for Scenario 1. The increase in ambiguity increases the Mean-Bias and the Ex-Post components of uncertainty, as well as the overall uncertainty. On the other hand, there is no change in either the perceived or the realized volatility, that is, the Ex-ante and Realized Risk components, respectively. This follows from the fact that: (i) the data generating process has not changed, and, thus, the realized variance (σ_z^2) has remained the same; and (ii) as Equation (2.18) suggests, the overall level of the ex-ante variance (σ_u^2) does not depend on n.

Panel B of Figure 2.10 shows the simulation results for the second

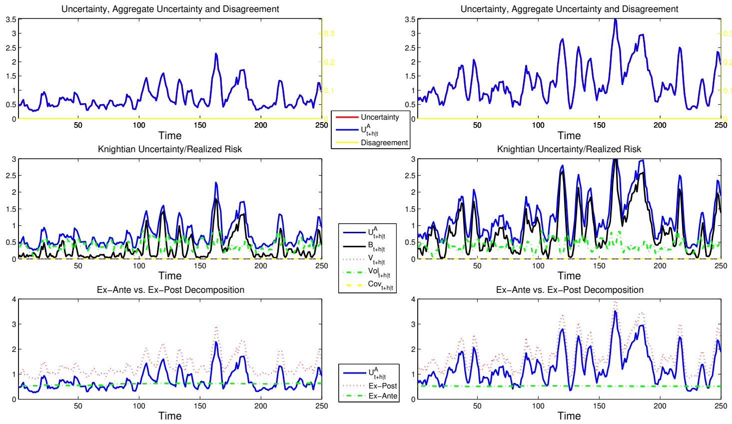

(a) Scenario 1: Ambiguity

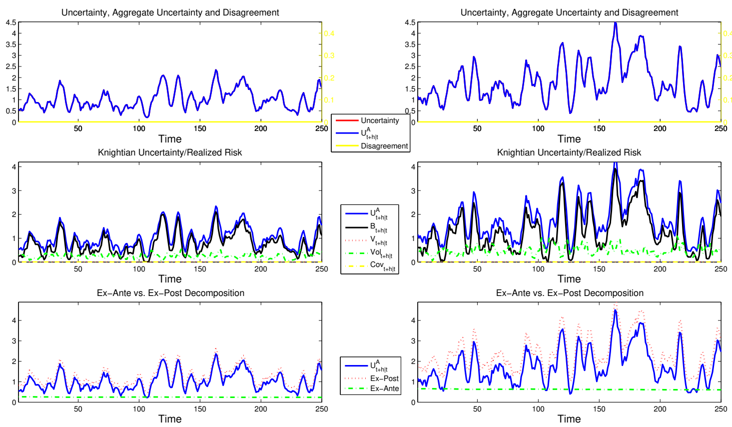

(b) Scenario 2: Increase in Risk with Ambiguity

FIGURE 2.10: Simulation Results

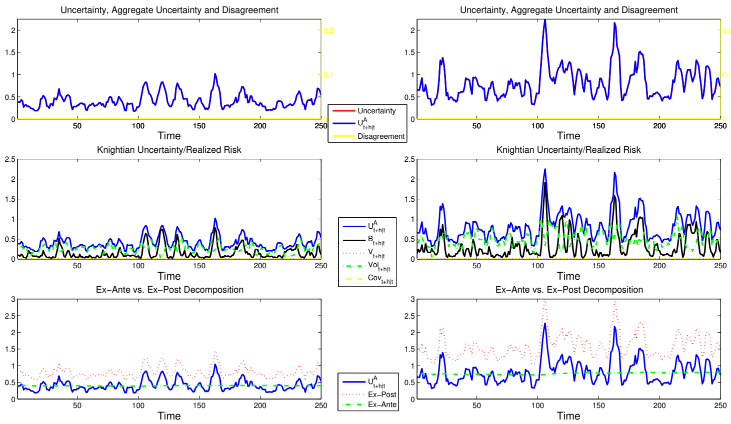

(a) Scenario 3: Increase in Risk, No Ambiguity

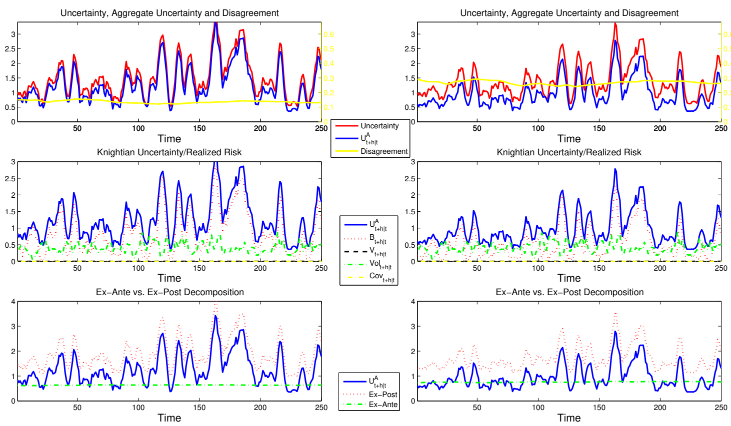

(b) Scenario 4: Increase in Dispersion of Ambiguity

Figure 2.11: Simulation Results, Continued

scenario. Here the increase in σ_u increases the measures of ex-ante (σ_u^2 itself) and ex-post risks. It is also important that there is feedback from risk to ambiguity. As discussed in the description of Scenario 2, both the mean ($\overline{a}$) and the variance (σ_a^2) of ambiguity are affected by the increase in the overall risk. Consequently, the overall measure of uncertainty increases due to both sources: increase in risk and increase in ambiguity.

Panel A of Figure 2.11 shows the dynamics of uncertainty and its components when there is an increase in risk in a model with no ambiguity. In this setup it appears that both ex-ante and ex-post components of uncertainty increase. However, this increase is proportional such that the average level of overall uncertainty increases due to the upward shift in ex-ante uncertainty and its volatility mimics that of ex-post uncertainty (in the right panel). Thus, comparing Panels B and C suggests that, in the presence of ambiguity, uncertainty increases proportionally more than the increase of risk.

Finally, Panel B of Figure 2.11 shows the dynamics when there is an increase in the cross-sectional dispersion of ambiguity while the overall level of ambiguity remains unchanged. Note that the component that is most largely affected by the increase in the cross-sectional dispersion in ambiguity is disagreement, as we would expect.

To summarize, the simulations show that the increase in ambiguity can increase the ex-post component, as well as the mean bias, thus resulting in an overall increase in uncertainty. The increase in the true volatility of the DGP increases both the realized volatility as well as the ex-ante volatility measures. However, the increase in the overall uncertainty affects the ex-post volatility and mean-bias as well. In the absence of ambiguity, the impact on the bias is negligible (it is more similar to noise), thus the increase in the aggregate uncertainty reflects the increase in the ex-post volatility. On the other hand, the increase in the ex-post uncertainty is twice as much the increase in the ex-ante uncertainty, such that the resulting measure of aggregate uncertainty still reflects the increase in the ex-ante uncertainty. Now, in the presence of ambiguity, on the other hand, the bias goes up and the ex-post uncertainty goes up proportionally more, such that the aggregate uncertainty reflects the increase in all sources of uncertainty.

These simulation results match the empirical findings. The model has a potential to generate an ex-ante uncertainty measure that is smoother than the realized variance and at the same time relatively volatile measures of bias and ex-post uncertainty. The simulations also suggest the existence of ambiguity in the empirical setup as the aggregate uncertainty does not move proportionally with the variance. In fact, the predominant sources of aggregate uncertainty are the Knightian measures.[21]

2.8 INFLATION UNCERTAINTY

The indicator of macroeconomic uncertainty that I have exposed in the previous section can be applied just as well to inflation forecasts, which are also available from the U.S. Survey of Professional Forecasters. This last section does precisely that.

Understanding inflation uncertainty is important for several reasons. High uncertainty about future inflation, possibly spurred by high inflation itself, may have effects on real variables (Ball, 1990). For example, Gürkaynak and Wright (2012) and Wright (2011) have argued that inflation uncertainty matters because it might help explain the behavior of bond risk premia, and therefore help economists understand why monetary policy differently affects short term rates (the instrument of monetary policy) and the long term rate (the rate that is of interest for investors and consumers). In fact, Wright (2011) has found a positive and strong relationship between long-term inflation uncertainty and

[21] Note that it is possible that the effect of the Knightian uncertainty is underestimated, since it is possible that the data generating process, thus realized volatility, can also change in response to ambiguity.

bond term premia in a large cross-section of countries. The important policy implication of Wright's (2011) findings is the possibility that eliminating long-run inflation uncertainty might facilitate the transmission of monetary policy to the economy. Also, D'Amico and Orphanides (2014) consider ex-ante measures of risk for inflation forecasting and Caporale et al. (2010) have shown that inflation uncertainty has decreased in the Euro area, possibly due to the fact that inflation decreased steadily since the beginning of the Euro.

Figure 2.12 depicts the measure of inflation uncertainty (Panel A) and its components (Panels B,C). Inflation uncertainty was high in the early 1980s, possibly due to oil price shocks, and decreased substantially afterwards; typically, it tends to be high around recessions. The behavior over time of uncertainty is very different from that of disagreement, which instead does not necessarily peaks around recession times. While the volatility component is pretty constant over time, the majority of the fluctuations in aggregate inflation uncertainty are associated with the bias component and the ex-post components; interestingly, ex-ante inflation uncertainty seems to have decreased monotonically since the early 1980s.

The empirical results suggest that the most effective policies to decrease inflation uncertainty are those that influence ex-post uncertainty. In other words, policies should aim at ensuring that ex-post realizations of inflation are in line with the average expected inflation (for example, by minimizing shocks to inflation), not those that decrease the agents' ex-ante uncertainty (i.e. not those that affect the agents' expectation formation process), although the latter can also be effective.

2.9 CONCLUSION

After having insisted on the importance of probabilistic assessments in the definition of economic uncertainty, I have exposed an alternative measure of uncertainty based on survey density forecasts that provide

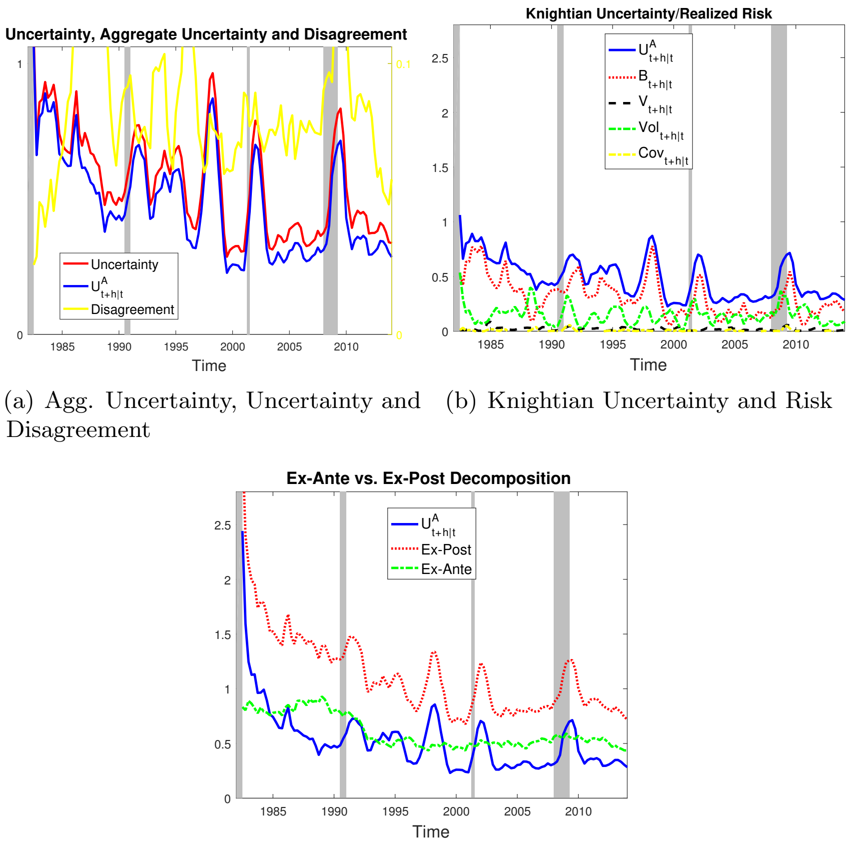

(a) Agg. Uncertainty, Uncertainty and Disagreement

(b) Knightian Uncertainty and Risk

(c) Ex-Ante and Ex-Post

Figure 2.12: Decomposing Inflation Uncertainty

such assessments. This new measure has the advantage that it can be used to decompose uncertainty into components that can help researchers understand what existing uncertainty indices relate to. In particular, this measure of uncertainty can be decomposed into aggregate uncertainty and disagreement, and aggregate uncertainty can itself be decomposed into Knightian uncertainty and realized risk. The latter inherently measure different phenomena, have specific business cycle dynamics and different macroeconomic impact. These sources of uncertainty resolve differently across prediction horizons, which is a new fact uncovered in these data.

The proposed uncertainty index is an ex-post measure of uncertainty – which is only half of the picture –, but it can be decomposed into a component that only reflects ex-ante uncertainty, which can be related to existing measures of uncertainty based on the inter-quantile spread of the forecast distribution, and a component that measures ex-post uncertainty. Some existing measures of uncertainty capture ex-ante uncertainty (such as existing measures of uncertainty based on policy uncertainty), while others capture ex-post uncertainty.

Finally, while an increase in overall uncertainty has recessionary effects, the effects of the various components of uncertainty differ. For example, disagreement is only a small portion of the overall uncertainty, and may both underestimate and lag the actual degree of uncertainty in the economy; thus it may not be a timely measure of uncertainty. In addition, both realized risk and Knightian uncertainty were important components of uncertainty over the last three decades, although the former sharply decreased as soon as the financial recession of 2007-2008 ended while the latter remained high even after the end of the crisis. This suggests that the high overall uncertainty that persisted after the end of the latest recession was mostly due to agents' being unable to assign the correct probability to the economic outcomes and disagreeing on them, rather than because risk was high. Simulation results from a stylized macroeconomic model suggest that the behavior of uncertainty and its components is largely reconcilable with a macroeconomic model with ambiguity. Ambiguity can be a source of its own in increasing the overall level of uncertainty; alternatively, it can also act as an amplifying mechanism for the increase in the level of risk.

The quantity of information that can be uncovered using density forecast from forecasters surveys suggests that there is possibly even more to discover by understanding how those forecasts were made to begin with, which is the subject of Chapters 3 and 4.

2.10 APPENDIX

Proofs of the Uncertainty Decompositions

The appendix provides the proofs for the results in this chapter. For simplicity of notation, the proofs are written for the unconditional expectation, but they all naturally hold in conditional form with the information set.

Proof of Equation (2.7)

The proof is a mechanical consequence of adding and subtracting the average forecast in the individual Brier scores averaged over the set of forecasters. Namely:

$$\begin{aligned} U &= E\left[\frac{1}{N}\sum_{s=1}^{N}\left[x_{t+h}(r) - P_{s,t+h|t}(r)\right]^2\right] \\ &= E\left[\frac{1}{N}\sum_{s=1}^{N}\left[x_{t+h}(r) - P_{t+h|t}(r) + P_{t+h|t}(r) - P_{s,t+h|t}(r)\right]^2\right]. \end{aligned}$$

Therefore:

$$\begin{aligned} U = E&\left[\frac{1}{N}\sum_{s=1}^{N}\left(x_{t+h}(r) - P_{t+h|t}(r)\right)^2\right] \\ &+ 2E\left[\frac{1}{N}\sum_{s=1}^{N}\left(x_{t+h}(r) - P_{t+h|t}(r)\right)\left(P_{t+h|t}(r) - P_{s,t+h|t}(r)\right)\right] \\ &+ E\left[\frac{1}{N}\sum_{s=1}^{N}\left(P_{t+h|t}(r) - P_{s,t+h|t}(r)\right)^2\right] \end{aligned}$$

The first term is the uncertainty of the average forecast. The other two terms need a bit of rewriting:

$$
\begin{aligned}
U &= U^A + 2E\left[\left(x_{t+h}\left(r\right) - P_{t+h|t}\left(r\right)\right)\frac{1}{N}\sum_{s=1}^{N}\left(P_{t+h|t}\left(r\right) - P_{s,t+h|t}\left(r\right)\right)\right] \\
&\quad + E\left[\frac{1}{N}\sum_{s=1}^{N}\left(P_{t+h|t}\left(r\right) - P_{s,t+h|t}\left(r\right)\right)^2\right] \\
&= U^A + 2E\left[\left(x_{t+h}\left(r\right) - P_{t+h|t}\left(r\right)\right)\left(P_{t+h|t}\left(r\right) - \frac{1}{N}\sum_{s=1}^{N}P_{s,t+h|t}\left(r\right)\right)\right] \\
&\quad + E\left[\frac{1}{N}\sum_{s=1}^{N}\left(P_{t+h|t}\left(r\right) - P_{s,t+h|t}\left(r\right)\right)^2\right] \\
&= E\left[\left(x_{t+h}\left(r\right) - P_{t+h|t}\left(r\right)\right)^2\right] + 0 + E\left[\frac{1}{N}\sum_{s=1}^{N}\left(P_{t+h|t}\left(r\right) - P_{s,t+h|t}\left(r\right)\right)^2\right].
\end{aligned}
$$

Proof of Equation (2.10)

$$
\begin{aligned}
u^{A}_{t+h}(r) &\equiv E\left[\left(x_{t+h}(r) - P_{t+h|t}(r)\right)^{2}\right] \\
&= E\left[\left(x_{t+h}(r) - E\left(x_{t+h}(r)\right) + E\left(x_{t+h}(r)\right) - P_{t+h|t}(r)\right)^{2}\right] \\
&= E\left[\left(x_{t+h}(r) - E\left(x_{t+h}(r)\right)\right)^{2}\right] + E\left[\left(E\left(x_{t+h}(r)\right) - P_{t+h|t}(r)\right)^{2}\right] \\
&\quad + E\left[2\left(x_{t+h}(r) - E\left(x_{t+h}(r)\right)\left(E\left(x_{t+h}(r)\right) - P_{t+h|t}(r)\right)\right)\right] \\
&= E\left(\left[P_{t+h|t}(r) - E\left(x_{t+h}(r)\right)\right]^{2}\right) + V\left(x_{t+h}(r)\right) \\
&\quad - 2\mathrm{Cov}(x_{t+h}(r)\, P_{t+h|t}(r)),
\end{aligned}
$$

where the last line follows from the fact that:

$$
\begin{aligned}
&E\left[2\left(x_{t+h}(r) - E\left(x_{t+h}(r)\right)\left(E\left(x_{t+h}(r)\right) - P_{t+h|t}(r)\right)\right)\right] \\
&= 2E\left[\left(x_{t+h}(r) - E\left(x_{t+h}(r)\right)\left(E\left(x_{t+h}(r)\right) - P_{t+h|t}(r)\right)\right)\right] \\
&= 2E\left[P_{t+h|t}(r)\, Ex_{t+h}(r) - x_{t+h}(r)\, P_{t+h|t}(r)\right] \\
&= 2\left[EP_{t+h|t}(r)\, Ex_{t+h}(r) - E(x_{t+h}(r)\, P_{t+h|t}(r))\right] \\
&= -2\mathrm{Cov}(x_{t+h}(r)\, P_{t+h|t}(r)).
\end{aligned}
$$

Furthermore, note that

$$
E\left[P_{t+h|t}(r) - E\left(x_{t+h}(r)\right)\right]^{2} = \left[E\left(P_{t+h|t}(r) - E\left(x_{t+h}(r)\right)\right]^{2}\right) + V(P_{t+h|t}(r)).
$$

This identity is found by adding and substracting $E(P_{t+h|t})$ in the squared term.

Proof of Equation (2.12)

Our measure of uncertainty is the negative of the CRPS (Gneiting and Raftery, 2007). Note that $CRPS(F, y_{t+h}) = -\int_{-\infty}^{\infty}(F(r) - 1\{y_{t+h} < r\})^2 \, \mathrm{d}r = -U_{t+h}^{A}$, where $F(r)$ is the aggregate predictive distribution. Let $G(r)$ denote the ideal – perfect forecast – distribution, i.e. $G(r) = 1\{y_{t+h} < r\}$; then by Lemma 2.2 of Baringhaus and Franz (2004), we have:

$$\begin{aligned} u_{t+h}^{A} &= \int_{-\infty}^{\infty}(F(r) - 1\{y_{t+h} < r\})^2 \, \mathrm{d}y \\ &= E|Y_{1,t+h} - y_{1,t+h}| - \frac{1}{2}E|Y_{1,t+h} - Y_{2,t+h}| - \frac{1}{2}E|y_{1,t+h} - y_{2,t+h}|, \end{aligned}$$

where $Y_{1,t+h}$ and $Y_{2,t+h}$ are i.i.d draws from F, while $y_{1,t+h}$ and $y_{2,t+h}$ are i.i.d. draws from $G(r)$, and both of these variables have finite expectations. Given Lemma 2.1 of Baringhaus and Franz (2004),

$$E|y_{1,t+h} - Y_{1,t+h}| = \int_{-\infty}^{\infty} F(r)(1 - G(r)) \, \mathrm{d}r + \int_{-\infty}^{\infty} G(r)(1 - F(r)) \, \mathrm{d}r.$$

Now for $y_{1,t+h}$ and $y_{2,t+h}$, we have:

$$\begin{aligned} E|y_{1,t+h} - y_{2,t+h}| &= 2\int_{-\infty}^{\infty} G(r)(1 - G(r)) \, \mathrm{d}r \\ &= 2\int_{-\infty}^{\infty} 1\{y_{t+h} < r\}(1 - 1\{y_{t+h} < r\}) \, \mathrm{d}r \\ &= 0, \end{aligned}$$

where the last equality follows from the fact that, for a particular value of r, either $1\{y_{t+h} < r\}$ or $1 - 1\{y_{t+h} < r\}$ will be zero, and, thus, the product will always equal zero. Therefore,

$$\begin{aligned} u_{t+h}^{A} &= \int_{-\infty}^{\infty}(F(r) - 1\{y_{t+h} < r\})^2 \, \mathrm{d}r && (2.19) \\ &= E|Y_{1,t+h} - y_{1,t+h}| - \frac{1}{2}E|Y_{1,t+h} - Y_{2,t+h}|. && (2.20) \end{aligned}$$

This means we can rewrite aggregate uncertainty as the sum of expected absolute distance measures of random variables coming from the predictive distribution, and that coming from the predictive distribution and the true distribution which generates the realization. If $F(r)$ is Gaussian, i.e. if $Y_{t+h} \sim^{\text{iid}} \mathcal{N}(\mu_{t+h}, \sigma^2_{t+h})$, then because the difference of iid normal random variables is normally distributed (in this case centered around zero with a variance of $2\sigma^2_{t+h}$), and the fact that the absolute value of a mean zero normal random variable has a half-normal distribution with mean $\frac{2\sigma_{t+h}}{\sqrt{\pi}}$, we have

$$\frac{1}{2}E|Y_{1,t+h} - Y_{2,t+h}| = \frac{\sigma_{t+h}}{\sqrt{\pi}}. \tag{2.21}$$

To obtain $E|Y_{t+h} - y_{t+h}|$, we use the properties of Dirac delta function. We denote the PDF of y_{t+h} by a Dirac delta function $\delta(r - y_{t+h})$. From the properties of the Dirac function, $E(y_{t+h}) = y_{t+h}$ and $V(y_{t+h}) = 0$. Then, $Y_{1,t+h} - y_{1,t+h} \sim N(\mu_{t+h} - y_{t+h}, \sigma^2_{t+h})$. By property of the folded normal distribution, we have:

$$\begin{aligned} E|Y_{t+h} - y_{t+h}| = \sigma_{t+h} 2\varphi \left(-\frac{\mu_{t+h} - y_{t+h}}{\sigma_{t+h}} \right) \dots \\ + (\mu_{t+h} - y_{1,t+h}) \left(1 - 2\Phi \left(-\frac{\mu_{t+h} - y_{t+h}}{\sigma_{t+h}} \right) \right). \end{aligned} \tag{2.22}$$

Substituting (2.22) and (2.21) into (2.20), and taking expectations with respect to Q we get the result:

$$U^A_{t+h} = E_Q(\text{Ex-Post}) + E_Q(\text{Ex-Ante}) \tag{2.23}$$

where:

$$\text{Ex-Post} = 2\sigma_{t+h}\phi\left(\frac{y_{t+h} - \mu_{t+h}}{\sigma_{t+h}}\right) + (y_{t+h} - \mu_{t+h})\left(2\Phi\left(\frac{y_{t+h} - \mu_{t+h}}{\sigma_{t+h}}\right) - 1\right)$$

$$\text{Ex-Ante} = -\frac{\sigma_{t+h}}{\sqrt{\pi}}$$

Reliability and Resolution Analysis

Note that an additional, interesting decomposition for $u^A_{t+h|t}(r)$ can be obtained following Murphy (1973):

$$u^A_{t+h|t}(r) \simeq REL_{t+h|t}(r) - RES_{t+h|t}(r) + V(x_{t+h}(r)) \qquad (2.24)$$

where:

- $REL_{t+h|t}(r) \equiv E\left(\left[P_{t+h|t}(r) - E\left(x_{t+h}(r)\,|P_{t+h|t}(r)\right)\right]^2\right)$ measures the reliability of the forecast and scores the calibration of the forecast. A forecast is said "reliable" when the observed frequency is consistent with the probabilistic forecast made for a given event. For instance, forecasts that predict a probability of recession of 30 percent will be reliable if the economy effectively enters a recession 30 percent of the time every time such a forecast is made. Hence, reliability measures the unconditional (un)biasedness of the probabilistic forecasts. Because the term is expressed as a squared error, the smaller the calibration error, the better (i.e., the lower) the Brier score.

- $RES_{t+h|t}(r) \equiv E\left(\left[E\left(x_{t+h}(r)\,|P_{t+h|t}(r)\right) - E(x_{t+h}(r))\right]^2\right)$ is the resolution, i.e. the average squared differential of the conditional and unconditional means of the observed outcomes. It captures the "decisiveness" of forecasts by comparing the forecast probability and the long-term average of the underlying process. The larger the term, the lower the Brier score.

As we show below, Eq. (2.24) holds up to an approximation error that involves within bin variation. The decomposition can be estimated as follows. Reliability is estimated as follows. For each t, determine which of the forecast bins $P_{t+h|t}(r)$ falls into. Let $\left\{P^{(k)}_{t+h|t}(r)\right\}$ be the collection of probabilities in the k-th bin and let $P^E_{t+h|t}(r)$ denote the unconditional expected value over the bin. We estimate $P^{E}_{t+h|t}(r)$ using a Uniform distribution over the bin, so that $P^{E}_{t+h|t}(r)$ is the midpoint of the bin.[22] In addition, let the number of probabilities in each bin be n_k. Let $\overline{x}_k$ be the average of the realizations conditional on the forecaster having made the probability forecast associated with the collection of probabilities in bin k, $\left\{P^{(k)}_{t+h|t}(r)\right\}$. Reliability is the average square calibration error, that is,

$$REL(r) = \frac{1}{T}\sum_{k=1}^{K} n_k \left(P^{E}_{t+h|t}(r) - \overline{x}_k(r)\right)^2. \tag{2.25}$$

Thus, reliability measures the squared deviation of the predicted probability from the observed outcome conditional probability of the event. This effectively tells the user how often (as a percentage) a forecast probability actually occurred. In theory, a perfect forecasting model will result in forecasts with a probability of $\alpha\%$ being consistent with the eventual outcome $\alpha\%$ of the time. Note that a forecast is reliable if the average square calibration error (REL) is small. Figure 2.13 provides intuition to understand reliability. The x-axis reports the forecast probability,[23] while the y-axis reports the observed relative frequency. A reliable forecast would be the 45-degree line, where the observed frequency of realizations equals the forecast probability; the data clearly show departures from reliability in our sample.

Resolution is the squared average difference between the conditional mean (given the forecast) and the unconditional mean: $RES(r) =$

[22]In the 3-terms decomposition that we discuss here, we abstract from within bin variance and within bin covariance; thus, the unconditional expected value over the bin is indeed the midpoint of the bin and all forecasts in the bin are imposed to be equal to the midpoint (so their average is also the midpoint). We derive a 5-term decomposition which includes within bin variance and within bin covariance (Stephenson, Coelho and Joliffe, 2008). In that case, the reliability will be calculated using the average forecast in the bin without imposing that all forecasts in the bin are equal. That is, $\overline{p}^{(k)}_{t+h|t}(r)$ (which is the average of the collection of probabilities in the k-th bin, $\left\{P^{(k)}_{t+h|t}(r)\right\}$), replaces $P^{E}_{t+h|t}(r)$ in eq. (2.25).

[23]The forecast probability is the mid-point of the bin in the forecast distribution.

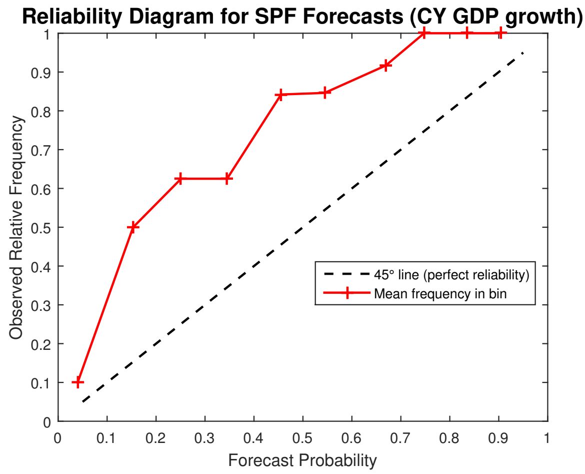

Figure 2.13: Reliability Diagram

Notes. The figure plots the reliability diagram for SPF forecasts of current year (CY) GDP growth.

$\frac{1}{T}\sum_{k=1}^{K} n_k \left(\overline{x}_k\left(r\right) - \overline{x}\left(r\right)\right)^2$. Note that good forecasts have high resolution.

Figure 2.14 shows the evolution of the components of the alternative decomposition over time.[24]

Proof of Equation (2.24)

In practice, the Murphy decomposition requires partitioning the range of forecasts – essentially, the [0,1] line – into K sub-segments. Let r be a number along the real line; let $\overline{p}^{(k)}$ denote the average probability in segment k;[25] and let n_k denote the number of forecast probabilities that

[24] Finally, note that the practical implementation of the Brier score involves "binning". Binning smooths the data and makes them less noisy, as larger bins limit the "sparseness" problem (Stephenson et al., 2008). Some information is lost, however, by approximating continuous probability densities with a discrete number of bins.

[25] Alternatively, one could consider $\overline{p}^{(k)}$ as the midpoint of the k-th segment

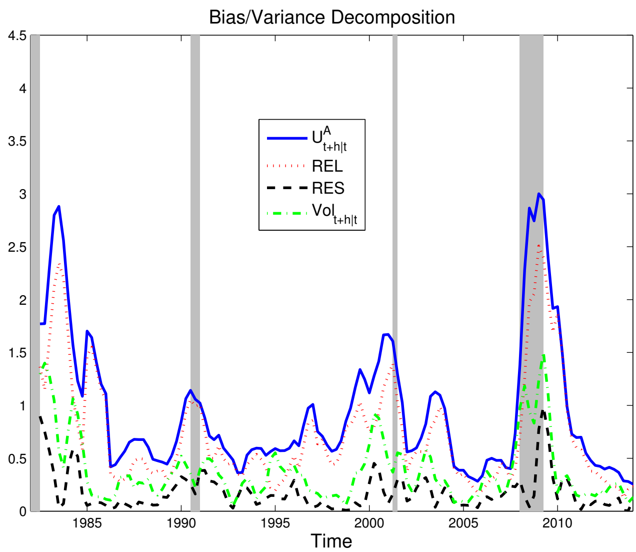

FIGURE 2.14: Aggregate Uncertainty, Reliability, Resolution and (Realized) Risk

Notes. The figure displays Aggregate Uncertainty, Reliability, Resolution and Realized Risk.

fall in the k-th sub-segment, for $k = 1, \dots, K$. Given all forecasts in the sample, the Brier score can be broken down as follows:

$$\frac{1}{T}\sum_{t=1}^{T}[x_{t+h}(r) - P_{t+h|t}(r)]^2 = \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x^{(j)}_{t+h}(r) - P^{(j)}_{t+h|t}(r)\right]^2$$

which further equals:

$$= \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x^{(j)}_{t+h}(r) - \overline{x}^{(k)}_{t+h}(r) + \overline{x}^{(k)}_{t+h}(r) - \overline{p}^{(k)}_{t+h|t}(r) + \overline{p}^{(k)}_{t+h|t}(r) - P^{(j)}_{t+h|t}(r)\right]^2$$

$$
\begin{aligned}
&= \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x^{(j)}_{t+h}(r) - \overline{x}^{(k)}_{t+h}(r)\right]^2 + \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[\overline{x}^{(k)}_{t+h}(r) - \overline{p}^{(k)}_{t+h|t}(r)\right]^2 \\
&+ \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[\overline{p}^{(k)}_{t+h|t}(r) - P^{(j)}_{t+h|t}(r)\right]^2 \\
&+ \frac{2}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x^{(j)}_{t+h}(r) - \overline{x}^{(k)}_{t+h}(r)\right]\left[\overline{x}^{(k)}_{t+h}(r) - \overline{p}^{(k)}_{t+h|t}(r)\right] \\
&+ \frac{2}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x^{(j)}_{t+h}(r) - \overline{x}^{(k)}_{t+h}(r)\right]\left[P^{(k)}_{t+h}(r) - P^{(j)}_{t+h|t}(r)\right] \\
&+ \frac{2}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[P^{(j)}_{t+h|t}(r) - \overline{p}^{(k)}_{t+h|t}(r)\right]\left[\overline{x}^{(k)}_{t+h}(r) - \overline{p}^{(k)}_{t+h|t}(r)\right] \\
&= \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x^{(j)}_{t+h}(r) - \overline{x}^{(k)}_{t+h}(r)\right]^2 + \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[\overline{x}^{(k)}_{t+h}(r) - \overline{p}^{(k)}_{t+h|t}(r)\right]^2 \\
&+ \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[\overline{p}^{(k)}_{t+h|t}(r) - P^{(j)}_{t+h|t}(r)\right]^2 \\
&+ \frac{2}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x^{(j)}_{t+h}(r) - \overline{x}^{(k)}_{t+h}(r)\right]\left[P^{(k)}_{t+h}(r) - P^{(j)}_{t+h|t}(r)\right].
\end{aligned}
$$

We can already recognize the reliability (REL) in the second term of this decomposition:

$$
\begin{aligned}
REL(r) &= \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[\overline{x}^{(k)}_{t+h}(r) - \overline{p}^{(k)}_{t+h|t}(r)\right]^2 \\
&= \frac{1}{T}\sum_{k=1}^{K} n_k \left[\overline{x}^{(k)}_{t+h}(r) - \overline{p}^{(k)}_{t+h|t}(r)\right]^2. \qquad (2.26)
\end{aligned}
$$

The first term can be expressed as follows:

$$\frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x_{t+h}^{(j)}(r)-\overline{x}_{t+h}^{(k)}(r)\right]^2 = \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x_{t+h}^{(j)}(r)-\overline{x}(r)+\overline{x}(r)-\overline{x}_{t+h}^{(k)}(r)\right]^2$$

$$= \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x_{t+h}^{(j)}(r)-\overline{x}(r)\right]^2$$

$$+ \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[\overline{x}(r)-\overline{x}_{t+h}^{(k)}(r)\right]^2$$

$$+ \frac{2}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x_{t+h}^{(j)}(r)-\overline{x}(r)\right]\left[\overline{x}(r)-\overline{x}_{t+h}^{(k)}(r)\right]$$

$$= \frac{1}{T}\sum_{t=1}^{T}\left[x_{t+h}(r)-\overline{x}(r)\right]^2$$

$$- \frac{1}{T}\sum_{k=1}^{K} n_k\left[\overline{x}(r)-\overline{x}_{t+h}^{(k)}(r)\right]^2$$

$$\equiv V\left(x_{t+h}\left(r\right)|\Im_{t-R}^{t}\right) - RES(r).$$

Note that because the outcome variable x is binary, the uncertainty term can be expressed as $V\left(x_{t+h}\left(r\right)|\Im_{t-R}^{t}\right) = \overline{x}(r)(1-\overline{x}(r))$. To summarize, we have decomposed the Brier score in the following way:

$$\frac{1}{T}\sum_{t=1}^{T}[x_{t+h}(r)-P_{t+h|t}(r)]^2 = V\left(x_{t+h}\left(r\right)|\Im_{t-R}^{t}\right) + REL(r) - RES(r)$$

$$+ \frac{1}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[\overline{p}_{t+h|t}^{(k)}(r)-P_{t+h|t}^{(j)}(r)\right]^2$$

$$+ \frac{2}{T}\sum_{k=1}^{K}\sum_{j=1}^{n_k}\left[x_{t+h}^{(j)}(r)-\overline{x}_{t+h}^{(k)}(r)\right]\left[P_{t+h}^{(k)}(r)-P_{t+h|t}^{(j)}(r)\right].$$

The last two terms measure the variance of forecasts within the sub-segments and the co-movement between forecasts within a segment and their corresponding outcomes. The decomposition therefore writes:

$$\frac{1}{T}\sum_{t=1}^{T}[x_{t+h}(r) - P_{t+h|t}(r)]^2 = V\left(x_{t+h}\left(r\right)|\Im_{t-R}^{t}\right) + REL(r) - RES(r) + WSV(r) + WSC(r).$$

Remark that the last two terms equal zero when all forecasts within the same segment are assumed identical. Because $WSV(r)$ and $WSC(r)$ are quantitatively very small in the data, we will work under the simpler decomposition:

$$\frac{1}{T}\sum_{t=1}^{T}[x_{t+h}(r) - P_{t+h|t}(r)]^2 \simeq V\left(x_{t+h}\left(r\right)|\Im_{t-R}^{t}\right) + REL(r) - RES(r),$$

as per the definitions.

Results For a Large Dimensional VAR

Variable	*Mnemonics*	*Description*
real GDP	GDPC96	Real Gross Domestic Product, 3 Decimal (Billions of Chained 2009 Dollars)
Employment	PAYEMS	All Employees: Total nonfarm (Thousands of Persons)
Real Consumption	PCECC96	Real Personal Consumption Expenditures (Billions of Chained 2009 Dollars)
PCE deflator	PCECTPI	Personal Consumption Expenditures: Chain-type Price Index (Index 2009=100)
real new order	AMDMNOx	Real Manufacturers' New Orders: Durable Goods (Millions of 2009 Dollars), deflated by Core PCE
real wage	AHETPIx	Real Average Hourly Earnings of Production and Nonsupervisory Employees: Total Private (2009 Dollars per Hour), deflated by Core PCE
hours	HOANBS	Nonfarm Business Sector: Hours of All Persons (Index 2009=100)
federal funds rate	FEDFUNDS	Effective Federal Funds Rate (Percent)
S&P 500 Index	S&P 500	S&P's Common Stock Price Index: Composite
M2	M2REALX	Real M2 Money Stock (Billions of 1982-84 Dollars)

Table 2.2: Description of Variables Included in the VAR

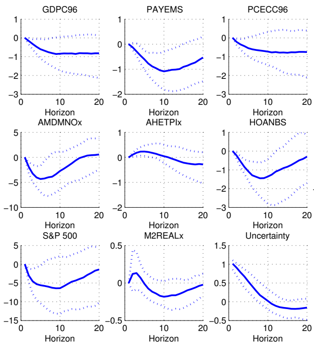

(a) Uncertainty

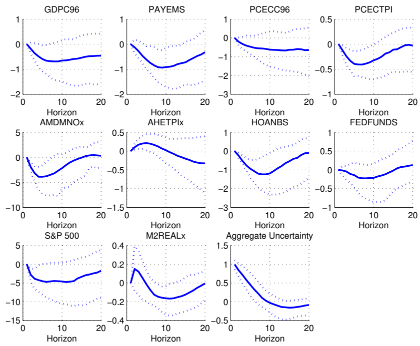

(b) Aggregate Uncertainty

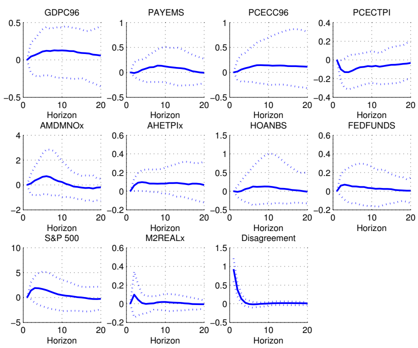

(c) Disagreement

Figure 2.15: Results for the Large VAR

Alternative Ex-Ante Uncertainty Measure

The ex-ante uncertainty, $\sigma_{t+h|t}/\pi$, can more generally estimated, for any predictive distribution, as:

$$\int_{-\infty}^{+\infty} E\left[\left(x_{t+h}(r) - P_{s,t+h|t}(t)\right)^2 |\Omega_{s,t}\right] \, \mathrm{d}r = \int_{-\infty}^{+\infty} P_{s,t+h|t}(t)(1 - P_{s,t+h|t}(t)) \, \mathrm{d}r;$$

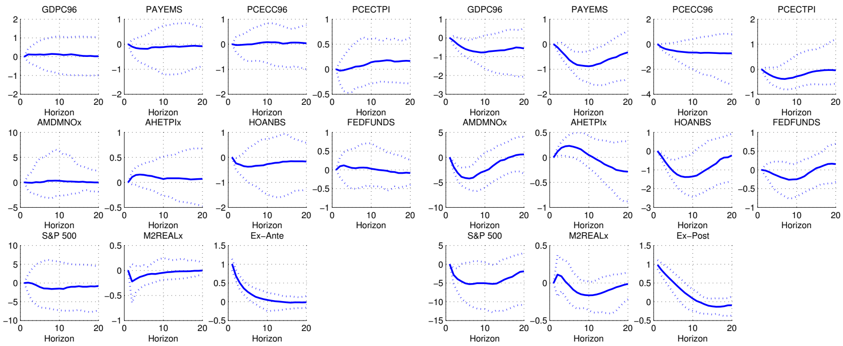

(a) Ex-Post (b) Ex-Ante

FIGURE 2.16: Results for the Large VAR

averaged across forecasters. Figure 2.17 shows indeed that they are the same object and behave very similarly.

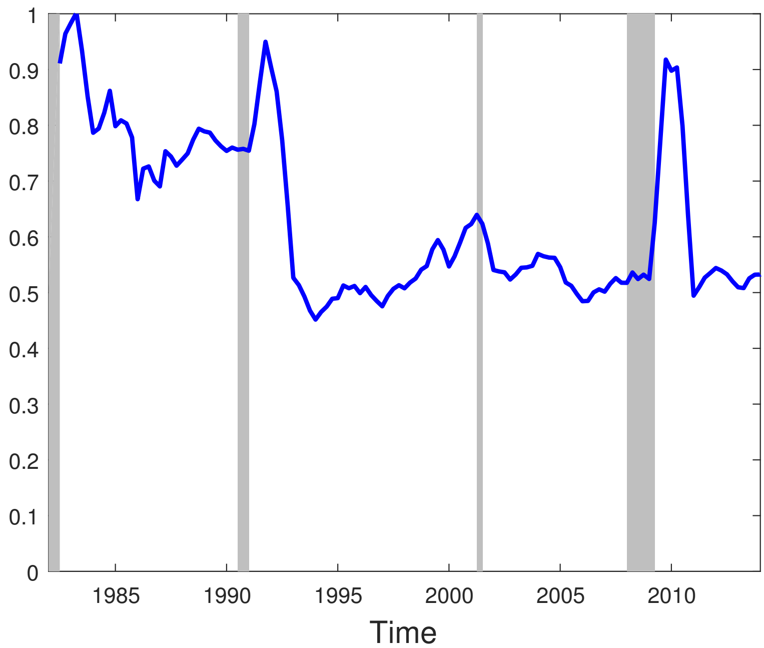

Figure 2.17: Alternative Uncertainty Measure

Estimation

The decomposition is estimated with its sample counterparts:

$$\widehat{U}_{t+h|t} = \int_{-\infty}^{+\infty} \widehat{u}_{t+h|t}\left(r\right) \, \mathrm{d}r, \;\; t = R, ..., T$$

where R is the size of the rolling window,

$$\begin{aligned} \widehat{u}_{t+h|t}\left(r\right) &= \frac{1}{R} \sum_{j=t-R+1}^{t} \frac{1}{N} \sum_{s=1}^{N} u_{s,j+h|j}\left(r\right) \\ &= \frac{1}{R} \sum_{j=t-R+1}^{t} \frac{1}{N} \sum_{s=1}^{N} \left[x_{t+h}\left(r\right) - p_{s,j+h|j}\left(r\right)\right]^2 \end{aligned}$$

and

$$\widehat{U}^A_{t+h|t} = \int_{-\infty}^{+\infty} \left(\overline{p}_{t+h|t}\left(r\right) - \overline{x}_{t+h}\left(r\right)\right)^2 \mathrm{d}r + \int_{-\infty}^{+\infty} \widehat{V}\left(p_{t+h|t}\left(r\right)\right) \mathrm{d}r$$
$$+ \int_{-\infty}^{+\infty} \widehat{\mathrm{Vol}}_{t+h|t}\left(r\right) \mathrm{d}r - 2\int_{-\infty}^{+\infty} \widehat{\mathrm{Cov}}(x_{t+h}\left(r\right), p_{t+h|t}\left(r\right)) \mathrm{d}r, \tag{2.27}$$

where the terms on the RHS of eq. (2.27) are as follows:

- $\overline{p}_{t+h|t}\left(r\right)$, $\overline{x}_{t+h}\left(r\right)$ are estimated by:

$$\frac{1}{R}\sum_{j=t-R+1}^{t} p_{j+h|j}\left(r\right),$$

  and

$$\frac{1}{R}\sum_{j=t-R+1}^{t} x_{j+h}\left(r\right);$$

- $\widehat{\mathrm{Vol}}_{t+h}\left(x_{t+h}\left(r\right)\right)$ is an estimate of the variance of $x_{t+h}\left(r\right)$, which is a binary variable, recursively over time:

$$\widehat{\mathrm{Vol}}_{t+h}\left(x_{t+h}\left(r\right)\right) = \overline{x}_{t+h}\left(1 - \overline{x}_{t+h}\right);$$

- $\widehat{V}_{t+h}\left(p_{t+h|t}\left(r\right)\right)$ is an estimate of the variance of $p_{t+h|t}\left(r\right)$ recursively over time:

$$\widehat{V}_{t+h}\left(p_{t+h|t}\left(r\right)\right) = \frac{1}{R}\sum_{j=t-R+1}^{t}\left(p_{j+h|j}\left(r\right) - \overline{p}_{t+h|t}\left(r\right)\right)^2$$

- $\widehat{\mathrm{Cov}}(x_{t+h}\left(r\right) p_{t+h|t}\left(r\right))$ is estimated as:

$$\widehat{\mathrm{Cov}}(x_{t+h}\left(r\right), p_{t+h|t}\left(r\right)) =$$
$$\frac{1}{R}\sum_{j=t-R+1}^{t}\left(p_{j+h|j}\left(r\right) - \overline{p}_{t+h|t}\left(r\right)\right)\left(x_{j+h}\left(r\right) - \overline{x}_{t+h}\left(r\right)\right)$$

While we do not need the Normality assumption to calculate the decomposition above, in practice we fit a Gaussian distribution to the predictive density. The main reason is to guarantee that the "Knightian uncertainty/(Realized) Risk" decomposition is consistent with the"Ex-ante"/"Ex-post", since the latter is valid only under Normality. Furthermore, in the empirical implementation we let $R = 4$, which amounts to calculating 4-quarter-moving average of the various components of uncertainty, and we proxy the indefinite integrals with definite ones by treating the extrema of either the realization or the bins as integral bounds.

Detailed Analysis of Uncertainty Across Forecast Horizons

Each plot in Figure 2.18 contains eight forecasts made in a given year: 1 quarter ahead, 2 quarters ahead, etc. Each density is then compared to the corresponding realization of GDP growth, depicted as a vertical line. Two things can be noted from those graphs. First, densities tend to get narrower at shorter horizons. That ís what one would expect based on our analysis: the shorter the horizon, the more concentrated the forecast will be. This illustrates why ex-ante uncertainty is lower at short horizons than at a longer horizons, which is what we found with our uncertainty measure. Second, since densities at longer horizons are less concentrated, the actual realizations may still end up well inside the predictive distribution and hence the ex-post error (in terms of likelihood) need not be greater than that of a concentrated, short-term forecast. To see this in detail, consider the examples for the following years:

– 1984: Long horizon forecasts were quite flat and in the end, the realization fell quite close to the center of the curve. On the other hand, the short term forecasts were concentrated and missed the realization substantially. Ex-post error is higher for short term horizons than for long term.
– 1995: This picture shows the opposite situation. Long-horizon forecasts missed the realization, but short-term forecasts hit the nail on the head. Ex-post error is lower at short horizons than at long horizons.
– 1992: Both long and short term horizon failed in predicting. Ex-post error should be about the same in both cases.

As one looks across diferent points in time, there are many more cases where the pictures look like the situation in 1984 than in 1995, which explains why, on average, the results show that ex-ante uncertainty decreases as the horizon decreases, but ex-post uncertainty increases.

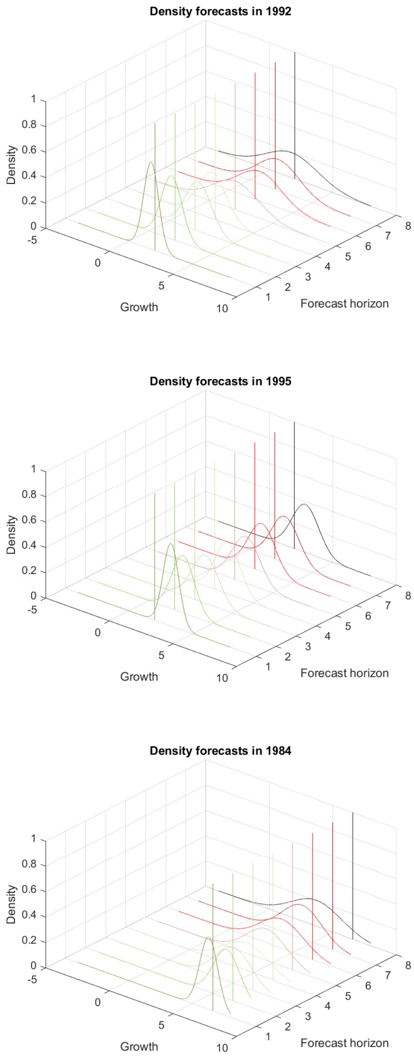

Figure 2.18: Examples of Predictive Densities and Realizations

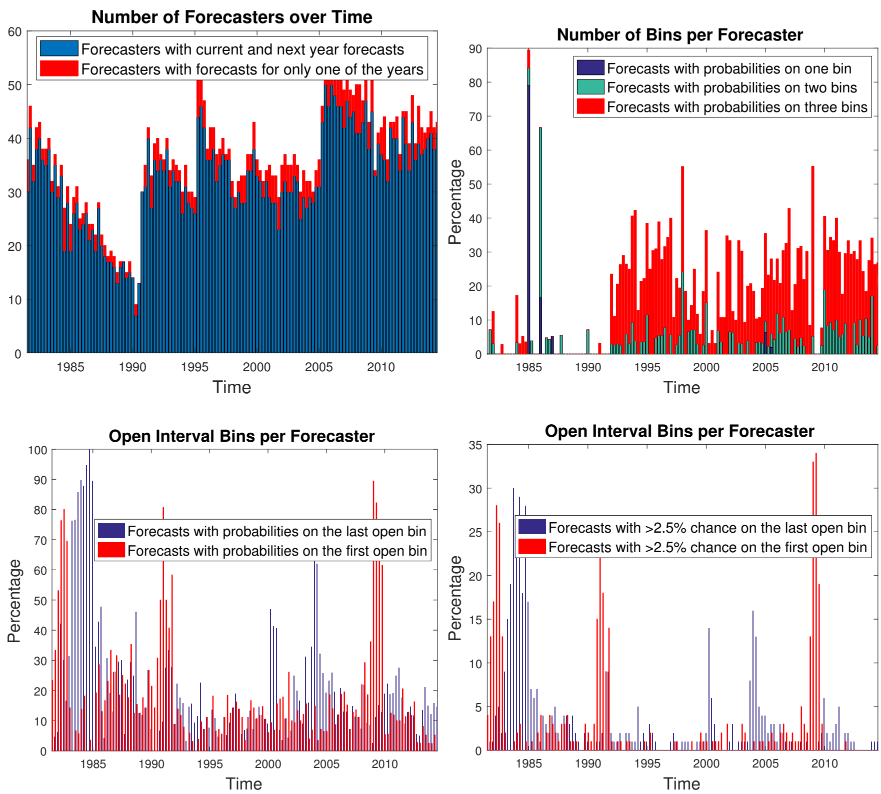

Figure 2.19: Descriptive Statistics on the SPF

3

INFORMATION AND UNCERTAINTY

My greatest concern was what to call it. I thought of calling it "information", but the word was overly used, so I decided to call it "uncertainty." When I discussed it with John von Neumann, he had a better idea. Von Neumann told me, "You should call it entropy, for two reasons. In the first place your uncertainty function has been used in statistical mechanics under that name, so it already has a name. In the second place, and more important, nobody knows what entropy really is, so in a debate you will always have the advantage."

— Claude Shannon, *Scientific American*, September 1971

3.1 Introduction

Claude Shannon's theory of information has had a tremendous impact on economic theory, yet little on the study of economic uncertainty as an object of its own. The rational inattention framework of Christopher Sims and the robustness ideas of Thomas Sargent both read like information theory between the lines. They inherently rely on the idea that agents have limited information processing capabilities and are therefore far from the perfect forecasters that the rational expectations hypothesis thought they were. Agents are "rationally inattentive"; they are aware of their limitations and optimize their decisions accordingly. Neither Sims nor Sargent, however, try to quantify uncertainty. This chapter does not discuss these ideas nor does it build on them; rather, it provides a simple framework to estimate uncertainty using entropic quantities on survey data.

Information theory was originally built with the purpose of better understanding noisy data and soon became a solution to the classical statistical dilemma. Statistical estimation always faces the choice of either imposing a specific a priori structure – which introduces arbitrary decisions in the estimation – or dealing with inherently under-determined problems – which leads to infinitely many solutions and interpretations. Under pre-imposed structures such as maximum likelihood estimation, information measures are often used as measures of the discrepancy between distributions, as goodness-of-fit measures and as other informative statistics for hypothesis testing, or for evaluating the informational content of the data. Notable examples include information criteria in econometrics, such as that of Akaike (1974) or Schwarz (1978) also known as the AIC and BIC criteria. These are, however, "objective" criteria used to skim through objective standardized econometric models. But they do not tell much about model uncertainty, nor about the uncertainty that belies the underlying quantity measured. In contrast, "maximum entropy" methods were developed as a decision tool to pick the most "informative" model, this time in the information theoretic sense, without putting too much arbitrary structure on the model. The principle of maximum entropy states that the probability distribution which best represents the current state of knowledge is the one with largest entropy. Even if information criteria and maximum entropy methods proceed from a very different logic, information is defined in a similar way through logarithms and plays a central role. For example, the Fisher "information" represents a degree of uncertainty for maximum likelihood procedures around the true value of the parameter. The Cramér-Rao bound gives an idea of the amount of incompressible risk carried by a statistical model under maximum likelihood estimation, is a direct function of the probabilistic "informativeness" of an estimator.[1] Whether one adheres to maximum entropy or maximum likelihood, distributions convey information that is tied to estimation uncertainty.

Knowing how much statistical theory used entropic measures as metrics of uncertainty, entropy was surprisingly never used to assess how much uncertainty probabilistic forecasts contained – in economics or else. And despite the clear link between probabilistic forecasts and information most forecast evaluation methods do not use entropy nor related quantities. Currently used methods include the Brier score (1950), ranked probability score (Epstein in 1969; Murphy in 1971), relative operating characteristics (Swets in 1973), or rank histograms (Anderson in 1996). Leung and North (1990) suggested that a relative entropy-type measure might be used as the basis of a skill score for deterministic forecasts. In rarer occasions entropy has been used in

[1]For an unbiased estimator, the Cramér-Rao (1946) states that:

$$R(\widehat{\theta}) \geqslant \frac{1}{I(\widehat{\theta})}; \tag{3.1}$$

where $I(\widehat{\theta})$ is the Fisher information of the estimator – i.e. the curvature of the log-likelihood around the true parameter value. A high curvature means a more efficient estimation. A result from Barron (1986) ties convergence in (relative) entropy and Fisher information in the context of the central limit theorem.

previous studies to quantify ensemble spread, for example by Stephenson and Doblas-Reyes (2000). Entropy was then suggested as a predictor of forecast skill, rather than as a measure of forecast skill. Rich and Tracy (2010) use some information theoretic measures on the survey of professional forecasters but their metrics are posited and not derived from theory. Because of that last point, their measures aren't related to one another to any way and fail in explaining the different features of commonly accepted proxies of uncertainty. Neither do they properly lay out the benefits of applying entropic measures on survey data.

This chapter aims at filling this gap and showing some possible applications in economics. With rationally inattentive agents or with *robust* decision making the concept of information entropy plays a central role but there is no measure of uncertainty *per se*. I start from forecasting theory and define optimal density forecasts to show how they relate to entropy and Shannon's theory of information. I use my results to develop a measure of uncertainty based on (survey) density forecasts, as opposed to forecast errors like it has been done in most of the existing literature. I then try to compare forecast-error type measures based on entropic quantities to the measures presented in Chapter 2 and try to uncover some principles linking economic uncertainty, forecast errors and disagreement.

3.2 Information-Theoretic Measures of Uncertainty

3.2.1 *Information, Entropy, Uncertainty and Forecasting*

Classical information science was founded by Claude Shannon with the aim of giving a mathematical structure to the concept of information transmission in the presence of noise. At the crux of information theory stands the concept of entropy which is a measure of the uncertainty of a random variable. More specifically, entropy is a measure of the amount of information required on the average to describe the random variable.[2] Little after Shannon's breakthroughs, Jaynes (1970) even went on to claim that the traditional mechanical view of economic systems might be inaccurate; he favors the thermodynamical view of the economy and with it the entropy methods that thermodynamics employs.

Without veering into heterodox economics, it can be shown that entropy and information actually lurk in the typical economic (density) forecasting problem. This section retains a framework close to that of Elliott and Timmermann (2016), but the problem is studied in much greater detail. Assume that agents want to figure out the odds of the states of nature in the next period, with their ultimate goal being to provide a predictive density over the states of nature. Nature itself is unknown to the forecaster, such that the probability density function of the data generating process is itself picked from a set of densities $\{f_y : y \in Y\}$ which is equipped with a σ-algebra and with a probability measure P_Y for a well-defined statistical model.[3] Forecasters seek to minimize the forecast error given a loss function $\mathcal{L}$:

$$\begin{aligned} R(f,\theta) &= E_{Y,Z}[\mathcal{L}(f(Z,\theta),Y,Z)] \\ &= \int_z \int_y \mathcal{L}(f(Z,\theta),y,z)\,\mathrm{d}P_{Y|z,\theta}(y)\,\mathrm{d}P_{Z|\theta}(z). \end{aligned} \tag{3.2}$$

In the case of density forecasts, a common choice for the loss function is

[2] This is a consequence of what is known as the asymptotic equipartition property that was described by Shannon in his original 1948 paper. Loosely, the theorem establishes that $nH(X)$ bits suffice on average to describe n independent and identically distributed random variables, or in other words, that a given distribution can be properly understood with at least $nH(X)$ bits of data.

[3] In a discrete case, one can think of nature picking a degenerate distribution at random, where all the mass is put on a given outcome. The mixture of such distributions produces a proper distribution whose probabilities are given by the probabilities of selecting a given degenerate distribution.

the Kullback–Leibler divergence:

$$\begin{aligned}\mathcal{K}(f_i, f_y) &= E_{f_i}\left[\ln\left(\frac{f_i}{f_y}\right)\right] \\ &= \int_t f_i(t)\ln\left(\frac{f_i(t)}{f_y(t)}\right)\mathrm{d}t\end{aligned}$$

The Kullback-Leibler distance possesses almost all properties of a distance safe that of symmetry. The choice of this ordering between the two distributions comes from the idea that forecasters will seek to minimize the error according their own odds, since they are trying to guess those of nature. Simplifying the writing by omitting the parameter vector θ and other covariates Z, the forecasting problem takes the form of a minimization of expected loss and essentially becomes:

$$\min_{f_i}\int_Y\int_t f_i(t)\ln\left(\frac{f_i(t)}{f_y(t)}\right)\mathrm{d}t\,\mathrm{d}P_Y \text{ such that } \int_t f_i(t)\,\mathrm{d}t \leqslant 1. \tag{3.3}$$

This is generalized version of a simpler problem known as the minimal cross entropy problem where the target distribution is known. This problem solves into:[4]

$$\begin{aligned}f_i(t) &= \mathrm{e}^{-1+\lambda+E(\ln f_y(t))} \\ &= \mathrm{e}^{-1+\lambda-H_t(Y)};\end{aligned} \tag{3.4}$$

where λ is the Lagrangean multiplier tied to the constraint of the problem and $H_t(Y)$ is the entropy of nature at event t such that:

$$H(Y) = -\int_Y\int_t \ln f_y(t)\,\mathrm{d}t\,\mathrm{d}P_Y(y) = \int_t H_t(Y)\,\mathrm{d}t. \tag{3.5}$$

Remark that with this solution $H(f_i) = H(Y)+C$ where C is a constant, meaning that the optimal forecast will be set in such a way that matches the entropy of nature. These derivations serve to motivate the estimation exercise that I conduct later in this chapter: Calculating the entropy of

[4] The proof can be found in the appendix.

a (supposedly optimal) density forecast helps uncover the underlying uncertainty of the process that is forecast.

An important point to make is that this problem is so far very loose. In practice, forecaster i may want to put some shape constraints on their forecast, which can be summarized by the following set of moment constraints:

$$G_i = \left\{ \gamma_{i,j}(\cdot) : \int \gamma_{i,j}(t) f_i(t)\, \mathrm{d}t = \mu_j ; j \in J_i \right\}; \tag{3.6}$$

where J_i is finite. This will change the solution to the optimal forecasting problem into:

$$f_i(t) = \mathrm{e}^{-1+\lambda+\sum_{j \in J_i} \lambda_{i,j}\gamma_{i,j}(t) - H_t(Y)} \tag{3.7}$$

where the $\{\lambda_{i,j}\}$ are the Lagrangean multipliers associated with the shape constraints of forecaster i. Consequently, $H(f_i) = H(Y) + C_i$, which allows for individual differences in uncertainty. When individual uncertainties are averaged out over the whole sample, this identity provides a condition for the consensus to arrive at the exact entropy of nature: All individual moment constraints "balance out", such that $\int C_i \,\mathrm{d}S(i) \simeq 0$, where S is a probability measure over the set of forecasters $\mathcal{S}$. It can be further be established that such a choice of density maximizes entropy over the set of constraints, see for example Cover and Thomas (2006) for a treatment in the context of the maximal entropy problem. Finally, remark that the "true" density that is being forecast may be the laws of nature that impose themselves to everyone or the idiosyncratic probabilities for someone's life events tomorrow.

Murphy (1993) argued that it is possible to distinguish three different dimensions of forecast "goodness": (i) Consistency – the correspondence between forecasts and judgments; (ii) Quality – the correspondence between forecasts and observations; and (iii) Value: incremental benefits of forecasts to users. Forecast consistency is an assumption implicitly made when using expectations conditional on the information set of the forecaster.[5] Building from the forecasting problem presented in this section, information theory can help us study forecast uncertainty from the point of view of quality and value at the same time. How it can do so is the object of the next section.

3.2.2 *Information-Theoretic Measures of Uncertainty*

3.2.2.1 Aggregate Entropy, Average Entropy, Disagreement

The Kullback-Leibler divergence used in the previous section showed that optimal density forecasts achieve the minimum cross entropy in such a way that matches the entropy of nature.[6] Using entropy isn't however, a new practice in the attempt to measure uncertainty. Rich and Tracy (2010) postulate empirical measures of disagreement and uncertainty using Shannon's entropy and comparing their results to the work of Wallis (2005). They do not show how uncertainty of the average forecast relates to disagreement and the average uncertainty, which is an important question when dealing with *macroeconomic* uncertainty. In particular, uncertainty measures should be able to disentangle the effects of individual divergences in guesses and the change in information in the consensus, which is simply an "aggregate density":

$$f^{\mathrm{A}} = \int_{\mathbb{S}} f_s \,\mathrm{d}S(s); \tag{3.8}$$

[5]Say that the forecast takes the form $y^*_{t+h} = E(y_{t+h}|\mathfrak{I}_t)$; Doob's lemma in probability theory tells us that y^*_{t+h} is a random variable that is $\sigma(\mathfrak{I}_t)$-measurable.

[6]It can be shown that maximal entropy (i.e. maximal uncertainty) is reached for a uniform distribution in a discrete world (and for the Gaussian distribution in the continuous one). This is consistent with what is often called the Principle of insufficient reason, corresponding to highest "uncertainty." Keynes (1921) strongly rejected the principle of insufficient reason as he argued that probabilities need not be "numerical" and judgments of uncertainty even less so; he very much agreed Knight (1921) in that sense. Until neuroscience takes a great leap, however, the data will always be numerical and the principle of insufficient reason is a good approximation for the highest uncertainty.

where the f_s stand for individual density forecasts, and forecasters are distributed over a set $\mathcal{S}$.[7] It can be easily checked that f^A is indeed a probability density function as long as the set of forecasters is normalized to have mass 1. Consider disagreement to be the average divergence of individual forecasts to the consensus:

$$\begin{aligned}
D &\equiv \int_{\mathcal{S}} \mathcal{K}(f_s, f^A)\,\mathrm{d}S(s) && (3.9)\\
&= \int_{\mathcal{S}}\int f_s(t)\ln(f_s(t)/f^A(t))\,\mathrm{d}t\,\mathrm{d}S(s)\\
&= \int_{\mathcal{S}}\int f_s(t)\ln(f_s(t))\,\mathrm{d}t\,\mathrm{d}S(s) - \int_{\mathcal{S}}\int f_s(t)\ln(f^A(t))\,\mathrm{d}t\,\mathrm{d}S(s)\\
&= -\int_{\mathcal{S}} H(s)\,\mathrm{d}S(s) - \int \ln(f^A(t))\int_{\mathcal{S}} f_s(t)\,\mathrm{d}S(s)\,\mathrm{d}t\\
&= -\int_{\mathcal{S}} H(s)\,\mathrm{d}S(s) + H^A && (3.10)
\end{aligned}$$

where $H(s) \equiv -\int f_s(t)\ln(f_s(t))\,\mathrm{d}t$ is the individual entropy (which measures uncertainty) for forecaster s, and $H^A \equiv \int f^A(t)\ln(f^A(t))\,\mathrm{d}t$ denote aggregate entropy (uncertainty). In essence, we have just shown that, with density forecasts:

$$\text{Aggregate Uncertainty} = \text{Uncertainty} + \text{Disagreement}. \qquad (3.11)$$

Note that this decomposition echoes what we had found with the CRPS in Chapter 2. Uncertainty can be high either when forecasters disagree a lot (but they are not individually uncertain), or when they are individually uncertain and don't disagree much about it (that is, when no one knows what happens next).

While the decomposition on its own does not uncover anything that wasn't known in Chapter 2, it presents several practical advantages.

[7] In the case of binned forecasts as in the Survey of Professional Forecasters, the aggregate forecast would be:

$$p^A = \{p_1^A, \ldots, p_K^A\} \text{ where } p_k^A = \frac{1}{N}\sum_{s\in\mathcal{S}} p_{s,k} \text{ and } \sum_k p_k^A = 1.$$

First, these quantities are completely agnostic. They do not rely on any model nor econometric estimation method; they merely need survey forecasts to be estimated and simply assume their consistency with the information set available to forecasters. Moreover, these measures tie into the limited information macroeconomic literature and can be very simply implemented in a macroeconomic model in which agents form density forecasts over the possible states of the world. Finally and most importantly, the decomposition is purely based on beliefs. Unlike the CRPS or most measures of uncertainty developed until now, the uncertainty does not take the form of a forecast error.[8]. In that sense, these quantities are all *ex-ante*. An alternative version of *ex-post* information theoretic measures is presented in the following section, while a simple example of these quantities in the case of Gaussian forecasts is presented in the Appendix.

3.2.2.2 Forecast Errors, Knightian Uncertainty and Relative Entropy

A common vision of uncertainty that, in a sense, also matches Knight's view, is that of a forecast mismatch. In probabilistic language, distribution languages do not match the distribution that nature seems to be working with, and agents are as uncertain as they are far from the "true distribution." With the same Kullback-Leibler distance as we used before:

$$E_{\text{Guess}} \ln \left(\frac{f_{\text{Guess}}}{f_{\text{True}}} \right) \tag{3.12}$$

This is essentially the forecast error as in the forecasting problem solved at the beginning of this chapter, and this has been the approach retained by most of the uncertainty measurement literature so far. Let f_s denote the (observed) density forecast of forecaster s, and define the uncertainty

[8]Baker et al. (2015) is another such measure, but much more costly in terms of estimation.

as a forecast error: $U_s \equiv \mathcal{K}(f_s, f^{\text{true}})$; then:

$$\begin{aligned} U_s &= \int f_s(t) \ln\left(\frac{f_s(t)}{f^{\text{true}}(t)}\right) \mathrm{d}t \\ &= \int f_s(t) \ln\left(\frac{f_s(t)}{f^A(t)}\right) \mathrm{d}t \ldots \\ &\qquad + \int f_s(t) \ln\left(\frac{f^A(t)}{f^{\text{true}}(t)}\right) \mathrm{d}t \end{aligned} \tag{3.13}$$

Integrating over the set of forecasters to get the total uncertainty for the panel:

$$U = \int_{\mathcal{S}} \mathcal{K}\left(f_s, f^A\right) \mathrm{d}S(s) + \int f^A(t) \ln\left(\frac{f^A(t)}{f^{\text{true}}(t)}\right) \mathrm{d}t \tag{3.14}$$

This last decomposition echoes Equation (3.11): Uncertainty = Aggregate Uncertainty + Disagreement. Also, note that for the average forecast uncertainty breaks down as such:

$$\int f^A(t) \ln\left(\frac{f^A(t)}{f^{\text{true}}(t)}\right) \mathrm{d}t = \int f^A(t) \ln f^A(t)\, \mathrm{d}t - \int f^A(t) \ln f^{\text{true}}(t)\, \mathrm{d}t; \tag{3.15}$$

which says mathematically that uncertainty decomposes into an ex-ante component – the entropy of f^A – and an ex-post term – the cross entropy between f^A and f^{true}. This is the exact same finding as in Chapter 2. In a perfect world where the randomness of nature is known with decent precision, this measure can be immediately estimated from survey forecasts. Finally, remark that such a measure makes sense only in the situation where the "true" density that is predicted is the same for everyone. Adrian and coauthors' (2016) entropies follow a similar logic but assume that the true distribution is the unconditional distribution of GDP growth estimated from the data. While this is not inherently a misguided approach, it comes with a number of drawbacks.

The issue is that the randomness of nature isn't known and such estimated densities are still very much window-dependent. Add the consideration that the underlying process may be non-stationary and one may really want to stray away from any assumption that estimated densities faithfully represent the true data generating process. Fortunately, forecasting practitioners have reverted to some simplifications to circumvent that issue. Brier (1950) used 0-1 forecast outcomes, that is, the true density puts all the weight on the event that actually occurred. Hersbach (2000) turns density forecasts into binary forecasts before integrating them over the whole support of the distribution to arrive to the continuous-rank probability score.[9] Either of these options is not feasible in this case because neither is well defined for the Kullback-Leibler distance. One possible solution is to use the Jensen-Shannon divergence, which unlike the Kullback-Leibler divergence, is symmetric, always well defined and bounded. The Jensen-Shannon divergence was introduced by Amari et al. (1987) and is defined, for two densities f^1 and f^2, as follows:

$$\mathfrak{J}(f^1, f^2) = H\left(\frac{1}{2}f^1 + \frac{1}{2}f^2\right) - \frac{1}{2}H\left(f^1\right) + \frac{1}{2}H\left(f^2\right); \tag{3.16}$$

where H denotes the Shannon entropy attached to the respective distributions. It can readily be seen that this divergence is symmetric – a property that had been forgone with the Kullback-Leibler divergence –, continuous and always properly defined, even for binary distributions. The assumption made here is that the "true" distribution of nature puts all the mass on one particular event, similar to that of Hersbach (2000).[10] With such definitions, the following decomposition holds:

$$\begin{aligned}\int \mathfrak{J}(f_s, f^{\text{true}})\, \mathrm{d}S(s) + \frac{1}{2}\int H(s)\, \mathrm{d}S(s) = \mathfrak{J}(f^A, f^{\text{true}}) + \frac{1}{2}H^A \ldots \\ + \int \left[H\left(\frac{1}{2}f_s + \frac{1}{2}f^{\text{true}}\right) - H\left(\frac{1}{2}f^A + \frac{1}{2}f^{\text{true}}\right)\right] \mathrm{d}S(s)\end{aligned} \tag{3.17}$$

[9] This is what is done in Chapter 2.

[10] The entropy of a distribution that puts all weight on one outcome is zero.

which, once the following substitution is made:

$$\frac{1}{2}H^A - \frac{1}{2}\int H(s)\,\mathrm{d}S(s) = \frac{1}{2}\int \mathcal{K}(f_s, f^A)\,\mathrm{d}S(s); \tag{3.18}$$

results in the following:

$$\int \mathcal{J}(f_s, f^{\text{true}})\,\mathrm{d}S(s) \simeq \mathcal{J}(f^A, f^{\text{true}}) + \frac{1}{2}\int \mathcal{K}(f_s, f^A)\,\mathrm{d}S(s) \tag{3.19}$$

All in all, that is tantamount to saying the same as before:

$$\text{Uncertainty} = \text{Aggregate Uncertainty} + \text{Disagreement}.$$

The last "noise" term (the difference in entropy) is bounded as follows:[11]

$$\left|\int \left[H\left(\frac{1}{2}f_s + \frac{1}{2}f^{\text{true}}\right) - H\left(\frac{1}{2}f^A + \frac{1}{2}f^{\text{true}}\right)\right]\,\mathrm{d}S(s)\right| \\ \leqslant -\frac{1}{2}E_S\left(\|f_s - f^A\|_1 \ln \frac{1}{2\kappa}\|f_s - f^A\|_1\right) \tag{3.20}$$

where κ is defined as:

$$\kappa \equiv \mathrm{e}^{\frac{1}{2}(1+\ln 2\pi\overline{\sigma}^2)};\ \text{for some } \overline{\sigma}^2 \geqslant \sup_{s\in\mathcal{S}}\int t^2\left(\frac{|f_s(t) - f^A(t)|}{\|f_s - f^A\|_1}\right)\mathrm{d}t \tag{3.21}$$

Because it is independent of the "true" density and involves only a comparison of the individual forecast against the aggregate, this bound justifies considering the "difference in entropy" term a "noise" of disagreement; it rewrites as a direct expected value of (a function of) the L^1 distance between the individual and consensus forecasts.[12] In particular, as $E_S\|f_s - f^A\|_1 \to 0$, that is, all individuals converge to a

[11] Under some regularity conditions detailed in the appendix together with the proof.

[12] Alternatively, one could also consider that the term in question is "the other half" of disagreement that is missing in the equation, since the average of Kullback-Leibler is affected with a factor of exactly 0.5.

consensus, the disagreement "noise" term disappears as well.[13] While the Jensen-Shannon doesn't have the same theoretical foundation and is not immediately apparent from the forecasting problem presented in the first section of this chapter,[14] the decomposition still mirrors the properties of aggregate uncertainty that were obtained in other approaches. More importantly, the Jensen-Shannon approach allows to extend information-theoretic measures to a feasible forecast-error based approach.

3.2.2.3 State-Dependent Uncertainty and Total Uncertainty

An important question when trying to capture the underlying uncertainty of density forecast is that of the extent to which it captures moments of the distribution. I have shown that the solution to the forecasting problem indeed factors in moment conditions through their associated multipliers. Large deviation theorists Donsker and Varadhan linked the relative entropy to (cumulant) moment generating functionals, but the intuition of how the overall entropy is linked to the conditional distributions implied by the division of the support. This sections presents a simple decomposition to break down the entropy of a distribution over its support.

To lighten the notation I will assume that the density forecast made since the beginning of this chapter represents the density of a random variable of interest X with density f with respect to the Lebesgue measure λ over $\mathbb{R}$. The forecast is made conditional to the information

[13] Jensen's inequality applied to $t \mapsto -t \ln t$ (which is concave) allows to rewrite the bound as:

$$\left| \int \left[H\left(\frac{1}{2} f_s + \frac{1}{2} f^{\text{true}}\right) - H\left(\frac{1}{2} f^A + \frac{1}{2} f^{\text{true}}\right) \right] \mathrm{d}S(s) \right| \\ \leqslant -\frac{1}{2} E_S\left(\|f_s - f^A\|_1\right) \ln \frac{1}{2\kappa} E_S\left(\|f_s - f^A\|_1\right) \to 0;$$

as $E_S \|f_s - f^A\|_1 \to 0$ – forecasters all converge to the same consensus.

[14] It nevertheless easy to check that the minimum "loss" is attained at the true density.

available at the beginning of the period, which I omit from the notation for simplicity. I show in the Appendix of this chapter that:

$$H(X) = P(X \in A)H(X|X \in A) + [1 - P(X \in A)]H(X|X \notin A)\dots \\ + \psi[P(X \in A)], \tag{3.22}$$

where $\psi[P(X \in A]$ denotes the binary entropy of the event $(X \in A) = \{\omega \in \Omega; X(\omega) \in A\}$. In common terms, the uncertainty of a given variable of interest can be decomposed into a weighted sum of the uncertainty of the variable in a given state and outside of it, to which the uncertainty around that state is added. One advantage of such a decomposition is to emphasize asymmetries in different areas (or "states" of the distribution forecast being made. A simple example will make this idea clearer. Imagine that GDP growth X can take 4 different values – x_1 and x_2 denote recessionary states, x_3 and x_4 denote expansionary states. Consider Situation 1 where the forecast predicts $P_X(x_1) = 0.2$; $P_X(x_2) = 0.2$; $P_X(x_3) = 0.3$ and $P_X(x_4) = 0.3$; and Situation 2 where $P_X(x_1) = 0.2$; $P_X(x_2) = 0.3$; $P_X(x_3) = 0.3$ and $P_X(x_4) = 0.2$. Both distribution forecasts have the same entropy since one is simply a rearrangement of the other. Note, however, that in the second case there is more uncertainty in either state (recession or expansion) than in the first case. Looking solely at entropy would not have been able to uncover this fact. More generally, as long as all quantities have been properly normalized[15], this decomposition helps understand where uncertainty changed along the distribution and where it did not. To be sure, such a decomposition makes sense only when the underlying density forecast is not perfectly symmetric around the states, say for example, with the normal distribution around its mean. It allows me to escape some simplifying assumption made in Chapter 2.

[15]See the estimation sections below for a proper explanation of the normalization of entropy when working in a discrete setting.

Finally, note that for a joint process (X,Y) the following is true:

$$\begin{aligned} H(X,Y) = P(Y \in A)H(X,Y|Y \in A) + [1 - P(Y \in A)]H(X,Y|Y \notin A) \\ + \psi[P(Y \in A)], \end{aligned} \tag{3.23}$$

which could be useful in looking at the regimes of growth in different states of, say, inflation. Obviously this can be estimated provided that one has joint predictive densities to feed in. While I restrict my estimation to the marginal densities provided in the Survey of Professional Forecasters, some ways of generating joint predictive densities from this survey exist, see for instance Odendahl (2017).

This decomposition is a generalization of what is commonly known as the "recursiveness" of the Shannon differential entropy. Such a decomposition has not been exposed nor used in the entropy literature and much less so in the economic uncertainty literature. The closest concept to this decomposition would be Adrian and coauthors' (2016) "downward" and "upward" entropy concepts, but they use estimated densities and focus about the distance between conditional and unconditional densities under and above the median. For the sake of comparison, let m denote the median of the distribution of X, that is, $P_X\{(-\infty, m)\} = 0.5$. My decomposition would read as follows:

$$H(X) = \frac{1}{2}H(X|X > m) + \frac{1}{2}H(X|X \leqslant m) + \ln 2. \tag{3.24}$$

In fact, the upside and downside entropies are related to the total (Kullback-Leibler) divergence as follows:

$$\mathcal{K}(\widehat{f}_{y_t|x_{t+h}}, \widehat{g}_{y_t}) = \frac{1}{2}\mathcal{L}_t^U + \frac{1}{2}\mathcal{L}_t^D + \ln 2, \tag{3.25}$$

although Adrian et al.' (2016) upward and downward measures are not linked to one another as such in the current version of the paper. For a density forecast conditional on the information available in the previous period, this decomposition highlights how total uncertainty related to uncertainty under and above the median. A direct application to survey forecasts will be conducted later in this chapter for introduction. Finally, note that the breakdown of entropy can be extended to any quantile, or more generally, to any partition of the support of X. In the case of growth, this means that distribution forecasts can be studied in and out of recessions, as opposed to simple quantiles.

3.3 ESTIMATING ENTROPY-BASED UNCERTAINTY MEASURES

3.3.1 *Job Market Uncertainty in the Survey of Economic Expectations*

I begin this illustrative section using data that is not purely macroeconomic nor comes from professional forecasters, and more importantly, with which typical mean-squared error based measures of uncertainty would be impossible to estimate. The Survey of Economic Expectations (SEE) was run between 1992 and 2002 (over 12 waves) by the University of Wisconsin Survey Center (UWSC). The goal of the SEE was to elicit probabilistic expectations of significant personal events, such as personal security, unemployment, insurance and income. While the methodology can be applied to all questions in the survey, I choose to focus on employment outcomes. Participants were asked the following questions:

- On job loss (All waves): "I would like you to think about your employment prospects over the next 12 months. What do you think is the percent change that you will lose your job during the next 12 months?"

- On finding as good a job (All waves): "If you were to lose your job during the next 12 months... What do you think is the percent chance (or chances out of 100) that the job you eventually find and accept would be at least as good as your current job, in terms of wages and benefits?

– On leaving job voluntarily (All waves): “What do you think is the percent change that you will leave your job voluntarily during the next 12 months?”

Note that these questions constitute a complete probability distribution over the following events: Losing one’s job and not finding as good a job in the following 12 months, losing one’s job and finding as good a job in the following 12 months, leaving one’s job voluntarily, and leaving things as they are.[16] However, the outcomes are individual and because the support is discrete only the “ex-ante” entropy measures presented in Section 3.2.2.1 are really relevant in this case.[17]

The entropy measure is estimated directly on the probabilities given in the survey. The entropy is normalized by $\ln(n)$, where n is the number of events considered, to make it a index bounded between zero and one.[18] The estimated job market uncertainty is presented in Figure 3.1. Uncertainty remained quite low over the 1992-2002 decade and seems to vary little except from a clear ”down then up” period towards the end of the sample period. In the latter, uncertainty remained stable while aggregate uncertainty went up together with disagreement. These data provide support to the idea that disagreement is a non-negligible part of uncertainty. The continuous line in the graph depicts the evolution of the importance of disagreement in overall uncertainty. In the case of job market uncertainty, disagreement represents between 20 and 25 percent of overall uncertainty. In Chapter 2, I had found that disagreement was a very small fraction of overall uncertainty; the finding that disagreement is high in the case of job market uncertainty begs

[16] The conditional probabilities are inferred from the survey data. Note that in a minority of cases people give a set of probabilities that sum to slightly over 1. I exclude these cases from the estimation.

[17] The state dependent entropy decomposition gives better insights on various quantiles of a continuous distribution than on a discrete distribution with only four outcomes.

[18] This is for consistency with the following sections where the number of bins changes and needs to be accounted for in the normalization.

the question of what drives disagreement in the Survey of Economic Expectations. A possible explanation for disagreement being such an important part of uncertainty is that agents are forecasting "individual" events, in the sense that individuals in the survey may have different underlying conditions that will shift their beliefs toward one direction or another. In what follows, I confirm that the level of studies may be an explanation for the uneven levels of uncertainty in the survey.

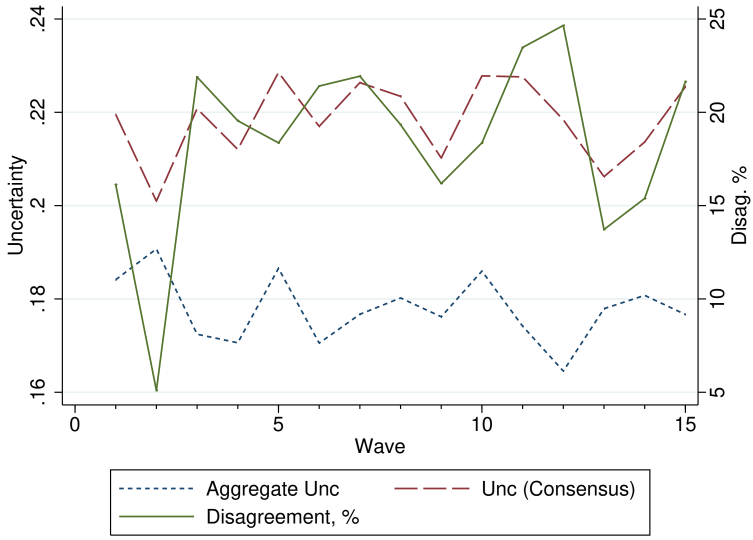

FIGURE 3.1: Job Market Uncertainty, Aggregate Uncertainty and Disagreement Estimated in the Survey of Economic Expectations

One advantage of survey data like the SEE is that they provide a number of socio-demographic characteristics of individuals together with the individuals' stated beliefs. An obvious question that is subject of a vivid debate nowadays is that of gender inequality on labor market outcomes and the resulting uncertainty. In the Survey of Economic Expectations, as far as job market uncertainty goes, women do not appear to face starkly higher uncertainty levels than men. A simple

t-test for difference in means yields a p-value of approximately 0.5. The data is pooled across different periods of time; the low variation in uncertainty over the sample period leads to believe that there isn't much time effect to account for.[19] This statement holds true when the level of education is accounted for – in particular, when jobs mostly occupied by women (for instance, nursing) are excluded. Therefore, the data do not seem to point at a significant difference between the two sexes in regards to job market uncertainty.

Finally, I consider the impact of education on the level of uncertainty. A sensible prior would be that higher levels of education confer more stable job outcomes and better prospects on the job market overall. To address this question, I run an analysis of variance (ANOVA) across the different education categories in the survey. The results barely meet the requirements for statistical significance with a p-value stopping dead on the 0.1 threshold, but given how simple the procedure is, they do indicate a trend. A box plot summarizing the differences across groups is presented in Figure 3.2. Higher levels of education seem to experience consistently lower levels of uncertainty over the sample period. While it may seem surprising that nurses have the lowest level of average uncertainty, it is important to remember that the United States has been facing a shortage of nurses for decades (Buchan and Aiken, 2008) and population ageing only contributes to make the situation more severe. Therefore, nurses are likely to enjoy a higher "bargaining" power on the job market and consequently lower labor uncertainty.

Beyond a skin-deep discussion on the determinant of uncertainty faced by survey respondents, the data point to some interesting facts. First, aside from signaling purposes, educational choices translate into a different set of beliefs later in life, not simply to initial outcomes after education as some claim. To be sure, in a signaling framework, agents choose education, reach a certain job; yet little is said about the

[19] As a robustness check, I ran the same exercises – this and the former – wave by wave and found similar results.

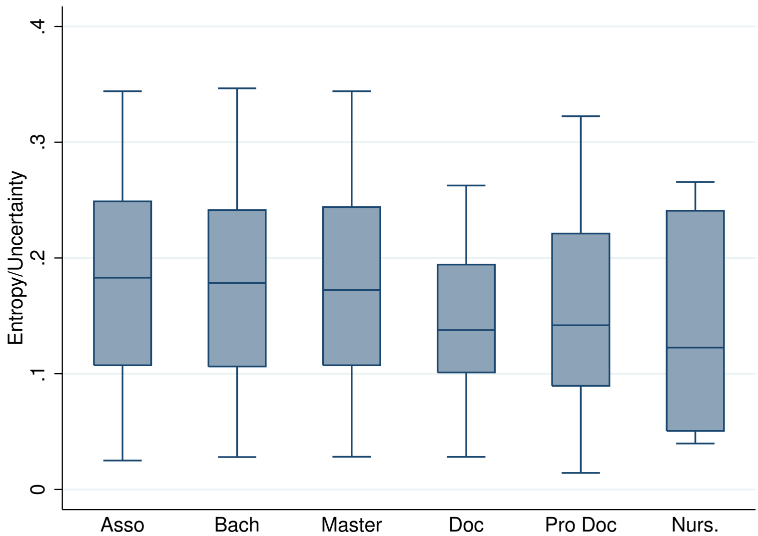

FIGURE 3.2: Mean Uncertainty by Education

sequential job searching uncertainty. The data of the SEE seems to show that the initial signaling choice also matters further down the road. Similarly, there is an argument to be made about agents internalizing their more or less volatile job market outcomes and formulating beliefs that agree with it. Finally, this fact gives credit to the idea that different social categories may need different levels of labor insurance for other reasons than the sheer difference in wage.

Provided that the data were made available to the public, such uncertainty measures could be estimated on more recent periods. For instance, the Michigan Survey of Consumer does ask some probabilistic forecast survey questions, but the data is not available in disaggregated form to the public and the questions are still binary for the most part. Similarly, the Manpower Employment Outlook Survey (MEOS)[20] could serve as a great dataset to estimate firm level job market uncertainty

[20]See https://www.manpowergroup.us/meos/.

and compare it to household job market uncertainty, but once again the data are not available to the general public. In particular, this data could help estimate firm-level job market uncertainty at two levels of disaggregation, following the decomposition for subgroups presented in the Appendix.

Entropy based measures go against the idea that uncertainty cannot be directly observed, but needs to be indirectly inferred from decisions that are influenced by uncertainty. The latter approach is defended by Ernst and Viegelahn (2014), and while that approach is not inherently wrong, it is more likely to result in an estimate of *ex-post* uncertainty, i.e. risk, than ex-ante Knightian uncertainty. In the context of labor market uncertainty, the determinants of economic decisions are likely to be a covariate of firm-level business uncertainty, which is only one dimension of job market uncertainty and imperfectly reflects the uncertainty that consumers face when making their decisions. Furthermore, entropy based measures are "agnostic" and do not rely on an assumed reduced form macroeconomic model to be estimated.

3.3.2 *Macroeconomic Uncertainty in the Survey of Professional Forecasters*

In this section, I estimate the information theoretic uncertainty measures I developed over the data of the Survey of Professional Forecasters – the same that I use in Chapter 2. The entropy measures are normalized to one by dividing the total entropy by $\ln n$ where n denotes the number of bins. This is to ensure consistency in the metric displayed, because the SPF changed the number of different scenario for growth outcomes a number of times over the sample period.[21] For the sake of comparison, I retain the same sample period as before.

[21] To be sure, if there are n outcomes, the maximum unnormalized entropy is $\ln n$ which is an increasing function of the number of bins. This was originally a requirement of Claude Shannon for his uncertainty measure, but in this case a finer grid need not mean a higher level of uncertainty.

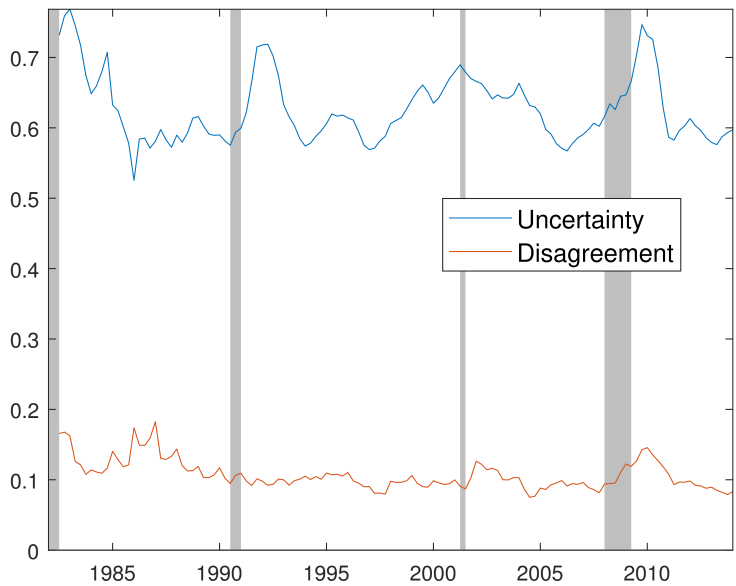

FIGURE 3.3: Uncertainty and Disagreement in the Survey of Professional Forecasters

Some very similar observations can readily be made: Uncertainty is definitely correlated with business cycles in that it increases after recessions. Disagreement is once again a small fraction of overall uncertainty – that of the aggregate forecast, in this case. In light of the previous example with job market outcomes, a possible explanation for this fact is the strong convergence in forecasts of practitioners. Indeed, it isn't unreasonable to expect that professional forecasters, who have studied similar courses and have been taught which models work and which do not, will provide similar forecasts. (This idea that professional forecasters may use a similar reduced set of models will give support to the measure of forecasting uncertainty that is presented in Chapter 4.) Finally, as claimed throughout this chapter, the entropy-based measure is very much like the ex-ante measure developed in the previous chapter. Figure 3.4 depicts the entropic uncertainty measure against the ex-ante and ex-post measures of uncertainty estimated in Chapter 2, and it can be seen that ex-ante uncertainty is almost a mirror of entropy in this case. Note that the apparent divergence between the ex-ante measure from Chapter 2 around 1992 has a simple "mechanical" explanation. In Chapter 2, the predictive densities are fitted to Gaussian densities. In 1992, the Survey of Professional Forecasters introduced a finer grid for the questionnaire, such that forecasters were now able to provide more "concentrated" forecasts than before – over fewer bins, the variance parameter of the fitted Gaussian density will likely appear larger than over a finer grid. Because the ex-ante measure of Chapter 2 is a direct function of the fitted variance, it sharply decreases at the moment where the questionnaire was changed, while the (normalized) entropy is robust to such (arbitrary) changes.

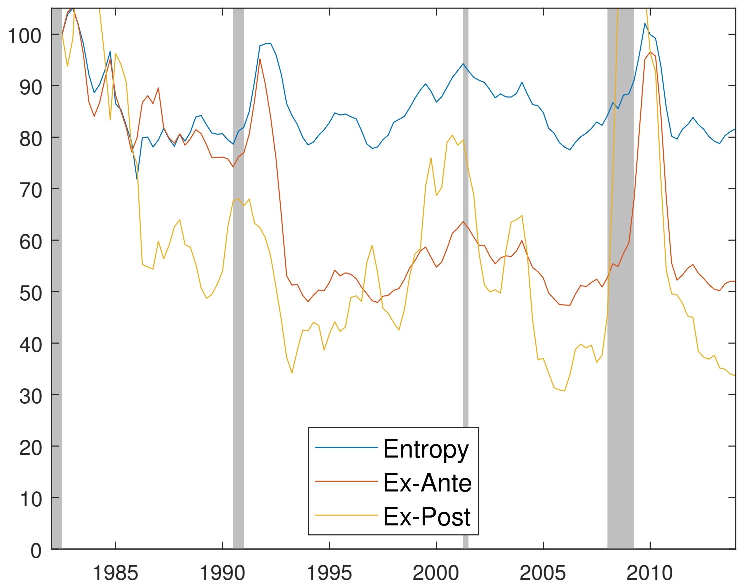

Figure 3.4: Entropic Uncertainty v. CRPS Based Uncertainty Measures

I now proceed to estimate the "forecast error" version of my information theoretic measures. Given the format of SPF forecasts, the outcome vector is a vector of 0s and a 1, where the 1 falls in the bin corresponding to the realized growth number. The estimated measure is presented in Figure 3.5. Some clear similarities with the uncertainty measures presented in Chapter 2 stand out. First, as it was to be expected, the introduction of a forecast error component in the uncertainty measures renders their behavior much more akin to that of ex-post uncertainty. Interestingly, the inclusion of the forecast outcome also introduced more variability in the uncertainty indices, even more so than in the ex-post and ex-ante CRPS based measures. This likely comes from the definition of the "distribution outcome" with the Jensen-Shannon distance. Here, the outcome is a vector of 0s and one 1 in the bin corresponding to the growth outcome for the period. In the measures presented in Chapter 2, the outcome was made "binary" through the setting of a threshold and compared to a binary forecast built from the complete forecast density. The squared error was then integrated over the real line to allow the threshold to cover the complete support of the distribution. To the extent that the outcome is harder to guess – since after all, there are more possibilities – it is natural to expect the forecast error to be greater. Finally, the two measures display a very similar finding than in the previous chapter in regards to the importance of disagreement in overall uncertainty. The disagreement seems to contribute extremely little to the overall variation of uncertainty, as the difference between the two lines of $\int \mathcal{J}(f_s, f^{\text{true}})\, \mathrm{d}S(s)$ and $\mathcal{J}(f^A, f^{\text{true}})$ is hardly noticeable. In this Jensen-Shannon decomposition, disagreement (including the noise term) represents a mere 4 percent of uncertainty on average. This hints at the more general idea that when considering uncertainty as a forecast error, disagreement becomes secondary. This comes in sharp contrast to the purely ex-ante aspect of uncertainty, where disagreement was a non-negligible part.

Finally, the conditional entropy decomposition can be estimated on the average forecast in the SPF over time. I consider the two

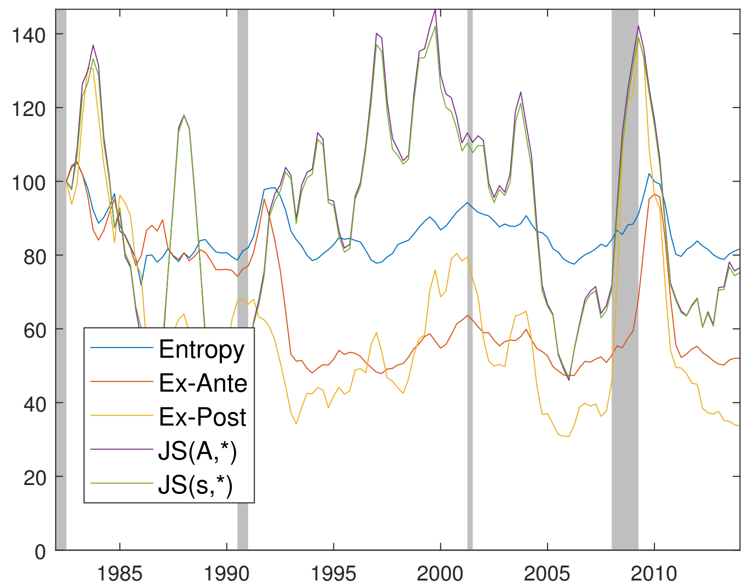

Figure 3.5: Entropic Uncertainty, CRPS Measures and Jensen-Shannon Uncertainty

states of interest to be recessions and expansions, that is, at each point in time, I compute the (conditional) uncertainty in the recessionary and expansionary scenarios for the average forecast.[22] Because the decomposition presented in Equation 3.22 is not normalized and I want to compare entropies on a level ground, I adopt the following normalized version of the decomposition:

$$H = p_{\mathrm{E}}H(E) + p_R H(R) + h(p_E, p_R), \tag{3.26}$$

where p_E (p_R) denotes the normalized probability of being in an expansion (recession), $H(E)$ ($H(R)$) denotes the normalized conditional

[22]It does not really make sense to average these decompositions calculated over individual forecasts unless the weights are kept the same for all forecasters, which is the case when the analysis is restricted to quantiles. In my case, the two different scenarios are expansions and recessions because they mark the movements of business cycles.

entropy in an expansion (recession), and $h(p_E, p_R)$ denotes the scaled entropy of the business cycle movements considered. The normalized probabilities are such that:

$$p_E \frac{\ln n}{\ln n_E} + p_R \frac{\ln n}{\ln n_R} = 1; \tag{3.27}$$

where n_E and n_R denote the number of bins in expansion and recession in the SPF forecast such that $n = n_E + n_R$. The entropies are computed with actual conditional probabilities and properly normalized to fall between 0 and 1. The scaled entropy of business cycle movements is the normalized entropy of the binary distribution of recessions and expansions multiplied by $\ln(2) \cdot \ln(n)^{-1}$. Note that the weights do not actually matter as they are very much arbitrary, but this is merely to keep track of the proper breakdown of uncertainty. The estimated uncertainties are presented in Figure 3.6.

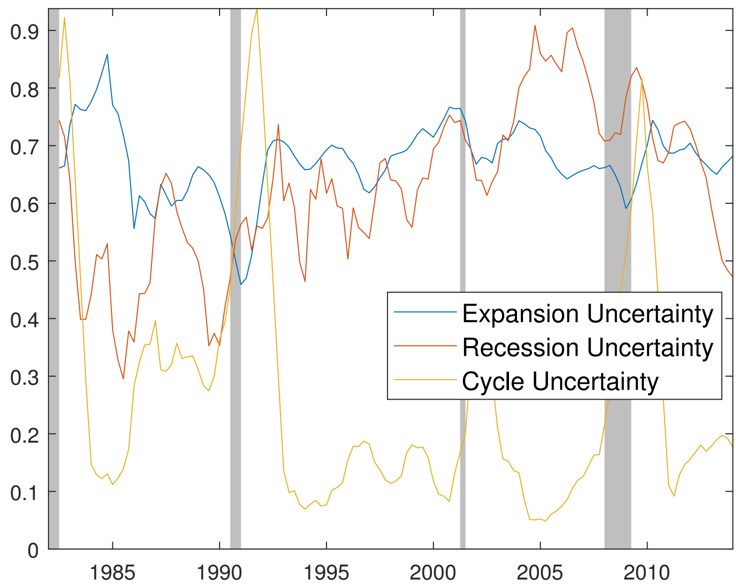

FIGURE 3.6: Conditional Uncertainties and Business Cycle Uncertainty

A few interesting facts stand out from the evolution of conditional entropies. First, business cycle uncertainty is a clear correlate of ex-ante uncertainty it was presented in the earlier part of this chapter and in Chaper 2. This hints at the fact that using probabilities of recessions and expansions, such as those provided by the Federal Reserve Board (see Chauvet, 1998) can provide a very "quick" indicator of macroeconomic uncertainty at very little computational cost. Second, it is clear that within either regime, uncertainty is quite high. Forecasters in the SPF do not provide extremely "clear cut" forecasts, even with the business cycle uncertainty is apparently at its lowest (see, for instance, after the dotcom bubble crash). Finally, uncertainties in either state show very little correlation. They do seem to cross and diverge in certain periods, but there is no pattern to be discerned with business cycles. Nevertheless, both are objects of interest because they exhibit movements that would not be visible by looking solely at the entropy of the whole forecast. In the first half of the 1980s, expansion uncertainty went very high, indicating that forecasters could not foresee the overheating that the Fed was trying to prevent. Similarly after the dotcom crash, business cycle uncertainty went down while recession uncertainty went close to its maximum, showing little visibility of forecasters in the negative scenarios. In short, the measures presented in this section bring some light on new uncertainty facts that existing measures of uncertainty overlook.

3.4 Conclusion

Because of its breakthrough in understanding noisy data, information theory and entropy methods had a great influence in theoretical economics and statistics. But so far little has been done to take these methods to the data in the actual measurement of economic uncertainty. This is precisely the void that this chapter aimed at filling.

Starting from an optimal forecasting problem, I have shown how and why entropy appears to be an adequate proxy of economic uncertainty.

From there, the decompositions presented either served to parallel what the mean-squared approach of economic uncertainty would consider as measures, or to extend them to situations that cannot be studied in the traditional framework. The consistency that was observed in this new approach gave credit to its validity, and at the same time shed light on the important distinction between ex-ante and ex-post uncertainty.

The choice between mean-squared errors and entropy to measure economic uncertainty seems to have limited consequences on the conclusions when ex-post proxies of uncertainty are considered. There is one situation, however, where entropy wins without a fight: when the outcomes – that we'd use in a mean-squared forecast error – are not available, which is often the case with microeconomic data. This chapter presented several options that can be used in this case, all of them employing entropy and proceeding from the same theoretical motivation. In that sense, information theory allows the estimation of Knightian *ex-ante* uncertainty in situations that were not accessible to the usual approach.

In the forecasting problem that we started from, an important assumption is that computing odds is a straightforward task. But in practice forecasters may be hard pressed to know which data generating process they should fit their data to. The difficulty in the choice of the proper statistical model represents a last layer of uncertainty that Chapter 4 explores.

3.5 APPENDIX

In this section, all densities and probability measures involved are assumed to be well defined on a probability space $(\Omega, \mathcal{A}, P)$ or on the real space equipped with the Borelian σ-algebra $(\mathbb{R}^m, \mathcal{B}(\mathbb{R}^m), P.)$, as needed.

A Simple Gaussian Forecasting Problem

This section presents a more concrete setting where forecasters form Gaussian density forecasts following based on parameter priors.[23] Here, the forecaster guesses the mean and variance to try to form a Gaussian density forecast.[24] Forecasters have different priors; each forecaster i has priors:[25]

$$\pi_i(\mu|\sigma^2) = \mathcal{N}(\overline{\mu}_i, \sigma^2) \text{ and } \pi_i(\sigma^2) = \mathrm{IG}(\ell_i, s_i);$$

where IG denotes the inverse Gamma distribution. Given those priors, the forecaster will choose (μ_i, σ_i^2) to solve:

$$\min \int_\mu \int_{\sigma^2} \mathcal{K}(f_i, f_y) \pi_i(\mu, \sigma^2) \, \mathrm{d}\mu \, \mathrm{d}\sigma^2. \tag{3.28}$$

Using the fact that the Kullback-Leibler divergence between two Gaussian densities is given by:

$$\mathcal{K}(f_i, f_y) = \ln \frac{\sigma^2}{\sigma_i^2} + \frac{\sigma_i^2 + (\mu_i - \mu)^2}{2\sigma^2} - \frac{1}{2};$$

[23]In the context of a classic linear regression where $y = x\beta + \varepsilon$, making a density forecast of the dependent variable given data x is tantamount to estimating the coefficients β and the variance of the error term σ^2.

[24]Under non-multicollinearity in a linear regression this is equivalent to guessing β. Here, the "priors" already represent the state of knowledge of the forecaster at that point in time.

[25]The distributions have been chosen to provide a parallel to the Bayesian regression literature.

the forecaster will choose the parameters to solve:

$$\min\left\{\int_\mu\int_{\sigma^2}\ln\frac{\sigma^2}{\sigma_i^2}\pi_i(\mu,\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2\right.$$
$$\left.+\int_\mu\int_{\sigma^2}\frac{\sigma_i^2+(\mu_i-\mu)^2}{2\sigma^2}\pi_i(\mu,\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2-\frac{1}{2}\right\}$$

Thanks to the assumptions on the priors and the Bayesian identity $\pi_i(\mu,\sigma^2)\propto\pi_i(\mu|\sigma^2)\pi_i(\sigma^2)$:

$$\begin{aligned}\int_\mu\int_{\sigma^2}\ln\frac{\sigma^2}{\sigma_i^2}\pi_i(\mu,\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2 &= \int_\mu\int_{\sigma^2}\ln\sigma^2\pi_i(\mu|\sigma^2)\pi_i(\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2\\ &\quad-\int_\mu\int_{\sigma^2}\ln\sigma_i^2\pi_i(\mu|\sigma^2)\pi_i(\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2\\ &= -\ln\sigma_i^2+\int_\mu\int_{\sigma^2}\ln\sigma^2\pi_i(\mu|\sigma^2)\pi_i(\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2\end{aligned}$$

The second integral is $E(\ln Z)$, and under the Inverse Gamma distribution this expected value is best found using the moment generating function of the inverse gamma distribution itself since $E(\mathrm{e}^{t\ln Z})=E(Z^t)$. It can be easily shown that:

$$E(Z^t)=\frac{(s_i)^t\Gamma(\ell_i-t)}{\Gamma(\ell_i)},$$

where Γ denotes Euler's Gamma function. Furthermore, the derivative at $t=0$ of the moment generating function is:

$$E(\ln Z)=\ln s_i-\frac{\Gamma'(\ell_i)}{\Gamma(\ell_i)}.$$

(The ratio in the identity above is also known as the digamma function.) Therefore:

$$\iint\ln\frac{\sigma_i^2}{\sigma^2}\pi_i(\mu,\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2=-\ln\sigma_i^2+\ln s_i-\frac{\Gamma'(\ell_i)}{\Gamma(\ell_i)}. \tag{3.29}$$

I now proceed to the second part of the objective function:

$$\begin{aligned}
\iint \frac{\sigma_i^2 + (\mu_i - \mu)^2}{2\sigma^2}\pi_i(\mu,\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2 &= \sigma_i^2 \iint \frac{1}{2\sigma^2}\pi_i(\mu,\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2 \\
&\quad + \iint \frac{(\mu_i-\mu)^2}{2\sigma^2}\pi_i(\mu,\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2 \\
&= \frac{\sigma_i^2}{2}E(Z^{-1}) \\
&\quad + \int \pi_i(\sigma^2)\frac{1}{2\sigma^2}\int (\mu_i-\mu)^2\pi_i(\mu|\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2 \\
&= \frac{\sigma_i^2}{2}E(Z^{-1}) + \frac{1}{2}\int \pi_i(\sigma^2)\frac{1}{\sigma^2}[(\mu_i-\overline{\mu}_i)^2+\sigma^2]\,\mathrm{d}\sigma^2 \\
&= \frac{\sigma_i^2}{2}E(Z^{-1}) + \frac{1}{2}\left((\mu_i-\overline{\mu})^2E(Z^{-1})+1\right).
\end{aligned}$$

For the Inverse Gamma distribution, the expected value can be computed directly $E(Z^{-1}) = \ell_i/s_i$. Substituting yields:

$$\int\int \frac{\sigma_i^2 + (\mu_i-\mu)^2}{2\sigma^2}\pi_i(\mu,\sigma^2)\,\mathrm{d}\mu\,\mathrm{d}\sigma^2 = \frac{\sigma_i^2}{2}\frac{\ell_i}{s_i} + \frac{1}{2}\left((\mu_i-\overline{\mu})^2\frac{\ell_i}{s_i}+1\right). \tag{3.30}$$

Putting it all together, the forecasters solve with respect to (μ_i, σ_i^2):

$$\min\left\{-\ln\sigma_i^2 + \ln s_i - \frac{\Gamma'(\ell_i)}{\Gamma(\ell_i)} + \frac{\sigma_i^2}{2}\frac{\ell_i}{s_i} + \frac{1}{2}(\mu_i-\overline{\mu}_i)^2\frac{\ell_i}{s_i}\right\} \tag{3.31}$$

The objective function is convex in (μ_i, σ_i^2) and the first order conditions for this problem make the forecaster set:

$$\sigma_i^{2*} = 2\frac{s_i}{\ell_i} \tag{3.32}$$

$$\mu_i^* = \overline{\mu}_i. \tag{3.33}$$

Remark that the variance is set at exactly twice the harmonic mean of the mean and the mode of the prior. Unless the location parameter of the distribution is known, however, the parameters are not identified. Note that in the most likely case where forecasters will have a similar location prior for the mean (since they all observe the same long term history of data), there will be no disagreement on the mean of the density forecast. Disagreement will be a function of the "ex-ante" uncertainty around the variance – the scale parameter of the Inverse Gamma distribution. This example supports the idea that disagreement is an imperfect proxy for uncertainty but should not be considered as a quantity completely orthogonal to it. In either case, the average entropy (Aggregate Uncertainty) is given by:

$$\int H(s)\,\mathrm{d}S(s) = \frac{1}{2} + \ln(\sqrt{2\pi}) + \frac{1}{S}\sum_{i=1}^{S} \ln\left(\sqrt{\frac{2s_i}{\ell_i}}\right); \tag{3.34}$$

which is an increasing function of the individual variances. There is no closed form for entropies of Gaussian mixtures[26] but some approximations exist, see for instance Huber et al. (2008).

Solution of the General Forecasting Problem

The forecasting problem is the following:

$$\min_{f_i} R = \int_Y \int_t f_i(t) \ln\left(\frac{f_i(t)}{f_y(t)}\right) \mathrm{d}t\,\mathrm{d}P_Y \text{ such that } \int_t f_i(t)\,\mathrm{d}t \leqslant 1. \tag{3.35}$$

The Kullback-Leibler being weakly convex, the problem is well defined. The Lagrangean of the problem is given by:

$$L(\{f_i\},\lambda) = \int_Y \int_t f_i(t) \ln\left(\frac{f_i(t)}{f_y(t)}\right) \mathrm{d}t\,\mathrm{d}P_Y - \lambda\left(\int_t f_i(t)\,\mathrm{d}t - 1\right). \tag{3.36}$$

The first order conditions of the problem are given by:

$$\frac{\partial L(\{f_i\},\lambda)}{\partial f_i(t)} = \int_Y \ln\left(\frac{f_i(t)}{f_y(t)}\right) \mathrm{d}P_Y + 1 - \lambda = 0; \tag{3.37}$$

[26]The average density can be seen as the density of an equal-weight mixture of Gaussians.

and the constraint. The conditions solve into:

$$f_i(t) = \mathrm{e}^{-1+\lambda+E(\ln f_y(t))}, \tag{3.38}$$

as stated.

A Bound on the Disagreement Term in the Jensen-Shannon Decomposition

Let $p \equiv \frac{1}{2}f_s + \frac{1}{2}f^{\text{true}}$ and $q \equiv \frac{1}{2}f^A + \frac{1}{2}f^{\text{true}}$. $|p-q| = \frac{1}{2}\left|f_s - f^A\right|$ and it follows that:

$$\int |p(t) - q(t)|\,\mathrm{d}t = \frac{1}{2}\int \left|f_s(t) - f^A(t)\right|\,\mathrm{d}t = \frac{1}{2}\int_{f^A \neq f_s} \left|f_s(t) - f^A(t)\right|\,\mathrm{d}t$$

Assuming that the integral term is less than 1:

$$\frac{1}{2}\|f_s - f^A\|_1 = \int |p(t) - q(t)|\,\mathrm{d}t \leqslant \frac{1}{2}.$$

Now, remark that the function $t \mapsto -t\ln t$ is concave, positive on $[0,1]$ and equals 0 at 0 and 1. Consider the chord of the function between t and $t+h$, where $h \leqslant 0.5$. The maximum absolute slope of the chord is at either end of it, meaning that for $0 \leqslant t \leqslant 1-h$:

$$|-t\ln t + (t+h)\ln(t+h)| \leqslant \max[-h\ln h, -(1-h)\ln(1-h)] = -h\ln h.$$

It follows that:

$$\begin{aligned}
|H(p) - H(q)| &= \left|\int -p(t)\ln p(t) + q(t)\ln q(t)\,\mathrm{d}t\right| \\
&\leqslant \int |-p(t)\ln p(t) + q(t)\ln q(t)|\,\mathrm{d}t \\
&\leqslant \int -|p(t) - q(t)|\ln|p(t) - q(t)|\,\mathrm{d}t
\end{aligned}$$

where the last inequality follows the convexity inequality established above. Finally, the last term is:

$$\begin{aligned}
&= -\frac{1}{2}\|f_s - f^A\|_1 \ln\frac{1}{2}\|f_s - f^A\|_1 + \frac{1}{2}\|f_s - f^A\|_1 H\left(\frac{|f_s - f^A|}{\|f_s - f^A\|}\right) \\
&\leqslant -\frac{1}{2}\|f_s - f^A\|_1 \ln\frac{1}{2}\|f_s - f^A\|_1 + \frac{1}{2}\|f_s - f^A\|_1 \frac{1}{2}(1 + \ln 2\pi\overline{\sigma}^2), \forall s
\end{aligned}$$

for some $\overline{\sigma}^2$, using the fact that in the continuous case the Gaussian distribution is the maximal entropy distribution. Note that:

$$\overline{\sigma}^2 = \sup_{s \in \mathcal{S}} \int t^2 \left(\frac{|f_s(t) - f^A(t)|}{\|f_s - f^A\|_1} \right) \mathrm{d}t$$

always exists (as long as all distributions admit second moments) and suffices. Let $\kappa \equiv \mathrm{e}^{0.5(1+\ln 2\pi\overline{\sigma}^2)}$, then:

$$|H(p) - H(q)| \leqslant -\frac{1}{2}\|f_s - f^A\|_1 \ln \frac{1}{2\kappa}\|f_s - f^A\|_1,$$

The inequality follows from taking expected values with respect to the distribution of forecasters on both sides.

A Decomposition of Uncertainty For Subgroups

Assume that the set of forecasters $\mathcal{S}$ contains G subgroups and that there exists a transition probability S_g for every subgroup g such that the aggregate forecast can be rewritten

$$f^A = \iint f_s \ln f_s \, \mathrm{d}S_g(s) \, \mathrm{d}G(g). \tag{3.39}$$

The aggregate forecast of subgroup g is naturally defined as:

$$f^{A_g} = \int f_s \ln f_s \, \mathrm{d}S_g(s). \tag{3.40}$$

The *intra-group* identity can be derived just like the simple case was:

$$\mathcal{D}_g = \text{Disagreement in } g = -\text{Uncertainty in } g + H^{A_g}. \tag{3.41}$$

Similarly, remark that:

$$\mathcal{K}(f^{A_g}, f^A) = -H^{A_g} - \int f^{A_g} \ln f^A. \tag{3.42}$$

Averaging this last identity over all subgroups;

$$\int \mathcal{K}(f^{A_g}, f^A) \, \mathrm{d}G(g) = -\int H^{A_g} \, \mathrm{d}G(g) + H^A; \tag{3.43}$$

and substituting the identity of intra-group disagreement integrated over g,

$$\int \mathcal{K}(f^{A_g}, f^A)\,\mathrm{d}G(g) = -\int \mathcal{D}_g\,\mathrm{d}G(g) - \iint H(s)\,\mathrm{d}S_g(s)\,\mathrm{d}G(g) + H^A. \tag{3.44}$$

All in all, this provides us a very similar decomposition where disagreement breaks down into inter and intra-group disagreement:

$$\begin{aligned}\text{Aggregate Uncertainty} = \text{Uncertainty} &+ \text{Intra-group Disagreement} \\ &+ \text{Inter-group Disagreement}.\end{aligned} \tag{3.45}$$

State-Dependent Entropy and Total Uncertainty

Consider a random variable X with distribution P_X and Radon-Nikodym density $\mathrm{d}P_X/\mathrm{d}\lambda = f_X$ with respect to the Lebesgue measure on $\mathbb{R}^m$; $m \geqslant 1$. Let A be a subset of the support of X such that $P(X \in A) > 0$. Let $\mathcal{B}(\mathbb{R}^m)$ denote the set of Borelian sets on $\mathbb{R}^m$; the Radon-Nikodym density of $P_{X|X\in A}$ is given by:

$$\begin{aligned}\forall C \in \mathcal{B}(\mathbb{R}^m) : P_{X|X\in A} &= \frac{P[X^{-1}(C) \cap X^{-1}(A)]}{P[X^{-1}(A)]} \\ &= \frac{1}{P[X^{-1}(A)]}\int_\Omega \mathbb{1}_{X^{-1}(A)\cap X^{-1}(C)}(\omega)\,\mathrm{d}P(\omega) \\ &= \frac{1}{P[X^{-1}(A)]}\int_\Omega \mathbb{1}_{X^{-1}(A)}(\omega)\mathbb{1}_{X^{-1}(C)}(\omega)\,\mathrm{d}P(\omega) \\ &= \int_C \frac{\mathbb{1}_A(x)}{P[X^{-1}(A)]}\,\mathrm{d}P_X(x);\end{aligned}$$

which means that $\mathrm{d}P_{X|X\in A}/\mathrm{d}P_X = \mathbb{1}_A(\cdot)/P[X \in A]$ and by extension: $\mathrm{d}P_{X|X\in A}/\mathrm{d}\lambda = \mathbb{1}_A(\cdot) f_X/P[X \in A]$. With that in mind, the Shannon entropy of X is given by:

$$
\begin{aligned}
H(X) &= -E[\ln f(X)] \\
&= -\int \ln f(x)\,\mathrm{d}P_X(x) \\
&= -\int_A \ln f(x)\,\mathrm{d}P_X(x) - \int_{\not{A}} \ln f(x)\,\mathrm{d}P_X(x) \\
&= -P(X \in A) \int_A \frac{\mathbb{1}_A(x)}{P(X \in A)} \ln \frac{f(x)\mathbb{1}_A(x)}{P(X \in A)}\,\mathrm{d}P_X(x) \ldots \\
&\quad - [1 - P(X \in A)] \int_{\not{A}} \frac{\mathbb{1}_{\not{A}}(x)}{1 - P(X \in A)} \ln \frac{f(x)\mathbb{1}_{\not{A}}(x)}{1 - P(X \in A)}\,\mathrm{d}P_X(x) \ldots \\
&\qquad + \psi[P(X \in A)] \\
&= -P(X \in A) \int_{\mathbb{R}^m} \ln f_{X|X\in A}(x)\,\mathrm{d}P_{X|X\in A}(x) \ldots \\
&\quad - [1 - P(X \in A)] \int_{\mathbb{R}^m} \ln f_{X|X\in A}(x)\,\mathrm{d}P_{X|X\in \not{A}}(x) + \psi[P(X \in A)];
\end{aligned}
$$

where ψ denotes the entropy function of a binary distribution with parameter p, that is, $\psi(p) = -p\ln(p) - (1-p)\ln(1-p)$.and the final result is achieved by recognizing the conditional entropies in the first two terms. Note that this method can be effectively applied to any partition of the space $X(\Omega)$ and to discrete distributions just as well by taking densities with respect to the counting measure. This decomposition is a generalization and extension of the recursiveness of entropy in the discrete case to the continuous case.

Extensions to Weighted Entropy

This section extends the previous framework where agents attribute different utility values to the outcomes that they are trying to forecast. Imagine that the forecaster puts a weight $w_i(t) > 0$ at event t to indicate the "importance" of event t relative to others. The weighted loss function and minimization problem become:

$$\min_{f_i} R = \int_Y \int_t w_i(t) f_i(t) \ln\left(\frac{f_i(t)}{f_y(t)}\right) \mathrm{d}t\, \mathrm{d}P_Y \text{ such that } \int_t f_i(t)\, \mathrm{d}t \leqslant 1. \tag{3.46}$$

The first order conditions of the new weighted problem are given by:

$$\frac{\partial L(\{f_i\}, \lambda)}{\partial f_i(t)} = \int_Y w_i(t) \ln\left(\frac{f_i(t)}{f_y(t)}\right) \mathrm{d}P_Y + w_i(t) - \lambda = 0; \tag{3.47}$$

and the constraint. The conditions solve into:

$$f_i(t) = \mathrm{e}^{-w_i(t) + \lambda / w_i(t) + E(\ln f_y(t))}, \tag{3.48}$$

as stated.

4

FORECASTING, UNCERTAINTY AND BUSINESS CYCLES

Taken together, [...] SPF panelists are quite flexible in their approach to forecasting. They use a combination of models in forming their expectations, rather than just one model. And, they vary their methods with the forecast horizon.

— Tom Stark, *"SPF Panelists' Forecasting Methods: A Note on the Aggregate Results of a Nov. 2009 Special Survey"*, 2013

4.1 INTRODUCTION

The overwhelming improvement in collecting information seen in the past two decades has made forecasting exercises at the same time easier and harder. The "big data" revolution supposedly allows forecasters to grasp a greater information set, which under the proper assumptions will always result in better predictions in a regression framework. But this slew of new potential predictors also leaves more room for misspecification and greater model uncertainty. The blurred lines between forecasting models makes economic policy decisions more complicated, hence the importance of quantifying how much one can trust their models.

Also known as model inadequacy, model bias, or model discrepancy, model uncertainty represents the lack of knowledge of the underlying law of nature governing the object being modeled. It is a measure of how accurately a mathematical model describes the true "laws of nature", notwithstanding that models can be approximations or simplifications of the reality they aim at describing. When economists try to capture the movement of economic aggregates with structural models, they know very well that the model itself is inaccurate since there always exists other "frictions" that are impossible to account for, either because the data are lacking *at the moment* or simply because those aren't measurable. And even if there is no unknown parameter in the model, a discrepancy is still expected between the model and true data generating process – the model itself, though, is as close as it gets to the laws of nature. In econometric analysis, model uncertainty represents a challenge because estimates will in most cases depend on the model estimated. Leamer and Leonard (1983) warn about the sensitivity of classical regression analysis to the selection of regressors. And only recently has macroeconomic theory started to account for model uncertainty to better describe the reality of agents' utility maximizing decisions, as seen in Hansen and Sargent's defense of the concept of robustness in macroeconomic models (Hansen and Sargent, 2010).

The push for more reliable models to describe economic aggregates truly began with the Great Recession when quite all models failed. The problem is that the complexity of identification increases exponentially as the number of variables and "shocks" grows, not to mention the absence of any guarantee of closing in on a realistic representation of the true data generating process.[1] What is more, these models involve – for the most part – few variables of interest and a vast number of potential explanatory controls. The perceived need to control for large numbers of potential confounds has led to over-specified models that decrease efficiency without necessarily guaranteeing the absence of omitted variable bias (Montgomery and Nyhan 2010). Because of the uncertainty around the truthfulness of the assumptions is indisputable, the common "hand waving" solution is to assert that reality is a convex combination of different representations of it. In rigorous economics, this takes the form of model averaging to formulate predictions, also known as optimal forecast combinations.

Forecast combinations date back to much further back than the Fed's panel of forecasters. Knight (1921) already considered a rudimentary form of model uncertainty in his definition of uncertainty. Even under what Knight calls true "uncertainty" agents estimate the veracity of their own model: "The business man himself not merely forms the best estimate he can of the outcome of his actions, but he is likely also to estimate the probability that his estimate is correct" (Knight, 1921, p. 226). In other words, agents facing uncertainty trust their forecasting models to a degree proportional to the probability of that model being a correct representation of nature. Tom Stark's words in epigraph of this chapter confirm Knight's predictions, at least for the Fed's forecasters. In practice, economists neither rely on a single model to formulate their predictions, nor do they seem to keep their set of models identical over time, mostly because they are aware of the risk of model misspecification, or have incomplete access to information, or trivially because the data do not remain available at all times. Second and most importantly, the exercise of forecasting at its deepest prone to uncertainty around the

[1] Romer (2016) describes the process of adding further nuts and bolts to a model as "putting lipstick on the pig."

"true" data generating process and the flexibility of forecasters aims at mitigating that risk by letting each model's likelihood be contingent. Finally and at a more philosophical level, the idea that there is such a thing as a "best" model is somehow vain in many situations of social science, where unequivocal proofs that one specification is the "true model" represent an impossible standard.

In this chapter, I try to estimate forecasting uncertainty by focusing on the level of trust – in the probabilistic sense – one may attribute to various forecasting models. I do so using Bayesian model averaging over a broad range of predictors and motivate my choice using statistical decision theory. Unlike survey forecasts whose merits I starkly defended in Chapter 2, measuring uncertainty solely with macroeconomic data gives a readily and "objective" – in that it does not depend on other people's estimated – proxy for what I argue to be another dimension of Knightian uncertainty.

4.2 Building a Measure of Model Uncertainty

4.2.1 *Optimal Forecast Combination and Model Uncertainty*

Model combination and weighting is a customary practice of forecasters and has been successfully applied in several areas of forecasting, in particular for GDP growth and inflation. By allowing the forecaster to eschew the challenge of selecting a single model, model combination precisely aims at circumventing uncertainty surrounding the model selection process. Claeskens and Hjort (2008) and Moral-Benito (2015) provide an excellent review of selection mechanisms, but I focus on model combination rather than model selection. Indeed, the latter may be undesirable for at least two reasons. First, the model selected might not be the one which is the closest to the true DGP since the definition of "best model" is up to the econometrician. Second, even if the model selected is the best one, it is rarely optimal to ignore the evidence from other "imperfect" models – at least in the sense of forecast risk.

In practice, forecast performance is one of many forecast combination strategies. Combination methods that are founded upon forecast errors are sometimes classified as frequentist model averaging. In the context of density forecasts, the complexity of evaluating forecast performance – what is the "realized density" and what is the mean squared forecast error of a density forecast? –, other approaches are preferable. Gánics (2018) proposes the Probability Integral Transform (PIT) as a tool for density forecast combinations. Another popular and general combination method is Bayesian model averaging, which given the definitions of uncertainty that we have been working with so far, might be the most natural framework to build a measure of model uncertainty. A more recent approach as presented in George and McCulloch (1993) and Raftery et al. (1997) is to include model uncertainty by assuming that there is uncertainty surrounding the population parameters *and* in the inclusion of certain predictors in a model. Such methods are called Bayesian model averaging (BMA). While the jury is still out on whether Bayesian model averaging and selection is preferable to their frequentist counterparts, I retain Bayesian model averaging for several reasons. First, these methods do seem to perform quite well in economics. In a thorough comparison exercise, Rossi and Sekhposyan (2014) find that forecasts obtained via Bayesian model averaging are the best calibrated for predicting output growth and inflation in the United States. Second, BMA has a very intuitive interpretation in statistical theory as I will present in greater detail below. Finally, it ties to the theory of uncertainty presented by Knight (1921).

For the sake of clarity, I remind the forecasting problem described in Chapter 3. Forecasters seek to minimize the following loss function:

$$
\begin{aligned}
R(f,\theta) &= E_{Y,Z}[\mathcal{L}(f(Z,\theta),Y,Z)] \\
&= \int_z \int_y \mathcal{L}(f(Z,\theta),y,z)\,\mathrm{d}P_{Y|z,\theta}(y)\,\mathrm{d}P_{Z|\theta}(z).
\end{aligned} \tag{4.1}
$$

I assume that the data generating process (DGP) of y is unknown in two dimensions: the (econometric model) form and the parameters for each model. For simplicity, I will not write out integrals with respect to the parameter vector θ, but all formulas and claims hold just the same. The statistical model is given by $(M, \mathcal{M}, P_M)$ where M is the set of models that forecasters know and can estimate, $\mathcal{M}$ is the smallest σ-algebra containing M and P_M a probability measure on M that is absolutely continuous with respect to the counting measure μ defined over $\mathbb{N}$. Up to a normalization, the Radon-Nikodym density of P_M with respect to μ is $\pi(m) \equiv \mathrm{d}P_M(m)/\,\mathrm{d}\mu(m)$. The decision problem can be rewritten in a Bayesian setting, integrating over the set of models M where $\mu(M) < \infty$:

$$\min_f \int_M R(f,m)\,\mathrm{d}P_M(m) = \int\limits_M \int\limits_Z \int\limits_Y \mathcal{L}(f(Z,m),y,z) \times \dots \\ \pi(m)\,\mathrm{d}P_{Y|z,m}(y)\,\mathrm{d}P_{Z|\theta}(z)\,\mathrm{d}\mu(m). \tag{4.2}$$

Note that this expression is general enough to allow for M to be a set of *classes* of models. Forgoing the exogenous covariates and replacing them with the information set available at that time, integrals can be rewritten as:[2]

$$\min_f \int_M R(f,m)\,\mathrm{d}P_M(m) = \int_M \int_Y \mathcal{L}(f,m,y)\pi(m|y,\Im)\,\mathrm{d}P_{Y|\Im}\,\mathrm{d}\mu(m). \tag{4.3}$$

A forecast f is a decision made on the space of possible outcomes $\mathcal{Y}$. In the case of a squared-error type of loss function the solution to the statistical decision problem is:

$$f^* = \int_M E(y|m,\Im)\pi(m|\Im)\,\mathrm{d}\mu(m) \tag{4.4}$$

[2]See Monfort (1982) for a proof.

which essentially means that the optimal forecast is the Bayesian average over all possible models. Furthermore, assuming that the set of possible models M is of finite measure, the equal weights forecast is strictly worse in terms of quadratic risk that the Bayesian weights forecast combination. In a discrete setting for a point forecast with M models and a squared error loss function, the optimal forecast combination is:

$$y_{t+h|t}^{\text{BMA}} \equiv \sum_{m=1}^{M} E(y_{t+h}|M_m, \Im_t)P(M_m|\Im_t) \tag{4.5}$$

The Appendix of this chapter provides an alternative non-probabilistic framework that reaches similar conclusions. By now it should be clear that classical forecasting decision theory naturally turns toward Bayesian averaging as optimal forecast combination.

The previous derivations taught us two features of the forecasting problem: under imperfect knowledge (of the underlying DGP) Bayesian forecast combination is optimal, and the most "uneducated guess" of putting equal weights on all models is the worst possible combination in the quadratic risk sense. Hence, it is natural to map forecast model uncertainty to how close one forecaster is to making the worst possible combination. Any distance or pseudo-distance fits the purpose and a natural candidate would be the Kullback-Leibler divergence. But because combination weights are not probability statements on the states of nature like I considered in Chapter 3 prefer to discard it in favor of a more traditional choice and define forecast model uncertainty as follows:

$$FMU \equiv \left[\sqrt{\int_M [\pi(m|\underline{y}) - \mu(M)^{-1}]^2 \, \mathrm{d}\mu(m)} \right]^{-1} . \tag{4.6}$$

Intuitively, this metric will go up whenever the optimal weights get close to equal weights and conversely go to zero when they stray away from them. Note that this measure is bounded below by $\sqrt{1 - M^{-1}}^{-1}$

in a discrete setting with M models.[3] This lower bound is attained for any "ultra-confident" guess where all weight is put on one model, which provides further intuitive support for this measure. It's also worth noting that in this measure need *not* make sense to estimate on survey forecasts. Say that every forecaster picks a different model; *ex-ante* model uncertainty would stand at the lower bound for each forecaster since each seems quite sure which model to pick. Model uncertainty averaged across the panel would be at its lowest. It is clear, however, that forecasters disagree on the right model to choose and hence some form of Knightian uncertainty is present.[4] As an aside, the choice of the Euclidian distance in the denominator is arbitrary but of little impact since I am considering finite dimensional model spaces. Weights vectors therefore are all in $\mathbb{R}^M$ – a Euclidian space on which all distances will be equivalent. In other words, any other distance would fall within a reasonable range around the one I chose. Another non-distance based approach would have been to take the entropy over the set of posterior probabilities, but because this would be somehow a repetition of Chapter 3 I prefer to take an alternative, less ideological, approach.[5]

4.2.2 *Data and Models Retained*

The data consist of measures of asset prices, real economic activity, wages, prices and money supply, similar to that of Stock and Watson (2003) from January 1959 through January 2011. The data are transformed to

[3]See the appendix of this chapter for a proof.

[4]This reinforces the idea that disagreement is a key component of Knightian uncertainty, as Chapter 3 concluded.

[5]Note that if we were to consider total uncertainty as in Equation (3.23), uncertainty would write as:

$$H(Y,M) = \sum_{i=1}^{M} H(Y|M = m_i)p_i + H(\{p_i\}); \text{ where } p_i = P(M = m_i).$$

This decomposition highlights two components of uncertainty: that of Y under the different DGPs that it is assumed to follow, and that of the overall model selection process.

eliminate stochastic or deterministic trends and seasonality, as in Rossi and Sekhposyan (2014). I use a fixed rolling window estimation scheme of 40 periods (10 years).

Marcellino (2008) provides an exhaustive list of forecasting models in economics without comparing their performance per se. Giannoni (2016) offers a restricted list that Fed forecasters allegedly use: a combination of (S)VAR, autoregressive, factor and DSGE models. Given the number of times the estimation has to be run, Bayesian estimation of a DSGE is not feasible considering the computational time[6] Hence, I choose to focus on autoregressive and factor models. More specifically, I use autoregressive distributed lag (ADL) models where I let the regressor change over the whole collection available in the data set, and factor models with factors being selected within the data set. While this choice may seem simplistic, there are at least two pitfalls to incorporating a broader set of forecasting models: nested models and over-fitting. ADL models are estimated through Bayesian least squares to account for parameter uncertainty and avoid the measure of model uncertainty to be driven by it. The ADL forecasting equation can be written as follows:

$$y_{t+h|t} = \beta_{k,0} + \beta_{k,1}(L)X_{t,k} + \beta_{k,2}(L)y_t + \varepsilon_{t+h}; \tag{4.7}$$

where the dependent variable is annualized GDP growth (using logarithmic growth rates), $X_{t,k}$ denote the k-th regressor in the database for $k = 1, \ldots, K$; $K = 32$. The error term is normally distributed: $\varepsilon_{t+h} \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. I consider one and four-quarter ahead forecasts. The lag polynomials $\beta_{;1}$ and $\beta_{;2}$ are left to have different lag orders for a given regressor and across regressors where lag orders are optimally selected using a Bayesian information criterion (BIC). The lag order is first selected for the autoregressive component and then augmented for the additional predictor. The variance of the error term is estimated and corrected for serial correlation with a Newey and West (1987) estimator.

[6] Furthermore, the relatively small time window might be insufficient for model estimation.

The variant of the ADL forecasting that I consider is the principal components model that augments the original DGP with factors extracted from the set of all regressors. This inclusion is motivated by two features of factor models. First, they have gained a lot of attention in the (economic) forecasting literature in the past few years, notably after Stock and Watson (2002) diffusion indices and Bernanke et al.'s (2004) Factor Augmented VAR or FAVAR. More specifically, I estimate the following static[7] factor model:

$$y_{t+h|t} = \beta_0 + \lambda \widehat{F}_t + \beta_2(L) y_t + u_{t+h} \tag{4.8}$$

where $\widehat{F}_t$ denote the estimated ℓ factors among the 32 predictors of the database; ℓ being set recursively to capture at least 60% of the total variation present in the predictors. (Overall, two to three factors end up selected for output growth or inflation). As in the ADL model the error term is normally distributed and its variance is estimated in a heteroskedasticity and autocorrelation robust manner.

This choice of models to estimate ties in to the work done by Wright (2011) with two salient differences: I do not try to forecast GDP growth for the sake of the exercise, and I average over a greater span of econometric models. The set I chose is a simple yet representative choice of practitioners. In the following section, I describe the estimation strategy for my measure of forecasting uncertainty.

4.2.3 *Econometric Methodology*

Traditionally, Bayesian model averaging uses the identity $\pi(m|y) \propto \pi(y|m)\pi(m)$ to estimate posterior weights. As a prior I assume that each model is of the same likelihood, that is: $\pi(m) = \mu(M)^{-1}\mathbb{1}_M$. In practice, this supports my claim that the most naive guess for a forecast combination is the prior. Because each model estimation spans over

[7] Bai and Ng (2007) explain that there is little gain in extending the framework to a dynamic factor model

certain parameter spaces and because of the choice of prior:

$$P(y|m) = \int_{\theta_m} \pi(y|m, \theta_m)\pi(m, \theta_m)\, \mathrm{d}\theta_m$$
$$\text{and } \pi(m_i|y) = \frac{\int_{\theta_{m_i}} \pi(y|m_i, \theta_{m_i})\pi(m_i)\, \mathrm{d}\theta_{m_i}}{\int_m \int_{\theta_m} \pi(y|m, \theta_m)\pi(m)\, \mathrm{d}\theta_m\, \mathrm{d}\mu m} \tag{4.9}$$

The estimation of these weights is usually done via Markov Chain Monte Carlo-type methods (MCMC), but in my case such a strategy is not feasible given the number of models that I consider. Instead, I use a result from Kass and Wasserman (1996) and Wasserman (2000) to approximate these weights. In the latter, the authors establish that:

$$\int_{\theta_{m_i}} \pi(y|m_i, \theta_{m_i})\pi(m_i)\, \mathrm{d}\theta_{m_i} = \widehat{\mu_i}(m_i|y)(1 + O_P(1))$$
$$\text{where } \log \widehat{\mu_i} = \log \pi(\widehat{m_i}|y) - \frac{1}{2} d_i \log n. \tag{4.10}$$

The last property gives credit to the idea that the estimated model can be used to give a reasonable approximation to the posterior of the model $\pi(m_i|y)$.[8] Under this approximation and with the assumption of a uniform prior on models:

$$\pi(m_i|y) \simeq \frac{\pi(y|m_i, \widehat{\theta}_{m_i})}{\int_m \pi(y|m, \widehat{\theta}_m)\, \mathrm{d}m} \tag{4.11}$$

From there the estimation strategy becomes clear. Once a set of acceptable forecasting models has been settled, one needs to estimate all models separately and compute the weights using the predicted parameters as in equation (4.11). In the case of a discrete and finite set of models, the integrals become discrete sums, hence:

$$P(M_m|y) \simeq \frac{P(y|M_m, \widehat{\theta}_{M_m})}{\sum_m P(y|M_m, \widehat{\theta}_{M_m})}. \tag{4.12}$$

[8]Note that $\widehat{\mu_i}$ need not converge to the true prior in probability, but Kass and Wasserman (1996) establish that the error of approximation does.

This can be readily computed from the estimated models according to the assumptions placed on the error terms (cf the previous section).

4.3 ESTIMATING FORECASTING MODEL UNCERTAINTY

4.3.1 *Model Uncertainty is Countercyclical*

The estimates of forecasting uncertainty for GDP growth at horizon $h = 1$ and $h = 4$ in Figures 4.1 and 4.2. Intuitively, forecasting uncertainty could have gone either down or up during a downturn. Economic recessions can be understood as unforeseen negative "shocks"[9] hitting the economy, meaning that economic predictions failed. All models stop working, become equally bad, and forecasters are hard pressed to prefer one to another. On the other hand, many practitioners consider that several predictors (such as real estate metrics) become particularly relevant during recessions and therefore forecasting in these periods is somehow "easier." In that case model uncertainty would actually fall during recessions. The data show that this second theory is not wrong for the forecasting framework that I consider. Business cycles and model uncertainty are indeed very much linked. The gray areas indicating NBER recessions over the sample period show that model uncertainty tends to increase sharply after recessions rather than decrease. This was particularly salient after the dotcom crisis or in the 1980s. Furthermore, forecasting uncertainty doesn't seem to change much with the forecast horizon, and nor does its dynamics. While some strand of the literature has found uncertainty to typically increase with the forecasting horizon, I have shown in Chapter 2 that this need not be true for all types of uncertainty. In particular, for Knightian uncertainty, forecast horizon seems to have little effect.

[9] The term "shock" does have a random connotation. To be sure, recessions always happen for a reason and they are, to a certain degree at least, possible to anticipate.

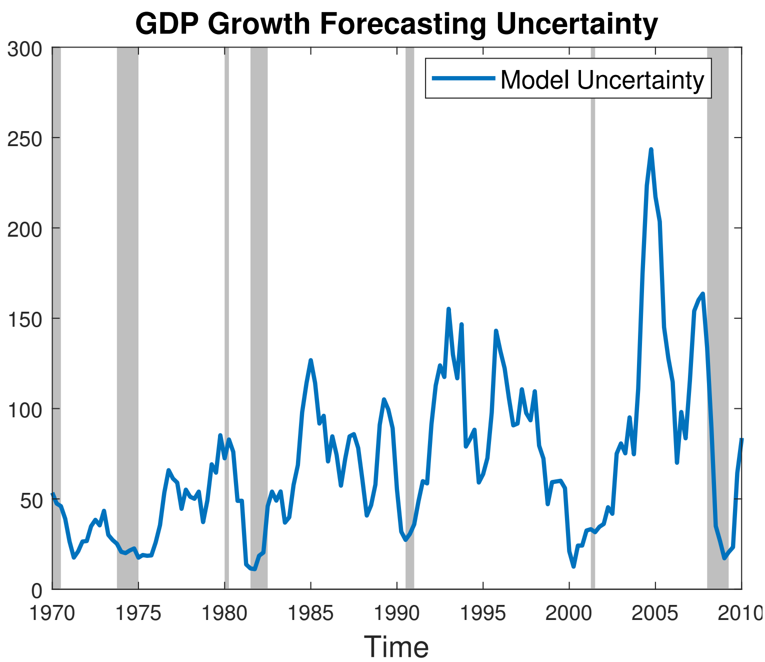

FIGURE 4.1: Forecasting Uncertainty for GDP, 1-quarter ahead forecasts

Coming back to the idea that some models may perform better during recessions, one could think that one or some model systematically outperform in tranquil times and fails during recessions, which would explain the pattern of model uncertainty that I observed. This could be true either in the broad sense of the class of models, or simply because some predictor always becomes relevant during downturns. This is not the case. A quick look at the Bayesian posterior weights as displayed in Figure 4.3 shows that there is no "star" model outperforming the rest. The factor model does seem to do quite well overall, even more so in the Great Recession. This is to be interpreted as the positive consequence of including more relevant information in the model, and has been documented in the literature. But at times some regressors to better and do worse in other periods.

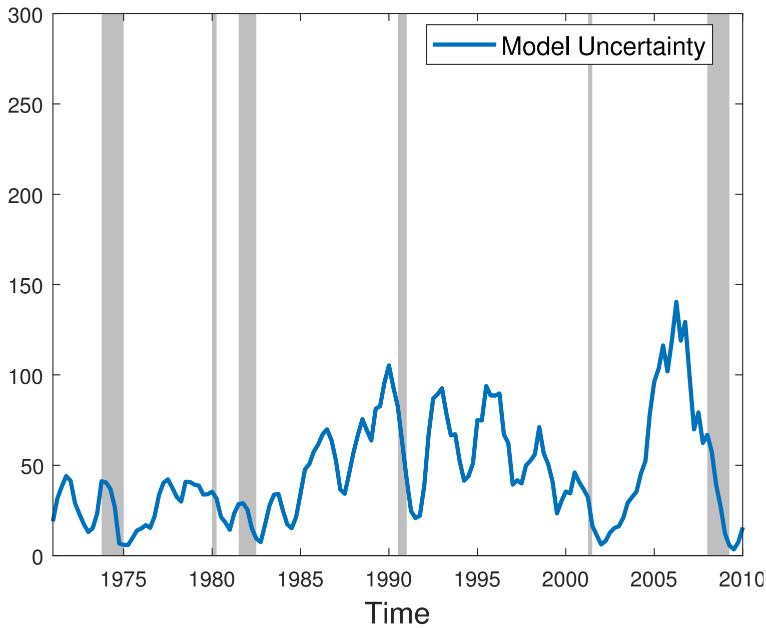

Figure 4.2: Forecasting Uncertainty for GDP, 4-quarter ahead forecasts

Another important issue was to know what type of uncertainty this measured "forecasting uncertainty" could fit in. Following Knight's typology, forecasting uncertainty would fall right under what is commonly coined Knightian uncertainty – not "radical" uncertainty nor risk. I claimed in the introduction that posterior weights nicely matched Knight's claim that under uncertainty agents would still try to estimate how likely their guesses were. (Only under radical uncertainty is such a judgment impossible to make.) Looking at how my measure correlates with the aggregates measured in Chapter 2 does however indicate that such a statement is partly false. In absolute forecasting uncertainty correlates the most with ex-post measures of uncertainty (about 0.4-0.5 v. 0.2-0.3 for ex-ante items). This most likely comes from the fact that model weights are estimated ex-post (since they depend on the

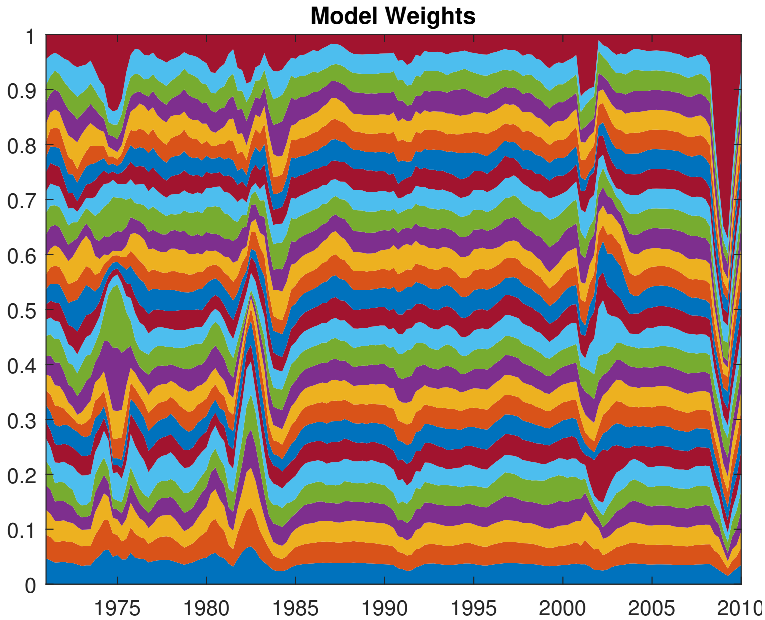

FIGURE 4.3: Posterior Weights, 1-quarter ahead forecasts

data) while Knight's probability statements do not depend on the past. In a sense, measured forecasting uncertainty behaves more like that of Jurado et al. (2015) (about 0.36 for $h = 1$). Note, however, that those correlations are quite small and give comfort in the idea that this measure is measuring a somehow different phenomenon.

Overall, the measured forecasting uncertainty confirms that forecasting isn't as trivial as some practitioners will assert – there is not one true model that always does better than the rest – and that it's even less so in periods of economic turmoils. This finding hasn't really been studied empirically nor theoretically since there are so few economic models using forecasts as an endogenous variable. The next section looks at a simple macroeconomic model to see if such findings could be replicated in simulations.

4.4 BAYESIAN MODEL AVERAGING IN A REDUCED FORM MACROECONOMIC MODEL

4.4.1 *Forecast Combination Equilibria*

I have briefly touched in the introduction to this dissertation that rational expectations let room for only a very specific type of uncertainty. It is therefore natural to turn to models that account for the bounded rationality of agents (forecasters) to model uncertainty in the economy. But letting go of rational expectations need not mean discarding all the conclusions that were reached while working with them. Perhaps the most underrated contribution of general equilibrium models with rational expectations was to show that macroeconomic aggregates are forward looking. Therefore, modeling forecasting uncertainty has to include a feedback effect of expectations on macroeconomic aggregates, in my case output.

The literature including these two features is very sparse. In fact, most models available to this date are versions of Branch and Evans (2007). Agents are modeled as econometricians that select the best forecast out of a menu of methods available to them using "dynamic predictor selection." This behavior imposed to the model generates more exotic economic fluctuations, including endogenous time varying volatility. It is unlikely, however, that agents behave according to model selection instead of model combination, which is a known practice among professional forecasters (cf. Tom Stark's quotation in epigraph of this chapter). What is more, some forecasts are directly published as forecast combinations, such as the Michigan Survey of Consumers and most Bloomberg forward looking data. Evans and Honkapohja (2013) stated that when abstracting from rational expectations, economic agents should be modeled to be "as smart as (good) economists", and it seems that good economists combine their forecasts rather than pick them. I use a variant of Gibbs's (2017) model where forecasts are optimally combined. The weight are derived minimizing the expected squared forecast error; and as a reference equal weights combinations represents the most naive and uninformed forecast combination.

4.4.2 *The Reduced Form Economy*

The simplistic economy I consider is described by a self-referential stochastic process driven by a vector of observable and unobservable exogenous shocks. The model takes the following form:

$$y_t = \mu + \alpha E_{t-1} y_t + \beta' x_{t-1} + \varepsilon_t; \tag{4.13}$$

where y_t denotes output, x_{t-1} is a vector of immediately observable exogenous shocks and ε is an unobservable, iid exogenous shock. This economy can be seen as a reduced form version the cobweb of Muth (1961) when $\alpha < 0$ and the aggregate supply model of Lucas (1973) for $0 < \alpha < 1$. Under rational expectations agents form expectations like econometricians would: they use linear regressions. In mathematical terms:

$$E_{t-1} y_t = \widehat{\beta}_i'(1, x_{t-1}). \tag{4.14}$$

The orthogonality of forecast errors to agents information sets constitutes the sole condition that pins down the rational expectation equilibrium.

Much like Gibbs (2017) I consider a world in which agents under-parameterize their models to form imperfect forecasts they'll wish to combine. There are as many misspecified models as observables shocks, meaning that each model's forecast is given by: $E_{t-1} y_{i,t} = \widehat{\beta}_i'(1, x_{i,t-1})$. The under-parameterized models represent the fact that forecasters, while they could observe the shocks, have limited access to data and limited information capacity (cf. Chapter 3). Following the cognitive consistency principle, agents choose to combine the k different forecasts to create a single forecast of output using a weighted sum approach:

$$E_{t-1} y_t = \sum_k \gamma_i E_{t-1} y_{i,t} = \gamma_i \widehat{\beta}_i'(1, x_{i,t-1}). \tag{4.15}$$

I work under the same forecast combination equilibrium as Gibbs (2017), meaning that agents' beliefs satisfy the following equilibrium conditions:

$$E_{t-1}(1, x_{i,t})[y_t - \widehat{\beta}_i'(1, x_{t-1})] = 0, \forall k. \tag{4.16}$$

These conditions require agents to consider only "econometrically coherent" models in equilibrium where the conditional expected forecast error of each model is zero. A Forecast Combination Equilibrium is reached when given weights $\{\gamma\}$ the equilibrium conditions are met.[10] I consider that all agents settle on the same combination strategy, which isn't too unrealistic given that my ultimate study population is professional forecasters.

Unlike Gibbs (2017) my forecasters combine weights according to Bayesian combination strategies. Despite the importance of forecast combination in the literature and among practitioners, macroeconomic models generally study agents that select, rather than combine forecasts. Some examples can be found with Brock and Hommes (1998), Branch and Evans (2007 and 2011), or Branch and McGough (2008). I show in the appendix that under the assumption that misspecified models cover the whole range of possibilities such combination strategies are optimal, which warrants the existence of a forecast combination equilibrium that coincides with rational expectations.[11] Shocks are independent and identically distributed to guarantee that agents can't forecast them. The misspecified regressions simply help implicitly figuring out the β coefficients.

4.4.3 *Simulating The Reduced Form Economy*

In this chapter, however, I worry less about model equilibria and use the framework solely for simulation purposes. I simulate the reduced form

[10] For an explicit algebraic version of those, simply replace $E_{t-1}y_t$ with its forecast combination expression of (4.15) in the equilibrium conditions. Also see the appendix of Gibbs (2017).

[11] Also called a "Fundamental Forecast Combination Equilibirum", see Gibbs (2017).

economy to see how forecast weights and the uncertainty measure that I built react with business cycles. I simulate the reduced form economy for 100 periods, 10,000 times with the ultimate goal of understanding how forecasting model uncertainty behaves closes to recessions (following the same definition as the NBER). I use a very standard calibration of the model and keep all coefficients to unity besides α, the strength of the feedback mechanism, that I let vary from zero to one. I try to match basic business cycle facts with an average growth of about 3% and a turning point (defined as two periods of consecutive decline) every ten years.

I then look at the importance of the self referential component of the reduced form economy. For that, I consider the measure of model uncertainty and run a regression on a dummy indicating a recession of the form: $m = a + b\mathbb{1}_{\text{Recession}} + error$, for different values of α ranging from zero to 0.9 (1 being off the table for stationarity reasons). I find, quite naturally, that the stronger the self referential component, the more model uncertainty correlates with recessions. Standard errors are estimated using Newey and West (1987) estimation procedures since the error term is very likely to be serially correlated. Figure 4.4 displays the estimated value and confidence interval in the simple regression that I run.

Though the interaction term is barely significant, the trend seems to be clear. Also note that the self referential component is still a minor determinant of the outcome variable. Consider the gradient of output at time t: $\nabla y_t = (\alpha, \beta, 1)$ (considering the trend to be deterministic.) To warrant non-explosiveness $\alpha < 1$, meaning that in the direction of, for instance, the steady-state the self referential component is still very much dominated by other shocks, and in particular always strictly dominated by the non-observable shocks.

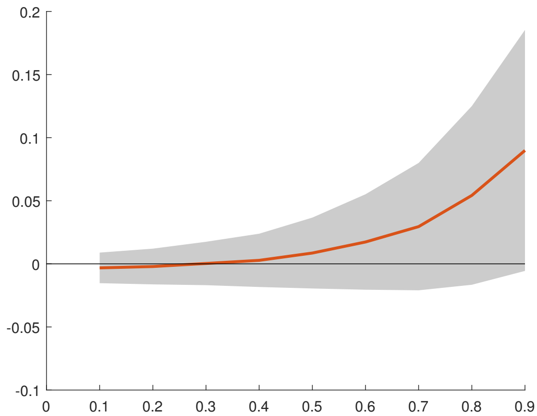

Figure 4.4: Reaction of Forecasting Model Uncertainty to Recessions

4.5 Conclusion

Knight's treatise on uncertainty has been interpreted in a great number of ways, but most agree that Knightian uncertainty denotes an incapacity to formulate a probability statement. Few, however, tried to center their attention on Knight's further observation that "probability estimates" were still given in situations of uncertainty, and that agents would further try to esteem the probability that their "estimates" got it right. This chapter used Bayesian econometric theory to try to measure such a feature of Knigthian uncertainty.

It appears that forecasting uncertainty, the inability to decide on a forecasting model, behaves much like other measures of Knightian uncertainty that I have presented in the previous chapters of this thesis. Recessions seem to represent the occasion for forecasters to re-think their paradigm, and at a much deeper level than the simple "best predictor" for the situation. There is not one model that performs better than the rest in a given situations, and if there were, everyone's using it would make it useless. All in all, this gives credit to the Lucas critique and the idea that historical forecasting is useful at best in hindsight.

Because the benefits of forecast combinations become clearer, a branch of economic theory has grown to study their effect in standardized theoretical economies. Simulating such economies under the forecast combination scheme that I retain here seems to confirm the conclusion that I have reached with the help of the data. Nevertheless, macroeconomic models including forecast combinations are few and far between and the conclusions that one can derive from their use are still to be considered with care.

Forecasting uncertainty represents the last piece of uncertainty that I wanted to estimate as it is a key metric to look at for economic policy-making. Because there seems to be a tight link between business cycle movements and uncertainty and Keynes had made this a central point in his plea for government intervention. The next and last chapter of this thesis tries to empirically challenge Keynes's hypothesis.

4.6 APPENDIX

An Alternative Framework

In the particular example where the loss function is the mean squared forecast error and if some structure is put on the MSE and correlations of MSEs, the problem of forecast combination can be solved analytically. Let y denote the variable to forecast, $f = (f_1, \ldots, f_m)$ the vector of (unbiased) forecasts and $e = \iota.y - f$ the vector of forecast errors with variance Σ_e. The forecast combination problem can be written as follows:

$$\min_{w'\iota=1} w'\Sigma_e w. \tag{4.17}$$

Setting a standard Lagrangean one can show that the solution to this problem is $w^* = (\iota'\Sigma_e^{-1}\iota)^{-1}\Sigma_e^{-1}\iota$. Furthermore, in the particular case where forecasts have the same MSE σ^2 and identical correlation ρ, it can be shown that the optimal forecast will put equal weight on all models. In other words, when model uncertainty is at its highest – no forecast performs better than the other and correlations do not allow proper diversification of model risk – the optimal forecast will be the equal weight forecast. I summarize this idea in the following proposition:

Proposition Let M unbiased forecast models in a vector f for the variable y be such that:

$$MSE(f) = \Sigma_e = \begin{pmatrix} \sigma^2 & \rho & \ldots & \ldots & \rho \\ \rho & \sigma^2 & \rho & \ldots & \rho \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \rho & \ldots & \ldots & \rho & \sigma^2 \end{pmatrix}; \tag{4.18}$$

then the optimal forecast combination puts equal weight on every forecast.

Proof The forecast combination problem amounts to finding w that solves:

$$\min_{w'\iota=1} w'\Sigma_e w. \tag{4.19}$$

Setting a standard Lagrangean one can show that the solution to this problem is $w^* = (\iota'\Sigma_e^{-1}\iota)^{-1}\Sigma_e^{-1}\iota$: $\mathcal{L} = 0.5w'\Sigma_e w - \mu(w'\iota - 1)$; differentiate with respect to w to get: $w = \mu\Sigma_e^{-1}\iota$. Plug this in the constraint to get the value of $\mu = (\iota'\Sigma_e^{-1}\iota)^{-1}$.

To get the explicit form of the weights, note that $\Sigma_e = (\sigma^2 - \rho)I_M + \rho J$ where J is a matrix of ones such that $J^2 = M.J$ As long as $\rho \neq \sigma^2$, Σ_e is invertible and its inverse takes the form $aI_M + bJ$ where:

$$a = \frac{1}{\sigma^2 - \rho}; \tag{4.20}$$

$$b = \frac{\rho}{(\rho - \sigma^2)(\sigma^2 + (M-1)\rho)}. \tag{4.21}$$

This allows us to compute:

$$\Sigma_e^{-1}\iota = \frac{1}{\sigma^2 + (M-1)\rho}; \tag{4.22}$$

$$\iota'\Sigma_e^{-1}\iota = \frac{M}{(\sigma^2 + (M-1)\rho)}; \tag{4.23}$$

which entails that $w = \frac{1}{M}\iota$ and finishes the proof. In other words, when all models are equal on the ground of forecast error and there is no reason to prefer one to another, the optimal combination is the most naive one.

Properties of the Forecasting Uncertainty Measure

In a discrete setting with M models the Forecast Uncertainty Measure is bounded below by $\sqrt{1 - M^{-1}}^{-1}$. To see that, remark that the identity:

$$\sum_{m=1}^{M} \left(p_m - \frac{1}{M}\right)^2 = \sum_{m=1}^{M} p_m^2 - \frac{1}{M} \tag{4.24}$$

holds true thanks to the fact that $\sum p_m = 1$. Because all weights are between zero and one, $\sum p_m^2 \leqslant \sum p_m = 1$ which finishes the proof. This lower bound is naturally reached in the "ultra-confident" guesses where all weight is put on one model. In the continuous case, the lower bound becomes zero; "ultra-confident" guesses become *even more* confident as each model is infinitesimal and betting on an infinitesimal class of models is tantamount to betting on an even of probability zero.

5

FISCAL MULTIPLIERS IN UNCERTAIN TIMES

Keynes' views on private spending and public spending are illustrative. We cannot leave investment decisions to businessmen in the private sector because they do not know enough about the future to make their "parting with liquidity" worthwhile. By contrast, when government spending on public works is under consideration, the debilitating uncertainty goes into remission. And the question of whether the spending is worthwhile somehow becomes irrelevant.

— Roger W. Garrison, *The Kaleidic World of Ludwig Lachmann*, 1987

5.1 introduction

The battle for government intervention resurfaces every time the economy seems to be put at risk. Such interventions can take various forms – changing regulations, modifying taxes, infrastructure spending, etc. – and economists are likely bound to keep on debating upon the scope and usefulness of these nearly forever.

One idea that most do agree upon is that government intervention need not be the same throughout the business cycles ebbs and flows. Which policy tools should be countercyclical, and to what extent they should be so, are still open questions. However, the dichotomy between recessions and expansions is reductive and may overlook other situations that warrant government intervention. Specifically, downturns and upswings might not be the only metric to look at. An completely unforeseen downturn may shift behaviors in a way that makes them more risk averse and bolster precautionary spending, thus reducing the consumption multiplier effect of government spending. The converse would be true for any situation where clear expectations for the future do not warrant more prudent consumption, be it during recessions or expansions. While it only recently gathered some interest, the question of the role of uncertainty has been the pet peeve of economic policy for more than seven decades. Standing at both ends of the economic spectrum, Keynes and Hayek had agreed on the importance of uncertainty in the determination of equilibrium outcomes, and made Frank Knight's theory of uncertainty a cornerstone of their own. But Keynes and Hayek evidently reached diametrically opposite conclusions, and even more so saliently on the question of state intervention. Keynes argued that the systematic incoordination caused by uncertain and irrational forces required correction through deliberate changes in public spending and taxation. Hayek believed that no central planning authority could gather nearly enough information to efficiently coordinate uncertain agents, thus axing every form of government intervention.

The confrontation between Keynes and Hayek's ideas is no news but little has been done to try to settle the debate. Some empirical work has been done to analyze the role of government under different regimes but not without limitations. In the case of uncertainty as with many other covariates, a crucial issue in the method is that economic downturns do tend to correlate with periods of heightened uncertainty in practice. This raises the issue of specious conclusions if one were to retain only an indicator of uncertainty and separating history according to its level. To a certain extent, all methods aimed at estimating regime switching model suffer from that flaw insofar as the indicator of regime changes may well be correlated with another series, rendering any causal analysis more a judgment of values than anything else.

Instead of attempting to solve an unsolvable debate, this chapter tries to fill part of a policy practice gap by assessing whether government spending is more effective in uncertain than in tranquil times. I do so using data over the whole 20th century in the United States, both on economic aggregates and acceptable proxies of economic uncertainty. I find that government intervention appears to effect more when perceivable economic uncertainty is high, while spending in tranquil times seems to bring insignificant benefits in terms of output. Because spuriousness easily sneaks in discussions on the impact of uncertainty, I propose a new way to make estimation tools robust to some stylized facts of uncertainty, in particular that measured uncertainty is highly counter-cyclical. This encouraging result seems, however, quite sensitive to the sample window retained. I consider this evidence of a possible structural break in the relation between fiscal multipliers and uncertainty, calling for a deeper investigation of the microeconomic channels it relies on, for which the data is unfortunately still lacking.

5.2 uncertainty and fiscal policy

5.2.1 *Lessons from Economic Modeling*

This section exposes some intuitive reasons why the effects of policy may depend on the degree of economic uncertainty. Interestingly, neither theory nor intuition provide an unequivocal answer.

The role of expectations in equilibrium outcomes is an old idea dating back to Keynes (1936) General Theory. Loosely summarizing, Keynes (1936) states that as perceived uncertainty increases, so does "liquidity preference", which conditions the size of the liquidity traps where he sees government intervention to be more potent. For Keynes, uncertainty begs government intervention. Akin to the idea of liquidity preference, firms may delay investment decisions during uncertain times, even after having observed a positive fiscal spending shock. The reason is that investment is (at least partially) irreversible, and firms may be unwilling to risk suffering the adjustment cost given the uncertain outlook. That is, economic uncertainty induces "wait-and-see" behaviors that could potentially dampen the effects of fiscal policy. Conversely, uncertainty can unveil growth options for firms, and have the opposite effect. By spurring "effective demand" (Keynes, 1936), government spending could potentially increase the value of these options, causing agents to "exercise" the option and invest. Fiscal stimuli can also send mixed signals and entail unclear effects during uncertain times. Indeed, a government spending shock during times of heightened uncertainty may confirm some pessimistic views, in turn producing a decline in consumption and activity. Or it could, as explained before, constitute a positive signal on "effective demand", justifying an increase in investment and output.

These ideas are apparent in a bare-bones macroeconomic model. The simplified model abstracts from private capital accumulation, habit formation in consumption, a fixed cost in production, wage stickiness, and price dispersion, following that of Sims and Wolff (2017). I do not consider formally the firms' optimization problem, taking the price markup as a measure of the overall level of distortion in the economy. This simplified economy can be summarized by the following conditions:

$$\begin{aligned} Y_t &= C_t + G_t + G_{I,t} \\ Y_t &= A_t K_{G,t}^{\varphi} N_t \\ w_t &= \mu_t^{-1} A_t K_{G,t}^{\varphi} \\ U_t &= u(C_t, G_t) - \ell(N_t) \\ \ell_N(N_t) &= u_C(C_t, G_t) w_t \\ K_{G,t+1} &= G_{I,t} + (1 - \delta_G) K_{G,t}. \end{aligned} \tag{5.1}$$

The penultimate condition is the optimality condition for the consumers supply of labor; the last condition is the accumulation dynamic for the government's capital stock (there is no private capital). I assume with little originality that $u_C > 0$, $u_{CC} < 0$, $u_{CG} \geqslant 0$, $\ell_N > 0$ and $\ell_{NN} > 0$. Taking differentials of these conditions in the vincinity of a steady state, the government consumption and investment multipliers can be expressed as follows:

$$\frac{\mathrm{d}Y_t}{\mathrm{d}G_t} = \frac{-u_{CC} + u_{CG}}{\ell_{NN} \frac{\mu}{(AK_G^{\varphi})^2} - u_{CC}} - \frac{u_C}{\ell_{NN} \frac{\mu}{(AK_G^{\varphi})^2} - u_{CC}} \frac{\mathrm{d}\mu_t/\mu}{\mathrm{d}G_t}; \tag{5.2}$$

$$\frac{\mathrm{d}Y_t}{\mathrm{d}G_{I,t}} = \frac{-u_{CC}}{\ell_{NN} \frac{\mu}{(AK_G^{\varphi})^2} - u_{CC}} - \frac{u_C}{\ell_{NN} \frac{\mu}{(AK_G^{\varphi})^2} - u_{CC}} \frac{\mathrm{d}\mu_t/\mu}{\mathrm{d}G_{I,t}}. \tag{5.3}$$

The takeaways from these formulas are that the response of the wage markup is critical.[1] The markup will typically fall after an increase in either type of government expenditure. Furthermore, the "uncertain" part about the response of the markup lies in the productivity shock, and any shift in its underlying variance will drive the sensitivity of the markup upward. The expression of the markup allows to claim that:

$$E\left(\frac{\mathrm{d}\mu_t/\mu}{\mathrm{d}G_{I,t}}\right) \overset{\text{sign}}{\propto} E\left(A_t\right). \tag{5.4}$$

[1]Note that in a frictionless economy, the second term in each equation equals zero and government multipliers are deterministic. In particular, the government investor multiplier is exactly one, that is, government investment is completely neutral.

From the government's standpoint, the *risk* associated in spending is proxied by the variance of the response of the markup to government spending, which with an argument akin to the one I just made is a multiple of the variance of the productivity process. Any increase in "uncertainty" – which is exactly the variance of A_t, at least in Bloom's (2009) terms – leads to higher government multipliers. Note, however, that the closer to the frictionless ideal where $\mu \sim 0$, the less that statement holds true. For a more general discussion around state dependent government multipliers in a medium scale DSGE framework, see Sims and Wolff (2017). Overall, there are many channels through which uncertainty may affect the efficiency of fiscal policy, most of which are likely to interact. Because these channels might well cancel each other out, this thesis takes a first step and empirically assesses whether uncertainty has an actual effect on fiscal multipliers.

5.2.2 *Current Research Frontier*

Keynes (1936) original intuition has been wanting a formal theoretical framework for decades. The importance of uncertainty in economic fluctuations is, however, a much more recent concern in economic modeling. In a partial equilibrium setting, Bloom (2009) showed that "uncertainty shocks" expanded firms' "inaction region", thus generating sharp recessions and recoveries. Christiano et al. (2014) assessed the importance of risk shocks in a dynamic stochastic general equilibrium (DSGE) framework, and concluded that "risk shocks" – shocks to the variance of capital returns – have the greatest contribution to fluctuations along the business cycle, trumping that of government spending shocks. Christiano et al. (2014) do not discuss, however, whether the impact of fiscal policy is mitigated, or augmented, in periods of higher "risk." Taschereau-Dumouchel et al. (2013) describe an economy in which multiple steady states of uncertainty exist. In particular, the economy endogenously creates "uncertainty traps", which are episodes of long-lasting recessions, where high uncertainty meets low activity, even though fundamentals are strong. While the authors do not address economic policy's potential recourse to exit such inefficient situations, Taschereau-Dumouchel et al. (2013) nonetheless mention that optimal policy interventions can be hamstrung by economic uncertainty, hinting at lower multipliers in these eras.[2] Finally, Saijo (2013) considers feedback effects: an endogenous rise in aggregate uncertainty hinders economic activity, generating more uncertainty in return, and so forth. None is said on fiscal stimulus, and whether it may or may not offset the "uncertainty multiplier."

Beyond the still unsettled role of uncertainty in generating fluctuations, there is nary a consensus on how it affects the efficiency of government intervention. Indeed, Keynes (1936)' intuition for "animal spirits" as a cause of economic fluctuations is still a rather controversial point in empirical work (Barsky and Sims, 2012). As regards fiscal multipliers, the empirical literature has extensively studied how they varied over the business cycle; yet there are few assessments of the efficiency of fiscal policy when the economy is more "uncertain." Without having created a proper consensus, the data seem to have confirmed the Keynesian intuition (Keynes, 1936) that government spending is more potent in dire times. For instance, fiscal multipliers have been found to be higher in liquidity traps (Farhi and Werning, 2012). Using data from the OECD countries, Auerbach and Gorodnichenko (2011) find evidence of larger multipliers during recessions, mirroring smaller multipliers in expansions. In a more general setting, Fazzari et al. (2013) document how fiscal policy can have non-linear effects in certain states of the business cycle; they find that fiscal multipliers are substantially higher when there is considerable "economic slack" (unused resources in the economy). Ramey and Zubairy (2014) investigate whether U.S. government spending multipliers differ according to the amount of slack

[2] Their definition of uncertainty – as the dispersion of the subjective beliefs about the fundamental – differs from that of the common uncertainty-driven business cycle literature. A key difference is that there can be uncertainty without volatility.

and whether interest rates are near the zero lower bound, using a new, extended quarterly dataset. Using a different estimation technique than that of Fazzari et al. (2013), Ramey and Zubairy (2014) find that the magnitude of multipliers does not seem to depend on the amount of slack in the economy, while the proximity to the zero lower bound has an unclear effect. (Note that their conclusions radically discord with those of Fazzari et al. (2013) on the state-dependence of multipliers.) With all that research in mind, there is still no answer provided to the question of fiscal multipliers in the face of higher economic uncertainty.

This chapter's question is closest to that of Bachmann and Sims (2012), who look at consumer and business confidence as a potential channel for fiscal policy to effect. Bachmann and Sims find that confidence rises following an increase in spending during periods of economic slack, rendering multipliers much larger then. The systematic response of confidence is irrelevant for the output multiplier during normal times, but critical in downturns. In unpublished work, Alloza (2014) reproduces the exercise of Bachmann and Sims, but tries to compare periods of high uncertainty – measured by indices of consumer confidence and financial volatility – to periods of low economic activity. Alloza concludes that multipliers are highest in periods of low uncertainty and expansion, contrary to the conclusions reached by a large strand the literature on state-dependent multipliers. In front of these existing papers, my approach stands out in at least three ways. First, I employ long term quarterly data over the 20th century. Second, the econometric methodology retained imposes little restrictions on the behavior of impulse response functions, contrary to the (structural) VAR methods used by the papers mentioned. Third, I allow for smoother transitions between states of high and low uncertainty, rather than arbitrarily splitting the sample into seasons of high and low uncertainty. Finally, in the state transitions, I factor in that uncertainty tends to correlate with recessions (Bloom, 2014) and propose a way to adjust for the state of the business cycle directly in the transition function. This affords me to substantially refine the conclusions previously reached by the literature.

5.3 METHODOLOGY

5.3.1 *Data*

Historical, quarterly macroeconomic data for the United States on GDP, the GDP deflator, government spending, population, and the unemployment rate, come from various sources described in the Appendix.[3] In addition, I use Ramey (2011) "news" variable reflecting changes in the expected present value of government spending in response to military events. The sample ranges from the first quarter of 1890 to the last quarter of 2010. To these data items, I append constructed measures of uncertainty. These include Baker et al. (2012) news-based economic policy uncertainty index (available from 1900), Jurado et al. (2015) composite uncertainty measure (available from the second quarter of 1960), the Consumer confidence index from the Michigan Survey of Consumers, and a proxy for quarterly stock market volatility based on stock prices collected by the Yale NYSE Research Project. The resulting series (off-trend deviations) is exposed in Figure 5.1. In the regressions, all aggregate variables are expressed in real terms and transformed to logarithms.

5.3.2 *Government Multipliers and Smooth Transition Local Projections*

The literature has considered several methods to estimate state-dependent effects, but without agreeing on one best way to do so. Auerbach and Gorodnichenko (2012) consider a "smooth transition" vector autoregressive model (STVAR) to compare the effects of fiscal policy in recessions and expansions in the United States. In an ensuing exercise, Auerbach and Gorodnichenko (2013) extend their study to a broader array of

[3]The extended macroeconomic data is essentially the same as that of HERE for the United States

OECD countries, this time using direct projection methods following Jordà (2005) seminal paper .

I opt for Jordà (2005) local projection method to estimate impulse response functions for a number of reasons. Using direct projections rather than the SVAR approach to estimate multipliers affords us to relax the orthogonality assumptions imposed by the SVAR method, which ultimately constrain the shape of the impulse response functions, especially in the long run. Because causality is still a widely open question on the subject, not having to posit any form of relationship between uncertainty and other sorts of shocks is, in a sense, preferable. Additionally, direct projections do not constrain how long the economy remains in a given state and whether the shock entails leaving the state. As we will see later on, the method obviously carries a number of minor shortcomings. Finally, a good share of the literature has used that method in the context of transition models; using local projections therefore makes more sense for the sake of comparison.

I also use the transformation recommended by Hall (2009) and Barro and Redlick (2011) to steer clear of issues related to units of measurement:

$$\frac{X_{t+h} - X_{t-1}}{Y_{t-1}} \simeq (\ln X_{t+h} - \ln X_{t-1}) \cdot \frac{X_{t-1}}{Y_{t-1}}; \tag{5.5}$$

where X denotes the aggregate variable whose reaction to government spending shocks are assessed. Formally, the state-dependent impulse responses are estimated with sequential regressions of the form:

$$\begin{aligned} z_{t+h} = \text{Quartic Trend} + I_{t-1}(\alpha_{\text{HU},h} + \Phi_{\text{HU},h}(L)y_{t-1} + \beta_{\text{HU},h}\eta_t) \\ + (1 - I_{t-1})(\alpha_{\text{LU},h} + \Phi_{\text{LU},h}(L)y_{t-1} + \beta_{\text{LU},h}\eta_t) + \varepsilon_{t+h}; \end{aligned} \tag{5.6}$$

where I_t is a variable that retains information on the state (here the economy being in "high uncertainty"); $\Phi_{\text{HU},h}$ and $\Phi_{\text{LU},h}$ are quartic lag polynomials, and ε_{t+h} is an error term. The responses of the dependent variable z_{t+h} h periods after the shock η are given by $\beta_{\text{HU},h}$

in the high-uncertainty state and $\beta_{\mathrm{LU},h}$ in the low-uncertainty state. In the sequential regressions, the standard errors are corrected for serial correlation using Newey and West (1987) autocorrelation robust estimator.

5.3.3 *Conditional Smooth Transition Models*

Smooth transition autoregressive (STAR) models were first proposed by Terasvirta and Anderson (1992) as a combination of the self-exciting threshold autoregressive (SETAR) and the exponential autoregressive (EAR) models. The smooth transitions allow the model to be interpreted in two ways. On the one hand, STAR models are similar to Markov switching models but allow for smoother, more realistic transitions. On the other hand, STAR models can be interpreted as multi-state models with a continuum of states Terasvirta (1996)

The main advantage in favor of STAR models is that changes in economic aggregates are influenced by changes in the behavior of many different agents and it is highly unlikely that all agents react simultaneously to a given economic signal. In financial markets, for example, with a large number of investors, each switching at different times (probably due to heterogeneous objectives), a smooth transition or a continuum of states between the extremes appears more realistic.

On the other hand, smooth transition models do suffer from a significant "design flaw" that applies to most regime switching models where the state transition is estimated with a single observable. Indeed, in most cases, the state transition variable is merely a proxy for the true state. This does not violate the underlying assumptions of a typical smooth transition model, but it does make the conclusion one draws from it spurious at best. For a more formal example, consider the following simplified situation. The two processes X and Y evolve structurally as follows:

$$\begin{aligned} X_t &= a_1 X_{t-1} + a_2 Y_{t-1} + \varepsilon_{X,t} \\ Y_t &= a_3 Y_{t-1} + a_4 X_{t-1} + \varepsilon_{Y,t}; \end{aligned} \tag{5.7}$$

where $\varepsilon_{X,t}$ and $\varepsilon_{Y,t}$ are uncorrelated shocks, exogenous to X and Y. X_t can be rewritten as follows:

$$X_t = a_1 X_{t-1} + \sum_{k=2}^{\infty} a_2 (a_3)^{k-2} a_4 X_{t-k} + a_2 \sum_{k=1}^{\infty} (a_3)^{t-k} \varepsilon_{Y,t-k} + \varepsilon_{X,t}; \tag{5.8}$$

which entails that X depends on past shocks to Y. In the context of (smooth) state transitions, this implies that the transition to a new state in the variable X may be the consequence of past shocks to the variable Y; therefore the transition function should control for the level of Y. Note that I have still assumed that shocks to X and Y were uncorrelated; in the context of economics and business cycles that assumption is likely to hold false for any representation that is estimated, even if the theoretical existence of such shocks gives support to that enterprise. In the case of uncertainty, whichever indicator is retained is very likely to be affected by other economic trends at play, such as recessions, financial market crashes, wars, etc. While these phenomena do correlate with uncertainty, they *are not* economic uncertainty as I have described it in the previous chapters or as the fathers of economic uncertainty theory would have defined it.

It is to this purpose that I propose an adjustment to the typical logistic STAR model to account for mutual influences among state variables. Consider the typical LSTAR model equation:

$$y_t = \mathcal{I}(s_t, \gamma)\Phi_1(L) y_{t-1} + (1 - \mathcal{I}(s_t, \gamma))\Phi_2(L) y_{t-1} + \varepsilon_t; \tag{5.9}$$

where $\mathcal{I}(s_t, \gamma) = 1 - (1 + \mathrm{e}^{-\gamma s_t})^{-1}$. The logistic state transition factor does not, in its current form, account for the potential mutual influence business cycle indicators. I propose an adjustment based on a result established in the Appendix of this chapter. Let $x = (x_0, x_1, \ldots, x_n)$ be a multidimensional logistic vector with parameter $\gamma = (\gamma_0, \gamma_1, \ldots, \gamma_n)$ such that the first line of $\Sigma = E(xx') = (\sigma_{i,j})$ contains at least one non-zero value other than $\sigma_{1,1}$. Let $x_{-0} = (x_1, \ldots, x_n)$. Then:

i. x_{-0} is a multidimensional logistic vector;

ii. The density of x is given by:

$$f(x) = \frac{(n+1)! \left(\prod_{\gamma_i} \gamma_i\right) \left(\mathrm{e}^{\sum_{x_i} \mathrm{e}^{-\gamma_i x_i}}\right)}{\left(1 + \sum_{x_i \in \{x_0, x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+2}}; \tag{5.10}$$

iii. The density of x_0 conditional on x_{-0} is given by:

$$f(x_0|x_{-0}) = \frac{\gamma_0 (n+1) \mathrm{e}^{-\gamma_0 x_0} \left(1 + \sum_{x_i \in \{x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+1}}{\left(1 + \sum_{x_i \in \{x_0, x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+2}}; \tag{5.11}$$

iv. The transition factor conditional on x_{-0} is given by:

$$\mathcal{T}^*(x_0|x_{-0}) = 1 - \frac{\left(1 + \sum_{x_i \in \{x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+1}}{\left(1 + \sum_{x_i \in \{x_0, x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+1}}. \tag{5.12}$$

This result allows me to build a "conditional" smooth transition model where the positive mutual influence between different business cycle aggregates is factored directly in the transition factor. To be sure, the model can be written down as follows:

$$y_t = \mathcal{T}^*(x_{0,t}, \gamma | x_{-0,t}) \Phi_1(L) y_{t-1} + (1 - \mathcal{T}^*(x_{0,t}, \gamma | x_{-0,t})) \Phi_2(L) y_{t-1} + \varepsilon_t; \tag{5.13}$$

where:

$$\mathcal{T}^*(x_{0,t}, \gamma | x_{-0,t}) = \frac{\left(1 + \sum_{x_i \in \{x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+1}}{\left(1 + \sum_{x_i \in \{x_0, x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+1}}. \tag{5.14}$$

I call $\mathcal{I}^*$ the *conditional* smooth transition factor, which accounts for the level of other covariates in the state transition function, and summarize the properties of the conditional model as follows:

i. The conditional factor $\mathcal{I}_t^*$ tends to one, that is, to the high state, if and only if the underlying index of interest x_0 tends to $+\infty$, i.e. the high state
ii. If the underlying index of interest x_0 tends to $-\infty$, i.e. the low state, the conditional factor $\mathcal{I}_t^*$ tends to zero, that is to the low state. It does so too when one of the other covariates tends to the low state.
iii. If $\{\varepsilon\}$ is a martingale difference sequence, local projection estimates of the impulse response functions are consistent (and robust to misspecification).

The limits of the new conditional term imply that the new factor does not over-weigh uncertainty and fits what one would like to have. First, if the state of interest goes to infinity (minus infinity), the indicator goes to 1 (to 0). Additionally, if all other covariates tend to infinity (that is, other metrics converge to high states), the indicator is still "weighed down" by the level of the state of interest. If all other state covariates tend to the low state, the indicator tends to that of the low state too, as one would expect under a careful choice of the covariates. On the other hand, the unconditional term would have converged towards the low state in the same way, so the conditional version is at worst a non-improvement – not a distortion from the previous method.

In contrast to the scant literature considering the effects of fiscal policy under uncertainty, I do not specify the state I_t factor as a step function and follow the procedure that I have described earlier in this section. It allows me to eschew the arbitrary classification of states with high or low uncertainty. Alloza (2014) considers states in which uncertainty is 1.65 standard deviations off its long term mean. My state transition variable of interest is s_t which denotes (normalized) deviations from the trend for a given measure of uncertainty.[4] I prefer this specification – the logistic STAR (LSTAR) model – to their exponential counterpart (ESTAR), where the transition function is given by $\mathcal{I}_t = 1 - \exp(-\mu s_t^2)$. It can readily be seen that this specification allows for "symmetric" states, where highs and lows are treated equally. In my case and with the data series that I retain, it seems natural that high uncertainty and low uncertainty should have different effects. Furthermore, it allows me to apply the previous lemma and improve upon the typical "unconditional" model and factor comovements of business cycles and uncertainty. To account for the fact that higher uncertainty usually comes with recessions (without assuming any causal direction), I specify $\mathcal{I}_t$ as a conditional probability as follows:

$$\mathcal{I}_t^* \equiv 1 - \frac{(1 + \mathrm{e}^{-\gamma g_t})^2}{(1 + \mathrm{e}^{-\mu s_t} + \mathrm{e}^{-\gamma g_t})^2}; \tag{5.15}$$

where g_t denotes a seven-quarter moving average of quarterly real GDP growth, as in Auerbach and Gorodnichenko (2012). The challenge of such specification – and also of the unconditional one – lies in the calibration of the parameters γ and μ, which is done to match business cycles facts (probability of recessions and observed high uncertainty periods.[5] This new specification factors in that uncertainty tends to rise during recessions, in an attempt to avoid observing spurious results. $\mathcal{I}_t^*$ is close to one in when uncertainty is high, but corrects for the fact that the economy might be in a recession and "penalizes" high uncertainty states where the economy is at the same time undergoing a recession – which is documented to affect the size of the multipliers. The deeper the recession for a given level of uncertainty, the lower the weight on that state. By comparing the results between the specifications with $\mathcal{I}_t$ and $\mathcal{I}_t^*$, it will then be clear whether uncertainty alone is justifies

[4] I transformed the series of deviations from the trend such that a large negative value indicates a reading of uncertainty high above the trend.

[5] $\gamma = 10$ and μ set to a similar value in the uncertainty transition function seem to work quite well.

government intervention, or if the effect appears to be more linked to recessions than uncertainty per se.

5.4 do uncertain times conjure higher multipliers?

5.4.1 *Estimated Responses to Government Spending Shocks*

For illustration purposes, I first present the responses to a government spending shock in a linear, non state-contingent setting. The results are exposed in Figure 5.3. In all the different specifications, I use the constructed historical stock market volatility as a proxy for economic uncertainty. The series of government spending shocks is that of Ramey (2011) so as to consider admittedly exogenous shocks. Other proxies will be considered in section 5.4.3. In short, basic linear models point towards positive effects of government spending, peaking at around 12 quarters after the shock. The estimated peak multiplier comes out at approximately 0.9.

Escaping from the linear world, I first run through a state-dependent estimation using I_t – the unconditional transition function – to distinguish between uncertain and tranquil times. The results of the estimation are presented in Figure 5.4. The somewhat oscillating impulse response functions are an unfortunate side effect of local projection methods. The results are however clear. In uncertain times, government spending appears to have a positive effect on output, fading to zero after roughly 4 years. On the flip side, these effects are hardly significant in tranquil times. The effects of government spending shock under different levels of uncertainty are no different than in confident times. Moreover, the effects of government spending shocks are essentially nil in both states.

Finally, to account for the fact that uncertainty rises during recessions, I use a conditional version of the probability of switching to a higher uncertainty state, as defined by I_t^* in the previous section. The estimated impulse response functions are presented in Figure 5.5. Interestingly, the estimated responses are very similar to those obtained in the unconditional model, although somehow a bit smoother in the low uncertainty state with conditional transition function. While the conditional transition function is not the nail in the coffin of spuriousness, it does comfort the conclusions reached in the unconditional model. The fact that the results are very similar to those of the unconditional model comes from the notable similarity between the unconditional and the conditional transitions, which can be observed in Figure 5.6. Note, however, that the conditional version accentuates non-contractionary periods, in particular that of the 1930s. The conditional transition strongly adjusts before the unconditional transition, then kicks back when uncertainty plateaus even though the economy picks up.

Besides that of the similarity between the two measures, another concern that I had was that despite the adjustment, I was still capturing the effects of recessions in these estimates. An informal comparison between the conditional transition factor and that of Auerbach and Gorodnichenko (2012) in Figure 5.6 shows that this is far from the truth. To comfort this fact, I replicated the Auerbach and Gorodnichenko (2012) results in the Appendix, using local projection methods for comparability. Multipliers appear higher in recessions, but the estimated responses have very different shapes. Interestingly, the differences between I_t^* and the Auerbach-Gorodnichenko state transition go both ways – although $I_t^* > I_t$ does appear more often than its contrapositive. This supports the idea that the effect estimated here is genuinely different.

5.4.2 *Fiscal Multipliers in Uncertain Times*

I now turn to actual fiscal multipliers. For state-dependent models, multipliers are derived from the estimated coefficients from both states, for GDP and government spending. I compute multipliers over three horizons: as the cumulative responses through two years, four years, and at the peaks of each response. Formally, the cumulative multiplier at horizon H $M_Y(H)$, and the peak multiplier $M_Y^{\infty}(H)$, are defined by:

$$M_Y(H) = \frac{\sum_{h=1}^{H} \beta_{\mathrm{S},h}^{Y}}{\sum_{h=1}^{H} \beta_{\mathrm{S},h}^{G}}; \tag{5.16}$$

$$M_Y^{\infty}(H) = \max_{h \leqslant H} \left\{ \frac{\beta_{\mathrm{S},h}^{Y}}{\beta_{\mathrm{S},h}^{G}} \right\}. \tag{5.17}$$

Multipliers derived from these formulas are presented in table 5.1. The comparative interpretations from the impulse response functions hold. An interesting observation is that the linear model appears to underestimate multipliers. Surprisingly, multipliers are much higher when I distinguish between low and high uncertainty levels. The estimated levels are consistent with the conclusions from the previous sections. Albeit they appear surprisingly high, multipliers in tranquil times are likely insignificant.

Table 5.1: Estimated fiscal multipliers for all models

Model	2-year cumul	4-year cumul	4-year peak
Linear	0.78	0.87	0.92
High σ (unconditional)	1.68	2.04	1.36
Low σ (unconditional)	0.86	1.25	0.93
High σ (conditional)	1.79	2.81	1.37
Low σ (conditional)	0.90	1.33	0.96
In recession (A-G)	0.78	0.87	0.93
In expansion (A-G)	0.90	0.56	n.m

Notes: Cumulative multipliers calculated according to the formula in equation (5.16), and peak multipliers according to the formula in (5.17). "n.m." stands for "not meaningful", which can happen when the quantities aren't statistically significant (for instance a denominator extremely close to zero, pushing the value of the multiplier). "A-G" stands for my replication of Auerbach and Gorodnichenko (2013) exercise.

My conclusions differ in some aspects with the existing literature. Alloza (2014) found that fiscal policy was more efficient during times of lower uncertainty and in expansions, contrary to what most of the empirical literature has uncovered. To account for his results, Alloza suggests that precautionary savings effects outweigh that of the hike in anticipated demand. Pessimism, which reportedly increases in response to government spending shocks, annihilates the positive effects of fiscal stimuli. A first remark is that Alloza's question is somehow misguided, as recessions tend to coincide with high uncertainty (whatever causal direction one is willing to assume). An crucial precaution to take on this question is to try to partial out the effects of economic recessions on multipliers before asking whether uncertainty affects the efficiency of fiscal policy. Here, I did not assume any more than the fact that uncertainty and recessions were a joint "random" process, and I derived a conditional version of the transition function. Another concern with Alloza (2014) study is that the state switching process I_t for expansions and recessions is ultimately as precise as the NBER's business cycle dating; while transitions between high and low levels of uncertainty are very much arbitrarily decided. In particular, the effects of fiscal policy under both regimes could be estimated with little precision if fiscal policy shocks are observed close to a change of regime that could be "misdated" (be it for recessions or uncertainty spikes). Although this argument does not really apply to state indicators based on the data, I believe that the current state of knowledge on uncertainty does not afford such a clear cut distinction.

5.4.3 *Robustness Checks*

I first evaluate the robustness of results to different versions of the data. Until now, the estimation has been carried out with the market volatility index constructed from the Yale NYSE History Project. I get very similar results by substituting Baker et al. (2012) measure of uncertainty to my initial proxy. The reason why I did not use this measure in the first place is that it essentially is a metric of policy uncertainty rather than aggregate economic uncertainty, for which stock market volatility is considered a more appropriate instrument (Bloom, 2014). In the same vein, I consider different forms of uncertainty detrending. A Hodrick-Prescott detrending method yields almost identical results in both the unconditional and conditional models. I keep quartic detrending for simplicity. Because they start much later than my main proxy for uncertainty, I cannot run the estimation with series from the Michigan Survey of Consumers nor with Jurado et al. (2015) data.

I also compare my results with a different series of government spending shocks. With the same data, I reproduce the estimation procedure of Blanchard and Perotti (2002), identifying shocks as the residuals in a structural VAR that imposes the exclusion restriction that government spending cannot react within one quarter to shocks to output and tax revenues. The results are qualitatively robust to the new data in the sense that impulse response functions have a very similar shape. The confidence bands, however, appear much wider than before. While the similar shapes of impulse responses are comforting, the broader confidence bands are somehow worrisome. It means that conclusions on the actual size of multipliers, namely those in section 5.4.2, need to be taken with caution.

An important verification to carry out was that of the sample window. Considering only the post-WWII era measurably modifies the results. The shape of the impulse responses in the high uncertainty state is roughly the same, but that of the low uncertainty state is starkly affected. The response of aggregates basically oscillates in that state, which I only partially attribute to local projections. The most likely culprit is that the way uncertainty influenced the transmission of government spending shocks has changed in the sample. The further away we push the beginning of the sample, the less significant the differences between the two states, whatever uncertainty proxy is retained. This puts a clear limit to the conclusions, as it now appears that the results were driven by the earlier part of the sample considered.

On a side note, another important robustness check would be to confront such a study to data from different countries. My conclusions are essentially confined to the United States. Even then, the sensibility to the sample window calls for an investigation of the structural break that generates it, and more importantly, to the elaboration of an estimation technique that is robust to such a change.

5.5 TAKEAWAYS FOR FISCAL MULTIPLIERS

Overall, the response of output (and other aggregates) to fiscal spending shocks appears sharply different in periods of high and low uncertainty. As highlighted by the literature, there could be many explanations for this fact. Nevertheless, the exercise appears highly reliant on the sample window retained. A policy lesson would be that existing proxies of uncertainty are not relevant metrics to consider for the efficiency of government stimuli. Pre-1945, the story would be that the "effective demand" shock may outweigh the negative signal sent by government intervention, or that the irreversibility of investment is generates quantitatively negligible differences from a frictionless equilibrium. Post-1945, there is disappointingly no story to be told, at least at this stage and with the measures of uncertainty available at the moment. Still, a major challenge remains to stop guessing the story and actually find out which effect dominates.

Of the possible extensions of this work, none is more important than understanding the mechanism though which uncertainty might condition the potency of fiscal stimuli, why it likely has changes and whether it may change again or back. Further trying to partial out the effects of recessions and expansions on economic uncertainty in this question, as well as in empirical exercises that involve uncertainty in general, only comes second by a nose. A formal test for non-linearity could give further support to the study. The linear model is indeed nested in the (L)STAR version retained (take $\mu \to \infty$), and a proper calibration of the transition function is paramount to consolidate the conclusions. The fact that there is no perceivable differences between the two regimes gives some support to this objection. Another possible specification would be to try to adapt multiple-regime LSTAR (MR-LSTAR) models to this question. MR-LSTAR models formally allow for more than two states, which, in the case of uncertainty, would consist in distinguishing between normal times, high-uncertainty states ("panic") and low-uncertainty states ("tranquility"). Given the noisiness of the current proxies of uncertainty available, this other version comes last to the other issues to fix here, such as that of endogeneity or spuriousness.

5.6 Conclusion

We began from the Keynesian assertion that the level of uncertainty may affect the efficiency of fiscal policy. From there, I estimated state-dependent responses to government spending shocks, and showed that the data pointed at higher fiscal multipliers in uncertain times. Importantly, these results were obtained by correcting for the fact that uncertainty rises during recessions, an issue that haunts many discussions on the effects of uncertainty on the economy.

Although using century-long series seemed like an improvement on the existing literature, the solidity of the results seemed to largely rely upon the sample window selection. Spending multipliers do appear higher in uncertain times over the whole course over the 20th century, but the data hardly support this conclusion if one considers only post-WWII history. To be sure, this suggests that a structural change occurred before the soaring second half of the century. The Great Moderation is the most likely culprit of that fact.

Rather than disappointing, I see this difficulty as leaving a challenging research agenda to understand the channels through which fiscal spending effects. A thin literature, and to some extent this chapter's exercise, seems to support uncertainty as a determinant of the efficiency of fiscal policy. Because there are countless ways uncertainty could matter, the next step will boil down to deciphering these links and sorting out the meaningful ones. The growing availability of microeconomic data will – without a doubt – make this possible soon enough.

5.7 APPENDIX

DATA SOURCES

GDP and GDP deflator

1947 – 2010: Quarterly data on chain-weighted real GDP, nominal GDP, and GDP deflator from BEA NIPA

1889 – 1946: Annual data from 1929 – 1946 from BEA NIPA

For 1889 – 1928, *Historical Statistics of the United States, Earliest Times to the Present: Millennial Edition*, edited by Susan B. Carter, Scott Sigmund Gartner, Michael R. Haines, Alan L. Olmstead, Richard Sutch, and Gavin Wright. New York: Cambridge University Press, 2006.

1889 – 1938: Quarterly data on real GNP and GNP deflator from Nathan Balke and Robert J. Gordon, The Estimation of Prewar Gross National Product: Methodology and New Evidence, Journal of Political Economy, 97, February 1989. Data available at: `http://www.nber.org/data/abc/`

Uncertainty

1900 - 2010: Bloom, Baker and Davis uncertainty measure, available at `www.policyuncertainty.com`

Properties of Multivariate Logistic Distributions

Multivariate logistic distributions were first introduced by Malik and Abraham (1973). The first part of the lemma announced is simply the definition of the multivariate logistic density function:

$$f(x) = \frac{(n+1)!\left(\prod_{\gamma_i} \gamma_i\right)\left(\mathrm{e}^{-\sum_{x_i} \gamma_i x_i}\right)}{\left(1 + \sum_{x_i \in \{x_0, x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+2}}.$$

It follows that the density of X_0 conditional on all following components is equal to:

$$f_{X_0|X_{-0}}(x) = \frac{f_X(x)}{f_{X_{-0}}(x_{-0})}.$$

The marginal density in the denominator can be readily obtained by integrating over x_0:

$$\int f_X(x)\,\mathrm{d}x_0 = n!\left(\prod_{\gamma_i \neq \gamma_0} \gamma_i\right)\left(\mathrm{e}^{-\sum_{-x_0} \gamma_i x_i}\right)\int \frac{(n+1)\gamma_0 \mathrm{e}^{-\gamma_0 x_0}}{\left(1 + \sum_{x_i \in \{x_0, x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+2}}\,\mathrm{d}x_0$$

The integral term can be computed analytically:

$$\begin{aligned}\int f_X(x)\,\mathrm{d}x_0 &= n!\left(\prod_{\gamma_i \neq \gamma_0} \gamma_i\right)\left(\mathrm{e}^{-\sum_{x_{-0}} \gamma_i x_i}\right)\left[\frac{-1}{\left(1 + \sum_{x_{-0}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+1}}\right]_{-\infty}^{+\infty} \\ &= \frac{n!\left(\prod_{\gamma_i \neq \gamma_0} \gamma_i\right)\left(\mathrm{e}^{-\sum_{x_{-0}} \gamma_i x_i}\right)}{\left(1 + \sum_{x_{-0}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+1}}.\end{aligned}$$

The conditional density now boils down to:

$$\begin{aligned}&\frac{(n+1)!\left(\prod_{\gamma_i} \gamma_i\right)\left(\mathrm{e}^{-\sum_{x_i} \gamma_i x_i}\right)}{\left(1 + \sum_{x_i \in \{x_0, x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+2}} \frac{\left(1 + \sum_{x_{-0}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+1}}{n!\left(\prod_{\gamma_i \neq \gamma_0} \gamma_i\right)\left(\mathrm{e}^{-\sum_{x_{-0}} \gamma_i x_i}\right)} \\ &= \frac{(n+1)\gamma_0 \mathrm{e}^{-\gamma_0 x_0}\left(1 + \sum_{x_{-0}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+1}}{\left(1 + \sum_{x_i \in \{x_0, x_{-0}\}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+2}}.\end{aligned}$$

The cumulative distribution function – i.e the “conditional transition factor” – is obtained by immediate integration for a set value of x_0:

$$F(x_0|X_{-0} = \int_{-\infty}^{x_0} \frac{(n+1)\gamma_0 \mathrm{e}^{-\gamma_0 t}\left(1+\sum_{x_{-0}} \mathrm{e}^{-\gamma_i x_i}\right)^{n+1}}{\left(1+\sum_{x_{-0}} \mathrm{e}^{-\gamma_i x_i} + \mathrm{e}^{-\gamma_0 t}\right)^{n+2}}.$$

Transition Functions for Two Joint Underlying Processes

Assume that the transition to different states is determined by a joint process (X, Y), where no assumption is placed on the dependence between X and Y. Let F_X and F_Y denote the transition functions for X and Y, respectively:

$$F_X(x) = \frac{\mathrm{e}^{-\mu x}}{1+\mathrm{e}^{-\mu x}}; \qquad F_Y(x) = \frac{\mathrm{e}^{-\gamma y}}{1+\mathrm{e}^{-\gamma y}}.$$

Those can be seen as survival functions for two logistic distributions. The two marginal distributions of (X, Y) are logistic, making (X, Y) a bivariate logistic distribution, with marginal densities f_x and f_Y defined by:

$$f_X(x) = \frac{\mu \mathrm{e}^{-\mu x}}{(1+\mathrm{e}^{-\mu x})^2}; \qquad f_Y(x) = \frac{\gamma \mathrm{e}^{-\gamma y}}{(1+\mathrm{e}^{-\gamma y})^2}.$$

It can be shown that the joint density and cumulative distribution functions are given by:

$$f_{(X,Y)}(x,y) = \frac{2\mu\gamma \mathrm{e}^{-\mu x - \gamma y}}{(1+\mathrm{e}^{-\mu x}+\mathrm{e}^{-\gamma y})^3}; \quad F_{(X,Y)}(x,y) = \frac{1}{1+\mathrm{e}^{-\mu x}+\mathrm{e}^{-\gamma y}}.$$

In this context, the conditional density of X can be computed as follows:

$$\begin{aligned} f_{X\,|\,Y}(x\,|\,y) &= \frac{f_{(X,Y)}(x,y)}{f_Y(y)} \\ &= \frac{2\mu(1+\mathrm{e}^{-\gamma y})^2 \mathrm{e}^{-\mu x}}{(1+\mathrm{e}^{-\gamma y}+\mathrm{e}^{-\mu x})^3}; \end{aligned} \tag{5.18}$$

implying the following form of the conditional cumulative distribution function:

$$F_{X \mid Y}(x \mid y) = \frac{(1+\mathrm{e}^{-\gamma y})^2}{(1+\mathrm{e}^{-\gamma y}+\mathrm{e}^{-\mu x})^2}; \tag{5.19}$$

with survival function:

$$S_{X \mid Y}(x \mid y) = 1 - \frac{(1+\mathrm{e}^{-\gamma y})^2}{(1+\mathrm{e}^{-\gamma y}+\mathrm{e}^{-\mu x})^2}. \tag{5.20}$$

RESULT: *With loose notation,*

$$\lim_{x \to -\infty} S_{X \mid Y}(x \mid y) = 1 \tag{5.21}$$

$$\lim_{y \to -\infty} S_{X \mid Y}(x \mid y) = 0 \tag{5.22}$$

Proof – (5.21) simply comes from $S_{X \mid Y}$ being a survival function. For (5.22), remark that the expression can be factored as follows:

$$S_{X \mid Y}(x \mid y) = 1 - \frac{1}{1 + \frac{2\mathrm{e}^{-\mu x}}{(1+\mathrm{e}^{-\gamma y})} + \frac{\mathrm{e}^{-2\mu x}}{(1+\mathrm{e}^{-\gamma y})^2}};$$

which goes to zero as y goes to $-\infty$.

FISCAL MULTIPLIERS IN RECESSIONS AND EXPANSIONS

I reproduced Auerbach and Gorodnichenko's exercise and compared the responses of GDP, government spending and tax receipts in recessions and in expansions.

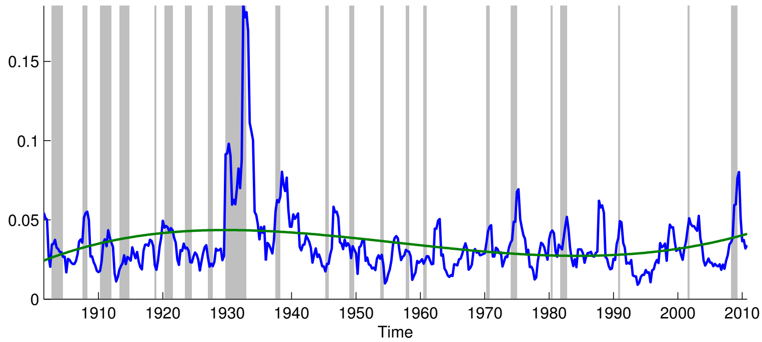

(a) Trailing twelve month stock return volatility

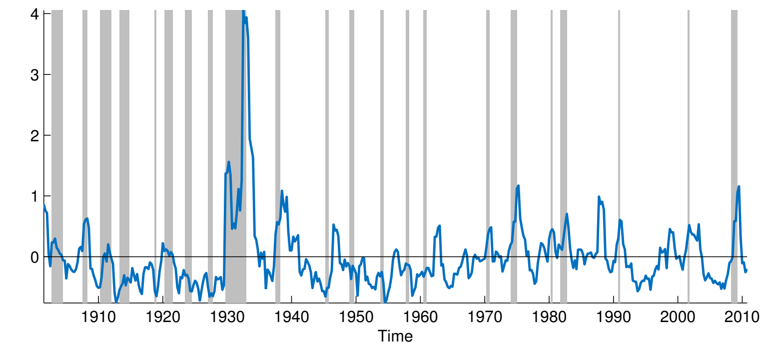

(b) Normalized deviations from the quartic trend in uncertainty

Figure 5.1: Historical data for uncertainty through the 20th century business cycles

Notes: The shaded gray bands denote NBER-dated recessions. The upper graph shows the uncertainty data over the 20th century, computed as the trailing-twelve month market volatility for each quarter. The lower graph displays deviations from the quartic trend. A reading high above zero indicates high economic uncertainty. (For ease and consistency of exposition in the transition function, I take the negative of this measure.)

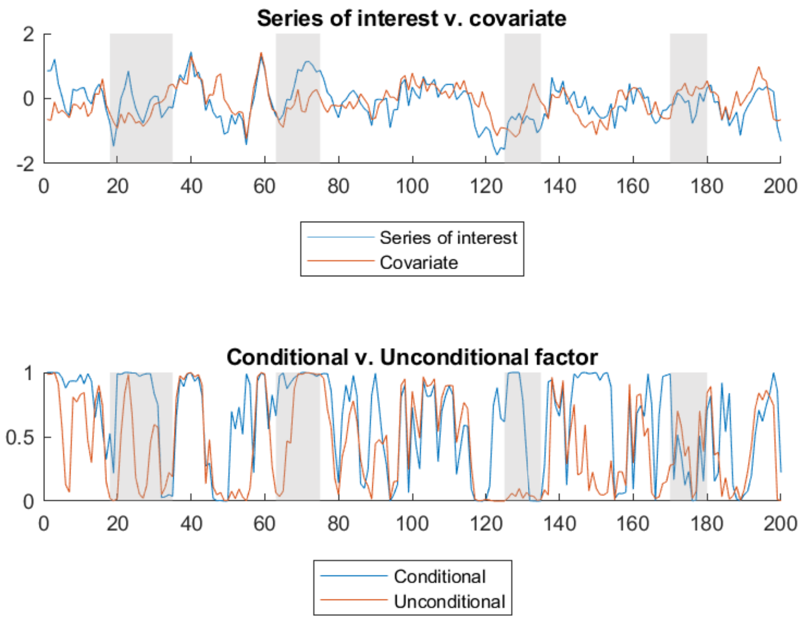

FIGURE 5.2: Comparison of the unconditional and conditional factors

Notes: The upper graph depicts two artificial time series. The series of interest is correlated ($\rho = 0.5$ to the second series) but still subject to an idiosyncratic shock. This is problematic when considering typical smooth transition models, as they are impossible to partial out the changes in the underlying covariate in the variable of interest. The second graph compares the usual unconditional factor (the orange line) to the conditional adjustment that I suggest. Some sample periods showing the importance of this distinction are highlighted. In the first period, the variable of interest persists and shoots up while the covariate goes down. The usual unconditional factor goes down to the low state, which is erroneous since the variable of interest keeps increasing to the high state. In the second period, the unconditional factor is "late" in identifying the high state. In the third period, the unconditional factor completely misrepresents what happens to the variable of interest: It stays in the low state while the variable of interest is clearly picking up, then goes up only when the covariate goes up and there is a slight fall in the variable of interest. In the fourth and last sample period, the unconditional factor remains higher than the conditional factor while the variable of interest is clearly headed towards the low state.

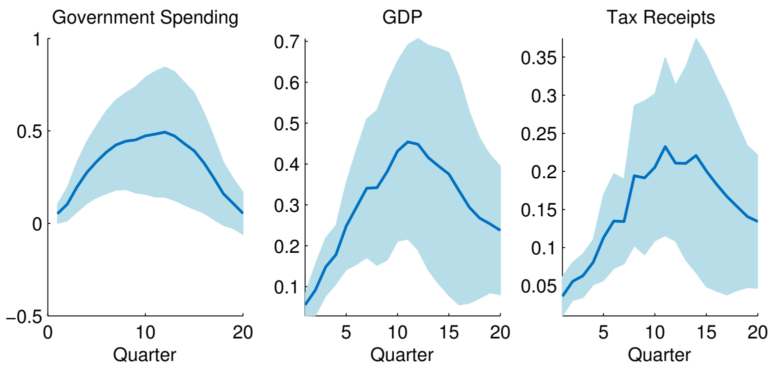

Figure 5.3: Impulse response in the linear model

Notes: Solid lines denote the estimated responses to a government shock; shaded areas around the lines represent 95 percent confidence bands (using Newey-West corrected standard errors).

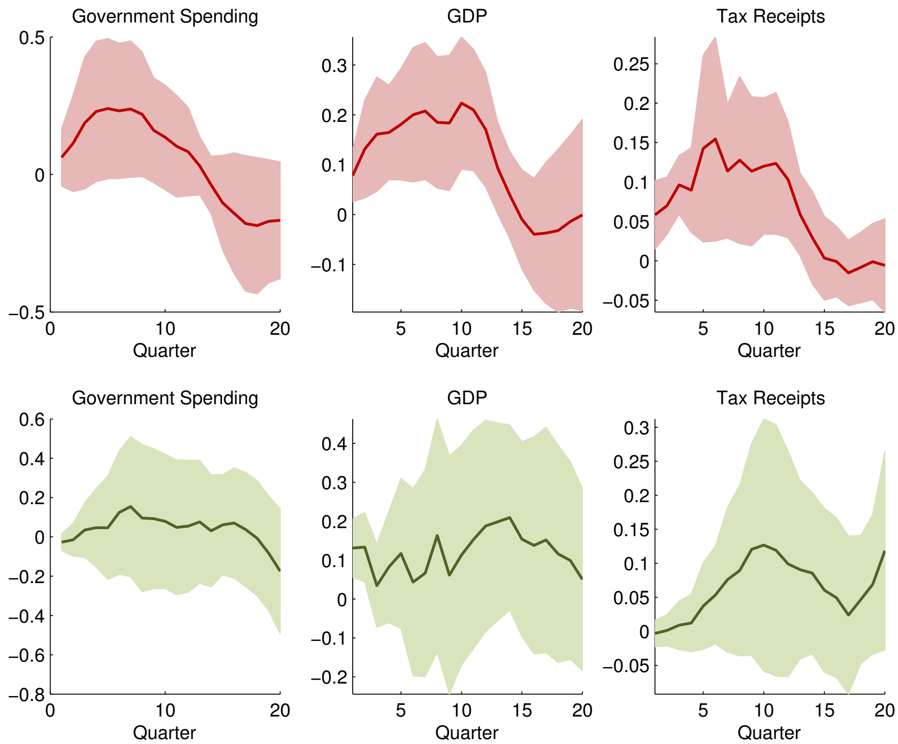

FIGURE 5.4: Impulse response in the unconditional, state dependent model

Notes: Solid lines within red areas denote the estimated responses to a government shock when uncertainty is high, while they denote the estimated responses for lower uncertainty levels within green areas. Shaded areas around the lines represent 95 percent confidence bands (using Newey-West corrected standard errors).

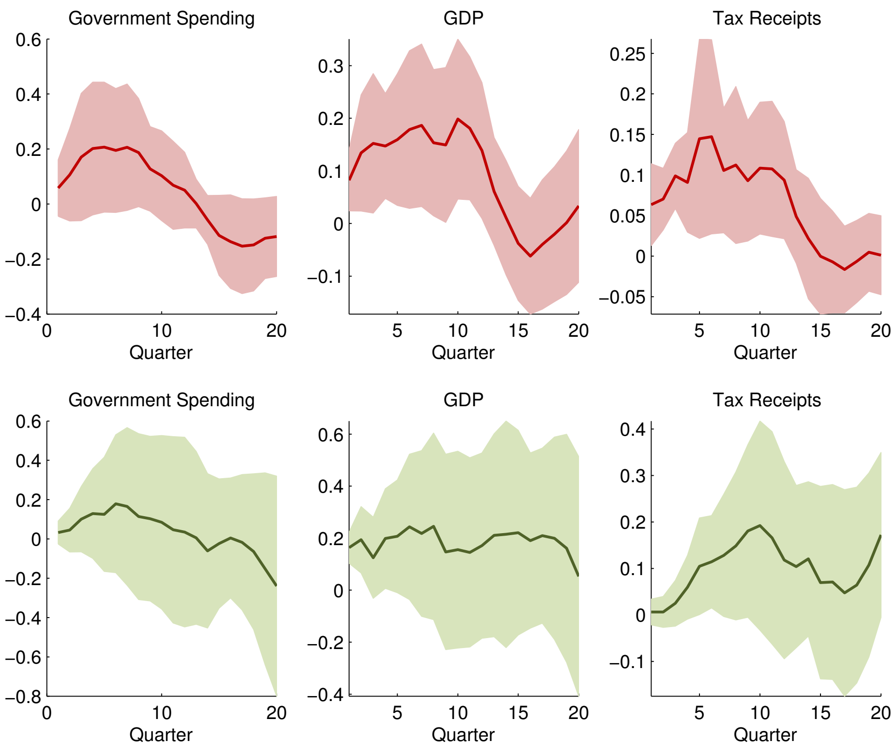

Figure 5.5: Impulse response in the conditional state-dependent model

Notes: Solid lines within red areas denote the estimated responses to a government shock when uncertainty is high, while they denote the estimated responses for lower uncertainty levels within green areas. Shaded areas around the lines represent 95 percent confidence bands (using Newey-West corrected standard errors).

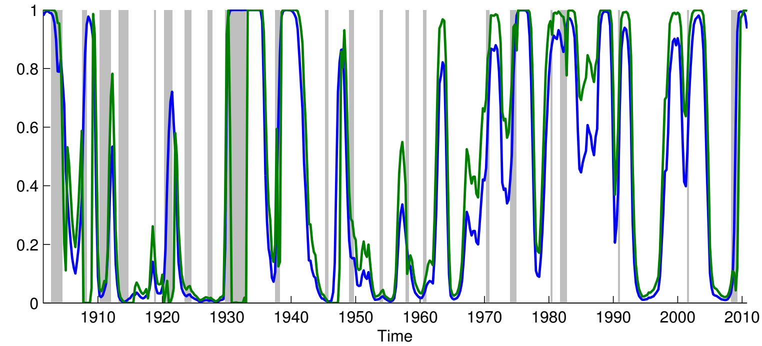

(a) Unconditional (blue line) v. conditional (green line) transitions

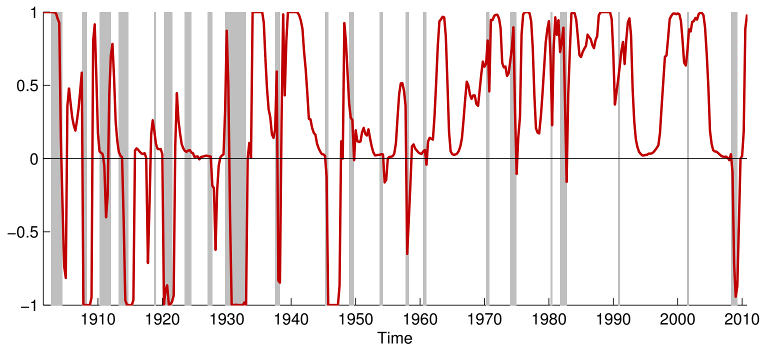

(b) Difference between conditional and Auerbach-Gorodnichenko transitions

Figure 5.6: Transition functions throughout recessions

Notes: In both graphs, the shaded gray bands denote NBER-dated recessions. The upper graph compares the unconditional transition function I_t (in blue) against the conditional version I_t^* (in green). The lower chart plots the difference between the conditional transition function I_t^* and Auerbach and Gorodnichenko (2012) transition ($F(z_t)$ in the original article.) In the bottom graph, periods in which the two measures coincide correspond to those when the line hits the zero axis.

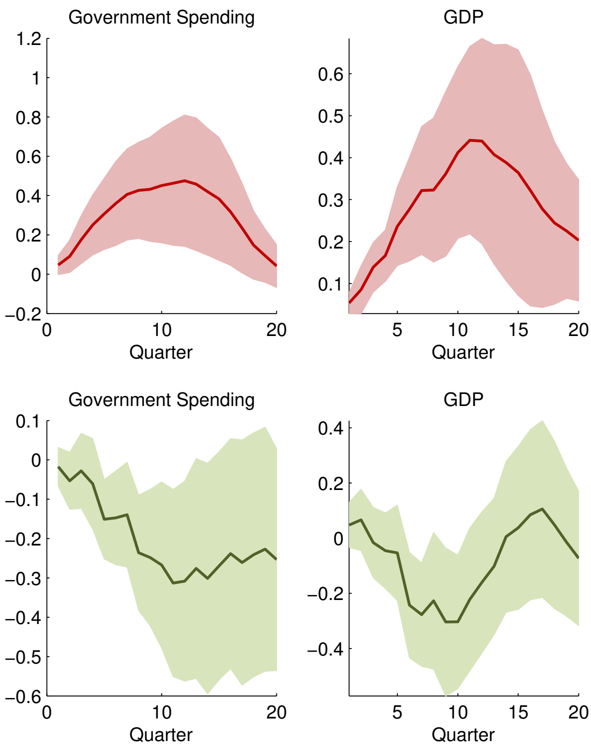

Figure 5.7: Impulse response in recessions and expansions

Notes: Solid lines within red areas denote the estimated responses to a government shock in recessions, while within green areas they denote the estimated responses in expansions. Shaded areas around the lines represent 95 percent confidence bands (using Newey-West corrected standard errors). These impulse response functions are estimated using the Jordà (2005) method, following Auerbach and Gorodnichenko (2013)

6

CONCLUDING REMARKS

Economics has progressively accepted uncertainty. From trying to abstract from it and attributing it to "sunspots" like William Stanley Jevons, economists started modeling it in a predictable way with rational expectations; only recently was uncertainty brought back to the research agenda and made an interesting object on its own. And there is a legitimate hope that the bounty of data that progresses of information technology and growing "connectedness" will eventually bridge economics and other disciplines such as psychology or neuroscience, which will open the door to a much more profound understanding of uncertainty faced by the *homo œconomicus*.

Until then, studying uncertainty means making the most of what economic data has to offer. Backed by historical definitions of uncertainty, this dissertation strongly hinges on the premise that probabilistic data would provide the most accurate approximation of true uncertainty – what is now better known as *Knightian* uncertainty. I have shown that applying the typical "forecast error" approach to probabilistic data allowed to isolate Knightian and non-Knightian uncertainty components that empirically behaved in accordance to what a macroeconomic model featuring Knightian uncertainty would lead to believe. I then turned to a theory of uncertainty that originated outside of economics – Claude Shannon's information theory – and shown how it not only applied to the study of economic uncertainty, but also how it could match the findings that I had uncovered in the fist chapter. Furthermore, the tools provided by Shannon's mathematical theory of uncertainty, and entropy above all else, do apply to a broader range of uncertainty situations, be it labor market uncertainty, security, meteorology, or pretty much any topic as long as the probabilistic forecast data are available. Applying my uncertainty measures also showed that disagreement was not to be left out of the uncertainty landscape, and perhaps with a leap of faith in generalizing, that uncertainty need not be a "representative agent" phenomenon. Because forecasting is at the root of any probability statement, I then turned to study uncertainty from the angle of forecast model uncertainty. Using a tool more commonly used for (optimally) combining predictions – Bayesian model averaging –, I built a measure of forecasting uncertainty and showed how it matched the movements of Knightian uncertainty as I had measured it in the previous chapters. While the scarcity of macroeconomic models featuring forecasting uncertainty calls for avoiding any peremptory conclusion, I was able to match the models' findings with my data. Finally, the ultimate chapter of this dissertation studied the fiscal policy implications of economic uncertainty. Because state dependent models in econometrics are far from robust to specious conclusions, I proposed a way of factoring in the correlation of uncertainty with other covariates to show that fiscal policy would likely have greater effects in uncertain times. Keynes' "animal spirits" hypothesis was insightful, after all. Whether and when uncertainty is a cause or an effect of policy is still an open question to which the current data are unlikely to provide an answer.

While measures of uncertainty like the ones I have presented may very well have little predictive value, they at least do show the extent of our failures in business cycle reversals. Be it for forecast errors, ambiguity or for model uncertainty, crises sadly brings much light on our inability to predict what comes next. In that sense, the uncertainty I tried to measure shows a striking pattern. Perhaps this fact should be remembered when some strive to find ways to "predict crises" or "understand cycles." Because uncertainty is so prevalent one might be tempted to ask the following questions: What does it imply for policy, and can or should it be reduced? Because the latter is loaded with political implications, I have only tried to contribute to the first. But there is no reason why some experiments in the lab should not be conducted to understand how much uncertainty is tolerable, productive, and when there is simply too much uncertainty to go forward. The latest financial crisis showed how uncertainty can, when it reaches unprecedented levels, completely paralyze the economy. There is likely a fine line to walk between letting agents be uncertain enough and letting them be too uncertain, and delineating this limit is to me one of potentially most insightful research agenda that stands unexplored.

Studying uncertainty does not entail trying to eschew it. Many a philosopher have explained how uncertainty is an inescapable part of being human. This dissertation opened with a quotation of Saint Augustine making uncertainty a sort of ontological proof, a *cogito* before Descartes's. This relatively long essay on uncertainty has hopefully shown that uncertainty is better understood than avoided. Sisyphus lived with the perpetual uncertainty that the boulder he relentlessly pushed up his hill would remain atop, and in Albert Camus's own words, *one must imagine Sisyphus happy.*

BIBLIOGRAPHY

Tobias Adrian, Nina Boyarchenko, and Domenico Giannone. Vulnerable Growth. Staff Reports 794, Federal Reserve Bank of New York, September 2016.

H. Akaike. A New Look At The Statistical Model Identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, Dec 1974.

Mario Alloza. Is Fiscal Policy More Effective in Uncertain Times or During Recessions? Unpublished manuscript, April 2014.

S.-I. Amari, O. E. Barndorff-Nielsen, R. E. Kass, S. L. Lauritzen, and C. R. Rao. Differential Geometry in Statistical Inference. *Lecture Notes-Monograph Series*, 10:i–240, 1987.

Jeffrey L. Anderson. A Method for Producing and Evaluating Probabilistic Forecasts from Ensemble Model Integrations. *Journal of Climate*, 9(7):1518–1530, 1996.

Kenneth J. Arrow and Gerard Debreu. Existence of an Equilibrium for a Competitive Economy. *Econometrica*, 22(3):265–290, July 1954.

Alan J. Auerbach and Yuriy Gorodnichenko. Fiscal Multipliers in Recession and Expansion. NBER Working Papers 17447, National Bureau of Economic Research, Inc, September 2011.

Alan J. Auerbach and Yuriy Gorodnichenko. Measuring the Output Responses to Fiscal Policy. *American Economic Journal: Economic Policy*, 4(2):1–27, May 2012.

Alan J. Auerbach and Yuriy Gorodnichenko. Output Spillovers from Fiscal Policy. *American Economic Review*, 103(3):141–46, May 2013.

Rüdiger Bachmann and Eric R. Sims. Confidence and the Transmission of Government Spending Shocks. *Journal of Monetary Economics*, 59 (3):235–249, 2012.

Jushan Bai and Serena Ng. Determining the Number of Primitive Shocks in Factor Models. 25:52–60, 02 2007.

Scott R. Baker, Nicholas Bloom, and Steven J. Davis. Policy Uncertainty: a New Indicator. CentrePiece - The Magazine for Economic Performance 362, Centre for Economic Performance, LSE, February 2012.

Scott R. Baker, Nicholas Bloom, and Steven J Davis. Measuring Economic Policy Uncertainty. CEPR Discussion Papers 10900, C.E.P.R. Discussion Papers, October 2015.

Laurence Ball. Why Does High Inflation Raise Inflation Uncertainty? NBER Working Papers 3224, National Bureau of Economic Research, Inc, January 1990.

Robert J. Barro and Charles J. Redlick. Macroeconomic Effects From Government Purchases and Taxes. *The Quarterly Journal of Economics*, 126(1):51–102, 2011.

Andrew R. Barron. Entropy and the Central Limit Theorem. *The Annals of Probability*, 14(1):336–342, January 1986.

Robert B. Barsky and Eric R. Sims. Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence. *American Economic Review*, 102(4):1343–77, 2012.

Ben S. Bernanke, Jean Boivin, and Piotr Eliasz. Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach. NBER Working Papers 10220, National Bureau of Economic Research, Inc, January 2004.

Olivier Blanchard and Roberto Perotti. An Empirical Characterization Of The Dynamic Effects Of Changes In Government Spending And Taxes On Output. *The Quarterly Journal of Economics*, 117(4): 1329–1368, November 2002.

Nicholas Bloom. The Impact of Uncertainty Shocks. *Econometrica*, 77 (3):623–685, 05 2009.

Nicholas Bloom. Fluctuations in Uncertainty. *Journal of Economic Perspectives*, 28(2):153–76, Spring 2014.

William A. Branch and George W. Evans. Model Uncertainty and Endogenous Volatility. *Review of Economic Dynamics*, 10(2):207 – 237, 2007.

William A. Branch and Bruce McGough. Replicator Dynamics in a Cobweb Model with Rationally Heterogeneous Expectations. *Journal of Economic Behavior & Organization*, 65(2):224–244, February 2008.

Glenn W. Brier. Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, 78(1):1–3, 1950.

William A. Brock and Cars H. Hommes. Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model. *Journal of Economic Dynamics and Control*, 22(8):1235 – 1274, 1998.

James Buchan and Linda Aiken. Solving Nursing Shortages: a Common Priority. *Journal of Clinical Nursing*, 17(24):3262–3268.

Albert Camus. *The Myth of Sisyphus and Other Essays.* Hamish Hamilton, 1955.

Guglielmo Maria Caporale, Paolo Paesani, and Luca Onorante. Inflation and Inflation Uncertainty in the Euro Area. Working Paper Series 1229, European Central Bank, July 2010.

G. Cardano. *The Book on Games of Chance.* Holt, Rinehart and Winston, New York, New York, NY (United States), 1953.

Marcelle Chauvet. An Econometric Characterization of Business Cycle Dynamics with Factor Structure and Regime Switching. *International Economic Review*, 39(4):969–996, 1998.

Lawrence J. Christiano, Roberto Motto, and Massimo Rostagno. Risk Shocks. *American Economic Review*, 104(1):27–65, 2014.

G. Claeskens and N.L. Hjort. *Model Selection and Model Averaging.* Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2008.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing).* Wiley-Interscience, New York, NY, USA, 2006.

H. Cramér. *Mathematical Methods of Statistics.* Princeton landmarks in mathematics and physics. Princeton University Press, 1946.

Stefania D'Amico and Athanasios Orphanides. Inflation Uncertainty and Disagreement in Bond Risk Premia. Working Paper Series WP-2014-24, Federal Reserve Bank of Chicago, January 2014.

G. Debreu. *Theory of Value: An Axiomatic Analysis of Economic Equilibrium.* Wiley, 1959.

G. Elliott and A. Timmermann. *Economic Forecasting.* Princeton University Press, 2016.

Daniel Ellsberg. Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.

Ekkehard Ernst and Christian Viegelahn. Hiring Uncertainty: A New Labour Market Indicator. Staff reports, International Labor Organization, February 2014.

G.W. Evans and S. Honkapohja. Learning as a Rational Foundation for Macroeconomics and Finance. In Roman Frydman and Edmund S. Phelps, editors, *Rethinking Expectations: The Way Forward For Macreconomics.* Princeton University Press, 2013.

Emmanuel Farhi and Iván Werning. Fiscal Multipliers: Liquidity Traps and Currency Unions. Working Paper 18381, National Bureau of Economic Research, September 2012.

Steven Fazzari, James Morley, and Irina Panovska. State-Dependent Effects of Fiscal Policy. Discussion Papers 2012-27A, School of Economics, The University of New South Wales, April 2013.

Laurent Ferrara and Pierre Guérin. What Are The Macroeconomic Effects of High-Frequency Uncertainty Shocks? Technical report, 2015.

Roger W. Garrison. The Kaleidic World of Ludwig Lachmann. *Critical Review*, 1(3):77–89, 1987.

Edward I. George and Robert E. McCulloch. Variable Selection via Gibbs Sampling. *Journal of the American Statistical Association*, 88 (423):881–889, 1993.

Domenico Giannone, Lucrezia Reichlin, and David Small. Nowcasting: The Real-Time Informational Content of Macroeconomic Data. *Journal of Monetary Economics*, 55(4):665–676, May 2008.

M. Giannoni. The Federal Reserve in the 21st Century: Models for Forecasting and Policy Analysis. Available on the New York Fed's website, 2016.

Christopher G. Gibbs. Forecast Combination, Non-Linear Dynamics, and the Macroeconomy. *Economic Theory*, 63(3):653–686, Mar 2017.

Itzhak Gilboa and David Schmeidler. Maxmin Expected Utility With Non-Unique Prior. *Journal of Mathematical Economics*, 18(2):141–153, April 1989.

Paolo Giordani and Paul Söderlind. Inflation Forecast Uncertainty. *European Economic Review*, 47(6):1037–1059, December 2003.

Tilmann Gneiting and Adrian E Raftery. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

Gergely Gánics. Optimal Density Forecast Combinations. Banco de España Working Paper 1751, Banco de España, January 2018.

Refet S. Gürkaynak and Jonathan H. Wright. Macroeconomics and the Term Structure. *Journal of Economic Literature*, 50(2):331–67, June 2012.

I. Hacking. *The Emergence of Probability*. Cambridge University Press, 1984.

Robert E. Hall. By How Much Does GDP Rise if the Government Buys More Output? NBER Working Papers 15496, National Bureau of Economic Research, Inc, November 2009.

Lars Peter Hansen and Thomas J. Sargent. Wanting Robustness in Macroeconomics. In Benjamin M. Friedman and Michael Woodford, editors, *Handbook of Monetary Economics*, volume 3 of *Handbook of Monetary Economics*, chapter 20, pages 1097–1157. Elsevier, 2010.

P. T. Harker. The U.S. Economy in 2017: Why Uncertainty Is the 'Biggest Risk'. Knowledge At Wharton, 2017.

Hans Hersbach. Decomposition of the Continuous Ranked Probability Score for Ensemble Prediction Systems. *Weather and Forecasting*, 15 (5):559–570, 2000.

M. F. Huber, T. Bailey, H. Durrant-Whyte, and U. D. Hanebeck. On Entropy Approximation for Gaussian Mixture Random Vectors. In *2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 181–188, Aug 2008.

Cosmin L. Ilut and Martin Schneider. Ambiguous Business Cycles. *American Economic Review*, 104(8):2368–99, August 2014.

E. T. Jaynes. How Should We Use Entropy In Economics? February 1970.

W. S. Jevons. *The Theory of Political Economy*. London: Macmillan and Co., 1888.

Òscar Jordà. Estimation and Inference of Impulse Responses by Local Projections. *The American Economic Review*, 95(1):pp. 161–182, 2005. ISSN 00028282.

Kyle Jurado, Sydney C. Ludvigson, and Serena Ng. Measuring Uncertainty. *American Economic Review*, 105(3):1177–1216, March 2015.

Robert E. Kass and Larry Wasserman. The Selection of Prior Distributions by Formal Rules. *Journal of the American Statistical Association*, 91(435):1343–1370, 1996.

John Maynard Keynes. *A Treatise on Probability*. London: Macmillan and Co, 1921.

John Maynard Keynes. *General Theory of Employment, Interest and Money*. Cambridge University Press, Cambridge, 1936.

Frank H. Knight. *Risk, Uncertainty, and Profit.* Houghton Mifflin Co., 1921.

Kajal Lahiri and Xuguang Sheng. Measuring Forecast Uncertainty by Disagreement: The Missing Link. *Journal of Applied Econometrics*, 25(4):514–538, 2010.

P.-S. Laplace. *A Philosophical Essay on Probabilities.* Dover, Mineola, NY (United States), 1953.

Edward Leamer and Herman B Leonard. Reporting the Fragility of Regression Estimates. *The Review of Economics and Statistics*, 65(2): 306–17, 1983.

Lai-Yung Leung and Gerald R. North. Information Theory and Climate Prediction. *Journal of Climate*, 3(1):5–14, 1990.

Jr Lucas, Robert E. Some International Evidence on Output-Inflation Tradeoffs. *American Economic Review*, 63(3):326–334, June 1973.

Henrick J. Malik and Bovas Abraham. Multivariate Logistic Distributions. *The Annals of Statistics*, 1(3):588–590, 1973.

Massimiliano Marcellino. A Linear Benchmark for Forecasting GDP Growth and Inflation? *Journal of Forecasting*, 27(4):305–340, 2008.

D. N. McCloskey. *The Rhetoric of Economic (2 ed.).* University of Wisconsin Press, 1998.

Michael W. McCracken, Michael T. Owyang, and Tatevik Sekhposyan. Real-Time Forecasting with a Large, Mixed Frequency, Bayesian VAR. Working Papers 2015-30, Federal Reserve Bank of St. Louis, October 2015.

A. Monfort. *Cours de Statistique Mathématique.* Collection "Économie et Statistiques Avancées". Economica, 1982.

Jacob M. Montgomery and Brendan Nyhan. Bayesian Model Averaging: Theoretical Developments and Practical Applications. *Political Analysis*, 18(2):245–270, 2010.

Enrique Moral-Benito. Model Averaging in Economics: An Overview. *Journal of Economic Surveys*, 29(1):46–75, 2015. ISSN 1467-6419.

A. H. Murphy. What Is a Good Forecast? An Essay on the Nature of Goodness in Weather Forecasting. *Weather and Forecasting*, 8: 281–293, June 1993.

Allan H. Murphy. A Note on the Ranked Probability Score. *Journal of Applied Meteorology*, 10(1):155–156, 1971.

John F. Muth. Rational Expectations and the Theory of Price Movements. *Econometrica*, 29(3):315–335, 1961.

Whitney K Newey and Kenneth D West. A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703–08, May 1987.

Florens Odendahl. Survey-Based Joint Density Forecasts. Working paper, October 2017.

Andrew J. Patton and Allan Timmermann. Why Do Forecasters Disagree? Lessons From The Term Structure Of Cross-Sectional Dispersion. *Journal of Monetary Economics*, 57(7):803–820, October 2010.

Adrian E. Raftery, David Madigan, and Jennifer A. Hoeting. Bayesian Model Averaging for Linear Regression Models. *Journal of the American Statistical Association*, 92(437):179–191, 1997.

Valerie A. Ramey. Identifying Government Spending Shocks: It's all in the Timing. *The Quarterly Journal of Economics*, 126(1):1–50, 2011.

Valerie A. Ramey and Sarah Zubairy. Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data. NBER Working Papers 20719, National Bureau of Economic Research, Inc, November 2014.

Robert Rich and Joseph Tracy. The Relationships among Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts. *The Review of Economics and Statistics*, 92(1):200–207, February 2010.

P. Romer. The Trouble With Macroeconomics. Available on Paul Romer's website, 2016.

Barbara Rossi and Tatevik Sekhposyan. Evaluating Predictive Densities of US Output Growth and Inflation in a Large Macroeconomic Data Set. *International Journal of Forecasting*, 30(3):662–682, 2014.

Barbara Rossi and Tatevik Sekhposyan. Macroeconomic Uncertainty Indices Based on Nowcast and Forecast Error Distributions. *American Economic Review*, 105(5):650–55, May 2015.

Barbara Rossi and Tatevik Sekhposyan. A Macroeconomic Uncertainty Index for the Euro Area. mimeo, 2016.

Epstein Edward S. Stochastic Dynamic Predictions. *Tellus*, 21(6): 739–759.

Hikaru Saijo. The Uncertainty Multiplier and Business Cycles. UTokyo Price Project Working Paper Series 016, University of Tokyo, Graduate School of Economics, November 2013.

J.-B. Say. *A Treatise on Political Economy.* Philadelphia: Claxton, Remson and Haffelfinger, translated from the 4th French edition, 1803.

Gideon Schwarz. Estimating the Dimension of a Model. *The Annals of Statistics*, 6(2):461–464, 1978.

Chiara Scotti. Surprise and Uncertainty Indexes: Real-Time Aggregation of Real-Activity Macro Surprises. Technical report, 2013.

Eric Sims and Jonathan Wolff. The State-Dependent Effects of Tax Shocks. Working paper, August 2017.

D. B. Stephenson, C. A. S. Coelho, and I. T. Jolliffe. Two Extra Components in the Brier Score Decomposition. *Weather and Forecasting*, 23(4):752–757, 2008.

D.B. Stephenson and F.J Doblas-Reyes. Statistical Methods for Interpreting Monte Carlo Ensemble Forecasts. *Tellus A*, 52(3):300–322, 2000.

James H. Stock and Mark W. Watson. Business Cycle Fluctuations in U.S. Macroeconomic Time Series. NBER Working Papers 6528, National Bureau of Economic Research, Inc, April 1998.

James H Stock and Mark W Watson. Macroeconomic Forecasting Using Diffusion Indexes. *Journal of Business & Economic Statistics*, 20(2): 147–162, 2002.

James H. Stock and Mark W. Watson. Forecasting Output and Inflation: The Role of Asset Prices. *Journal of Economic Literature*, 41(3): 788–829, September 2003.

John A. Swets. The Relative Operating Characteristic in Psychology. *Science*, 182(4116):990–1000, 1973.

Mathieu Taschereau-Dumouchel, Edouard Schaal, and Pablo Fajgelbaum. Uncertainty Traps. 2013 Meeting Papers 677, Society for Economic Dynamics, 2013.

T. Terasvirta and H. M. Anderson. Characterizing Nonlinearities in Business Cycles Using Smooth Transition Autoregressive Models. *Journal of Applied Econometrics*, 7:pp. S119–S136, 1992.

Timo Terasvirta. Modelling Economic Relationships with Smooth Transition Regressions. SSE/EFI Working Paper Series in Economics and Finance 131, Stockholm School of Economics, November 1996.

Kenneth F. Wallis. Combining Density and Interval Forecasts: A Modest Proposal. *Oxford Bulletin of Economics and Statistics*, 67(s1):983–994, December 2005.

Larry Wasserman. Bayesian Model Selection and Model Averaging. *Journal of Mathematical Psychology*, 44(1):92 – 107, 2000.

Jonathan H. Wright. Term Premia and Inflation Uncertainty: Empirical Evidence from an International Panel Dataset. *American Economic Review*, 101(4):1514–34, June 2011.

Victor Zarnowitz and Louis A Lambros. Consensus and Uncertainty in Economic Prediction. *Journal of Political Economy*, 95(3):591–621, June 1987.