Chapter One

Why Study Information Choice?

The developed world is becoming increasingly a knowledge-driven economy. Fewer and fewer workers are involved in producing goods. Instead, much of value added comes from activities such as consulting, forecasting, and financial analysis. Even traditional firms devote significant resources to activities such as managerial decision making, price-setting, and evaluating potential investments, each of which involves acquiring, processing, and synthesizing information. Most macroeconomic models focus on goods production. Similarly, most portfolio and asset-pricing theories derive the investment decisions that maximize investors' utility. Only a small body of research tells us about the information acquisition that regulates these production and investment decisions, even as the amount of resources devoted to information-related activities grows ever larger. This book examines macroeconomics and finance models where people choose what information they want to know.

Every expectation, mean, variance, and covariance, every moment of every random variable, is conditional on some information set. Typically, we think of that information set as consisting of all the past realizations of the variable. But what if people do not know the entire history of all the realizations? Or, alternatively, what if they have information in addition to the history? Which information set people have will typically change every moment of the random variable. The information does not affect what the future realizations of the random variable will be. It changes what people know about those realizations. Moments of a random variable summarize what someone knows about it.

Since means, variances, and covariances appear all over economics and finance, how people evaluate these moments affects how they behave in any environment where a random variable matters. The effect of random changes in productivity in a business cycle model, of changes in asset valuation in a portfolio problem, of shocks to endowments in a consumption/savings problem, of money-supply shocks in a price-setting problem, and of changes in the state in a coordination game all depend on what information people know.

This book describes both a narrow and a broad research area. It focuses narrowly on predicting what information agents have and how that information affects aggregate outcomes. It presents a tool kit for writing applied theories: it does not explore theory so much for its own sake but explores theories that provide explanations for the phenomena we see in the world around us. The applications are wide-ranging, from asset pricing to monetary economics to international economics to business cycles. This book covers both the mathematical structures necessary to work in this area and ways of thinking about the modeling choices that arise. My hope is that at

the end, the reader will be at the frontier of applied theory research in information choice and might be inspired to contribute to it.

1.1 TYPES OF LEARNING MODELS

The literature on learning is expansive. Before delving into the material, let's review what kinds of subjects are covered here and what is omitted.

Learning is often used to refer to a literature in which agents do not use Bayes' law to form their expectations. One example is adaptive least-squares learning, where agents behave as econometricians, trying to discover the optimal linear forecasting rule for the next period's state. Evans and Honkapohja (2001) offer an exhaustive treatment of this literature. Information frictions also feature prominently in the literature on model misspecification. In these models, agents do not know the true model of the economy and choose actions that would produce good outcomes even if their model of the economy is not quite right (Hansen and Sargent 2003). This book focuses exclusively on Bayesian learning. It considers only environments where agents update their information sets using Bayes' law. They know the true model of the economy and are uncertain only about which realization of the state will be drawn by nature. Agents in this class of models are not uncertain about the distribution of outcomes from which nature draws that state. In other words, they have rational expectations. This focus creates a distinction between the study of the process through which agents learn and the information they acquire and learn from. It allows for a deeper understanding of what information agents observe by making simple assumptions about how learning takes place.

Among models with Bayesian learning, there are models of passive learning and models of active learning. In passive-learning models, agents are endowed with signals and/or learn as an unintended consequence of observing prices and quantities. One set of examples is models where information is exogenous. Information may be an endowment (Morris and Shin 1998), or it may arrive stochastically (Mankiw and Reis 2002). Even when information is endogenous, agents can still be passive learners. For example, information could be produced as a by-product of other activity, or information could be conveyed by market prices. This is still passive learning because agents are not exercising any control over the information they observe.

The other way people acquire information is intentionally, by choice. Acquiring information by choice is also called active learning. This choice might involve purchasing information, choosing how to allocate limited attention, or choosing an action, taking into account the information it will generate. Such models go beyond explaining the consequences of having information; they also predict what information agents will choose to have. Active learning is starting to play a more prominent role in macroeconomics and finance. In macroeconomics, it has been used to reexamine consumption savings problems (Sims 1998), price-setting frictions (Maćkowiak and Wiederholt 2009b; Reis 2006), and business-cycle dynamics (Veldkamp and Wolfers 2007). In finance, active learning has a long tradition in investment-allocation models (Grossman and Stiglitz 1980; Hellwig 1980) and has

more recently been used in dynamic asset-pricing theories (Peng and Xiong 2006), models of mutual funds (Garcia and Vanden 2005; Kacperczyk, Van Nieuwerburgh, and Veldkamp 2010), and models of decentralized trade (Golosov, Lorenzoni, and Tsyvinski 2008).

The vast majority of models in dynamic macroeconomics and finance still employ passive learning. Beliefs change when agents observe new outcomes or when new signals arrive because the model assumes they do. While such models clarify the role of information in aggregate outcomes, they do not tell us what information we may or may not observe. Models of active learning complement this literature by predicting what information agents choose to observe. Because an active-learning model can predict information sets on the basis of observable features of the economic environment, pairing it with a passive-learning model where information predicts observable outcomes results in a model where observables predict observables. That is the kind of model that is empirically testable. If the goal is to write applied theories that explain observed phenomena, we need a testable theory to know if the proposed explanation is correct. Therefore, while this book devotes substantial attention to passive-learning models, which form the bulk of the literature, it systematically complements this coverage with a discussion of how information choice affects the predictions.

Bayesian learning models are also used prolifically in other literatures. While this book focuses on applications of interest to researchers in macroeconomics and macro-finance, others cover related topics: Vives (2008) and Brunnermeier (2001) focus on market microstructure and information aggregation in rational expectations markets; Chamley (2004) explores models of herding where agents move sequentially and learn from each other's actions. Each of these volumes would be a good complement to the material presented here.

1.2 THEMES THAT RUN THROUGH THE BOOK

Theme 1: Information choice bridges the gap between rational and behavioral approaches. The first theme is about the place information-based models occupy in the macroeconomics and finance literature. We look to incomplete-information models in situations where the standard full-information, rational models fail to explain some feature of the data. Since these are the same set of circumstances that motivate much of the work in behavioral economics, modeling incomplete information is an alternative to the behavioral economics approach. Both literatures take a small step away from the fully informed, ex post optimal decision making that characterizes classical macroeconomics and finance. This step is useful because adding incomplete information allows the standard models to explain a broader set of phenomena such as portfolio underdiversification, asset-price bubbles, inertia in asset allocation, and sticky prices.

The line between incomplete information and behavioral assumptions can be further blurred because some approaches to modeling incomplete information—for example, information-processing constraints—are a form of bounded rationality. Yet, there is a fundamental distinction between the two approaches.

At its core, information choice is a rational choice. Agents in information-choice models treat their information constraints just as agents in a standard model treat their budget constraints. Rather than attacking tenets of the rational framework, information choice seeks to extend it by enlarging the set of choice variables. The advantage of this approach is that requiring information sets to be a solution to a rational-choice problem disciplines the forms of information asymmetry one can consider in a given environment.

Theme 2: Information is different from physical goods because it is expensive to discover and cheap to replicate. Because information is expensive to discover and cheap to replicate, the more copies of it are sold, the lower the cost per copy. In other words, information has increasing returns to scale. That property often produces extreme solutions and complementarities that make the predictions of models with information choice fundamentally different from those without it (Wilson 1975).

The fact that the economics of producing information is so different from the economics of producing physical goods allows information-choice models to explain phenomena that pose a challenge to standard theory. This idea shows up in the discussion of information choice and investment, where investors value information about a larger asset more because one piece of information can be used to evaluate every dollar of the investment. It reappears in the section on information markets, where information has a market price that decreases in the quantity of information sold because it is free to replicate. That makes observing information that others observe a cost-effective strategy.

Theme 3: Correlated information or coordination motive? Many works in economics and finance study situations where many agents take similar actions at the same time. Examples include bank runs, speculative attacks, market bubbles, and fads. The most common explanation is the existence of a *coordination motive*: people want to behave similarly because acting like others directly increases their utility. Yet, in some settings, people appear to be acting as if they had a coordination motive when no such motive is apparent. An alternative explanation is that people observe *correlated information*, which leads them to behave similarly. One example is the theory of herds, which is based on the following parable (taken from Banerjee 1992). Suppose diners arrive sequentially at two adjacent restaurants, A and B. The first and second diners believe A to be the better restaurant. So, both dine there. The third diner believes that B is the better restaurant but sees no one there. He infers the two diners in restaurant A must have been told that A was better and chooses A over B. All subsequent diners make similar calculations and dine at restaurant A. Thus, sequential, publicly observable actions provide one reason that people could have similar information sets. An alternative reason why people may have similar information is based on information choice. In the language of the parable, all the diners might go to restaurant B because a reviewer on the evening news recommends it. A diner could acquire private information by trying each restaurant. But watching the evening news and seeing the information that everyone else sees is less costly.

Theme 4: The effect of introducing information choice depends on the strategic motives in actions and on whether the information being chosen is public or private. This theme governs the organization of the book. The first half considers settings with strategic complementarity, while the second half focuses on substitutability. Each half examines public and private information choice.

One example of the interaction between strategic motives and information choice is when the combination of coordination motives and heterogeneous (private) information creates powerful inertia. Because information is imperfect, people do not know what the state of the world is and cannot adjust their actions precisely whenever that state changes. Furthermore, people do not know what the average action of others is because they do not know what information others have. If information accumulates over time, all people know more about what happened in the past than about the present state of the world. Thus, to be better coordinated, people put more weight on past than on current signals. Being unresponsive to new information creates inertia in actions. They also choose to acquire little new information because the old information is much more useful. Since actions can only respond to the changes in the state that people know about, delayed learning slows reactions. This delay creates even more inertia. Researchers in monetary economics use this mechanism to generate price stickiness.

However, when people want to coordinate their actions and have access to public information, the result is typically multiple equilibria. Such an economy may exhibit very little inertia because changes in expectations could cause switches from one equilibrium to another.

In other settings, people do not want to coordinate but instead prefer to take actions that are different from those of other people. Such strategic substitutability makes people want to act based on information that others do not have. They prefer private information. Since information about recent events is scarcer, people give more weight to recent signals. Instead of inertia, this kind of model generates highly volatile actions. Since asset markets are settings where investors prefer to buy assets that others do not demand because their price is lower, this mechanism provides a rationale for the volatility of investment behavior and, in turn, the volatility of asset prices.

Even though people with strategic substitutability in actions prefer private information, they may acquire public information if it is substantially cheaper than private information. Because information is expensive to discover and cheap to replicate, public or partially public information that can be sold in volume is less expensive to produce and therefore to purchase. As explained in theme 3, people who observe common information typically make similar decisions. Even though people's strategic motives dictate they should take actions that differ, their similar information sets may lead them to take similar actions. The result is coordinated actions.

The following table summarizes the key model outcomes from the four types of models described in the four preceding paragraphs. Each has either complementarity or substitutability in actions. In addition, each model allows agents to choose either public or private signals. For example, models with complementarity and private information choice typically produce inertia in actions, while models with

substitutability and private information typically generate actions that are dispersed and volatile.

Model Outcome	Private Signals	Public Signals
Complementarity	Inertia	Multiple equilibria
Substitutability	Dispersion and volatility	Coordinated actions

Theme 5: Information choice facilitates empirical testing of information-based theories. Writing models that predict what agents know is a way to get around the problem that information sets are not directly observable. If the theory predicts information sets on the basis of observable variables and then uses those information sets to predict actions of agents that are also observable, then the theory begins and ends with observable variables. A theory that begins and ends with observable variables is empirically verifiable. Therefore, an entire chapter is devoted to various ways of empirically evaluating information-based theories. It draws together ideas about what kinds of measurable fundamentals determine information sets, suggests proxy variables that others have used to measure information, and collects stylized facts that support a range of asymmetric information theories.

1.3 ORGANIZATION OF THE BOOK

This book is meant to be used as a guide or reference for researchers or as a textbook for a second-year Ph.D. course in economics or finance. One of the unusual features of this book is that it touches on many different topics, and yet it is not meant as an introductory text to macroeconomics, monetary economics, or finance. The reason for this approach is that each of these fields has developed insights about endogenous information that can and should be applied to other fields. Yet, since communication across research communities is often limited, opportunities for doing cutting-edge research are being lost because those who know the tools are not familiar with many of their potential applications and vice versa.

As a result, the text is filled with suggestions for how ideas in one field might be applied to another. This feature should make it a worthwhile read for graduate students searching for a dissertation topic. It also draws connections between literatures by illustrating, for example, that the driving force in a monetary-policy model is the same as that in a portfolio-choice model. These connections can help established researchers familiar with one area to quickly have a strong intuitive grasp of the logic behind the models in another area. This feature makes the chapters on monetary non-neutrality and business cycles perfectly appropriate for a class on asset pricing and the portfolio-choice chapter appropriate for a class on macroeconomics. The ideas about how information choice can generate inertia in prices

are just as relevant for explaining inertia in portfolio choices and the mechanism by which news about future productivity affects current output can be used as the basis of an asset-pricing model. Likewise, ideas about how agents choose to trade financial securities can be used to explain puzzling features of cross-country trade in goods and services. By drawing together insights from various corners of macroeconomics and finance, I hope this book might advance work in endogenous information by illuminating the possibilities for new applications of existing tools. Ultimately, such applications could generate new insights about the role of information in the aggregate economy.

Chapters 2 and 3 contain essential mathematical tools for understanding the material in chapters 4 through 9. Chapters 4 and 5 consider a strategic game with many players. Stripping away the details of any particular application makes the general themes that run through the later models more transparent. Chapter 4 shows that when agents can choose information to acquire before playing a strategic game, complementarity (substitutability) in actions typically generates complementarity (substitutability) in information choices. It also explains why heterogeneous information can deliver unique predictions out of a model with coordination motives where common knowledge would predict multiple equilibrium outcomes. Chapter 5 illustrates how changing the amount of public information agents know affects welfare differently than does changing the amount of private information.

Chapter 6 uses models from monetary economics to show how the combination of coordination motives in price-setting and heterogeneous information creates strong inertia in actions. When information choice is added to this setting, the coordination motive in information strengthens this inertia. Chapter 7 shows how the decision to invest in risky assets exhibits substitutability instead of complementarity. As a result, agents want to acquire information that is different from what other agents know. This different information leads them to hold different portfolios. Chapter 8 explores a setting where agents choose risky investments. But rather than focusing on the interaction between strategic motives in actions and information choices, it introduces the idea of increasing returns in information. Neither the monetary nor the investment models considered a production economy. Introducing an active role for information into such a general-equilibrium setting creates new challenges. Chapter 9 explores what model features allow information to act as an aggregate shock and create realistic business cycle fluctuations or realistic asset-price dynamics.

Each of the applied theory chapters (6 through 9) concludes with ideas about how its tools or insights might be applied to answer other questions. Each particular project may or may not be viable. The goal is to illustrate how the main concepts might be applied more broadly so that the reader might be better able to formulate his or her own research idea.

Of course, a work of applied theory is rarely considered complete without some empirical support for its hypothesis. Therefore, chapter 10 is devoted to the subject of testing information-based models. Chapter 11 offers some concluding thoughts.

As information technology makes it possible to access a vast array of knowledge, constraints on what one can and cannot observe become less relevant. Yet, having

access to information is not the same as knowing it. Just as you can buy a textbook and not learn its material, agents with access to information might not really acquire it in their working memories. But putting information in our working memories—learning—is a choice. In other words, information constraints are systematically being replaced with information choices. It is such choices that this book seeks to understand.