Econometric Learning

Emi Nakamura and Jón Steinsson
Harvard University

September 18, 2003

1 Introduction

In some situations, the appropriate model of one's environment is quite clear—for example, for a flip of a coin or a throw of a die. But this is not the case for most economic situations. Consider a worker who is trying to predict whether his firm will lay him off next month or a bond trader trying to determine how short term interest rates will evolve over the next year. In these cases it is far from obvious what the true model of the situation is.

Savage (1954) argued that in cases like these, agents could be modelled as maximizing utility given a subjective probability distribution over states of the world. Savage points out that it is simple to elicit an agent's subjective beliefs by offering the agent a sequence of gambles. Consider, for example, the famous example of determining someone's subjective probability of a rotten egg by asking how much he would pay for a bet that pays off if the next egg from a carton turns out to be rotten.

Most modern economic analyses impose the added assumption that the subjective beliefs of the agents in the model coincide with the objective probability distributions in the model. This is for instance a feature of Nash equilibria and rational expectations equilibria. One problem with this assumption is that it endows the agents in the economic model with a great deal of knowledge—knowledge that one doesn't necessarily have when one is trying to make an omelet, let alone forecast bond prices.

Savage's framework is more general in that it recognizes that there is always an answer to the question of how much you would bet on an uncertain event, even when the objective probabilities are difficult to fathom. However, Savage's framework leaves as a black box the source of these subjective beliefs. Savage is right that the question of how much an agent is willing to bet on a rotten egg always has an answer. But surely, some answers are more rational than others.

The assumption that agents are endowed with a complete probability distribution over states of the world, whether objective or subjective, also focuses attention away from a major aspect of real-life desicion making, namely model choice and model building. In most important economic situations a great deal of effort is expendend by economic agents trying to develop appropriate models. This is for instance the task economists spend most of their time thinking about. It is therefore curious that the models economists write down don't assume that the agents in the model have to engage in any model building activity.

A primary criticism of the literature that has tried to provide alternatives to the idea that "agents know the model" has been its arbitrariness. If not rational expectation then what? Sims (1980) cites the "wilderness" of alternatives once one deviates from rational expectations. In this paper, we argue that the wilderness is less forboding than it seems if we take as a guideline for how the agents in the model should behave how applied economists behave in building and estimating empirical models.

Take, for example, the work of Campbell and Viceira (2001). They develop a model of portfolio allocation with rational expectations. The first step is to estimate the stochastic processes for bond returns and inflation. Thus, while the rational expectations agents in the model are simply provided with a stochastic model, Campbell and Viceira must first build and estimate a model based on a dataset. Though significant disagreements persist on how to do empirical work, a great deal of structure can nevertheless be placed on how Campbell and Viceira proceed in building the term structure model. It is not an unfathomable jungle.

In our framework, agents are endowed only with a dataset of past observations. They develop a model of the world based on this dataset. The idea is to make the agents behavior approximate as closely as possible what applied economists, such as Campbell and Viceira, would have done if they were in the position of the agents at any given point in time. The agents then maximize expected utility given the model they select.

Thus, the paper establishes a link between the modelling of rational agents and best-practice methods in econometrics and computer science. [1] The approach is no more arbitrary than rational expectations since the agents in the model carry out an empirical process that the applied economist would have implemented implicitly in the rational expectations framework. The framework introduced in this paper simply makes the model choice problem faced by applied economists explicit and assumes that it is carried out by the agents being modelled.

Agents are assumed to have preferences over different models of the world. These preferences are motivated by a couple of broad lessons from the scientific literature on model choice. First, all else equal, agents prefer models that fit the existing data better. Second, all else equal, agents prefer simpler models.

The preference for simplicity may arise for a number of reasons. First, we assume that more complex models harder to learn and use. Second, following Willaim of Ockham many have argued that a preference for simplicity should be taken as a primitive of science. Zellner et al. (2001)

[1] Hansen and Sargent (2000) also make this connection, though they focus on the issue of model uncertainty surrounding an approximating model, rather than on the issue of how the approximating model is chosen in the first place.

provides a wide-ranging discussion of the role of simplicity in scientific inference. A number of commonly used econometric model selection criteria can be thought of as implementing "Occam's razor", in the context of nested model choice.

The idea that the standard modelling paradigm in economics assumes that the agents in the model are much more knowledgable about the world they live in than seems realistis is of course not new. In the macroeconomics literature, for instance, there has been quite a bit of discontent with this aspect of rational expectations models. Evans and Honkapohja (2001, pp. 12–13) make the following argument:

> The rational expectations approach presupposes that economic agents have a great deal of knowledge about the economy. In empirical work economists, who postulate rational expectations, do not themselves know the parameter values and must estimate them econometrically. It appreas more natural to assume that the agents in the economy face the same limitations on knowledge about the economy. This suggests that a more plausible view of rationality is that the agents act like statisticians of econometricians when doing the forecasting about the future state of the economy.

Sargent (1993, pp. 21–22) also makes this argument. There has been quite a large literature within macroeconomics, under the heading of learning, that seeks to implement this view. A prominent example from this literature is Sargent's (2002) work on least-squares learning. Evans and Honkapohja (2001) present a thorough overview of this literature.

There is also a large literature on learning in the game theory literature. Fudenberg and Levine (1999) present an overview of this literature. In the game theory learning literature, agents are often assumed to follow the technique of "fictitious play" in which the probabilities of various outcomes equal the historical frequency of their occurence. Fictitious play can be viewed as an extremely simple approach to model choice.

The approach to inference that we advocate in this paper may be viewed as a complex learning algorithm. Yet, our approach to learning differs in spirit from much of the learning literature. While the learning literature largely investigates learning rules that are intentionally sub-rational, our objective is to model the behavior of agents aiming at rationality given highly limited information.[2] One advantage of our approach is that the implications of the model are then useful from a

[2] Within the learning literature the work of Evans and Ramey (1992) and Brock and Hommes (1997) is more closely related to our approach. Like in our framework, they model rational agents who face costs of forming forecasts. However, they assume that the agents in the model can form rational expectation if they are willing to bear a fixed cost. We assume that agents have much more limited knowledge.

normative, as well as positive perspective.

Even the work on econometrics-based learning algorithms places considerably more restrictions on learning than would be appropriate for an econometrician. For example, in least-squares learning, the agents always update their expectations using OLS regression on the same set of variables. They do not switch models, even if circumstances change drastically. Similarly, in the macroeconomic learning algorithms discussed in Evans and Honkapohja (2001), the agents are assumed to update on parameter values alone.

In contrast the computer science literature and the econometrics literature have sought to develop methods for optimal inference. A prominent example in the computer science literature is Q-learning, e.g., Kearns and Vazirani (1994). In the econometrics literature, well-known model choice criteria such as the Akaike information criterion and the Bayesian information criterion are often used to select the number of variables in a model. These criteria trade off increased complexity of the models—more variables—for improved fit, as do the agents in our model. In the context of nested model choice, a penalty for complexity is necessary to prevent maximum likelihood from always selecting the model with the most variables. Other work in the econometric literature investigates the issue of model choice in broader classes of models.[3] Stylistically, one of the main differences between the computer science and econometrics literatures is that the former often implicitly assumes a deterministic data generating process, whereas stochastic data generating processes are the norm in the latter.

One paper that makes the connection between the econometrics model choice literature and economic behavior is Hong and Stein (2003). However, unlike our agents, Hong and Stein's agents do not face costs of complexity. In addition, Hong and Stein's agents consider a highly limited set of models—which does not include the true model, whereas our agents consider a wide variety of models. Thus, once again, our paper differs from Hong and Stein's work in that our agents our intended to be as rational as possible given the infomational constraints and costs of model production, whereas Hong and Stein focus on the sub-rational behavioral implications of their model selection framework.

The framework we adopt emphasizes model production as an important and costly activity. In many circumstances it is natural to think of the agents as building their own models. In other cases, however, it is useful to assume that models are produced by a "consulting" sector and then bought by agents in a market. The "market for models" framework implies a number of analytical

[3] Reference forecasting model selection stuff

results.

First, to the extent that it is equally costly, in terms of complexity, to improve the fit of the model to all the phenomena the agent is trying to model (a condition that will not always hold), more "important" phenomena are modelled more closely than less important phenomena. (We make this statement precise in Section 2).) This result follows simply from the agent's desire to allocate his modelling resources efficiently. This result is closely related to Fryer's (2003) work showing that in categorical models, minority behavior is likely to be poorly modelled, although our framework also identifies cases in which this result will not hold. A related observation is that the models created in our framework tend to be both simpler and more cohesive than they would be if model creation were costless. The reason is that, since modelling is costly, agents prefer models that are simple to use, and that can be applied to a broad range of phenomena.

Second, the dynamics with econometric learning diverge most from the dynamics with rational expectations following regime changes. In general, learning algorithms yield considerably more gradual responses to regime changes, as a number of authors have noted.[4] Our framework goes further by allowing for the possibility that more knowledgable agents who *know* about the regime change can intervene in the market for models to affect the behavior of less knowledgable agents.

Finally, the framework of the "market for models" can help us to understand the ideas of planning costs and inability to plan. We can think of an agent who plans poorly as someone who– for one reason or another– fails to buy the appropriate model for decisionmaking. For example, he may have a poor understanding of the model's value, or he may have a particularly high cost of using complex models. Subsidizing the agent's acquisition of a model may have a large effect on his modelling ability if it helps him to see the value of other models, producing a cascade in his acquisition of models. (???)

The market for models framework can help us to understand the activities of a large number of agents in the economy since private consultants, firm managers, the government, and even academics are all involved, in one way or another, in the process of model production.

An agent's decision to buy a particular model often affects other agents, since the choice of model may affect his decisions about where to work, what to buy, and so on. As a consequence, some agents have an interest in influencing other agents' decisions about which models to buy. In other words, there are often externalities to individuals adopting a better—but more complex—models of their economic environment.

[4]See Evans and Honkapohja (2001) for a survey of this work.

Particularly when large agents, such as firms and government, are affected by someone's choice of model, these agents may attempt to "manage the model" adopted by other agents in the market. In this situation, the competitive equilibrium does not arise. Rather, the decisions of the small agents about which models to buy are affected by the large agents' efforts to subsidize particular models.

For example, a firm has a great interest in managing its workers model of its labor management practices– since this model plays a large role in determining worker morale. Similarly, the US government closely manages the model adopted by other countries about its foreign policy responses to various types of hostile actions by other countries.

In Section 4.4, we apply our general framework to two specific examples. We discuss the experience of Hewlett-Packard in recreating the "HP way" in the 1980's, and the experience of Japanese firms in creating and exporting Japanese labor management practices.

When there are not incentives (or opportunities) for large agents to intervene in the market for models, the externalities associated with particular models can lead to market failure. For example, as we discuss in Sections 2.2 and 4, there are often positive externalities associated with individuals adopting a more realistic—but more complex—models of their economic environment. The potential for market failures is an important justification for government intervention in the market for models. Indeed, many government institutions, such as the "small business bureau" (???) and career counselling play some role in subsidizing peoples' acquisition of models.

There is an intimate relationship between the issues discussed in this paper and information economics. Information economics is about what happens when data is missing, but the model is known– so how to interpret any data that one acquires is a matter of statistical inference. In contrast, in our framework, the model must also be inferred from the available data. A "known" model can often be seen as the limiting outcome of observing a great deal of empirical evidence in support of a particular model.

This paper is also related to Sims (2000) work on information processing constraints. In Sims' view, there is an exogenous constraint on the amount of information that an agent is able to process. In our view, the source of the information constraint is, to a large extent, the availability of models that the agent has available to interpret the information. Improvements in the agent's modelling capabilities (perhaps brought on by reductions in the cost of modelling) are likely to improve the agent's ability to process information.

2 Econometric Learning

Consider an agent who's optimal actions depend on a set of unknown phenomena that the agent must make predictions about. The most popular way to model such a situation in economics is to assume that the agent either "knows the model" and has rational expectations or that he at least has a subjective probability distribution over states of the world and updates these beliefs using Bayes rule. As we discussed in the introdution these two approaches leave as a black box the issue of how the agent aquired this knowledge or these beliefs. The main alternative to this approach is to assume that the agents use past data to learn about their environment. The learning literature, however, usually makes fairly restrictive assumptions about how the agent learns. For example, in the macroeconomics literature, a common assumption is that the agents estimate the unknown phenomena using OLS regressions on a specific set of variables. In the game theory learning literature a common assumption is that the agents learn according to "fictitious play"—whereby the agent assumes that the distribution of the phenomenon equals the unconditional empirical distribution.

A major justification for the least squares learning approach is that it seems more reasonable to assume that agents behave like econometricians than to assume that they know the model or have a detailed subjective prior. However, the assumption of least squares learning misses crucial aspects of the behavior of real life statisticians and econometricians. In particular, it ignores the model choice problem. Rather than endowing the agents with one particular learning model to use, we will assume that agents are faced with a variety of possible models and have to make a decision about what model to use.

We assume that the agent has access to a set of relevant data. Although, as in the information economics literature, the agent probably does not have access to *all* the relevant data. In order to make use of the data, the agent must build a model that connects the observable data with the phenomena that concern him.

The agent is faced with a set of models from which to choose. In principle, this set encompasses all possible models. However, in practice, the set will be somewhat limited, as in empirical economics applications. The agent has preferences over the set of models.

These preferences depend on the models' past performances, but also on their complexity. We will motivate these assumptions below. The agent also has preferences over a set of endogenous variables Y_t, the unknown phenomena. The agent's preferences may be represented by the following utility function:

$$U_t(Y_t) + V(m_t, \eta_t). \tag{1}$$

Here m_t denotes the complexity level of the model the agent chooses to purchase at time t and η_t are past prediction errors for the endogenous variables Y given the model the agent chooses to buy at time t. We discuss measures of complexity in more detail below. The agent knows his preferences, he knows the set of actions a_t available to him and, as mentioned above, he has access to a set of data X_t. The agent does not, however, know how the endogenous variables evolve or how his actions affect their evolution. The evolution of the endogenous variables is given by $Y_t(a_t, X_t, Z_t)$, where Z_t is a set of unobservable variables. X_t may include past values of Y_t, but does not include past values of a_t. The agent builds models that describe this relationship. The agent enters each period with knowledge of the model he purchased in the previous period.

Our assumptions about the function V are guided by the large literature in econometrics, computer science, physics and philosophy on the model choice problem. Drawing on this literature, we assume that V is increasing in past predictive success. This assumption is based on the scientific view that the structure of the world we live in exibit some degree of continuity. Our assumption that V is decreasing in the complexity of the model is motivated by another well established principle of science, Occam's razor. The Occam's razor principle asserts that the simplest explanation of a phenomena is best. The penalty on complexity is also meant to capture the fact that models are more difficult to *use and develop* the more complicated they are since people have limited mental resources. This cost of complexity is likely to be partially determined by a person's general level of education, suggesting that public education is one means by which the type of model the agents choose may be altered.

These preferences over models encompass a variety of different model choice paradigms in the econometrics and computer science literature. In the context of nested linear regression models, common model selection criteria such as the Akaike information criterion and the Bayesian information criterion have the form,

$$MSE + p(T), \tag{2}$$

where MSE is the model's mean squared error when it is estimated on past data and $p(T)$ is a penalty for the model's complexity. According to these model selection approaches, the model is selected which minimizes this criterion. These model selection criteria are generally justified on asymptotic grounds.[5] In this case, the penalty for complexity serves to counteract the problematic

[5] One unsatisfying aspect of the asyptotic justifications for model choice criteria is that the cost of overspecifying

feature of maximum likelihood (ML) that ML methods always select the more complicated model in a nested model choice problem.

More generally, the idea that one should trade off goodness of fit with the complexity of a model is encapsulated in Rissanen (1987)'s concept of "minimum description length" (MDL). According to MDL, one should select the model which "encodes" the data most efficiently, where the description of the data consists of two parts: the model itself and deviations from the model (the error term). MDL makes use of "code- length" type measures of complexity that are popular in parts of the computer science and mathematics literature.[6]

However, citeasnounhansenandyu2003 discuss how MDL can be viewed as a framework for understanding many of the model selection and estimation approaches in statistics and econometrics.[7] (Add more here about how to compute MDL?) In the econometrics literature, the relevant measure of complexity often turns out to be the number of variables in a linear model.Zellner, Keuzenkamp, and McAleer (2001b) provides an overview of a variety of other approaches to measuring complexity that arise in the econometrics, statistics and computer science literatures.

Whether a particular specification of preferences over goodness of fit and complexity are actually optimal—in the sense of selecting the best model in the model class—depends on the particular situation in which the agent finds himself. There is no "model-free" way of finding the optimal preferences over models—for any model selection criteria, an environment can be devised in which the criteria will perform badly. Howevever, a great deal of effort has been expended in the computer science and econometrics literature in trying to answer this question for "reasonably" specified data-generating processes. Our objective here is simply to argue that a variety of broadly defined principles of model selection can be used as primitives in models of economic behavior.

The first thing the agent must do in each period is to choose a model of the environment he lives in that he will use in making other decisions. The agent chooses a model by minimizing the function $V(m(\eta), \eta)$, where $m(\eta)$ is the minimum attainable level of complexity given η. The first order conditions of this choice problem are

$$V_m(m(\eta), \eta) m_{\eta_j}(\eta) + V_{\eta_j}(m(\eta), \eta) = 0. \tag{3}$$

In this expression, subscripts on functions denote partial derivatives. The first term on left hand

the model goes to zero as the sample size goes to infinity.

[6]The codelength approach is based loosely on the theoretical approach of Kolmogorov complexity or algorithmic complexity. See Li and Vitanyi (1997) for a discussion of Kolmogorov complexity. Algorithmic complexity is closely related to the concept of Shannon information. See Cover and Thomas (1991) for a discussion of this linkage.

[7]See Hansen and Yu (2003).

side represents the utility cost of having to increase the complexity of the model in order to achieve a lower level of η_j. The seond term represents the direct utility benefit of lowering the level of η_j.

We would like to be able to capture the intuitive idea that some pheonomena are more important to the agent than others, e.g., phenomena relating to the situation at the firm the agent works for are plausibly more important to the agents wellbeing than the economy wide unemployment rate. In this framework the importance of a particular phenomena may be measured by the partial derivative $V_{\eta_j}(m(\eta^p), \eta^p)$.

It is clear from equation (3) that the level of detail with which a particular phenomena will be model depends not only on the importance of that phenomena but also on how easily it can be modelled, $m_{\eta_j}(\eta)$. As a consequence more important phenomena may not be modelled more carefully than less important phenomena. However, if two phenomena are equally hard to model, the more important one will be modelled in a more detailed way. It is easy to see from equation (3) that at the optimum

$$\frac{V_{\eta_j}(m(\eta^p), \eta^p)}{V_{\eta_i}(m(\eta^p), \eta^p)} = \frac{m_{\eta_j}(\eta)}{m_{\eta_i}(\eta)}. \tag{4}$$

One might think that the agent should be able to derive the optimal shape of the function V from the function U. This is however only the case if the agent knows the relationship between his actions a_t and the endogenous variables Y_t that determine his utility. For example, suppose the agent knows that his utility depends only on his prediction errors in predicting the phenomena Y_t. In this case, the agent can estimate the gains from adopting a more complex model simply by estimating the associated reduction in prediction errors. (Of course, it is another big leap to assume that the agent has a way of estimating the gains associated with adopting the more complex model.)

In contrast, if the agent's actions go beyond forecasting, the benefits to an improved model are much harder to predict. For example, consider the problem of a central bank in deciding whether to develop a new macro-model for analyzing monetary policy. In developing the model, the central bank may discover that it has unexpected powers—as in the discovery of the Phillips curve—or far less power than it thought—as in the discovery of the expectations-augmented Phillips curve. As a consequence, there is no objective way to decide how much research should be done.

Thus, in the majority of situations, it seems reasonable to think of V as determined by the agent's inherent "curiousity" about determinants of the phenomena he observes.[8] This is the

[8] Of course, a model could be developed to estimate the effectiveness of past research efforts. This type of modelling effort might be expected to have a gradual effect on the agent's "curiousity" preference.

difference between scientists and backward people. Of course, in many situations, there are likely to be reasonable restrictions on preferences. For example, if a phenomenon has never been encountered then modelling this phenomenon should probably not play a large role in V.

The agent may be indifferent between more than one model for a given phenomenon. In this situation, he will also be indifferent between actions which maximize expected utility for either stochastic model. This property of our framework is one way of understanding the Ellsburg paradox, as we discuss below.

A realistic extension to this model would be to make models durable goods. In this case, agents would build up a stock of models over their lifetimes. Thus, the agents' ability to plan would gradually improve as they got older. Since the model production costs would only have to be paid when a model was first acquired, agents would want to use models already purchased, all else equal, to avoid purchasing a new model. This type of bias is investigated in a corporate finance setting by Hong and Stein (2003).

Another modification to the model that seems reasonable from the perspective of the psychology and behavioral economics literature would be for agents to prefer models that paint a rosy picture of the future. This could be implemented in our framework by allowing V to depend on expected future utility as well as past empirical success and complexity. Brunnermeier and Parker (2003) develop a model in which the preference for optimistic models is the *only* determinant of the agent's choice of subjective probabilities. In contrast, in our framework, the preference for optimistic models would always be tempered by a preference for realism and simplicity. As a consequence, the agent would not choose arbitrarily optimistic models, even if no adverse consequences would arise from such wildly optimistic thinking.

2.1 Simple Applications of Model Selection Framework

1. Standardized Tests and Occam's Razor[9]

Standardized exams often ask the student to complete a sequence like $2, 4, 8, ...$ Although, there are clearly multiple "models" that could have produced this sequence, only one answer is correct– 16. This answer is associated with the simplest model (2^n) associated with the sequence.

2. Ellsburg Paradox

In the classic Ellsberg Paradox, an agent is offered the choice of randomly choosing a ball from one of the two urns filled with red and black balls. Urn A has a known fraction p of red balls; while

[9] This example is taken from Li and Vitanyi (1997).

Urn B has an unknown fraction.

The paradox is that, subjects choose to draw from urn A regardless of whether the reward is offered for a red ball or a black ball. This outcome is inconsistent with the subjects having a fixed prior over the number of red balls in urn B, since, in this case, the agents should favor Urn B in one of the two prize situations.

This result arises more naturally in our framework, since the agent is likely to be indifferent between probabilistic models for the number of red balls in Urn B of the form "red with probability q". (The models have equal complexity and no data has been observed.) As a consequence, the agent is predicted to be indifferent between the two urns. In this situation, factors that would otherwise be of second-order importance, such as uncertainty aversion, dominate the agent's choice of urns.[10]

3. A Fair Coin?[11]

Suppose that one is flipping a coin that is known to be fair. An analogy can be made between this type of model and many models in economics. In this situation, one should be no more surprised by a string of 100 heads than any other (equally likely) string of heads and tails.

Yet, there are few situations in which it would seem truly rational for the agent flipping the coin not to question whether the coin being flipped was actually fair. This example underscores the importance of model choice in the rational decisionmaking process.

4. Unawareness

In their seminal paper Dekel, Lipman, and Rustichini (1998) pose the problem of how to model unawareness in the standard state-space framework. The problem is that in the standard state-space framework, the agent may not know something, but he is perfectly informed about the areas in which he is uninformed. This situation does not seem to correspond to our intuitive concept of unawareness.

In our framework, the natural way of understanding unawareness of states is as states that exist but have not been incorporated into the agent's model. For example, the agent is likely to be unaware of states that have small effects on observable variables, or which have not occurred in the past.

5. Money Illusion

A longstanding literature in economics suggests that people respond to nominal rather than real

[10] Gilboa and Schmeidler (1989) make a similar point regarding the Ellsberg paradox, though they focus on providing an axiomatic foundation for a maxmin decision rule when multiple priors are possible.???

[11] This example is from Li and Vitanyi (1997).

prices of goods. In periods of low inflation, "money illusion" could be an optimizing response to the costs associated with including an additional variable (inflation) in one's model of expenditure. (???)

2.2 Example: The Cobweb model revisited

An interesting, yet relatively simple, example of a situation for which our model of behavior is relevant is the market for agricultural goods that take time to breed or grow, such as the market for hogs or grain. It was in this context that Muth (1961) introduced rational expectations to the economics literature. An earlier literature had analysed such markets using various forms of adaptive expectations (see, e.g., Goodwin, 1947, and Nevlove, 1958). These models were refered to as Cobweb models since they predicted cyclical fluctuations in equilibrium prices. Since the publication of Muth's seminal paper, models of these markets have become a standard example in papers about expectation formation (see, e.g., Cyert and DeGroot (1974), Decanio (1979), Bray and Savin (1986), Brock and Hommes (1997)).

To be specific we will refer to the market we analyze in this example as the market for hogs. Demand in this market is given by

$$q_t^d = -a_1 p_t + a_2 p_t^c - a_3 u_{t-1} + \epsilon_t^d, \tag{5}$$

where q_t^d denotes the quantity of hogs demanded, p_t denotes the price of hogs, p_t^c denotes the price of chicken, u_t referes to some aggregate variable that may affect the demand for hogs, e.g. the unemployment rate, and ϵ_t^d is a demand disturbance. All variables are written as log deviations from a steady state. Supply in the hog market is given by

$$q_t^s = b_1 p_{t-1|t} + b_2 f_t + \epsilon_t^s, \tag{6}$$

where q_t^s denotes quantity supplied, $p_{t-1|t}$ denotes farmers' time $t-1$ predictions about the price of hogs at time t, f_t denotes the price of feed and ϵ_t^s is a supply disturbance. Finally, we assume that the hog market clears each time period, i.e. $q_t^d = q_t^s = q_t$. We have assumed that there are several observable disturbances to demand and supply in the model as well as the regular unobservable disturbances. We do this in order to make the inference that the farmers must undertake interesting and moderatly complex.

Since farmers' supply depends on past predictions about the current price, we must make assumptions about how these predictions are made. First, we assume that the only knowledge the farmers have about the market is past data on q_t, p_t, p_t^c, u_t, f_t and possibly other variables that the farmers think may be relevant. Importantly, the farmers have no knowledge of the structure of the market.

Second, we must make assumptions about the class of models that the farmers consider using to form their forecast. In principle we would like to consider *all* models, or as large a class of models as possible. However, to start with, we limit the class of models considered in this example to all possible linear models involving the time $t-1$ variables. We also assume that the farmers never choose a model with as more variables than the number of observations available.[12]

Third, we must specify the farmers' preferences over the different models they consider. These preferences are given by the following loss function

$$L = \log\left(1 + \frac{1}{T}\sum_{t=0}^{T}(p_t - p_{t-1|t})^2\right) + \alpha^T \lambda \frac{\log T}{T} K + \kappa K, \tag{7}$$

where K is the number of parameters in the model being evaluated and α, Λ and κ are parameters. The first term formalized the farmers' preference for models that have produced small prediction errors in the past. The second and third term penalize complex models. We use the number of variables as our measure of complexity in this example. The second term is meant to capture the farmers' preference for parsimonious models. The weight on K in this term falls as the data set grows since the farmers are less and less weary about overfitting the model the more data they have observed. The third term, for which the weight on K does not fall as the data set grows, captures the farmers' costs of buying and using the models. It also captures the farmers' long run preference for parsimony. That is, their preference for parsimonious models that persists even when the farmers have observed a large amount of data about their environment. We will discuss several rationals for this preference below.

We assume that the demand and supply shifters follow AR(1) processes. We also make specific assumptions about all the parameters in the model and about the shocks driving the demand and supply shifters. For the most part, the exact assumptions we make are not important. We therefore won't detail all these assumptions here. The discussion below will highlight the features of the model that seem to be important.

In the first period of our simulations, we assume that the agents have only observed two periods in which the market is in steady state. The obvious model to use to forecast the price in this case is just the constant model $p_{t-1|t} = 0$. The farmers then observe the market equilibrium in

[12]One can view this restriction as an additional feature of the farmers' preference for simplicity.

the next period. Based on this information, the farmers make predictions about the next period's price. They estimate the various possible models using OLS, and choose the model that produced the smallest value of the loss function (7). As the number of observations increases, they consider models with more and more variables. After six periods the farmers can, in principle, fit all linear models with any combination of the variables in their data set.

Changes in the farmers' forecasts have large effects on the relationship between the equilibrium price and the other observable variables. When the farmers switch from simpler to more complex models, in the early periods of the model, the result is large prediction errors since the changes in the model lead to changes in the data generating process. Moreover, once the farmers have used a number of different models, their prediction task is even more difficult since the data are generated by more than one model. However, there is a large range of parameters for which the hog market eventually converges to a rational expectations equilibrium. In this sense, our results for the hog cycle model parallel those of Evans and Honkapohja (2001).[13]

The extent to which the farmers are reluctant to include a large number of variables in their models is dictated by the values of α, λ and κ. If we assume very high values of these parameters, the farmers always forecast using the trivial model $p_{t-1|t} = 0$. The lower the values of these parameters the more complex the models that the farmers choose.[14]

The penalty parameters have important effects on the dynamics of prices in the hog market. On the one hand, when the penalty for choosing complex models is "too" small, the farmers quickly switch from forecasting using $p_{t-1|t} = 0$ to more complex models. However, the more complex models are highly innacurate initially since they are estimated using very little data. This chain of events often leads prices in the market to spiral rapidly away from the rational expectations equilibrium, producing a hog price "bubble". Moreover, when the parameter κ in the farmers' loss functions is zero, the farmers soon begin to use all the variables in their data set to forecast the price, even variables that have week– and likely spurious– correlations with the price.

On the other hand, when the penalty for complexity is "too" large, the farmers forecast using excessively simple models for a long time (???).

Evidently, there is an "optimal" value of κ in the sense of maximizing consumer welfare in the hog market by avoiding price bubbles and other types of instability. This type of optimality is not

[13]See Section 2.3 for a further discussion of how our approach relates to the macroeconomic learning literature.

[14]The appropriate values for these parameters are not provided by econometrics-based justifications for model selection criteria, since these justifications rely on asymptotic arguments in which the precise values of the penalties are irrelevant, so long as they are finite.???

generally considered in the econometrics literature, since statistical models typically do not allow for feedback effects between the econometrician and the data generating process. Indeed, the value κ that is optimal in this sense probabably differs from the value of κ that minimizes the mean squared forecasting error– the usual criterion in the econometrics literature. (???Add more here).

One reason for the large effect of the farmers' model choice on the data generating process is that the farmers are identical. As a consequence, when they switch forecasting models, they all switch at once. Yet, in real life, there is likely to be much more heterogeneity in the approach that agents use to forecasting. After all, most econometricians probably do not believe that the publication of their paper will have a major effect on the data generating process either by a) convincing everyone of a new theory or b) because everyone will independently derive the same theory at once. This holds even when there is considerable agreement that expectations play an important role in the data generating process, as for many macroeconomic variables.

One source of this type of heterogeneity is probably different penalties placed by agents on the complexity term. For example, this type of heterogeneity might arise from agents having different costs of using complex models, for example, arising from different levels of education. Another source of heterogeneity is probably that agents have access to different data sets. Yet other sources of heterogeneity are likely to be price, wealth, and marginal utility of consumption effects associated with the market for models discussed in Section 3.

[ADD MORE ANALYTICAL RESULTS FROM HOG CYCLE MODEL ABOUT HETEROGENEITY???] [It would be interesting to see if the level of exogenous noise affects the size of the prediction errors in the case when agents aren't that heterogenious.]

An interesting augmentation to the hog cycle model presented so far is to analyze what happens when the agents are allowed to consider the possibility of a structural break. Structural breaks are an important part of peoples' intuitive forecasting approach. For example, most economists agree that there was a regime change in US monetary policy around the time of the Volcker disiniflation. However, structural breaks are not allowed for in most learning models, even those with more complicated learning algorithms such as neural networks.[15]

The standard econometric analysis of structural breaks can be incorporated into the econometric learning framework developed so far.[16] First let us consider the case of a single structural break. The time period of the structural break is identified as the period k when the value of the following

[15] However, setting the "gain" parameter at a constant level is one way of accounting for frequent structural breaks in standard macroeconomic learning models. (See Evans and Honkapohja (2001).

[16] See Stock (1994) for an overview of the econometrics literature about testing for structural breaks.

statistic is maximized,

$$F_T(\frac{k}{T}) = \frac{SSR_{1,T} - (SSR_{1,k} + SSR_{k+1,T})}{\Omega}, \tag{8}$$

where $\Omega = (SSR_{1,k} + SSR_{k+1,T})/(T - 2dim(X_t)$. The Quandt Likelihood Ratio (1960) statistic can be shown to have a stable distribution depending only on the position of the structural break. (???)

This distribution can be used to perform a statistical test for whether the structural break occurred. The test takes the form,

$$F_T(\frac{k}{T}) > \alpha, \tag{9}$$

which can be rewritten as,

$$SSR_{1,T} - (SSR_{1,k} + SSR_{k+1,T} + \alpha * \Omega) > 0. \tag{10}$$

Thus, the standard test for structural breaks can easily be incorporated into our framework if the penalty for the increased complexity associated with the structural break model is assumed to be $\alpha * \Omega$.

A similar approach can be used in the case of multiple structural breaks.[17]

[???The changes in the agents' model selection already create considerable non-stationarity in the hog cycle data. Thus, it would be interesting to consider the effects of allowing for structural breaks even in the standard model. It might also be interesting to consider the effects of the augmented model set given some other type of regime change.]

2.3 Feedback Effects and Structural Breaks

A number of features of the econometric learning framework are evident from the hog cycle example.

First, when the agent's expectations significantly affect the dynamics of the phenomenon he is trying to predict (or if an aggregate of identical agents plays this role) then the feedback effects can make the dynamics of the phenomenon considerably more difficult to predict.

The effect is essentially a variant of the Lucas critique. The revisions in the agent's model affect his actions. These actions, in turn, affect the data generating process of the phenomenon he is trying to predict; and these effects invalidate the model he has built using past data. In contrast, the standard case in the econometrics literature assumes that the forecasting approach has no effect on the data generating process.

[17]See Bai and Perron (1998) for a discussion of the econometric theory associated with testing for multiple structural breaks.

In the macroeconomics literature, a number of results have been derived regarding the stability of equilibria given these feedback effects. Evans and Honkapohja (2001) derive conditions for stability for a fairly broad specification of the learning process. However, our approach does not fall into the Evans and Honkapohja (2001) framework. We assume that agents use data to update their choice of models, as well as to estimate parameters; while Evans and Honkapohja (2001) consider rules that broadly ressemble least squares learning. Another difference is that in our framework, agents' initial expectations are likely to be far from rational expectations (due to the paucity of the available data); whereas, Evans and Honkapohja (2001) focus on the dynamics of models in the neighborhood of the rational expectations equilibirum.

Second, as with other types of learning models, econometric learners respond considerably differently to regime changes than agents with rational expectations. In the econometric learning framework, agents select models that have performed well in the past, all else equal. However, almost by definition, models that performed well in the past cease to work well following a structural break.

In contrast, agents with rational expectations are assumed to know immediately when a regime change has occurred– for example, a change in the management policy of a firm, or in the monetary policy of the central bank. This assumption seems unrealistic. Moreover, it is difficult to accomodate the intuitive idea of "credibility problems" in the rational expectations framework, while credibility is a prime issue for actual central banks.

One novel feature of the econometric learning framework is that it allows agents to consider models with structural breaks, as we discuss in hog cycle example in Section 2.2. In contrast, other types of learning approaches generally do not allow for this type of non-stationarity. For example, least squares learners continue to update their expectations according to OLS even in the face of large discontinuous change in the observable data.

Structural change is sometimes accounted for in macroeconomics models by adjusting the "gain" parameter– intuitively, the discount factor for new data.[18] However, adjusting the gain parameter is more appropriate for dealing with small, frequent structural changes than large infrequent changes such as leadership changes at a central bank or a firm.

Similarly, econometric learners are also likely to have difficulty forecasting phenomena that have a short history.

[18] See Evans and Honkapohja (2001) or Sargent (2002) for a discussion of the gain parameter.

2.4 Testing the Consumption Euler Equation

In a classic paper, Hall (1978) shows that in the standard life-cycle conssumption model, the marginal utility of consumption is a random walk with trend. Thus, according to theory, variables such as income and wealth do not provide any additional information useful for forecasting consumption, beyond the information provided by current consumption. The intuition for this result is that any new information causes the agent to immediately adjust his consumption, encapsulating the new information in the current level of consumption.

Hall (1978) tests this result by estiming a linear regression model for consumption as a function of lagged consumption and other variables, namely disposable income and stock market returns. Hall finds that while the income variable is insignificant in the forecasting equation, the stock market variable is significant. Hall concludes that the results of the income test support the lifecycle model, while the results of the stock market test suggest that some modification to the model is necessary.

Rational expectations play an important role in Hall's approach to testing the consumption Euler equation since the consumers are assumed to anticipate the correlation between consumption and stock market returns in 1945, even though the correlations are estimated for the period 1945-1977.[19]

The finance literature has tended to be rather skeptical of the assumption that correlations involving stock returns could reasonably haven been anticipated in advance. Cochrane (2001) writes,

> Here is the fundamental issue: Was it clear to people in 1947 (0r 1871, or whentever one starts the sample) and throughout the period that the average return on stocks would be 8% greater than that of bonds, subject only to the 16% year-to-year variation?... I do not think it *was* obvious in 1947 that the United States would not slip back into depression, or another world war, but would instead experience a half century of economic growth and stock returns never seen in human history. Eight percent seems like an extremely– maybe even irrationally– exuberant expectation for stock returns as of 1947, or 1871. (You can ask the same question, by the way, about value effects, market timing, or other puzzles we try to explain....)

In this section, we revisit Hall's tests of the consumption Euler equation. However, we assume that the agents are econometric learners. Thus, they rely only on models that seem justified by

[19] Mankiw, Reis, and Wolfers (2003) is a more recent example of this type of testing approach in the macroeconomics literature.

the data they have observed so far. In order to avoid the unit root problems inherent in Hall's analysis, we assume that the agents are attempting to forecast changes in consumption using a linear regression model. They consider models including some combination of the following variables: several lags of consumption growth, lagged changes in disposable income and real stock market returns. Table 1 shows that...

[ADD RESULTS HERE]

3 The Market for Models

In many circumstances it is natural to think of the agents as building their own models. In other cases, however, it is useful to assume that the models are developed by a sector of the economy that specializes in model building. For example, firms often hire consulting firms such as Accenture to develop models of the potential marketability of their products.

In this section, we assume that models are produced by a "consulting" sector and then "bought" by agents in the "real" sector. We separate the costs of model production from the costs of model use, and the penalty for complexity associated with scientific model choice. This formula allows us to think about prices and subsidies in the "market for models".

We assume a simple competitive market structure for the consulting sector. There exist a large number of identical consulting firms. There also exist a large number of agents. The agents are heterogeneous in that each one of them needs models to explain different phenomena – his/her personal circumstances. The consulting firms must therefore build specific models for each customer. Since there are many agents and many firms, neither the agents nor the firms have any market power. The equilibrium price of the models will therfore equal the firm's cost of producing them. As in the home production case, we assume that the cost of producing a model is increasing in the complexity of the model.

In each period the agent seeks to maximize expected utility, where the expectations are formed using the model the agent brings into the period. The part of this optimization problem that has to do with how the agent optimally trades one action off against another is standard. The novel part of the agent's problem is his decision about which model to buy and how much resources to spend on buying new models. The agent will make these decisions in the following way: First, for each level of past performance, the agent will only consider the model with lowest complexity. The agent will then optimally choose the level of past performance of the model, trading better past performance against the cost of a more complex model as well as the direct utility loss from complexity. The agents optimal choice is given by the solution to the following set of equations:

$$V_{\eta_j}(m(\eta^p), \eta^p) = p_m(m(\eta^p))m_{\eta_j}(\eta^p)\Lambda - V_m(m(\eta^p), \eta^p)m_{\eta_j}(\eta^p), \tag{11}$$

where $m(\eta^p)$ denotes the lowest complexity level that attains the chosen level of past errors and Λ denotes the agents beliefs about marginal utility of income given the model he brought into the period. Subscripts denote partial derivatives. The term on the left hand side of this equation represents the direct marginal utility improvement that results from choosing lower past errors. The first term on the right hand side represents the utility loss due to lower consumption that results from having to buy a more complex model in order to be able to attain a lower level of past errors. The second term on the right hand side represents the direct utility loss from increasing the complexity of the model in order to be able to attain a lower level of past errors.

It is clear from equation (11) that agents who live in a complex world may optimally choose not to form predictions using a model that is complex enough that is makes full use of all information in the agents data set. This will be the case if there exists another model that is sufficiently simpler but still produces reasonably good predictions. More complex models will always perform better in terms of past errors. However, this does not mean that the agent will use the most complex model, since he bears the cost of purchasing the model and of using it. But also, because he is wary that complex models may be "fitting the noise" in the past data and may therefore not perform well in the future (Occam's razor).

The bias toward simpler models and models that results from the costs of building complex models can only be justified on scientific grounds to the extent that the penalty for complexity that is typically argued for on scientific grounds is too low.

It is also clear from equation (11) that agents will spend more resources on modelling important phenomena, i.e. phenomena with high $V_{\eta_j}(m(\eta^p), \eta^p)$.

3.1 Comparative Statics

We can investigate comparative statics in the market for models using the budget line and indifference curve. The indifference curve and budget line can be drawn in V-C space, where V is the utility derived from modelling a particular phenomenon and C is consumption. Since the agent divides his resources between modelling and consumption, C is related inversely to the price of the model and complexity. The choice of models for all other phenomena held fixed in this analysis.

The indifference curve in V-C space reflects the tradeoff the agent perceives between consumption and the model selction criterion. Since there are diminishing returns to both higher V and higher consumption, the indifference curve has the usual concave shape.

In the case of the budget line, while more complex models cost more (and thus yield lower C), they do not necessarily yield higher values of V. As a consequence, the budget line does not always have the usual downward sloping linear shape.[20] Figures 1 and 2 depict the type of shapes the budget line is likely to have in $V - C$ space.

We assume that there is free disposal of consumption so the budget set is convex. We also assume that for any given level of complexity, a model with an arbitrarily low level of V can be found (for example by using a totally unrelated model). These restrictions ensure that the budget line is a 1-1 mapping.

Case 1: Budget Line with High Returns to Complexity

Figure 1a depicts the case where increases in the complexity of the model (which correspond to decreases in consumption) lead to higher values of V. In this case, the penalty imposed by V for higher complexity is more than compensated for by better fit of the model. For example, in the simple case where complexity is simply the number of variables in a linear model, this is the case where adding an additional variable, up to a certain point, always yields an increase in V. Note, however, that the plateau in the budget line reflects the point where V cannot be increased further.

Case 2: Budget Line with Low Returns to Complexity

Figure 1b depicts the case where the simplest model (which corresponds to the highest value of C) is best. In this case, adding additional complexity to the model never yields any increases in V. As we discuss below, the optimizing choice of C and V in this case is always a corner solution.

Case 3: Budget Line with Varying Returns to Complexity

Figure 2 depicts what we think of as the "typical" shape of the budget line in the market for models. There is a range over which increases in complexity yield higher values of V so the budget line is downward sloping. However, there are also significant ranges where higher complexity does not yield higher values of V (the flat spots to the left and right of the graph).

Now let us consider the consider the comparative statics results associated with modifying various parameters.

First, consider the effects of changing the cost of model production on the agent's model choice.

[20] In addition, in practice, the complexity of a model often increases in discrete jumps– for example, with the addition of new variables. As a consequence, the budget line may also be non-differentiable at other points. We ignore these issues here.???

A decrease in the cost of model production causes the downward sloping portion of the budget line to pivot outward. As a consequence, if the initial solution occurs where the budget line is downward-sloping, the decrease in cost yields a more complex model.

However, if the initial solution is a corner solution, then the change in the cost of model production has no effect on the choice of model. Figures 1 and 2 show that corner solutions arise in two main cases: either the initial solution is the simplest possible (which usually occurs if the simplest model is fairly good) or the initial solution has the highest possible level of V. Indeed, if the simplest model (with zero cost of complexity) is best as in Figure 1a, the change in the cost of model production leaves the budget line unchanged.

Second, consider the effects of changing the agent's value for modelling, or "curiousity"– that is, the extent to which an agent is willing to trade off consumption for an improved model. An increase in the agent's curiousity leads to a flatter indifference curve, reflecting a more favorable tradeoff between V and C.

If the budget constraint is downward sloping in the region of the initial equilibrium, then a flatter indifference curve yields a more complex model, with lower C and higher V. Once again, a flatter indifference curve does not yield a more complex model if the initial solution is a corner solution.

A simpilar line of reasoning applies to changes in the relative importance of various phenomena. Consider the simple case in which V can be written as the sum of the model selection criteria for several different phenomena, and the models for the different phenomena are completely independent. That is, for phenomena $i = 1...n$,

$$V = a_1 V_1(m_{1,t}, \eta^p_{1,t}) + a_2 V_2(m_{2,t}, \eta^p_{2,t}) + ... + a_n V_n(m_{n,t}, \eta^p_{n,t}), \tag{12}$$

where each $V_i(m_{i,t}, \eta^p_{i,t})$ reflects the quality of model i, $m_{i,t}$ is the complexity of model i and $\eta^p_{n,t})$ is the prediction error associated with model i. In this case, an increase in the relative importance of phenomenon i can be expressed as an increase in the value of a_i.

In this case, the same line of argument we used in analysing the effects of differences in the agent's curiousity applies to differences in the relative importance of various phenomena. In particular, an increase in a_i leads to an increase in the complexity of the model associated with phenomenon i and V_i if the initial solution lies on a downward sloping segment of the budget line. This observation is closely related to Fryer and Jackson (2003)'s result that, in a categorical model with a fixed number of categories, outlyer phenomena are modelled with coarser categories.

One way of defining rational expectations would be as the expectations associated with the model with the highest value of V. This would be the model chosen if the costs of model production were zero (though V still incorporates the costs of model use).

We can think of an agent as "model-constrained" if he is not using the model with the highest value of V. Model-constraints are one way of understanding the intuitive idea that agents have difficulty planning. Model constraints are likely to be a problem when the best model of a phenomenon is quite complicated. Model constraints will *not* be a problem when the simplest model is best, since the corner solution is always selected. (See Figure 1a).

If an agent is model-constrained then subsidizing the costs of complexity may have an impact on his choice of model, through the price effects discussed above.

3.2 The Black-Scholes Formula and Options Pricing

An interesting example of how the introduction of a new model effected the workings of a market is the introduction of the Black-Scholes formula in options pricing. Before the introduction of the Black-Scholes (B-S) formula, formal options markets did not exist, largely because financial traders did not know how to price, and more importantly, how to hedge, options. The introduction of the Black-Scholes formula in 1973 had an immense and rapid effect on options markets. Formal options markets developed rapidly after the introduction of the B-S formula.

Moreover, only months after the publication of the B-S formula, Texas Instruments introduced a calculator specifically desigened to price options using the B-S formula. No doubt, the new gadget significantly reduced the cognitive costs of using the B-S formula to options traders.

In introducing the B-S formula, Black and Scholes (1973) note that Before the 1970's there were systematic deviations between options prices and the B-S formula. However, the authors argue that these deviations were associated with the large transactions costs resulting trading options in the absence of formal markets. (See Black and Scholes (1973)).

From the early 1970's until the crash of 1987 the B-S model fit the data remarkably well. In particular, there was very little relationship between implied volatility and the extent to which the option was in the money – i.e. very little "volatily smile"– during this period. (See Bates (2000).)[21] Indeed, some researchers argue that the B-S formula fit *suspiciously* well during the

[21] The "volatility smile" refers to the characteristic relationship between the implied volatility of a derivative and the extent to which the derivative is in the money. The volatility estimates should be identical according to the basic B-S approach.

1970's, given its simplistic assumptions. Indeed, Cherian and Jarrow (1998) argue that the B-S model is a self-fulfilling prophecy.(See Cherian and Jarrow (1998).???)

However, the behavior of options prices changed dramatically following the crash of 1987. After the crash, options prices developed considerable volatility smiles and smirks, as discussed by Bates (2000).

Bates (2000) argues that the dramatic change in option price behavior can be explained by a a change in investors' assessment of the S&P 500 stochastic process. In particular, Bates argues that the crash raised investors perceptions of the probability of future stock market crashes, a perception that Bates notes was validated somewhat by further stock market drops of 5-8% on January 11 1988 and October 13 1989. If this interpretation is correct, the Black-Scholes case provides a dramatic example of investors adopting a new and more complex model of their environment, with important implications for the derivatives market.

4 Managing the Model

When the payoffs to an action are unknown, the model an agent buys can play a large role in determining whether a particular action seems appealing. For example, a worker has more incentive to work hard if he believes that he will be rewarded with high wages and job security in the future.

This type of complementarity between actions and models provide an incentive for large agents such as firms and governments to try to "manage the model" adopted by other agents. For example, a firm has an interest in convincing its workers that it will reward hard work in the tuture. Similarly, the US government has a large interest in convincing foreign leaders that the US will retaliate with nuclear weapons if it is attacked with nuclear weapons. Moreover, the OPEC members have a large interest in convincing fellow members that any country that deviates from OPEC pricing will be punished with a price war.

Indeed, many of the activities of firms' managers can be associated with "managing the model" adopted by the firm's employees, customers, and business partners. There are two main ways in which a large agent can try to "manage" the model adopted by other agents.

First, if the agent is model-constrained, it is often possible to alter the agent's choice of model by subsidizing the production of models.

1. The large agent can subsidize the production of a particular model, or of model production in general.

In a practical situation, this might be achieved by, for example, offering free education about a new model. For example, firm managers spend a great deal of time outlining their company philosopy to business partners and workers. Similarly, US diplomats make great efforts to communicate the US foreign policy doctrines– i.e. the Bush doctrine, the Monroe doctrine etc.– to foreign officials. None of this would be necessary if the parties involved were operating in the context of a known model.

Model subsidies do not, however, give the firm full control over the model adopted by the workers. While the costs of producing the model (paid to the consulting firm) can be subsidized, the costs of using the model cannot be subsidized. As a consequence, even if the cost of a structural model is reduced to zero, agents may still be unwilling to accept a more complicated model if it does not fit the data significantly better.

Moreover, model subsidies may not affect the agent's choice of model if he is at a corner solution– for example, if he is already using the model with the highest V or the simplest model. (See Section 3.1).

Notice that the incentives to subsidize a model may exist even if the model is false. For example, it is in the firm's interest for the workers to believe they will be rewarded for hard work, and therefore work hard, whether or not the belief is true.

When the large agent also has some control of the data observed by the other agents, there is another avenue available for managing the model.

2. The large agent can manage the "dataset" used by other agents.

In some cases, the large agent is actually involved in generating the data. For example, a firm has control of its layoff policy; the US has control of its foreign policy. In these cases, the large agent can manage the model in part by adjusting the data generating process.[22]

In other cases, the large agent controls the dataset through some form of information control. One approach is to limit the agents' access to data. For example, the North Korean government clearly believes that severe limits on their population's access to information play an important role in upholding their model of government. (North Korea also subsidizes its governmental approach considerably by providing considerably "education" about the government's intentions and approach.)

Alternatively, the large agent can manage other agents' expectations by choosing to interact with only those agents who have been exposed to (or shielded from) particular information. For

[22]This view of managing the model is closely related to Fryer and Jackson (2003)'s concept of "image management".

example, prosecutors try to avoid having people on juries who have had bad experiences with the police or childhood trauma. We discuss another example of this type of behavior by Japanese firms in Section 4.4.

The issue of "managing the dataset" to affect other agents' expectations also arises in the context of a known model. For example, Holmstrom (1999) analyzes a setting in which workers work particularly hard at the start of their career since they know that good performance at this stage has a particularly large impact on their employer's perception of their talents. (In this case, the worker is managing the model adopted by the firm.)

However, the analysis is even more interesting in the context of an unknown model. First, modifying the dataset can have a particularly large effect on the agent's expectations when the model is unknown since since the agent has the option of switching models.

Second, in order to analyze the effects of managing the dataset in a known model, one must first specify a stochastic model and prior. For example, Holmstrom (1999) specifies a prior over types of workers and a stochastic process for the workers' performance. But, as we discuss in the introduciton, what would be a reasonable stochastic model and prior depend on the existing data. The framework of an unknown model provides a connection between the existing data and the stochastic model and prior that is common to both the researcher and the agents in the model.

4.1 Example of Market Failure

The choice of a model often has externality effects as I discuss above. As a consequence, the competitive equilibrium in the market for models may produce a Pareto loss.

For example, consider the following interpretation of the repeated game prisoner's dilemna. Player 1 is a firm. Player 2 is a worker. The firm can either pay a high or low wage, and the worker can exert high or low effort. For simplicity, we assume that the firm has perfect information about the game, while the worker must form a model about the firm's behavior.

As a consequence, we have the following payoff matrix.

...............High Wage..........Low Wage

High Effort.....(10,10).............(9,15)

Low Effort......(15,9)...............(2,2)

The Nash equilibrium of the static game would be low wage and low effort. However, the cooperative outcome clearly Pareto dominates the Nash equilibrium of the static game. The cooperative outcome can be achieved if both the firm and the worker have strategies that reward their opponents in future periods for playing the cooperative strategy (via the folk theorem).

Suppose the exogenous variables considered by the workers includes the past wages paid by the firm, and the matrix of payoffs for each period. Suppose also that the worker and firm have been playing the static Nash equilibrium (low effort, low wages) for the past 100 periods.

Let us focus on two possible models. The first model is a statistical model,

$$w = 10. \tag{13}$$

This model fits the existing data perfectly. Moreover, it is likely to be judged to be very simple according to any reasonable complexity criterion.

The second model is the standard game theoretic structural model. This model predicts that the outcome of the game is a Nash equilibrium, which could be either the cooperative or the static non-cooperative outcome. This model is a function of the payoffs and is considerably more complicated than the statistical model.

Evidently, both models fit the data, but the structural model is more complex. In this case, whether the worker is willing to "buy" the structural game theory model depends on how useful it is in explaining other important phenomena. Moreover, he worker's choice of model is crucial to the outcome of the game since the cooperative equilibrium is unachievable if the worker adopts the statistical model.

Suppose the worker chooses not to buy the structural model due to a small difference in complexity between the structural and statistical models. In this case, the externality effects associated with model choice cause a Pareto loss since the agent fails to take into consideration the potential effects of his model choice on the firm's behavior.

4.2 The Role of the Government in Model Production

A number of government agencies are involved in model production. Consider, for example, the research arm of the Federal Reserve Banks, the Japanese business association (see Section 4.4), or government funding of research universities. In addition, public schooling probably reduces the costs of using, as well as producing, complex models.

There are two main justifications for government subsidies to model production.

First, there is the externalities argument. In some cases, the private sector has a great interest in producing good models– for example, in the area of financial research. However, in other cases, the private sector has little interest in subsidizing model production, or even has an interest in subsidizing false models.

Indeed, in a number of areas, academics have made a significant contribution by developing complex models where none would have otherwise been developed. Campbell and Viceira argue that academics have an important role to play in educating the public about the appropriate models of long-term investment. Vissing-Jorgenson investigates what kind of model seems appropriate for entrepreneurial returns– and finds that government encouragement of entrepreneurial investment is probably misplaced.

When there are externalities to model choice, the public benefits from using a more complex model may outstrip the private benefits of acquiring the more expensive model, as we discuss above. In this case, government subsidies to model production can yield a Pareto improvement.

Second, models are often useful to more than one agent. Moreover, models are generally non-rival and non-excludable objects. In these cases, government intervention (or regulation) is necessary to ensure production of the relevant models. The government also plays an important role in ensuring the production of non-excludable models by enforcing intellectual property laws.

Both of these motivations also exist in the standard information economics case where the agent is deciding whether or not to acquire a piece of information.

4.3 Data Mining

The analogy between the problem of an economic agent in an unknown model and an econometrician reminds us of a persistent problem in econometrics: data mining. Roughly speaking, the data mining problem refers to cases in which the usual statistics must be modified to account for the fact that the researcher doing the analysis had an interest in arriving at a particular conclusion.[23]

In particular, if the data were subjected to some form of pretest before the main analysis was conducted– for example, only data fitting a particular hypothesis was considered– then the pretest should be accounted for in the statistics. The data management approach of only providing the agent access to data supporting a particular hypothesis is just such a pretest. Thus, the usual model selection criteria are not appropriate when this type of approach is being used by the large agent. Rather, an optimizing agent should make use of modified model selection criteria to the extent that he is aware of the tampering with the dataset.

In contrast, the data mining issue is not a problem when the other approaches are used to

[23]See xx and xx for a discussion of the data mining problem in the econometrics literature.

manage the model. An agent's decision to subsidize a particular model should not affect the criterion used to judge the model, though the reduced cost of purchasing the model should be taken into consideration. Similarly, a change in the data generating process should not affect the model selection criterion. However, the decision to subsidize a model may provide new information, for example about the incentives of the large agent.

4.4 Real-world Examples of Managing the Model

In this section, we discuss two real-world examples of firms' attempts at "managing the model" adopted by their employees.

4.4.1 Japanese Management Practices: Nihonshiki

A number of the examples we have discussed so far focus on the implications of model choice for achieving cooperative equilibria. In a number of ways, Japanese firms and workers seem to have achieved a high-return cooperative equilibrium to the labor management problem in recent years.[24] Japanese workers generally enjoy a high level of job security and large bonuses that depend both on employee marit and firm profits. In return, Japanese workers tend to be more hardworking, and have much lower quit-rates than their American counterparts. Moreover, Japanese workers generally react less adversely to new technologies than their American counterparts, since they are less concerned that the technological improvements will provoke layoffs.[25]

However, what are today known as Japanese labor management practices were actually introduced to Japan relatively recently– in the 1950's and 1960's. Before that time, the labor management practices of Japanese firms ressembled current American labor management practices. For example, large-scale layoffs were relatively common. In particular, prior to the 1920's or 1930's, it has been documented that Japanese workers were not especially loyal, hard-working or dedicated.[26] Moreover, Japanese quit rates were similar to those in the US before WWII. Following many bitter labor disputes and strikes resulting from oppressive management, union bashing, low wages and long working hours, firms undertook a conscious effort to institute a new style of labor management. What we currently regard as contemporary Japanese labor management practices came to

[24] However, some of these practices have begun to change over the past decade as a consequence of Japan's prolonged recession.

[25] See Hashimoto (1990) for a comparative study of Japanese and American labor management practices.

[26] See Fruin (1983) or Nakamura (1993).

be widely adopted only after the 1950's.[27]

Japanese industry associations of various sorts, as well as other economy-wide management organizations played a large role in introducing a more cooperative type of labor management approach to Japan, by organizing meetings and disseminating information about the new management approach. These organizations also helped to strategically coordinate the switch to modern Japanese management practices among a number of large Japanese firms. Japanese management practices then diffused through the economy to small and medium-sized firms as well.

Government organizations also played an important role in the transformation. In particular, the Japan Productivity Center, with the support of the Ministry of International Trade and Industry, played an important role in disseminating information about the new management practices. The Japan Productivity Center was established in 1955 jointly by organizations representing all sectors of the Japanese economy including firms, government and academia. The center was founded for the purpose of improving Japan's productivity.[28]

In the context of our earlier discussion in Section 4, the efforts of these organizations to educate firms and workers can be understood as as subsidies toward the cooperative model. Moreover, the organizations' efforts to coordinate the switch to the new labor management practices among Japanese firms was, among other things, an attempt to avoid any firm being an outlier in its labor management policies. As we discuss in Section 2, outlier phenomena are often poorly modelled when there are costs of complexity.(???)

Nevertheless, the Japan's transition to the new cooperative equilibrium was a gradual process. Confidence in the system gradually increased as Japanese firms developed a track record for honoring the no-layoff policy. The gradual transition suggests that Japanese workers were learning about the new system, much like the farmers in our hog cycle example in Section 2.2.

The story continues when Japanese auto manufacturers tried to export Japanese management practices to America. In the 1980's Japanese auto manufacturers began producing cars in the US. The Japanese firms also wanted to transplant Japanese labor management practices to their US plants.

However, the Japanese labor management approach was clearly an outlier relative to the approaches of other American auto manufacturers. The Big Three auto manufacturers had frequent layoffs, high absentee rates and frequent strikes relative to the Japanese firms.

[27] See Taira (1970), Tsurumi (1978), Saxonhouse (1976), and Aoki (1988) pp. 189-191.

[28] For a discussion of the role of the Japan Productivity Sector in affecting Japanese industrial relations policy see Hart and Kawasaki (1999).

Since American workers who were employed by Japanese plants were completely unfamiliar with the Japanese labor management model, Japanese firms spent an enormous amount of resources trying to educate American workers and managers about the Japanese labor management model. In addition to holding lectures and conferences on the Japanese labor management approach, new American workers were often brought to Japan, at considerable cost, to observe the workings of Japanese plants.

Japanese firms also invested a considerable about of conscious effort in trying to manage the datasets of their employees. Japanese auto manufacturers intentionally located their factories in "greenfields" communities in the south or the midwest, far away from existing car manufacturing plants, in order to avoid hiring workers with prior experience with US auto manufacturers. In general, Japanese auto manufacturers' actually preferred to hire new highschool graduates with no labor force experience at all, who would start with a blank slate in learning about Japanese employment practices.

4.4.2 Hewlett-Packard and the HP Way

Hewlett-Packard (HP) is an example of another case in which a firm managed for a period to maintain a successful cooperative equilibrium with its workers.[29] HP's labor management approach was characterized in the 1970's by a set of corporate objectives, business values and management practices that had been passed down since the company's founding, and were known as the "HP way". HP management put a great deal of effort into disseminating the HP way to employees. Employees were exposed to a costant stream of discussion about the HP-way.

Job security was viewed by many employees as an integral part of the HP-way. Until the 1980's, HP had a no-layoff policy– although employees could be laid off for poor performance, HP never laid people off because of problems with overall economic conditions. During an economic downturn in the 1970's, HP avoided layoffs by instituting a pay cut and a shorter work-week. The HP-way was a source of much loyalty and satisfaction on the part of HP employees.

However, in the 1980's, economic conditions caused HP's management to decide that a reduction in HP's workforce was necessary. In order to achieve this goal, HP introduced downsizing and redeployment. In the downsizing initiative, HP offered employees voluntary retirement packages and severence incentives. In the redeployment initiative, Employees were given 3 months to find a

[29] The details about HP in this section are based on two Harvard Business School cases: "Human Resources at Hewlett-Packard A and B" (Case numbers 9-495-051 and 9-495-052).

new job within HP; and if they failed to find a new job themselves, HP found them a new job– but the new job could involve relocating anywhere in the country, and taking a pay cut.[30]

As a consequence of the new employment policies, many HP employees said that the HP way was dead. HP lost a number of its best workers, and the morale and loyalty of its workers dropped immensely.

HP responded to the crisis with a picture that has since become well-known in MBA classrooms. The picture is of three concentric circles, with the center circle representing HP's "core values" and the outer circles reflecting more ephemeral "corporate objectives" and "strategies and practices" respectively. HP's management argued that while "employment security" was a core value that could be expected to remain constant over time, "job security" was not. Moreover, even employment security could not be expected to remain constant in times of extreme economic duress.

HP subsequently spent considerable resources trying to educate workers about the "new HP-way". HP's efforts can be understood as efforts to subsidize the workers' acquisition of a new model (just as HP had spent considerable resources in trying to educate workers about the old HP-way). HP wanted workers to keep track of "employment probabilities", for which HP still had an unblemished track record, rather than "job loss probabilities".

HP's efforts to convince its employees about the new HP-way were not, however, entirely successful. Many workers rejected the new HP-way model in favor of a model with a structural break, with HP becoming a poor employer in the 1990's. Many of these workers left the company.

In our view, one reason why workers did not immediately buy the "new HP-way" was that it was significantly more complicated than the old HP-way. Moreover, unlike in Japan, the "new HP-way" was an outlier model.

In principle, HP could have introduced the concept of the "new HP-way" from the start– but this would have been a considerably more complicated model of HP's management approach, and the additional complications seemed unnecessary in the 1970's. Thus, HP's experience illustrates clearly the tradeoff between a complex, costly structural model and a simple statistical model: the latter is much cheaper to disseminate, but may lead to significant "inertia" if parameters of the structural model change.

To the extent that HP was serious about its new labor management policies (which it appears to have been) the failure to convince workers about the new HP-way model produced a significant

[30] Interestingly, this change in HP's labor management policy moved HP closer to the Japanese labor management model, which has always allowed for reassignments to new tasks or locations (sometimes in other countries).

Pareto loss.

5 Conclusion

Much of the economic analysis of decision-making focuses on how agents should maximize expected utility in the context of a known stochastic model. Yet, much of what makes decision-making hard in everyday life, as well as in economic policymaking, is developing the appropriate model— rather than optimizing once the model has been agreed upon.

In this paper, we provide a general framework for thinking about this aspect of economic behavior. We emphasize three main points. First, model choice is an important aspect of rationality when the true model is unknown. Second, modelling is a costly activity, and changes in the costs of modelling have important effects on peoples' decisions about which model to use. Third, one's choice of model often has externality effects on other people. As a consequence of these externalities, firms and governments often have an incentive to intervene in the model formation process.

The problem of writing down rules for rational inference in an unknown model is certainly a difficult one. Indeed, an infalliable approach is probably unattainable. Nevertheless, analyzing the large literature in econometrics and computer science on the model choice problem can help us to at least choose rules of inference that are not blatantly ridiculous. In this sense, we follow in the footsteps of Savage (1954) and Hansen and Sargent (2000) in suggesting that microtheory and econometrics are not far apart as they might seem from the familiar dichotomy between empirical and theoretical research.

References

Aoki, M. (1988): *Information, Incentives and Bargaining in the Japanese Economy*. Cambridge University Press, Cambridge and New York.

Bai, J., and P. Perron (1998): "Estimating and Testing Models with Multiple Structural Changes," *Econometrica*, 66(1), 47–78.

Bates, D. S. (2000): "Post-'87 crash fears in the S&P 500 futures option market," *Journal of Econometrics*, 94, 181–238.

Black, F., and M. Scholes (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economics*, 81(3), 637–654.

Bray, M. M., and N. E. Savin (1986): "Rational Expectations Equilibria, Learning and Model Specification," *Econometrica*, 54(5), 1129–1160.

Brock, W. A., and C. H. Hommes (1997): "A Rational Route to Randomness," *Econometrica*, 65(5), 1059–1095.

Brunnermeier, M. K., and J. A. Parker (2003): "Optimal Expectations," Working Paper, Princeton University.

Campbell, J. Y., and L. M. Viceira (2001): "Who Should Buy Long-Term Bonds," *American Economic Review*, 91(1), 99–127.

Cherian, J., and R. A. Jarrow (1998): "Options Markets, Self-fulfilling Prophecies, and Implied Volatilities," *Review of Derivatives Research*, 2, 5–37.

Cochrane, J. H. (2001): *Asset Pricing*. Princeton University Press, Princeton, NJ.

Cover, T. M., and J. A. Thomas (1991): *Elements of Information Theory*. John Wiley & Sons, New York.

Cyert, R. M., and M. H. DeGroot (1974): "Rational Expectations and Bayesian Analysis," *Journal of Political Economics*, 82(3), 521–536.

Decanio, S. J. (1979): "Rational Expectations and Learning from Experience," *Quarterly Journal of Economics*, 93(1), 47–57.

Dekel, E., B. L. Lipman, and A. Rustichini (1998): "Standard State-Space Models Preclude Unawareness," *Econometrica*, 66(1), 159–173.

Evans, G. W., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.

Evans, G. W., and G. Ramey (1992): "Expectation Calculation and Macroeconomic Dynamics," *American Economic Review*, 82(1), 207–224.

Fruin, W. M. (1983): *Kikkoman*. Harvard University Press, Cambridge.

Fryer, R. G., and M. O. Jackson (2003): "Categorical Cognition: A Psychological Model of Categories and Identification in Decision Making.," unpublishe, University of Chicago.

Fudenberg, D., and D. K. Levine (1999): *The Theory of Learning in Games*. MIT Press, Cambridge, Mass.

Gilboa, I., and D. Schmeidler (1989): "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

Hall, R. E. (1978): "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economics*, 86(6), 971–987.

Hansen, L. P., and T. J. Sargent (2000): "Wanting Robustness in Macroeconomics," unpublished, University of Chicago.

Hansen, M. H., and B. Yu (2003): "Model Selection and the Principle of Minimum Description Length," Working Paper.

Hart, R. A., and S. Kawasaki (1999): *Work and Pay in Japan*. Cambridge University Press, Cambridge and New York.

Hashimoto, M. (1990): *The Japanese Labor Market in a Comparative Perspective with the United States*. Upjohn Institute, Kalamazoo, MI.

Holmstrom, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169–182, Special Issue: Contracts.

Hong, H., and J. C. Stein (2003): "Simple Forecasts and Paradigm Shifts," Working Paper, Princeton University.

Kearns, M. J., and U. V. Vazirani (1994): *An Introduction to Computational Learning Theory*. MIT Press, Cambridge, Mass.

Li, M., and P. Vitanyi (1997): *An Introduction to Kolmogorov Complexity and Its Applications (Graduate Texts in Computer Science)*. Springer Verlag, ???

Mankiw, N. G., R. Reis, and J. Wolfers (2003): "Disagreement about Inflation Expectations," NBER Working Paper 9796.

Nakamura, M. (1993): "Japanese industrial relations in an international business environment," *North American Journal of Economics and Finance*, 4, 225–251.

Rissanen, J. (1987): "Stochastic Complexity," *Journal of the Royal Statistical Society. Series B (Methodological)*, 49(3), 223–239.

Sargent, T. J. (1993): *Bounded Rationality in Macroeconomics*. Oxford University Press, Oxford, U.K.

——— (2002): *The Conquestion of American Inflation*. Princeton University Press, Princeton, New Jersey.

Savage, L. J. (1954): *The Foundations of Statistics*. John Wiley & Sons, New York, New York.

Saxonhouse, G. (1976): *Country Girls and Communication among Competitors in the Japanese Cotton-spinning Industry*. University of California Press, Berkley.

Sims, C. A. (1980): "Macroeconomics and Reality," *Econometrica*, 48(1), 1–48.

Stock, J. H. (1994): "Unit Roots, Structural Breaks and Trends," in *Handbook of Econometrics*, ed. by R. F. Engle, and D. L. McFadden, vol. IV, pp. 2740–2841. Elsevier Science B.V.

Taira, K. (1970): *Economic Development and the Labor Market in Japan*. Columbia University Press, New York.

Tsurumi, Y. (1978): *Japanese Business: A research guide with Annotated Bibliography*. Praeger, New York.

Zellner, A., H. A. Keuzenkamp, and M. McAleer (eds.) (2001a): *Simplicity, Inference and Modelling*. Cambridge University Press, Cambridge, U.K.

——— (eds.) (2001b): *Simplicity, Inference and Modelling*. Cambridge University Press, Cambridge, UK.

Figure 1a: Budget Line with High Returns to Complexity

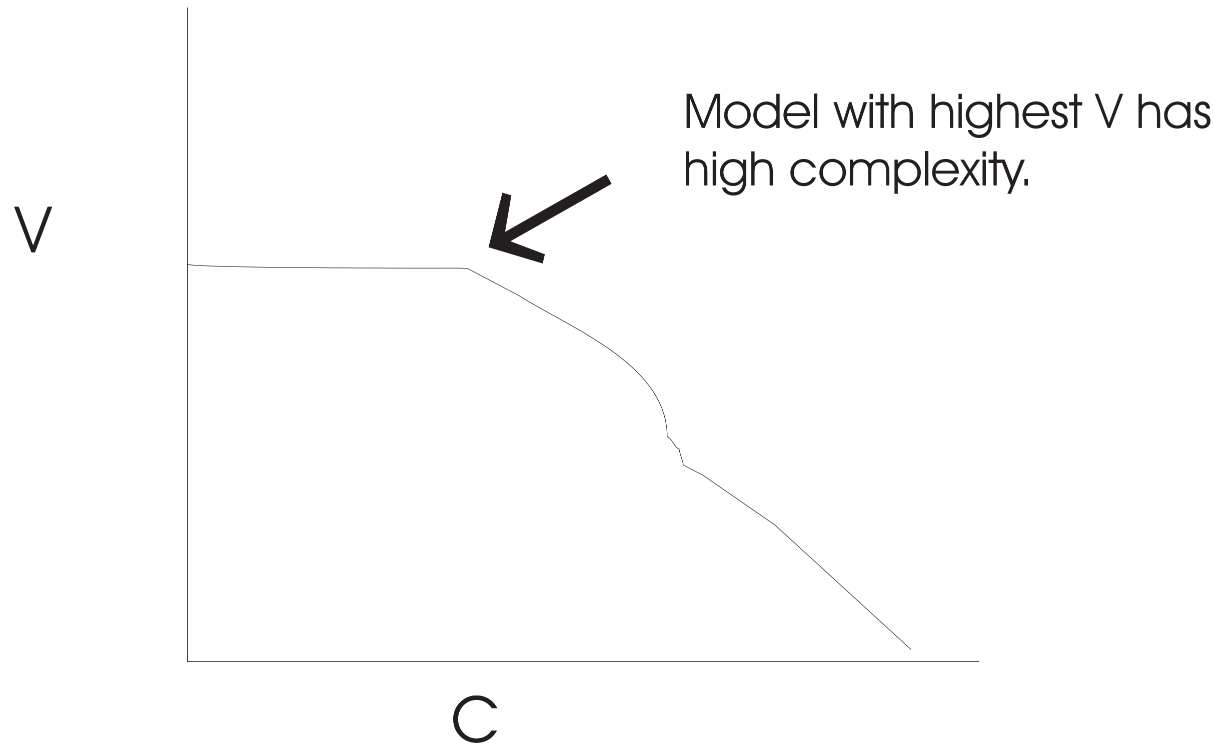

Figure 1b: Budget Line with Low Returns to Complexity

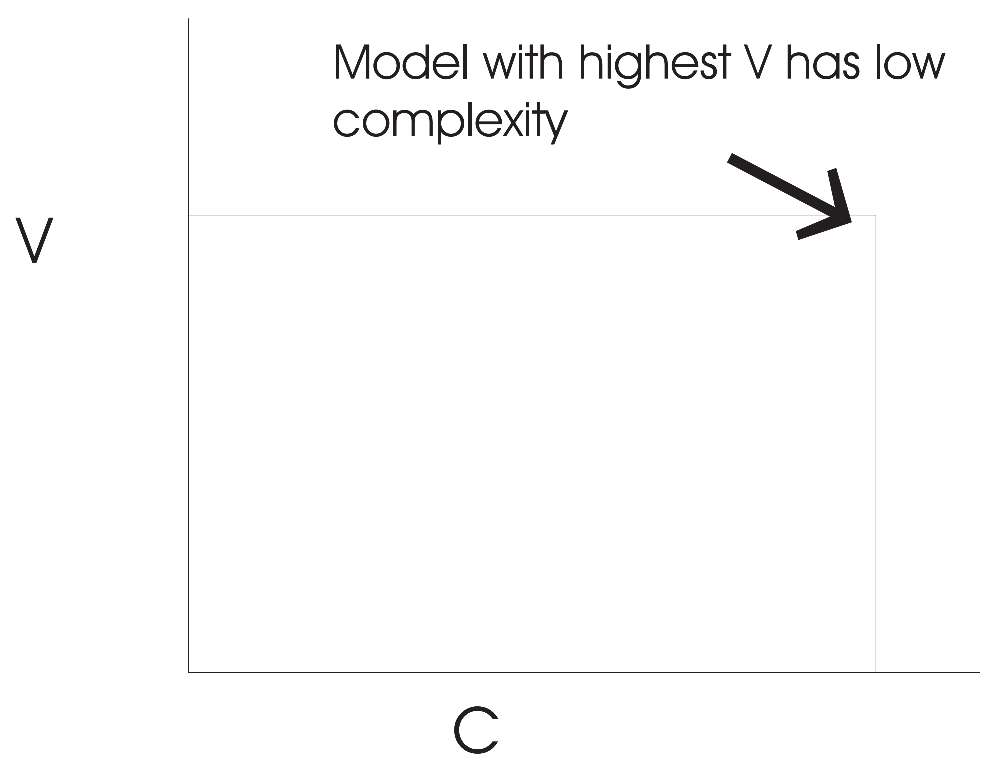

Figure 2: Budget Line with Varying Returns to Complexity