

Ran Spiegler

Bayesian networks and boundedly rational expectations

Discussion paper

Original citation:

Spiegler, Ran (2014) *Bayesian networks and boundedly rational expectations*. CFM discussion paper series, CFM-DP2014-17. Centre For Macroeconomics , London, UK.

Originally available from [Centre For Macroeconomics](http://www.cfm.ac.uk/)

This paper has benefitted from ESRC grant no. ES/L003031/1

This version available at: <http://eprints.lse.ac.uk/57994/>

Available in LSE Research Online: July 2014

© 2014 The Author

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

Bayesian Networks and Boundedly Rational Expectations*

Ran Spiegler[†]

July 7, 2014

Abstract

I present a framework for analyzing decision makers with an imperfect understanding of their environment's correlation structure. The decision maker faces an objective multivariate probability distribution (his own action is one of the random variables). He is characterized by a directed acyclic graph over the set of variables. His subjective belief filters the objective distribution through his graph, via the factorization formula for Bayesian networks. This belief distortion implies that the decision maker's long-run behavior may affect his perception of the consequences of his actions. Accordingly, I define a "personal equilibrium" notion of optimal choices. I show how recent models of boundedly rational expectations (as well as new ones, e.g. reverse causality) can be subsumed into this framework as special cases. Some general properties of the Bayesian-network representation of subjective beliefs are presented, as well as a "missing data" foundation.

*This paper has benefitted from ESRC grant no. ES/L003031/1. I am grateful to Noga Alon, Simon Byrne, Philip Dawid, Kfir Eliaz, Erik Eyster, Philippe Jehiel, Michael Woodford and seminar participants at Tel Aviv University and the 2014 BRIC workshop, for helpful conversations and comments.

[†]Tel Aviv University and University College London. URL: <http://www.tau.ac.il/~rani>. E-mail: rani@post.tau.ac.il.

1 Introduction

The rational-expectations postulate is a basic building block in the vast majority of economic models. It means that agents in an economic model have a *perfect* understanding of the statistical regularities that characterize equilibrium in the model - in particular, the structure of *correlations* among the relevant economic variables. Over the years, economists have become increasingly interested in modeling systematic departures from this extreme assumption. In this paper, I propose a rather general approach to modeling decision makers with an *imperfect* understanding of correlation structures, based on the borrowed concept of *Bayesian networks*.

I focus on choices by an individual decision maker (DM), whose environment is defined by an objective probability distribution p over $x = (x_0, x_1, \dots, x_n)$. One way of writing down p employs the standard chain rule

$$p(x) = p(x_0)p(x_1 | x_0)p(x_2 | x_0, x_1) \cdots p(x_n | x_0, \dots, x_{n-1}) \quad (1)$$

where the enumeration of the variables is arbitrary. This basic formula suggests a natural way of capturing an imperfect understanding of the correlation structure of p : *pick an enumeration and selectively omit conditioned variables*. This implies the following representation of subjective beliefs. The DM is characterized by an asymmetric, acyclic binary relation R over the set $N = \{0, 1, \dots, n\}$. Given that the objective distribution is p , the DM's subjective belief is given by the factorization formula

$$p_R(x) = \prod_{i=0}^n p(x_i | x_{R(i)}) \quad (2)$$

where $R(i) = \{j \in N \mid jRi\}$.

When R is any linear ordering over N , we recover the standard chain rule. This is the case of correct beliefs, i.e. "rational expectations". When R is empty, $p_R(x) = p(x_0)p(x_1) \cdots p(x_n)$ - i.e., p_R is the product of the marginals

of p over the $n + 1$ random variables. The more complete the relation R , the more thorough the DM's understanding of the underlying correlation structure - and his expectations are intuitively "more rational".

The binary relation R and the set of distributions representable by (2) define what is known as a *Bayesian network*. This concept was introduced by statisticians and Artificial-Intelligence researchers, and has become ubiquitous in the AI literature, as it provides a useful platform for studying algorithmic aspects of reasoning about uncertainty (for textbook treatments, see Cowell, Dawid, Lauritzen and Spiegelhalter (1999) or Koller and Friedman (2009)). In particular, Pearl (1988,2000) advocated the view of R as a *causal structure* that underlies the probability distribution p , and used the graph-theoretic representation of R to visualize causal relations and systematize reasoning about causality. I adopt the graph-theoretic terminology of this literature, and refer to R as a *directed acyclic graph* (DAG).¹

This paper puts Bayesian networks to novel use, as a representation of imperfect understanding of correlation structures, which is then integrated into a "personal equilibrium" model of individual decision making under "boundedly rational expectations". Suppose that $x_0 = a$ describes the DM's own action, $x_1 = t$ describes his signal, and denote $y = (x_2, \dots, x_n)$. An objective distribution p is a "*personal equilibrium*" if the DM's behavior, given by the conditional probabilities $p(a \mid t)$, maximizes his expected payoff $u(a, t, y)$ w.r.t the subjective conditional distribution $p_R(y \mid a, t)$. Because some actions may be played with zero probability in equilibrium, the definition of personal equilibrium involves "trembling hand perfection", conventionally capturing the idea that the DM almost never plays subjectively sub-optimal actions.²

¹Graphical probabilistic models have been introduced into economics in other contexts: as a way of representing games and facilitating computation of Nash equilibrium (see Kearns, Littman and Singh (2001) and Koller and Milch (2003)), or as a way of discussing causality in econometric models (see White and Chalak (2009)).

²The term "personal equilibrium" was introduced by Koszegi (2009) in the context of decision making with reference-dependent preferences.

Example 1.1: Medication

Consider a DM who considers fighting a disease by taking a medication that has adverse side effects. The DM's choice and his state of health are statistically independent, yet they share a common consequence: the level of some chemical in the DM's blood goes up when his medical condition improves, *or* when he takes the medication. Let x_0, x_1, x_2 denote the DM's consumption decision, state of health and chemical level, respectively. The objective distribution p is thus consistent with a "true DAG" $0 \rightarrow 2 \leftarrow 1$. Suppose that the chemical level as such is payoff-irrelevant. Then, a DM with rational expectations would choose not to take the medication.

Now assume that our DM's subjective DAG is $0 \rightarrow 2 \rightarrow 1$; that is, he inverts the causal link between the chemical level and his health. The DM will choose x_0 to maximize

$$\sum_{x_1} \sum_{x_2} p(x_2 \mid x_0) p(x_1 \mid x_2) u(x_0, x_2)$$

Given the true underlying p , x_1 is generally *not* independent of x_0 conditional on x_2 . Therefore, the conditional probability $p(x_1 \mid x_2)$ may be affected by the equilibrium distribution over x_0 . This is why the DM's choice is fundamentally an *equilibrium* decision: his evaluation of any choice of x_0 effectively takes his own long-run consumption as given.

When the disutility from the medication's side effects is not too large, there is a personal equilibrium in which the DM *takes the medication*. In such an equilibrium, the DM observes (possibly as a result of "trembles") the positive correlation between the chemical level and his health. He correctly grasps that if he quits taking the medication, the chemical level will go down. However, the subjective causal link $x_2 \rightarrow x_1$ means that the DM erroneously concludes that a lower chemical level will aggravate his disease. As a result, he prefers to continue taking the medication.

The Bayesian-network approach to modeling decision making under boundedly rational expectations has a number of merits:

Modeling subjective perceptions of causality. As Example 1.1 illustrates, the DM's DAG can be interpreted as a subjective (and possibly mis-perceived) model of causal relations in his environment. This "Pearlian" interpretation enables us to study implications of incorrect causality perceptions on equilibrium behavior in economic models.

Unification. In Section 3, I demonstrate that several existing models of boundedly rational expectations emerge as special cases captured by distinct graphical representations. When the objective distribution is consistent with a "true DAG", the systematic departures from rational expectations embodied in such models can be captured by simple operations on the true graph (removing, inverting or reorienting links). This unification uncovers connections among existing concepts, suggests new ones and leads to results that hold for general classes of models of boundedly rational expectations (see Section 5).

From solution concept to type. The literature has typically presented models of boundedly rational expectations in the form of *solution concepts* in a class of games. The Bayesian-network approach reduces such notions to *types* of individual agents, and this expands their scope of applicability. First, a DM represented by a DAG can be incorporated in other classes of models (e.g. competitive markets). Second, the representation opens the door for richer comparative statics (see Section 5.2) and greater heterogeneity of agents in economic models. Third, the fact that the representation is not tied to a particular interactive model simplifies analysis. Finally, the representation may facilitate the study of "high-order" boundedly rational expectations (when one of the variables in N is another agent's DAG).

Structured belief heterogeneity. It is instructive to compare the Bayesian-network approach to the traditional notion of subjective priors. A DM with

a subjective prior has a fixed belief, which is *independent* of the objective distribution p . In contrast, a DM represented by a DAG has a subjective prior that *changes systematically* with p . An economic model in which agents are characterized by distinct graphs exhibits "structured belief heterogeneity", because the agents' subjective beliefs are different deterministic transformations of the same objective distribution p . In particular, when p is consistent with the empty DAG, all agents' beliefs will coincide with p .

However, since the Bayesian-network representation of beliefs combines objective and subjective elements, it gives rise to an important question: how does the DM manage to fit p into his subjective graph? In Section 4 I provide a foundation for the Bayesian-network representation, in terms of a *naive procedure for extrapolating from limited datasets*. The idea is that the DM receives partial feedback about p , in the form of a large dataset with "*missing values*"; he employs an iterative procedure for imputing the missing values. I show that when the support of the missing-data process satisfies a condition known in the Bayesian-network literature as the "running intersection property", the procedure generates a "completed" dataset in which the frequencies have an essentially unique DAG representation, which the DAG is restricted to be *perfect*. When the running intersection property is violated, the procedure will be aborted for some objective distribution. Thus, although for most of the paper I treat the DM's graph as a fundamental characteristic, the missing-data foundation is an integral part of this approach. It has the advantage of imposing additional structure on Bayesian-network representations, and possibly inspiring more general ones.

Related works on models of boundedly rational expectations

Non-rational expectations have been explored from many different points of view (e.g., see Evans and Honkapohja (2001) and Woodford (2013) for macroeconomic models; the notion of "restricted perceptions equilibrium" is particularly relevant). In this partial literature review, I emphasize attempts to model *imperfect understanding of correlations*, as these have inspired the

present study. Osborne and Rubinstein (1998) studied games with players who falsely believe their own actions influence the opponents' behavior, as a result of naive extrapolation from small samples. Eyster and Rabin (2005) examined Bayesian games in which each player may underestimate (exaggerate) the correlation between his opponents' actions and their signals (his own signal). Esponda (2008) and Esponda and Pouzo (2012,2014b) focused on situations in which agents neglect the effect of their own actions on the observed distribution over consequences. A key development in this literature involved "coarse reasoning". Piccione and Rubinstein (2003), Jehiel (2005), Jehiel and Koessler (2008) and Eyster and Piccione (2013) assumed that agents' beliefs are measurable w.r.t a *coarse* representation of the set of contingencies (by omitting variables from their subjective model or by clumping contingencies into an "analogy class").³

The Bayesian-network approach is related to the view of a boundedly rational agent as an econometrician working with a misspecified model, as in Bray (1982), Cho, Sargent and Williams (2002), or Rabin and Vayanos (2010). More recently, Esponda and Pouzo (2014a) extended this tradition by formulating a rather general model of static games, in which each player has a set of subjective prior beliefs over the states of Nature (possibly with an incorrect support). The player receives partial feedback about the equilibrium distribution (similar in spirit to the "missing data" model in this paper). In equilibrium, each player's subjective belief is the closest (in terms of relative entropy) to the objective distribution, among the beliefs in the set of possible posterior beliefs over consequences defined by his set of priors and his feedback. Esponda and Pouzo justify their solution concept as a steady state of a dynamic Bayesian learning model.

³The representation (2) can be interpreted in terms of Jehiel's (2005) concept of analogy-based expectations, applied to a *fictitious* extensive-form game: the DM believes that he is part of an $(n + 1)$ -player extensive game, in which each player $i \in N$ moves once according to a *fixed* order of moves given by a linear ordering R^* that extends R . All player- i histories with the same moves by the players in $R(i)$ form an analogy class.

2 The Decision Model

Let $X = X_0 \times \cdots \times X_n$ be a finite set of *states*, where $n \geq 2$. Denote $N = \{0, \dots, n\}$. For every $M \subseteq N$, denote $x_M = (x_k)_{k \in M}$. The set $X_0 = A$ is the set of feasible *actions* available to a decision maker (DM). Thus, the DM's action is part of the definition of a state. I use x_0 or a interchangeably to denote an action. The set $T = X_1 \times \cdots \times X_l$, $l < n$, is the set of signals that the DM may receive (note that each $t \in T$ has l components), such that $x_{\{1, \dots, l\}}$ and t can be used interchangeably to denote a signal. Denote $Y = X_{l+1} \times \cdots \times X_n$. Thus, a state x can be written as the triple (a, t, y) . Let $p \in \Delta(X)$ be an *objective* probability distribution over states.

Beliefs

To capture limited understanding of the correlation structure of p , I introduce a new primitive. Let R be an asymmetric and acyclic binary relation over N , and define $R(k) = \{j \mid jRk\}$. I refer to R as a *directed acyclic graph* (DAG), and use \tilde{R} to denote the undirected version (or "skeleton") of R - that is, $i\tilde{R}j$ if and only if iRj or jRi . For any p and R , define p_R by the factorization formula (2). The notation p_R is a convenient short-hand for a *mapping* from $\Delta(X)$ to itself. This mapping assigns a subjective distribution to every objective distribution p , such that $p_R(x)$ is the subjective probability of x when the objective probability of x is $p(x)$. We say that a distribution p is consistent with a DAG R if $p_R(x) = p(x)$ for every x .

Recall that when R is a linear ordering, the enumeration of the state variables is irrelevant for the representation (2). This means that all linear orderings are equivalent, as far as the representation is concerned. In general, a given probability distribution can admit multiple DAG representations, a consequence of the basic identity $p(x, y) = p(x)p(y \mid x) = p(y)p(x \mid y)$.

Definition 1 *Two DAGs R and Q are **equivalent** if $p_R = p_Q$ (that is, $p_R(x) = p_Q(x)$ for every x) for every $p \in \Delta(X)$.*

Verma and Pearl (1991) provided a characterization of equivalent DAGs which will be useful in the sequel. We say that R and Q have *the same v -structure* if for every nodes i, j, k for which iRk , jRk , iQk and jQk , it is the case that $i\tilde{R}j$ if and only if $i\tilde{Q}j$. In other words, if both R and Q treat i and j as causes of k , then R establishes a causal link between i and j if and only if Q does so, too.

Proposition 1 (Verma and Pearl (1991)) *Two DAGs R and Q are equivalent if and only if they have the same skeleton and the same v -structure.*

To illustrate this result, the DAGs $0 \rightarrow 2 \leftarrow 1$ and $0 \rightarrow 2 \rightarrow 1$ have identical skeletons but different v -structures. Therefore, these DAGs generate different subjective beliefs for some objective distribution p . In contrast, the DAGs $0 \rightarrow 2 \rightarrow 1$ and $0 \leftarrow 2 \leftarrow 1$ are equivalent because they have the same skeleton and the same (vacuous) v -structure.

Comment: Conditional probabilities. Subjective conditional probabilities are calculated as usual. In particular, for any p , the DM's subjective distribution over y conditional on a, t is

$$p_R(y \mid a, t) = \frac{p_R(a, t, y)}{\sum_{y'} p_R(a, t, y')}$$

as long as $p(a, t) > 0$. Note that at first glance, the representation (2) seems to contain potentially ill-defined conditional probabilities. Suppose that $j \in R(i)$ and that $p(x_j \mid x_{R(j)}) = 0$ for some value of x_j . Then, the term $p(x_i \mid x_{R(i)})$ is ill-defined. However, this does not pose any problem: when we sum over all x in the calculation of p_R , we simply exclude zero-probability realizations of x_j in the summation, such that the problematic term does not appear.

Decisions

Let us now turn to decision making under the Bayesian-network representation of subjective beliefs. Our DM, characterized by a particular R , is an expected utility maximizer, with a vNM utility function $u : X \rightarrow \mathbb{R}$. The objective distribution p is interpreted as a long-run (steady-state) joint distribution over all the relevant variables, including the DM's own action. We will require the DM's long-run behavior (given by $(p(a \mid t))_{a,t}$) to be optimal w.r.t his subjective belief, which is a systematic distortion of p given by (2). Because of this distortion, the DM's long-run behavior can influence his perception of the implications of his actions, hence beliefs and decisions cannot be separated; optimal behavior is fundamentally an *equilibrium* notion.⁴

The need for an equilibrium model of individual choice also requires us to take off-equilibrium actions into account. Consider a distribution p with full support on T . We say that p' is a *perturbation* of p if $p'(t) \equiv p(t)$, $p'(y \mid a, t) \equiv p(y \mid a, t)$, and $p'(a \mid t) > 0$ for all a, t . A perturbation fixes every aspect of p except for the DM's behavior, such that every action is played with positive probability for every signal.

Definition 2 *A distribution $p \in \Delta(X)$ with full support on $A \times T$ is an ε -personal equilibrium if*

$$a \in \arg \max_{a'} \sum_y p_R(y \mid a', t) u(a', t, y)$$

for every a, t for which $p(a \mid t) > \varepsilon$.

Definition 3 *A distribution $p^* \in \Delta(X)$ with full support on T is a **personal equilibrium** if there exist a sequence $\varepsilon^k \rightarrow 0$ and a sequence $p^k \rightarrow p^*$ of perturbations of p^* , such for every k , p^k is a ε^k -personal equilibrium.*

⁴This feature, which is impossible under rational expectations, was pointed out in earlier works - explicitly in Esponda (2008), and implicitly in Piccione and Rubinstein's (2003) discussion of bounded recall and Jehiel's (2005) Centipede-Game example.

Thus, in equilibrium, the DM takes actions that maximize his subjective expected utility given his information. The expectation is taken w.r.t the (perturbed) equilibrium distribution filtered through the DM's DAG. In a ε -personal equilibrium, the DM plays actions that are sub-optimal w.r.t his subjective belief with probability ε at most. Personal equilibrium essentially requires that such actions are *almost never* taken.

Proposition 2 *For any $(p(t))_t$ and $(p(y \mid a, t))_{a,t,y}$, there exists $((p(a \mid t))_{a,t})$ such that $p = ((p(t))_t, ((p(a \mid t))_{a,t}, (p(y \mid a, t))_{a,t,y})$ is a personal equilibrium.*

As we shall see in Section 3, "pure" personal equilibria (where for every t there is a such that $p(a \mid t) = 1$) need not exist, unlike the case of expected utility maximization w.r.t rational expectations.

3 Illustrations

This section demonstrates how the Bayesian-network framework accommodates a variety of systematic departures from rational expectations. Most of the examples will involve three or four variables, and the DM will be uninformed ($l = 0$). In each example, p will be consistent with some "true DAG" having a particular conditional-independence structure; the DM's subjective DAG will be obtained from the true DAG by simple operations: *removing, inverting or reorienting links*. Different belief biases can thus be captured by different operations on the true DAG. In particular, a given subjective DAG can capture different biases, depending on its exact relation to the true DAG.

If the DM's subjective DAG differs from the true DAG only by *adding links*, this can never lead to biased beliefs, because any belief that is consistent with a DAG is also consistent with another DAG that contains it. Thus, *over-estimation* of correlations among variables is not a bias that the current framework can capture.

3.1 Inverting Links: Reverse Causality

Let $n = 2$. Interpret x_1 as a state of Nature, and x_2 as a consequence of the two other variables. The DM is uninformed about x_1 at the time he makes his choice; x_0 and x_1 are statistically independent. The objective distribution can be written as

$$p(x_0, x_1, x_2) = p(x_0)p(x_1)p(x_2 \mid x_0, x_1)$$

which means that p is consistent with the "true DAG"

$$R : 0 \rightarrow 2 \leftarrow 1$$

A rational DM will choose x_0 to maximize

$$\sum_{x_1} \sum_{x_2} p_R(x_1, x_2 \mid x_0) u(x_0, x_1, x_2) = \sum_{x_1} \sum_{x_2} p(x_1) p(x_2 \mid x_0, x_1) u(x_0, x_1, x_2)$$

Suppose that the DM's subjective graph is

$$RC : 0 \rightarrow 2 \rightarrow 1$$

such that

$$p_{RC}(x_0, x_1, x_2) = p(x_0) p(x_2 \mid x_0) p(x_1 \mid x_2)$$

Thus, relative to the true DAG R , the DM *inverts the direction of the causal link between x_1 and x_2* . The DM chooses x_0 to maximize

$$\sum_{x_1} \sum_{x_2} p_{RC}(x_1, x_2 \mid x_0) u(x_0, x_1, x_2) = \sum_{x_1} \sum_{x_2} p(x_2 \mid x_0) p(x_1 \mid x_2) u(x_0, x_1, x_2)$$

Under RC , the notion of personal equilibrium cannot be reduced to straightforward maximization. The reason is that the conditional probability $p(x_1 \mid x_2)$ involves summing over the DM's equilibrium actions. If

$p(x_0)$ were to change, so could $p(x_1 \mid x_2)$, and so could the DM's optimal action. This means that the DM effectively takes his equilibrium behavior as given when choosing his action, hence the equilibrium aspect of his choice is not redundant.

Example 3.1: Fiscal policy

Let us impose additional structure on this setting. The DM is a government. All three variables take values in $\{0, 1\}$; $x_1 = 0$ (1) means that the rate of real growth is low (high); and $x_2 = 0$ (1) means that the level of public debt is low (high). The action $a = 0$ (1) represents fiscal austerity (expansion). The objective distribution p satisfies $p(x_1 = 1) = \frac{1}{2}$ (independently of x_0 , in line with the true DAG R), and x_2 is a deterministic function of x_0, x_1 given by $x_2 = x_0(1 - x_1)$. Thus, high public debt results from a combination of fiscal expansion and low growth.

The government's preferences are given by $u(x_0, x_1) = x_0 + kx_1$, where $k > 0$. Thus, the government cares about social programs (which it can provide via fiscal expansion) and growth; it does not care about debt per se. It follows that under rational expectations, the government will choose $x_0 = 1$, because growth is independent of fiscal policy. The government's subjective DAG is RC . The nature of its departure from rational expectations is that it believes fluctuations in public debt cause fluctuations in real growth, whereas the true causal link is in the opposite direction.

This structure of p implies the following conditional probabilities. Denote $p(x_0 = 1) = \beta$. Then,

$$\begin{aligned} p(x_2 = 0 \mid x_0 = 0) &= 1 \\ p(x_2 = 0 \mid x_0 = 1) &= \frac{1}{2} \\ p(x_1 = 1 \mid x_2 = 1) &= 0 \\ p(x_1 = 1 \mid x_2 = 0) &= \frac{\frac{1}{2}(1 - \beta) + \frac{1}{2}\beta}{1 - \beta + \frac{1}{2}\beta} = \frac{1}{2 - \beta} \end{aligned}$$

Let us first show that the action $x_0 = 1$ is inconsistent with personal equilibrium whenever $k > 2$. Assume the contrary, i.e. $\beta = 1$. The government's evaluation of the action $x_0 = 1$ is

$$\begin{aligned} & 1 + k \cdot [p(x_2 = 0 \mid x_0 = 1)p(x_1 = 1 \mid x_2 = 0) \\ & \quad + p(x_2 = 1 \mid x_0 = 1)p(x_1 = 1 \mid x_2 = 1)] \\ & = 1 + k \cdot \left[\frac{1}{2} \cdot \frac{1}{2-1} + \frac{1}{2} \cdot 0 \right] = 1 + \frac{1}{2}k \end{aligned}$$

If the government deviates to $x_0 = 0$, it expects to earn

$$\begin{aligned} & 0 + k \cdot [p(x_2 = 0 \mid x_0 = 0) \cdot p(x_1 = 1 \mid x_2 = 0) \\ & \quad + p(x_2 = 1 \mid x_0 = 0) \cdot p(x_1 = 1 \mid x_2 = 1)] \\ & = k \end{aligned}$$

hence the deviation is profitable whenever $k > 2$.

Thus, although expansion is the unique optimal action under rational expectations, it is inconsistent with personal equilibrium when the government cares enough about growth and when its subjective causal model is given by RC . The intuition is simple: if the government plays $a = 1$ in the putative equilibrium, it observes a negative correlation between debt and growth. It correctly grasps the effect of fiscal policy on debt. And since it misperceives debt-growth causality, it erroneously believes that the low debt resulting from austerity will lead to high growth. This is not the usual logic of self-confirming expectations. The government's reasoning does *not* rest on out-of-equilibrium beliefs, but rather on a misperception of the statistical regularities that characterize the observed equilibrium behavior.

Using similar analysis, it can be shown that when $k \geq 4$, playing the action $x_0 = 0$ with probability one is consistent with personal equilibrium. When $k \in (2, 4)$, personal equilibrium must be "mixed" (in the sense that $\beta \in (0, 1)$). Other specifications of the DM's preferences could lead to multiple

personal equilibria. Of course, these effects are impossible under conventional expected-utility maximization.

Medication example revisited. Example 1.1 can be cast in the same terms as the fiscal policy example: $x_1 = 0$ (1) indicates poor (good) health; $x_2 = 0$ (1) indicates a high (low) chemical level; and $x_0 = 0$ (1) indicates that the DM takes (avoids) the medication. The true DAG means that the DM's chemical level is low if (and only if) he avoids the drug and suffers from the disease. Taking the medication is analogous to austerity in the fiscal-policy story.

3.2 Removing Links: Coarse Reasoning

Models in which agents' beliefs are a coarse representation of the true underlying distribution have been prominent in the literature on boundedly rational expectations (see Piccione and Rubinstein (2003), Jehiel (2005), Jehiel and Koessler (2008), Eyster and Piccione (2013)). To capture coarse reasoning, let $n = 4$ and suppose that p is consistent with the following true DAG R :

$$\begin{array}{ccccc} 2 & \rightarrow & 1 & \rightarrow & 0 \\ \downarrow & \searrow & \uparrow & & \\ 4 & \leftarrow & 3 & & \end{array} \tag{3}$$

where the nodes 0 and 1 represent the DM's action and signal, respectively; 2 and 3 jointly represent the state of Nature; and 4 represents some "dependent variable" (an opponent's action in a game, a product price in a market setting). Every p that is consistent with R has the property that x_0 is independent of (x_2, x_3, x_4) conditional on x_1 .

Now suppose that the DM's subjective DAG, denoted C , differs from R only by removing the link $4 \leftarrow 3$. This represents a coarse perception of the causes of x_4 : while the true DAG admits both x_2 and x_3 as direct causes of x_4 , the modified graph C admits only x_2 . Using the terminology of Jehiel (2005), x_2 is the "analogy class" to which the state of Nature (x_2, x_3)

belongs. In general, coarse reasoning can be captured by omitting links in the true DAG that involve "exogenous" variables, namely variables that are independent of the DM's action conditional on his signal.

To illustrate the effect of coarse reasoning on the DM's beliefs and actions, consider the simpler, four-variable case of an uninformed DM, which is captured graphically by removing the node 1 and its links from both R and C . Then,

$$\begin{aligned} p_R(x) &= p(x_0)p(x_2)p(x_3 | x_2)p(x_4 | x_2, x_3) \\ p_C(x) &= p(x_0)p(x_2)p(x_3 | x_2)p(x_4 | x_2) \end{aligned} \quad (4)$$

The DM chooses x_0 to maximize

$$\sum_{x_2} \sum_{x_3} \sum_{x_4} p_C(x_2, x_3, x_4 | x_0) u(x) = \sum_{x_2} \sum_{x_3} p(x_2, x_3) \sum_{x_4} p(x_4 | x_2) u(x)$$

This is precisely the individual behavior that Jehiel and Koessler (2008) presume as part of the solution concept Analogy-Based Expectations Equilibrium in the context of games with incomplete information.⁵

By the structure of R , the term $p(x_4 | x_2)$ is independent of x_0 . As a result, the DM does not need to take his own behavior as given when choosing how to act - in other words, personal equilibrium under C is reduced to maximization.

Example 3.2: Bilateral trade

The following example, based on an adverse-selection experiment by Bazerman and Samuelson (1985), has developed semi-canonical status in recent literature on boundedly rational expectations (including Eyster and Rabin (2005), Jehiel and Koessler (2008) and Esponda (2008) - see Spiegler (2011,

⁵The present model does not address environments with an explicit time dimension, and therefore does not subsume analogy-based expectations in extensive-form games, as defined by Jehiel (2005).

Ch. 8) for a pedagogical exposition). I recast it in terms of the Bayesian-network framework.

A seller holds an object. Our DM is an uninformed buyer who bids x_0 for the object; $x_1 \in [0, 1]$ is the seller's valuation; and $x_2 \in [0, 1]$ is his ask price. Trade takes place whenever $x_0 \geq x_2$. The transaction price conditional on trade is x_0 . Thus, if the seller does not play weakly dominated strategies, then $p(x_2 = x_1) = 1$. The buyer's payoff is $u(x) = \mathbf{1}(x_2 \leq x_0) \cdot (x_1 + g - x_0)$, where $g > 0$ is the gain from trade.

The true distribution p is consistent with the DAG

$$0 \quad 2 \leftarrow 1$$

A buyer with rational expectations will choose x_0 to maximize

$$\begin{aligned} \sum_{x_1} p(x_1) \sum_{x_2} p(x_2 \mid x_1) \mathbf{1}(x_2 \leq x_0) (x_1 + g - x_0) \\ = p(x_1 \leq x_0) \left[\sum_{x_1} p(x_1 \mid x_1 \leq x_0) x_1 + g - x_0 \right] \end{aligned}$$

Now suppose that the buyer's subjective DAG is the empty graph - i.e., it removes the link $2 \leftarrow 1$ from the true DAG, thus displaying "coarse reasoning". Then, he chooses x_0 to maximize

$$\begin{aligned} \sum_{x_1} p(x_1) \sum_{x_2} p(x_2) \mathbf{1}(x_2 \leq x_0) \cdot (x_1 + g - x_0) \\ = p(x_1 \leq x_0) \left[\sum_{x_1} p(x_1) x_1 + g - x_0 \right] \end{aligned}$$

Thus, the buyer has a correct understanding of the seller's average behavior, but he fails to take into account its correlation with the seller's private information.

3.3 Disconnecting Nodes: Missing Variables

In the previous sub-section, we saw how coarse reasoning can be captured by a simple operation on the true DAG, namely omitting a link between variables which are independent of the DM’s action (conditional on his signal). Let us now consider a more elaborate operation: severing *all the links* of a payoff-irrelevant variable which is potentially correlated with the DM’s action. Thus, the DM’s subjective causal model is misspecified in the sense that it effectively omits a variable, and this omission potentially biases the DM’s perception of the consequences of his action.

For instance, let $n = 3$, and suppose that the true DAG is

$$\begin{array}{ccccc} 0 & \rightarrow & 1 & \rightarrow & 3 \\ & \searrow & & \nearrow & \\ & & 2 & & \end{array} \quad (5)$$

Let u be purely a function of x_1 and x_3 . Under rational expectations, the DM chooses x_0 to maximize

$$\sum_{x_1} \sum_{x_2} p(x_1, x_2 \mid x_0) \sum_{x_3} p(x_3 \mid x_1, x_2) u(x_1, x_3)$$

Suppose that the DM's subjective DAG, denoted MV , is

$$0 \rightarrow 1 \rightarrow 3 \rightarrow 2 \rightarrow 0 \quad (6)$$

Thus, MV differs from the true DAG by disconnecting 2 from all other nodes. The DM will choose x_0 to maximize

$$\sum_{x_2} p(x_2) \sum_{x_1} p(x_1 | x_0) \sum_{x_3} p(x_3 | x_1) u(x_1, x_3) \quad (7)$$

which is equal to

$$\sum_{x_1} p(x_1 | x_0) \sum_{x_3} \sum_{x'_0} \sum_{x'_2} p(x'_0 | x_1) p(x'_2 | x'_0) p(x_3 | x_1, x'_2) u(x_1, x_3)$$

It is clear from this expression that as in the case of reverse causality - but unlike the case of coarse reasoning - the DM's own equilibrium behavior given by $p(x'_0)$ affects his evaluation of actions, hence personal equilibrium cannot be reduced to simple maximization.

Example 3.3: Central banking

This example translates an argument by Sargent (1999) into the Bayesian-network framework. The DM is a central banker, and x_0 is a non-negative real number that represents monetary policy (say, inflation targeting). The true DAG is given by (5). Let x_1 represent actual inflation, and suppose that $x_1 = x_0 + \eta$, where η is an independently distributed variable with mean zero. Let x_2 represent the public's inflation expectations, which are formed *after* the public observes the central banker's choice of x_0 . In the spirit of Lucas (1976), the public has "rational expectations", such that $x_2 = E(x_1 | x_0) = x_0$. The variable x_3 , which represents unemployment, is given by $x_3 = x_3^* - k(x_1 - x_2) + \varepsilon$, where x_3^* is the "natural unemployment rate", $k > 0$ is a constant, and ε is an independently distributed variable with mean zero. Thus, the true relation between monetary and real variables is given by an "expectational Phillips curve", where only unanticipated inflation affects unemployment. In contrast, the central banker's subjective graph MV , given by (6), can be interpreted as if it postulates a "traditional" Phillips curve that omits the public's inflation expectations.

The central banker's preferences are given by $u(x_1, x_3) = -[(x_3)^2 + (x_1)^2]$. Under rational expectations, the central banker would choose x_0 to minimize

$$E[(x_3)^2 + (x_1)^2 | x_0] = E(x_3^* - k\eta + \varepsilon)^2 + E(x_0 + \eta)^2$$

The optimal policy is thus clearly $x_0 = 0$. This is the familiar Kydland-
Prescott "commitment policy".

Now consider the central banker's behavior under *MV*. Suppose that in personal equilibrium, $p(x_0^*) = 1$ for some x_0^* . Then, the central banker believes that $x_2 = x_0^*$ with probability one, independently of x_0 and for all x_1 . Applying (7), we obtain that the central banker chooses x_0 to minimize

$$E[x_3^* - k(x_1 - x_0^*) + \varepsilon]^2 + E(x_0 + \eta)^2 = E[x_3^* - k(x_0 + \eta - x_0^*) + \varepsilon]^2 + E(x_0 + \eta)^2$$

The optimal policy x_0^{**} is given by the first-order condition,

$$-2k(x_3^* - k(x_0^{**} - x_0^*)) + 2x_0^{**} = 0$$

In order for x_0^{**} to be consistent with the guessed personal equilibrium, we need to have

$$x_0^{**} = x_0^* = \frac{1}{kx_3^*}$$

This is the Kydland-Prescott "no-commitment policy". Thus, when the central banker omits expected inflation from his subjective causal model, he behaves as if he cannot commit to an inflation target, effectively taking the public's expectations as given. This is essentially the insight of Sargent (1999) - originally formulated in the context of an elaborate dynamic learning model - recast in the simpler, static language of Bayesian networks.

3.3.1 Retrospective Choice

This sub-section discusses a different missing-variable distortion of a true DAG. Let $n = 3$, suppose that the true DAG is

$$\begin{array}{ccccc} 0 & \rightarrow & 2 & \leftarrow & 1 \\ & & & \searrow & \downarrow \\ & & & & 3 \end{array} \tag{8}$$

where x_1 is payoff-irrelevant. Under rational expectations, the DM chooses x_0 to maximize

$$\sum_{x_1} \sum_{x_2} p(x_1) p(x_2 \mid x_0, x_1) \sum_{x_3} p(x_3 \mid x_1, x_2) u(x)$$

Suppose that the DM's subjective DAG, denoted EP , omits all the links from 1, such that EP consists of an isolated node 1 and the connected component $0 \rightarrow 2 \rightarrow 3$. Then, The DM chooses x_0 to maximize

$$\begin{aligned} \sum_{x_1} p(x_1) \sum_{x_2} p(x_2 \mid x_0) \sum_{x_3} p(x_3 \mid x_2) u(x) \\ = \sum_{x_2} p(x_2 \mid x_0) \sum_{x_3} p(x_3 \mid x_2) u(x) \end{aligned}$$

To see why this expression departs from rational expectations, note first that each of the terms $p(x_2 \mid x_0)$ and $p(x_3 \mid x_2)$ implicitly involves a summation over x_1 , as if the effects of x_1 on x_2 and x_3 are independent - whereas in reality they may be correlated. Moreover, the term $p(x_3 \mid x_2)$ implicitly involves summing over all possible values of x_0 , effectively taking the DM's equilibrium behavior as given. It follows that personal equilibrium under EP cannot be reduced to maximization.

This example is abstracted from Esponda (2008) and Esponda and Pouzo (2014b), and captures a mode of reasoning which Esponda and Pouzo call "retrospective choice". It is also close in spirit to the choice procedure underlying Osborne and Rubinstein (1998). Suppose that x_3 is a deterministic function of x_1, x_2 , and u is purely a function of x_0, x_3 (such that x_1 is not directly payoff-relevant). Imagine that the DM reviews a large historical database induced by the equilibrium distribution p . For every action x_0 that he considers, he first examines its possible consequences x_2 ; for every such consequence he calculates a payoff-relevant summary statistic x_3 ; then he aggregates over all possible consequences, according to their weights in the

database conditional on the contemplated action x_0 . This chain of reasoning, action \rightarrow consequence \rightarrow payoff, is an intuitive way of evaluating actions. Yet, it leads to a biased evaluation (affected by the DM's own equilibrium behavior) because it omits the "latent" variable x_1 .⁶

Bilateral trade revisited

Consider the following reformulation of Example 3.2, which translates an example due to Esponda (2008) into the current framework. The variables x_0 and x_1 are defined as before. Let $x_2 \in \{0, 1\}$ indicate whether trade takes place, and assume $x_2 = 1$ if and only if $x_0 \geq x_1$. Let $x_3 = x_1 + g$ be the buyer's valuation of the object, and let $u(x_0, x_2, x_3) = x_2(x_3 - x_0)$. The objective distribution obeys the DAG given by (8), and the buyer's subjective DAG is *EP*. The buyer chooses x_0 to maximize

$$\begin{aligned} p(x_2 = 1 \mid x_0) \sum_{x_3} p(x_3 \mid x_2 = 1) u(x_0, x_2, x_3) \\ = p(x_1 \leq x_0) \left[\left(\sum_{x'_0} \sum_{x'_1} p(x'_0) p(x'_1 \mid x'_1 \leq x'_0) x'_1 \right) + g - x_0 \right] \end{aligned}$$

If in equilibrium an action x_0^* is played with probability one, then

$$x_0^* \in \arg \max_{x_0} p(x_1 \leq x_0) \left[\left(\sum_{x'_1} p(x'_1 \mid x'_1 \leq x_0^*) x'_1 \right) + g - x_0 \right]$$

Thus, the DM chooses his bid as if the object's expected quality is given by its empirical distribution conditional on trade taking place, according to his own equilibrium action. He fails to perceive that this distribution would change if he changed his bid.

⁶Specifically, Esponda and Pouzo (2012) analyze a voting model with asymmetric information, in which each voter evaluates each candidate by looking at his own average historical payoff when the candidate was elected (conditioning on his current signal).

Relation to reverse causality

Assume that $u(x_0, x_2, x_3) = x_3$ and $x_3 = v(x_1, x_2)$. That is, x_3 records the DM's payoff itself, and it is a deterministic function of x_1 and x_2 . Then,

$$\begin{aligned}
p_{EP}(x_1, x_2, x_3 \mid x_0) u(x_0, x_2, x_3) &= \sum_{x_2} p(x_2 \mid x_0) \sum_{x_3} p(x_3 \mid x_2) x_3 \\
&= \sum_{x_2} p(x_2 \mid x_0) \sum_{x_1} p(x_1 \mid x_2) \sum_{x_3} p(x_3 \mid x_1, x_2) x_3 \\
&= \sum_{x_1} \sum_{x_2} p(x_2 \mid x_0) p(x_1 \mid x_2) v(x_1, x_2) \\
&= p_{RC}(x_1, x_2 \mid x_0) v(x_1, x_2)
\end{aligned}$$

where RC is the DAG over $\{0, 1, 2\}$ given by $0 \rightarrow 2 \rightarrow 1$. Thus, two apparently different biases turn out to be equivalent in this environment: inverting the causal link between the state of Nature and a consequence, and omission of the state of Nature from a causal model that includes the DM's payoff as a distinct state variable. This illustrates how the Bayesian-network approach may help clarifying the connection between different types of departures from rational expectations.

3.4 Reorienting Links: False Attribution

In this sub-section I briefly discuss another class of belief biases that can be captured by another simple operation on the true DAG, namely changing the *origin* of a link. These can be called "attribution biases".

Illusion of control

Let $n = 2$. The DM is informed of the realization of x_1 when he chooses x_0 - i.e., x_1 is the DM's signal. Accordingly, the true DAG is $0 \leftarrow 1 \rightarrow 2$. The DM's subjective DAG is $1 \rightarrow 0 \rightarrow 2$. Thus, although x_1 is the only true cause of x_2 , the DM believes that his own action is the sole direct cause of

x_2 . This captures an "illusion of control", namely an exaggerated perception of the impact of one's own actions (see Langer (1975)). The DM will choose x_0 to maximize

$$\sum_{x_2} p(x_2 \mid x_0) u(x)$$

Since both x_0 and x_2 are actually caused by x_1 , the two variables are correlated and the DM will misperceive this correlation as a causal relation. The term $p(x_2 \mid x_0)$ is not invariant to $p(x_0 \mid x_1)$, hence personal equilibrium cannot be reduced to maximization.

Hindsight bias

Let $n = 4$, and assume that x_1 represents the (informed) DM's signal. The true DAG is $0 \leftarrow 1 \leftarrow 2 \rightarrow 3 \rightarrow 4$, where x_2 represents the state of Nature, and x_3 and x_4 represent another agent's signal and action, respectively. The DM's subjective DAG is

$$\begin{array}{ccccccc} 0 & \leftarrow & 1 & \rightarrow & 2 & \rightarrow & 3 \\ & & \downarrow & & & & \\ & & 4 & & & & \end{array}$$

This captures the phenomenon known as *hindsight bias* (or "information projection bias"), namely, the false belief that other people share your knowledge (see Fischhoff (1975) for a pioneering experimental study, and Madarasz (2012) for a Bayesian-game solution concept that captures information projection bias).

Analogy-based reasoning and attribution errors

Ettinger and Jehiel (2010) proposed that analogy-based reasoning can accommodate attribution errors. To illustrate this point, suppose that the true DAG is $0 \rightarrow 3 \leftarrow 1 \rightarrow 2$, where x_1 is the state of Nature and x_3 is a consequence of x_0 and x_1 . The variable x_2 is a deterministic, coarse function of x_1 . The DM's subjective DAG is $0 \rightarrow 3 \leftarrow 2 \leftarrow 1$ - that is, the DM attributes x_3

to the "analogy class" of the state of Nature. This observation demonstrates that the bias that a given DAG captures depends on its relation to the true DAG, as well as on the model's framing (in this example, the state of Nature and the analogy class are represented by different variables, whereas in Section 3.2 the state of Nature was defined by a collection of variables, one of which was the analogy class).

4 A "Missing Data" Foundation

What is the origin of R ? If it represents an entirely subjective causal model, how does the DM manage to apply it systematically to the objective distribution? On one hand, he should be exposed to enough data to enable him to pin down the conditional probabilities $p(x_i \mid x_{R(i)})$; yet on the other hand, if the data is rich enough, the DM may end up rejecting his subjective causal model.⁷

One story is that the DM has access to a rich database consisting of infinitely many observations of p . He poses a sequence of $n + 1$ questions to the database, where question i is: "What is the distribution over x_i conditional on $x_{R(i)}$?" (Variations on this question might involve running a regression or generating a plot graph.) The DM then pastes together the answers to these questions, by taking the product of the conditional distributions he obtains as answers to his queries. The reason the DM poses these particular questions may be that he is confident of the subjective causal model described by R , and therefore he does not look for correlations that are ruled out by R . Alternatively, the questions need not reflect an explicit prior causal model, but rather an intuitive way of reasoning about the data, as in the "retrospective choice" example of Section 3.3.1. According to this interpretation, the essence of the DM's bounded rationality is that he "asks the wrong ques-

⁷For an analysis of the problem of learning a graphical model from a machine-learning perspective, see Koski and Noble (2009, Ch. 6).

tions".

In this section I explore in detail a different foundation for R , based on the idea that the DM receives *incomplete data* about the objective distribution, and employs some procedure for *extrapolating his belief* from the data. To motivate our discussion, recall the setting of Section 3.1: $n = 2$, and x_0, x_1, x_2 denote the DM's action, the state of Nature and the consequence, respectively. Imagine that the DM's data consists of a very large sample. However, in each observation, either the value of x_0 or the value of x_1 are missing, according to some random independent process. In other words, the consequence is always observed, but only one of its two causes can ever be observed. The DM has thus effectively observed two joint distributions, $p(x_0, x_2)$ and $p(x_1, x_2)$, but he lacks direct evidence about the correlation between x_0 and x_2 . Our DM wishes to extrapolate from this data in order to form a subjective joint distribution over the three variables.

The DM's dataset can be visualized as a large *spreadsheet*, where each variable is represented by a different column, and each observation corresponds to a row. The spreadsheet has many missing values, generated according to the above process. Our DM attempts to fill the missing cells and "*rectangularize the spreadsheet*", such that the frequencies in the completed spreadsheet will serve as his subjective distribution. The following is a natural extrapolation procedure, which is motivated by this image. When the DM faces an observation of x_0, x_2 , he imputes the missing value of x_1 using the joint distribution $p(x_1, x_2)$ - specifically, by taking a random draw from $p(x_1 | x_2)$. Similarly, when he observes x_1, x_2 , he imputes the missing value of x_0 by taking a random draw from $p(x_0 | x_2)$.

Using this imputation procedure, the DM can turn a dataset with missing values into a complete dataset. The imputed distribution in the first part of the dataset (where values of x_1 were originally missing) is $p(x_0, x_2)p(x_1 | x_2)$; and the imputed distribution in the second part of the dataset (where values of x_0 were originally missing) is $p(x_1, x_2)p(x_0 | x_2)$. The two distributions are

identical, as both can be written as $p(x_0)p(x_2 \mid x_0)p(x_1 \mid x_2)$. The DM will adopt this "constructed" distribution as his subjective belief. This belief is consistent with the DAG $0 \rightarrow 2 \rightarrow 1$. This example shows that the DAG representation can be a consequence of a natural procedure for extrapolating from a limited dataset.

An iterative imputation procedure

I now define a general procedure for extrapolating a belief from a dataset with missing values, which *iterates* the imputation procedure of the motivating example. The DM's limited feedback consists of a collection \mathcal{S} of $m > 1$ non-empty subsets $S \subset N$. For convenience, I assume that \mathcal{S} is a *cover* of N , and that there exist no $S, S' \in \mathcal{S}$ such that $S \subset S'$. Each S represents an infinitely large sample consisting of independent draws of x_S from p , such that he effectively learns the marginal of the objective distribution p over x_S . (In general, the description of the DM's limited feedback would specify the frequency of each S in the dataset. However, this detail is omitted here because it is irrelevant for the present exercise).

The procedure consists of precisely $m - 1$ rounds. In each round $k = 1, \dots, m - 1$ of the procedure, the DM executes the following steps:

Step 1: The initial condition of the k -th round is a pair (B^{k-1}, p^{k-1}) , where $B^{k-1} \subset N$ and $p^{k-1} \in \Delta(B^{k-1})$. In particular, B^0 is an arbitrary member of \mathcal{S} , and p^0 coincides with p over B^0 , i.e. $p^0(x_{B^0}) \equiv p(x_{B^0})$.

Step 2: Select a set $S^k \in \arg \max_{S \in \mathcal{S} - \{B^0, \dots, S^{k-1}\}} |S \cap B^{k-1}|$. Define $B^k = B^{k-1} \cup S^k$.

Step 3: Define two distributions over X_{B^k} :

$$\begin{aligned} p_1^k(x_{B^k}) &= p^k(x_{B^{k-1}}) \cdot p(x_{S^k} \mid x_{S^k \cap B^{k-1}}) \\ p_2^k(x_{B^k}) &= p(x_{S^k}) \cdot p^{k-1}(x_{B^{k-1}} \mid x_{S^k \cap B^{k-1}}) \end{aligned}$$

If the two distributions coincide, define $p^k \equiv p_1^k \equiv p_2^k$ and continue to Step 4. Otherwise, *abort the procedure*.

Step 4: If $k = m - 1$, the *procedure is terminated and $p^{m-1} \in \Delta(X)$ is the DM's final belief*. If $k < m - 1$, switch to round $k + 1$ and return to Step 1.

Discussion

The idea behind this procedure is that the DM gradually completes his limited dataset by imputing missing values according to the correlations he observes. By the end of round $k - 1$, the DM has "rectangularized" the part of the spreadsheet in which the observed sets of variables were B^0, S^1, \dots, S^{k-1} - that is, he has transformed the observations of the variable sets B^0, S^1, \dots, S^{k-1} into a joint distribution over $X_{B^{k-1}}$. In round k , the DM looks for a new variable set S^k having maximal overlap with B^{k-1} , and he exploits the observed correlation among the variables in $B^{k-1} \cap S^k$ in order to extrapolate the distributions of $x_{B^{k-1}}$ and x_{S^k} conditional on $x_{B^{k-1} \cap S^k}$. The rationale for the "maximal overlap" criterion invoked in Step 2 of the iterative procedure is that the DM tries to make the most of the observed correlations. The DM terminates the procedure when he has "rectangularized" and completed his entire spreadsheet.

I assume that the DM aborts the procedure in round k whenever the imputation method leads to a discrepancy between the two auxiliary distributions p_1^k and p_2^k . Alternatively, one could assume that the DM simply takes some weighted average between the two, and obtain results in the same spirit. I adopt the "abort" variation mainly for simplicity, but it also reflects the idea that when extrapolations from different pieces of the DM's dataset lead to contradictory conclusions, the DM may abandon this method of extrapolation in favor of another one.

Illustration

Let $N = \{1, 2, 3, 4\}$ and $\mathcal{S} = \{\{1, 3\}, \{1, 2\}, \{2, 4\}\}$. This means that the DM has effectively learned the joint distributions $p(x_1, x_3)$, $p(x_1, x_2)$ and

$p(x_2, x_4)$. Select the initial condition to be $B^0 = \{1, 3\}$. The only legitimate continuation is $S^1 = \{1, 2\}$, such that $B^1 = \{1, 2, 3\}$. Imputing the missing values of x_3 (x_2) in the observations of x_1, x_2 (x_1, x_3) is done in round 1 exactly as in the motivating example. Thus, after the first round, the DM has replaced the original observations of x_1, x_2 and x_1, x_3 with "manufactured observations" of the triple x_1, x_2, x_3 , and the joint distribution over these variables in the manufactured dataset is

$$p^1(x_1, x_2, x_3) = p(x_3)p(x_1 | x_3)p(x_2 | x_1)$$

In the second and final round, $S^2 = \{2, 4\}$, such that $B^2 = N$, and

$$\begin{aligned} p_2^2(x_1, x_2, x_3, x_4) &= p(x_2, x_4) \cdot p^1(x_1, x_3 | x_2) \\ &= p(x_2, x_4) \frac{p_{\{1,2,3\}}^1(x_1, x_2, x_3)}{p_{\{1,2,3\}}^1(x_2)} \\ &= p(x_2, x_4) \frac{p(x_3)p(x_1 | x_3)p(x_2 | x_1)}{\sum_{x'_1} \sum_{x'_3} p(x'_3)p(x'_1 | x'_3)p(x_2 | x'_1)} \\ &= p(x_2, x_4) \frac{p(x_2)p(x_1 | x_2)p(x_3 | x_1)}{p(x_2) \sum_{x'_1} p(x'_1 | x_2) \sum_{x'_3} p(x'_3 | x'_1)} \\ &= p(x_2, x_4)p(x_1 | x_2)p(x_3 | x_1) \\ &= p(x_3)p(x_1 | x_3)p(x_2 | x_1)p(x_4 | x_2) \\ &= p^1(x_1, x_2, x_3) \cdot p(x_4 | x_2) \\ &= p_1^2(x_1, x_2, x_3, x_4) \end{aligned}$$

hence this is the expression for $p^2(x_1, x_2, x_3, x_4)$. The distribution given by the completed dataset is thus consistent with the DAG $3 \rightarrow 1 \rightarrow 2 \rightarrow 4$. Note that if the DM ignored the "maximal overlap" criterion and picked $S^1 = \{2, 4\}$ in the first round, the procedure would terminate immediately; the DM's resulting belief would be $p(x_1, x_3)p(x_2, x_4)$, thus ignoring the correlation structure given by his observations of x_1, x_2 . The insistence on the

maximal-overlap criterion ensures that the DM extracts more information from observed correlations.

The main result

In the preceding example, \mathcal{S} had the property that its elements could be ordered such that the intersection between any set along the sequence and the union of its predecessors would be weakly contained in one of these predecessors. (This property was trivially satisfied in the motivating example, where \mathcal{S} consisted of only two sets). This "*running intersection property*", as it is known in the Bayesian-network literature (see Cowell, Dawid, Lauritzen and Spiegelhalter (1999), p. 54), turns out to allow the iterative imputation procedure to culminate in a final belief that has an essentially unique Bayesian-network representation.

Definition 4 *A sequence of sets S_1, \dots, S_m satisfies the **running intersection property** (RIP) if for every $k = 2, \dots, m$, $S_k \cap (\cup_{i < k} S_i) \subseteq S_j$ for some $j < k$. We will say that the set \mathcal{S} satisfies RIP^* if its elements can be ordered in a sequence that satisfies RIP.*

Before stating our main result, we shall need a few conventional graph-theoretic definitions. A *clique* of a DAG R is a set of nodes A such that $i\tilde{R}j$ for every distinct $i, j \in A$. A clique in R is *maximal* if there is no clique in R that contains it. A clique A in R is *ancestral* if $R(i) \subset A$ for every $i \in A$. Finally, we will say that a DAG R is *perfect* if $R(i)$ is a clique for every $i \in N$. In a perfect DAG, all direct causes of a variable are causally linked themselves.

Remark 1 *Perfect DAGs that share the same set of maximal cliques are equivalent in the sense of Definition 1.*

This is a direct consequence of our analysis of equivalent DAGs in Section 2. By definition, perfect graphs have vacuous v -structures. Therefore, by Proposition 1, perfect graphs that share the same skeleton - which is given by the set of maximal cliques - are equivalent.

Proposition 3 *(i) If \mathcal{S} satisfies RIP^* , there exists a DAG R such that for every p , the iterative imputation procedure is never aborted and produces a final belief p^{m-1} that has a Bayesian-network representation given by (2). (ii) Conversely, if \mathcal{S} violates RIP^* , there exists p such that the procedure is necessarily aborted in some round.*

Thus, datasets that satisfy RIP^* can be extrapolated - via the iterative imputation procedure - into a belief that has a DAG representation. Moreover, the DAG must be perfect and uniquely pinned down up to the multiplicity allowed by Remark 1 - regardless of the arbitrary selections that the procedure involves along the way (e.g. the initial condition B^0). The result also means that any Bayesian-network representation (2) in which R is perfect has a "missing data" foundation, where \mathcal{S} is in the set of maximal cliques in R . On the other hand, if a dataset violates RIP^* , the procedure will be aborted for some objective distribution. For instance, let $N = \{1, 2, 3, 4\}$ and $\mathcal{S} = \{\{1, 2\}, \{2, 3\}, \{1, 3, 4\}\}$. Note that \mathcal{S} violates RIP^* . For any initial conditions of the iterative procedure, there is a distribution it will be aborted in the second round for some objective distribution.

Remark 1 means that in a perfect DAG, the direction of causal links is not identified: if R is perfect and iRj , there is an equivalent perfect DAG Q for which jQi . It follows that our imputation procedure cannot pin down the direction of any causal link. This is consistent with the idea that causality cannot be inferred from purely statistical data. There are natural economic models in which the objective distribution is consistent with an imperfect DAG - e.g., the "true DAG" in Example 3.3, given by (5). Proposition 3

thus uncovers a sense in which the true conditional-independence structure in Example 3.3 cannot be extrapolated from limited datasets.

Comments on RIP

RIP* states that the sets in \mathcal{S} can be ordered in *some sequence* that will satisfy RIP. However, the sequence in which the procedure considers the sets in \mathcal{S} is governed by the maximal-overlap criterion of Step 2. A priori, such a sequence need not satisfy RIP. A key step in the proof is a recent theorem due to Alon (2014), which assures that it will.

A natural variant on the imputation procedure would terminate it in the earliest round k for which $B^k = N$. In other words, the DM would stop as soon as his "edited" spreadsheet has an infinite set of rows for which no cell has a missing value, instead of trying to "rectangularize the *entire* spreadsheet". When RIP* is satisfied, this variation would not make any difference (in particular, by the assumption that \mathcal{S} does not include sets that contain one another, the alternative procedure would terminate in exactly $m - 1$ rounds). When RIP* is violated, it is possible that under the alternative termination criterion, the imputation procedure will never be aborted, and culminate in a belief with a DAG representation. However, the DAG will not be essentially unique, as it will depend on arbitrary selections such as the initial condition B^0 . For instance, let $N = \{1, 2, 3\}$ and $\mathcal{S} = \{\{1, 2\}, \{1, 3\}, \{2, 3\}\}$. The alternative procedure will terminate after one round, leading to a belief with a perfect DAG representation $i \rightarrow j \rightarrow k$, where any permutation of i, j, k is possible.

In the Bayesian-network literature, perfect DAGs graphs are relevant because they enable efficient algorithms for computing Bayesian updating. A common practice is to transform the DAG that represents a given statistical environment into a (skeleton of a) perfect DAG (by linking all parents of any given node, and eliminating the directionality of all links). RIP is then observed as a property of the set of maximal cliques in the "transformed" graph. This observation links the role of RIP in the Bayesian-network literature and

in the present paper, because when the iterative imputation procedure is not aborted, it implies that the DAG in the representation of subjective beliefs is perfect.

5 Strategic Rationality

In Section 3 we saw that personal equilibrium is reducible to maximization for some DAGs, whereas for others the notion of equilibrium is fundamental. Let us attempt to systematize this distinction. For simplicity, assume that the DM is uninformed ($l = 0$) and that $R(0) = \emptyset$ (the DM believes his action is not directly caused by any other variable).

Definition 5 *A DAG R is **strategically rational** w.r.t a set of probability distributions P if for every pair of distributions $p, q \in P$ and every a, y , $p(y \mid a) = q(y \mid a)$ implies $p_R(y \mid a) = q_R(y \mid a)$.*

Strategic rationality requires that if we modify the underlying objective distribution p by changing $p(a)$ - without changing the stochastic mapping from a to y - the DM's perception of this mapping should remain unchanged as well. When R is strategically rational, we can rewrite the definition of personal equilibrium as a maximization problem, because $p_R(y \mid a)$ is independent of the DM's equilibrium behavior given by $p(a)$. When strategic rationality is violated, we need to take $p(a)$ as given when calculating $p_R(y \mid a)$, and therefore the notion of personal equilibrium is indispensable.

Proposition 4 *Let $l = 0$. A DAG R satisfying $R(0) = \emptyset$ is strategically rational w.r.t a set of probability distributions P if and only if $0 \notin R(i)$ implies $p(y_i \mid y_{R(i) \cup \{0\}}) = p(y_i \mid y_{R(i)})$ for every $p \in P$, $i = 1, \dots, n$.*

Thus, a necessary and sufficient condition for strategic rationality is that if the DM's subjective DAG R omits a as an explanatory variable of some y_i , then it must be the case that for every distribution in P , y_i is independent of a conditional on $y_{R(i)}$. To illustrate this result, revisit the reverse causality example of Section 3.1. The DM's subjective DAG is $0 \rightarrow 2 \rightarrow 1$, i.e. $x_0 = a$ is omitted as an explanatory variable of x_1 . And yet, according to the true DAG $0 \rightarrow 2 \leftarrow 1$, x_1 is *not* independent of x_0 conditional on x_2 . In contrast, consider the case of fully coarse reasoning in Example 3.2, where the true DAG is $0 \rightarrow 2 \leftarrow 1$. The DM's subjective DAG is the empty graph, hence it assumes that neither x_0 nor x_2 are direct causes of x_1 . Since the true DAG has the property that x_1 is independent of x_0 , the condition for strategic rationality is satisfied.⁸

Simple strategically rational representations

The following is a simple special case of strategic rationality, which will serve us in the remaining two sub-sections. Suppose all variables in y are independent of a . Then, w.l.o.g we can remove the DM's action from his subjective causal model, and redefine R as the subgraph over $\{1, \dots, n\}$ induced by the original DAG over $\{0, 1, \dots, n\}$. Since the DM is uninformed, $y = (x_1, \dots, x_n)$. Define

$$p_R(y) = \prod_{i=1}^n p(x_i \mid x_{R(i)}) \quad (9)$$

The DM's chooses a to maximize

$$\sum_y p_R(y) u(a, y) \quad (10)$$

⁸When P is the set of all distributions that are representable by (2) w.r.t some DAG, the condition for strategic rationality can be described in terms of the structure of R^* (using the concept of *d-separation* - see any textbook on Bayesian networks, e.g. Pearl (2000, Ch. 1.2)). Since this concept requires preliminary definitions and makes no further appearance in this paper, I do not define it here.

Thus, the DM chooses his action to maximize his expected payoff w.r.t his distorted belief p_R over y .

5.1 Rationalizable Predictions

Because the DM in our model has an imperfect understanding of correlation structures, his prediction of some variable x_i may be sub-optimal even when he is fully informed of x_{-i} . By comparison, the traditional source of imperfect predictions in economics models is limited information. This raises the following question: are the predictions of a DM with complete information but incomplete understanding of the correlation structure distinguishable from the predictions of a DM with complete understanding and incomplete information?

Let R be a DAG over $\{1, \dots, n\}$. The DM tries to predict x_i given complete information of x_{-i} . The DM's posterior over x_i is thus $p_R(x_i \mid x_{-i})$. In contrast, consider a DM with rational expectations, who is only informed of $x_{A-\{i\}}$, where A is some subset of $\{1, \dots, n\}$. Such a DM's posterior over x_i would be $p(x_i \mid x_{A-\{i\}})$.

Definition 6 *A DAG R induces **rationalizable predictions** if there exists $A \subseteq \{1, \dots, n\}$ such that $p_R(x_i \mid x_{-i}) = p(x_i \mid x_{A-\{i\}})$ for every p and every i .*

When R is a linear ordering, it induces rationalizable predictions, where $A = \{1, \dots, n\}$. To take the other extreme, when R is empty, $p_R(x_1, \dots, x_n) = p(x_1) \cdots p(x_n)$, in which case the DM's prediction of x_i is $p_R(x_i \mid x_{-i}) = p(x_i)$, hence it is rationalizable by $A = \emptyset$.

Proposition 5 *A DAG R induces rationalizable predictions if and only if it is empty or a linear ordering.*

Thus, unless R is empty or a linear ordering, predictions on the basis of an imperfect understanding of the correlation structure are distinguishable from predictions based on rational expectations and imperfect information.

5.2 Performance-Based Rationality Ranking

As mentioned in the Introduction, a more complete graph corresponds in some intuitive sense to "more rational" expectations. Is this intuition consistent with ranking graphs in terms of the expected performance they lead to?

Definition 7 *Let R, R' be two DAGs over $\{1, \dots, n\}$ that are not equivalent in the sense of Definition 1. We say that R is **more rational** than R' if for every p, u, a, a' , the pair of inequalities*

$$\begin{aligned} \sum_y p_R(y)u(a, y) &> \sum_y p_R(y)u(a', y) \\ \sum_y p_{R'}(y)u(a', y) &> \sum_y p_{R'}(y)u(a, y) \end{aligned}$$

implies

$$\sum_y p(y)u(a, y) > \sum_y p(y)u(a', y)$$

That is, if R ranks a above a' and R' ranks a' over a , then the rational-expectations ranking of the two actions necessarily sides with R . Clearly, if R is a linear ordering and R' is not, the property holds trivially. The question is whether there exist non-equivalent R, R' that are not linear orderings, and yet one is more rational than the other.

Proposition 6 *Let R, R' be two non-equivalent DAGs that are not linear orderings. Then, neither DAG is more rational than the other.*

Thus, DMs characterized by different DAGs cannot be unambiguously ranked in terms of their expected performance, unless exactly one of these DAGs is a linear ordering.⁹ Note, however, that the ranking criterion adopted in this sub-section is very strong, because the ranking is required to be consistent across *all* u, p . It would be interesting to investigate weaker criteria that are based on restricted domains of payoff functions and objective distributions.

Mixed DAG representations

The simple representation (9) can be extended, by allowing the DM to be characterized by a *probability distribution* λ over DAGs R , such that his subjective belief (given the objective distribution p) over y is

$$p_\lambda(y) = \sum_R \lambda(R) p_R(y) \quad (11)$$

This fits cases in which the DM has uncertainty regarding the causal relations in his environment, thus capturing gradations of belief biases such as omitting or inverting links.

Definition 7 is extendible to such mixed DAG representations. Let λ^* be a distribution that assigns probability one to linear orderings. Consider two distributions λ, λ' that satisfy $\lambda = \alpha\lambda^* + (1 - \alpha)\lambda'$, where $\alpha \in (0, 1)$. It is easy to see that λ is more rational than λ' . Thus, mixed representations are also useful because they enable us to order types according to a performance-based rationality ranking.

Example 5.1: Partial cursedness (Eyster and Rabin (2005))

Let $y = \{x_1, x_2, x_3\}$. Consider the "coarse" DAG $1 \rightarrow 2 \rightarrow 3$. The true DAG is an extension of the coarse DAG into a linear ordering. Suppose that the DM's type is characterized by a mixture, denoted ER , that assigns

⁹Eyster and Piccione (2013) made an observation in the same spirit in the context of their model of competitive asset markets in which traders hold diversely coarse theories.

probability δ to the true DAG and probability $1 - \delta$ to the coarse DAG. Then,

$$p_{ER}(x_1, x_2, x_3) = p(x_1, x_2)[\delta p(x_3 \mid x_1, x_2) + (1 - \delta)p(x_3 \mid x_1)]$$

where $1 - \delta$ measures the degree of the DM's underestimation of the correlation between x_3 and x_2 . This is the representation of "partial cursedness" due to Eyster and Rabin (2005), for the simple case of an uninformed decision maker. Thus, partially cursed DMs can be unambiguously ranked in terms of their performance: a DM type with a higher δ is "more rational". This conclusion can be extended to the case of partially informed DMs.

6 Concluding Remarks

Part of the appeal of the Bayesian-network approach to modeling boundedly rational expectations is that it establishes links with fields that are currently far from microeconomic theory. For instance, in Section 3 I showed how the formalism can capture fallacies of statistical inference - mistaking correlation for causation, relying on observed correlations to evaluate interventions, ignoring latent variables, etc. Statistics and econometrics teachers pour gallons of sweat to "cure" students of such fallacies. Indeed, the impetus behind part of the Bayesian-network literature (Pearl (2000)) is to systematize *correct* reasoning about identification of causal relations (see Heckman and Pinto (2013), Pearl (2013) and Dawid (2014), for recent reflections on this role of graphical probabilistic models). Instead, I used the formalism *positively*, to model precisely the kind of systematic errors that statisticians warn us against.

The framework also alludes to the (somewhat abandoned) tradition that linked bounded rationality and artificial intelligence, which dates back to Herbert Simon's works (see Rubinstein (1993,1998), Radner and van Zandt (2001), Cho (1995) or Jehiel and Samet (2005)). Bayesian networks are useful for machine learning because they provide a platform for efficient computa-

tion of probabilistic inferences; in particular, perfect DAGs with relatively small cliques are appealing in this respect. This suggests an interesting direction for future research. Just as the finite-automata formalism enabled game theorists to incorporate complexity considerations into players' choices of strategies (as in Rubinstein (1986)) or beliefs (as in Eliaz (2003) and Spiegel (2004)), the Bayesian-network representation of subjective beliefs may enable researchers to introduce complexity of statistical inferences into positive economic modeling.

References

- [1] Alon, N. (2014), "A Note on the Running Intersection Property," Tel Aviv University, mimeo.
- [2] Bazerman, M. and W. Samuelson (1985), "Negotiation under the Winner's Curse," *Research in Experimental Economics*, 3, 105-138.
- [3] Bray, M. (1982), "Learning, Estimation, and the Stability of Rational Expectations," *Journal of Economic Theory*, 26, 318-339.
- [4] Cho, I. (1995), "Perceptrons Play the Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 67, 266-284.
- [5] Cho, I., N. Williams and T. Sargent (2002), "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1-40.
- [6] Cowell, R., P. Dawid, S. Lauritzen and D. Spiegelhalter (1999), *Probabilistic Networks and Expert Systems*, Springer, London.
- [7] Dawid, P. (2014), "Statistical Causality from a Decision-Theoretic Perspective," mimeo.

- [8] Eliaz, K. (2003), “Nash Equilibrium when Players Account for the Complexity of their Forecasts,” *Games and Economic Behavior*, 44, 286-310.
- [9] Esponda, I. (2008), “Behavioral Equilibrium in Economies with Adverse Selection,” *The American Economic Review*, 98, 1269-1291.
- [10] Esponda, I. and D. Pouzo (2012), “Learning Foundation and Equilibrium Selection in Voting Environments with Private Information,” *mimeo*.
- [11] Esponda, I. and D. Pouzo (2014a), “An Equilibrium Framework for Modeling Bounded Rationality,” *mimeo*.
- [12] Esponda, I. and D. Pouzo (2014b), “Conditional Retrospective Voting in Large Elections,” *mimeo*.
- [13] Ettinger, D. and P. Jehiel (2010), “A Theory of Deception,” *American Economic Journal: Microeconomics* 2, 1-20.
- [14] Evans, G. and S. Honkapohja (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press.
- [15] Eyster, E. and M. Piccione (2013), “An Approach to Asset Pricing Under Incomplete and Diverse Perceptions,” *Econometrica*, 81, 1483-1506.
- [16] Eyster, E. and M. Rabin (2005), “Cursed Equilibrium,” *Econometrica*, 73, 1623-1672.
- [17] Fischhoff, B. (1975), “Hindsight is not Equal to Foresight: The Effect of Outcome Knowledge on Judgment under Uncertainty,” *Journal of Experimental Psychology: Human Perception and Performance* 1 (1975), 288-299.
- [18] Heckman, J. and Pinto, R. (2013), “Causal analysis after Haavelmo,” IZA Discussion Paper, No. 7628.

- [19] Jehiel, P. (2005), “Analogy-Based Expectation Equilibrium,” *Journal of Economic Theory*, 123, 81-104.
- [20] Jehiel, P. and F. Koessler (2008), “Revisiting Games of Incomplete Information with Analogy-Based Expectations,” *Games and Economic Behavior*, 62, 533-557.
- [21] Jehiel, P. and D. Samet (2005), “Learning to Play Games in Extensive Form by Valuation,” *Journal of Economic Theory*, 124, 129-148.
- [22] Kearns, M., M. Littman and S. Singh (2001), “Graphical models for game theory,” *Proceedings of the Seventeenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc.
- [23] Koller, D. and N. Friedman (2009), *Probabilistic Graphical Models: Principles and Techniques*, MIT Press, London.
- [24] Koller, D. and B. Milch (2003), “Multi-Agent Influence Diagrams for Representing and Solving Games,” *Games and Economic Behavior* 45, 181-221.
- [25] Koski, T. and J. Noble (2009), *Bayesian Networks: An Introduction*, John Wiley & Sons, Chichester.
- [26] Köszegi, B. (2010), “Utility from Anticipation and Personal Equilibrium,” *Economic Theory*, 44, 415-444.
- [27] Langer, E. (1975), “The illusion of control,” *Journal of personality and social psychology* 32, 311-328.
- [28] Lucas, R. (1976), “Econometric policy evaluation: A critique,” *Carnegie-Rochester conference series on public policy*, Vol. 1. North-Holland.
- [29] Madarasz, K. (2012), “Information projection: Model and Applications,” *Review of Economic Studies* 79, 961-985.

- [30] Osborne, M. and A. Rubinstein (1998), “Games with Procedurally Rational Players,” *American Economic Review*, 88, 834-849.
- [31] Pearl, J. (1988), *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo.
- [32] Pearl, J. (2000), *Causality: Models, Reasoning and Inference*, Cambridge University Press, Cambridge.
- [33] Pearl, J. (2013), “Reflections on Heckman and Pinto’s ‘Causal analysis after Haavelmo’,” mimeo.
- [34] Piccione, M. and A. Rubinstein (2003), “Modeling the Economic Interaction of Agents with Diverse Abilities to Recognize Equilibrium Patterns,” *Journal of the European Economic Association*, 1, 212-223.
- [35] Rabin, M. and D. Vayanos (2010), “The gambler’s and Hot-Hand Fallacies: Theory and Applications,” *Review of Economic Studies*, 77, 730-778.
- [36] Radner, R. and T. van Zandt (2001), “Real-Time Decentralized Information Processing and Returns to Scale,” *Economic Theory* 17, 497-544.
- [37] Rubinstein, A. (1986), “Finite Automata Play the Repeated Prisoner’s Dilemma,” *Journal of Economic Theory*, 39, 83-96.
- [38] Rubinstein, A. (1993), “On Price Recognition and Computational Complexity in a Monopolistic Model,” *Journal of Political Economy*, 101, 473-484.
- [39] Rubinstein, A. (1998), *Modeling Bounded Rationality*, MIT Press, Cambridge, Massachusetts.
- [40] Sargent, T. (2002), *The conquest of American inflation*, Princeton University Press.

- [41] Spiegel, R. (2004), “Simplicity of Beliefs and Delay Tactics in a Concession Game,” *Games and Economic Behavior*, 47, 200-220.
- [42] Spiegel, R. (2011), *Bounded Rationality and Industrial Organization*, Oxford University Press, Oxford.
- [43] Verma, T. and J. Pearl (1991), “Equivalence and Synthesis of Causal Models,” *Uncertainty in Artificial Intelligence*, 6, 255-268.
- [44] White, H. and K. Chalak (2009), “Settable systems: an extension of Pearl’s causal model with optimization, equilibrium, and learning,” *Journal of Machine Learning Research* 10, 1759-1799.
- [45] Woodford, M. (2013), “Macroeconomic Analysis without the Rational Expectations Hypothesis,” *Annual Review of Economics*, forthcoming.

Appendix: Proofs

Proposition 2

Fix p . The set of ε -perturbations of p is compact and convex. For a fixed $\varepsilon \in (0, 1)$, let Q^ε be the set of profiles of conditional distributions $(p(a \mid t)_{a \in A})_{t \in T}$ such that $p(a \mid t) \geq \varepsilon$ for every a, t . Define

$$BR(p) = \arg \max_{q \in Q^\varepsilon} \sum_t p(t) \sum_a q(a \mid t) \sum_y p_R(y \mid a, t) u(a, t, y)$$

If p^* is a ε -personal equilibrium, then $(p^*(a \mid t)_{a \in A})_{t \in T} \in BR(p^*)$. Because $p_R(y \mid a, t)$ is continuous function in p , BR is continuous as well. Also, the target function in the definition of BR is linear in q , hence $BR(p)$ is a convex set. Since the set Q^ε is compact and convex, BR has a fixed point, by Kakutani’s theorem. Therefore, a ε -personal equilibrium p^ε exists for any $\varepsilon > 0$. By standard arguments, there is a convergent sequence of ε -personal equilibria.

Proposition 3

Part (i). I begin by stating a recent result due to Noga Alon. A sequence of sets S^0, S^1, \dots, S^K is **expansive** if for every $k \geq 1$, $|S^k \cap (\cup_{j < k} S^j)| \geq |S^i \cap (\cup_{j < k} S^j)|$ for all $i > k$.

Theorem 2 (Alon (2014)) *Suppose that \mathcal{S} satisfies RIP*. Then, every expansive ordering of \mathcal{S} satisfies RIP.*

Suppose that \mathcal{S} satisfies RIP*. Consider the sequence of sets S^k that are introduced in round k until the procedure is aborted/terminated in some round K . By Step 2 of the procedure, the sequence B^0, S^1, \dots, S^K is expansive. By Theorem 2, it satisfies RIP. Thus, for every round $k \geq 1$ of the iterative imputation procedure, $S^k \cap B^{k-1}$ is weakly contained in some $S^j \in \{B^0, \dots, S^{k-1}\}$. My task is to show that for every p , the procedure terminates in round $m-1$, such that the final belief p^{m-1} is given by (2), where R is some perfect DAG over N whose set of cliques is \mathcal{S} .

I will now prove by induction on k that p^k is consistent with a perfect DAG R^k . Let $k = 1$. Recall that $B^1 = B^0 \cup S^1$. Because \mathcal{S} is a cover of N and does not include sets that contain one another, $S^1 - B^0$ and $B^0 - S^1$ are non-empty. The auxiliary beliefs p_1^1 and p_2^1 defined over X_{B^1} are given by

$$\begin{aligned} p_1^1(x_{B^1}) &= p(x_{B^0})p(x_{S^1-B^0} \mid x_{S^1 \cap B^0}) \\ p_2^1(x_{B^1}) &= p(x_{S^1})p(x_{B^0-S^1} \mid x_{S^1 \cap B^0}) \end{aligned}$$

By the basic rules of conditional probability, we have

$$p(x_{B^0})p(x_{S^1-B^0} \mid x_{S^1 \cap B^0}) = p(x_{S^1})p(x_{B^0-S^1} \mid x_{S^1 \cap B^0})$$

and therefore p^1 is consistent with a DAG R^1 defined over B^1 , where $i \tilde{R} j$ if and only if $i, j \in B^0$ or $i, j \in S^1$. For convenience, impose the following

direction of links: $j \not\rightarrow i$ for every $i \in B^0, j \in S^1$. The DAG R^1 has two cliques, B^0 and S^1 .

Consider the initial condition (B^{k-1}, p^{k-1}) of the k^{th} round, $k > 1$. The inductive hypothesis is that p^{k-1} has a Bayesian-network representation with a perfect DAG R^{k-1} over B^{k-1} , where the set of maximal cliques of R^{k-1} is $\{B^0, S^1, \dots, S^{k-1}\}$. By RIP, $S^k \cap B^{k-1}$ is weakly contained in one of the sets B^0, S^1, \dots, S^{k-1} . By the assumption that no set in \mathcal{S} contains one another, $B^{k-1} - S^k$ and $S^k - B^{k-1}$ are non-empty. The auxiliary beliefs p_1^k and p_2^k over $B^k = B^{k-1} \cup S^k$ are given by

$$\begin{aligned} p_1^k(x_{B^k}) &= p^{k-1}(x_{B^{k-1}})p(x_{S^k-B^{k-1}} \mid x_{S^k \cap B^{k-1}}) \\ p_2^k(x_{B^k}) &= p(x_{S^k})p^{k-1}(x_{B^{k-1}-S^k} \mid x_{S^k \cap B^{k-1}}) \end{aligned} \quad (12)$$

By the inductive step, p^{k-1} has a Bayesian-network representation with a perfect DAG over B^{k-1} . Therefore, the term

$$p^{k-1}(x_{B^{k-1}})p(x_{S^k-B^{k-1}} \mid x_{S^k \cap B^{k-1}}) \quad (13)$$

is a Bayesian-network representation with a perfect DAG R^k over $B^{k-1} \cup S^k$, where S^k is a maximal clique in R^k , and for every $i \in B^{k-1}, j \in S^k$, $i \tilde{R} j$ only if i or j are in $S^k \cap B^{k-1}$.

The expressions in (12) can be rewritten as

$$\begin{aligned} p_1^k(x_{B^k}) &= p(x_{S^k})p^{k-1}(x_{B^{k-1}}) \cdot \frac{1}{p(x_{S^k \cap B^{k-1}})} \\ p_2^k(x_{B^k}) &= p(x_{S^k})p^{k-1}(x_{B^{k-1}}) \cdot \frac{1}{p^{k-1}(x_{S^k \cap B^{k-1}})} \end{aligned} \quad (14)$$

I will show that $p(x_{S^k \cap B^{k-1}}) = p^{k-1}(x_{S^k \cap B^{k-1}})$. I use the simplifying notation $S^k \cap B^{k-1} = V$, $B^{k-1} = M$, $R^{k-1} = R$. Thus, for every $M' \subset B^{k-1}$, $p^{k-1}(x_{M'})$

is denoted $p_R(x_{M'})$. Observe that

$$p_R(x_V) = \sum_{x'} p_R(x_V, x'_{M-V})$$

We have established that V is weakly contained in some clique C of R . Remark 1 implies that we can take C to be an ancestral clique w.l.o.g. Therefore,

$$p_R(x_V, x'_{M-V}) = p(x_V) \cdot p(x'_{C-V} \mid x_V) \cdot \prod_{i \in M-C} p(x'_i \mid x_{R(i) \cap V}, x'_{R(i)-V})$$

When we sum this expression over x' , we obtain

$$p(x_V) \sum_{x'_{C-V}} p(x'_{C-V} \mid x_V) \sum_{x'_{M-C}} \left(\prod_{i \in M-C} p(x'_i \mid x_{R(i) \cap V}, x'_{R(i) \cap (C-V)}, x'_{R(i)-C}) \right)$$

which is equal to $p(x_V)$. We have thus established that $p(x_{S^k \cap B^{k-1}}) = p^{k-1}(x_{S^k \cap B^{k-1}})$.

Therefore, the procedure is not aborted in round k , and p^k is given by (13). We conclude that the procedure necessarily terminates in round $m-1$. By construction, the set of maximal cliques in R^{m-1} is \mathcal{S} , which means that R^{m-1} is unique in the sense of Remark 1.

Part (ii). Suppose that \mathcal{S} violates RIP*. Note that This means N contains at least three elements. Assume that the procedure is never aborted. Then, it terminates in round $m-1$, for any p . Recall that the sequence B^0, S^1 trivially satisfies RIP. Let $k > 1$ be the earliest round for which $S^k \cap B^{k-1}$ is *not* weakly contained in any of the sets B^0, S^1, \dots, S^{k-1} . By the sufficiency result, $p^{k-1} \in \Delta(X_{B^{k-1}})$ is consistent with a perfect DAG R over B^{k-1} , where the set of maximal cliques of R is $\{B^0, S^1, \dots, S^{k-1}\}$.

As in the case of Part (i), denote $S^k \cap B^{k-1} = V$, $B^{k-1} = M$. We have established that V contains at least two elements, denoted 1 and 2, which do not belong to the same clique. Let C_i denote the maximal clique in R that

contains the node i , $i = 1, 2$. Moreover, since R is perfect, 1 allows us to set C_1 to be ancestral in R , w.l.o.g. For the same reason, we can set w.l.o.g $V \cap C_1$ to be R -maximal in C_1 , and $V \cap C_2$ to be R -minimal in C_2 , such that for each $i = 1, 2$, i is the R -maximal element in $V \cap C_i$. Let us establish the following properties: (i) there exists an element in $C_2 - V$, denoted 3, such that $3R2$; (ii) $1 \not R 2$. Property (i) follows directly from the assumption that no set in \mathcal{S} contains one another, as well as from the designation of $V \cap C_2$ to be R -minimal in C_2 . As to property (ii), assume that $1R2$. Then, since C_2 is a clique in which $V \cap C_2$ is R -minimal, $jR2$ for every $j \in C_2 - V$. And since R is perfect, it follows that $1Rj$ for every $j \in C_2 - V$, which would mean that the set $\{1, 2\} \cup (C_2 - V)$ is a clique, contradicting the fact that 1 and 2 do not belong to the same clique. Note that since C_1 is ancestral, is impossible that $3R1$. We have thus established the existence of two nodes $1, 2 \in V$ and a third node $3 \notin V$, such that $2R3$, $1 \not R 3$ and $3 \not R 1$.

The auxiliary distributions p_1^k and p_2^k are given by (14). I now show that there exists p such that $p(x_{S^k \cap B^{k-1}}) \neq p^{k-1}(x_{S^k \cap B^{k-1}})$, which would mean that the procedure aborts in round k . To use our simplified notation, we need to show that $p_R(x_V) \neq p(x_V)$ for some p . We can assume w.l.o.g that p has the feature that every variable in $N - \{1, 2, 3\}$ is distributed independently of all other variables. Then, $p_R(x_V)$ can be written as

$$p_R(x_V) = p(x_{V-\{1,2\}}) \cdot p(x_1) \cdot \sum_{x'_3} p(x'_3 \mid x_{R(3) \cap \{1\}}) p(x_2 \mid x'_3)$$

whereas $p(x_V)$ can be written as

$$p(x_V) = p(x_{V-\{1,2\}}) \cdot p(x_1) \cdot \sum_{x'_3} p(x'_3 \mid x_{R(3) \cap \{1\}}) p(x_2 \mid x'_3, x_1)$$

Set p such that $p(x_2 \mid x'_3) \neq p(x_2 \mid x'_3, x_1)$ for some x_1, x_2, x'_3 . Then, $p_R(x_V) \neq p(x_V)$, such that the procedure will abort in round k for this p , a contradiction.

Proposition 4

Recall that when $l = 0$, we denote $a = x_0$, $y = (x_1, \dots, x_n)$. A DM with some DAG R satisfying $R(0) = \emptyset$ chooses a to maximize

$$\begin{aligned}
 p_R(y \mid a) &= p_R(x_1, \dots, x_n \mid x_0) \\
 &= \frac{p(x_0) \cdot \prod_{i=1}^n p(x_i \mid x_{R(i)})}{\sum_{x'_1, \dots, x'_n} p(x_0) \cdot \prod_{i=1}^n p(x'_i \mid x_{R(i) \cap \{0\}}, x'_{R(i) - \{0\}})} \\
 &= \frac{\prod_{i=1}^n p(x_i \mid x_{R(i)})}{\sum_{x'_1, \dots, x'_n} \prod_{i=1}^n p(x'_i \mid x_{R(i) \cap \{0\}}, x'_{R(i) - \{0\}})}
 \end{aligned}$$

Recall that we are considering a modification of p that change $p(a)$ for some a , while leaving $p(y \mid a)$ intact for all a, y . Therefore, if $0 \in R(i)$, the terms $p(x_i \mid x_{R(i)})$ and $p(x'_i \mid x_{R(i) \cap \{0\}}, x'_{R(i) - \{0\}})$ are unchanged; while if $0 \notin R(i)$, these terms can be written as follows:

$$p(x'_i \mid x_{R(i) \cap \{0\}}, x'_{R(i) - \{0\}}) = p(x'_i \mid x'_{R(i)}) = \sum_{x''_0} p(x''_0) p(x'_i \mid x''_0, x'_{R(i)})$$

Recall that $p(x'_i \mid x''_0, x'_{R(i)})$ is unchanged. If this probability is not constant in x''_0 , we can find a modification of $p(x'_0)$ for some values of x'_0 such that the expression for $p_R(y \mid a)$ will change. In contrast, if the probability is constant in x''_0 (i.e., $p(x'_i \mid x''_0, x'_{R(i)}) = p(x'_i \mid x'_{R(i)})$), the expression for $p_R(y \mid a)$ is necessarily unchanged.

Proposition 5

We already established the "if" part before the statement of the result. Let us turn to the "only if" part. Assume that R is neither empty nor a linear

ordering. Therefore, $n \geq 3$.

Suppose first that R is intransitive. Then, there are three variables, indexed w.l.o.g 1, 2, 3, such that $1R2$, $2R3$ and $1 \not R 3$. Let us restrict attention to distributions p for which $p(x_{\{1,2,3\}} \mid x_{-\{1,2,3\}})$ is constant - i.e., the variables x_1, x_2, x_3 are independent of all other variables. This means that all our formulas for conditional probabilities can be written as if the only variables are x_1, x_2, x_3 . In particular, $p_R(x_2 \mid x_{-2}) = p_R(x_2 \mid x_1, x_3)$ and $p(x_2 \mid x_{A-\{2\}}) = p(x_2 \mid x_{A \cap \{1,3\}})$. Thus:

$$\begin{aligned} p_R(x_2 \mid x_1, x_3) &= \frac{p(x_1)p(x_2 \mid x_1)p(x_3 \mid x_2)}{\sum_{x'_2} p(x_1)p(x'_2 \mid x_1)p(x_3 \mid x'_2)} = \frac{p(x_2 \mid x_1)p(x_3 \mid x_2)}{\sum_{x'_2} p(x'_2 \mid x_1)p(x_3 \mid x'_2)} \\ &= p(x_2 \mid x_1) \cdot \frac{1}{\sum_{x'_2} p(x'_2 \mid x_1) \left(\frac{p(x_3 \mid x'_2)}{p(x_3 \mid x_2)} \right)} \\ &= p(x_3 \mid x_2) \cdot \frac{1}{\sum_{x'_2} p(x_3 \mid x'_2) \left(\frac{p(x'_2 \mid x_1)}{p(x_2 \mid x_1)} \right)} \end{aligned}$$

The only possibilities for A we need to consider are $\{1\}$, $\{3\}$ and $\{1, 3\}$. It is clear that we can find p with non-constant $p(x_2 \mid x_1)$ and non-constant $p(x_3 \mid x_2)$, such that the denominators in the last two expressions are different from 1, which implies $p_R(x_2 \mid x_1, x_3) \neq p(x_2 \mid x_1)$ and $p_R(x_2 \mid x_1, x_3) \neq p(x_3 \mid x_2)$. Finally, we need to compare $p_R(x_2 \mid x_1, x_3)$ to

$$p(x_2 \mid x_1, x_3) = \frac{p(x_1)p(x_2 \mid x_1)p(x_3 \mid x_1, x_2)}{\sum_{x'_2} p(x_1)p(x'_2 \mid x_1)p(x_3 \mid x_1, x'_2)} = \frac{p(x_2 \mid x_1)p(x_3 \mid x_1, x_2)}{\sum_{x'_2} p(x'_2 \mid x_1)p(x_3 \mid x_1, x'_2)}$$

Once again, it is clear we can find p for which $p(x_3 \mid x_1, x_2)$ does not coincide with $p(x_3 \mid x_2)$, such that the two expressions are different.

Now suppose that R is transitive, but not anti-symmetric. Then, $\{1, \dots, n\}$ can be partitioned into $K \geq 2$ subsets $\{A_1, \dots, A_K\}$ such that R is complete and transitive when restricted to any given A_k . Let $A(i)$ denote the partition cell to which i belongs. Take i, j such that $A(i) \neq A(j)$ and $|A(i)| \geq 2$. Since

R is not empty, there must exist such i, j . Then, $p_R(x_i \mid x_{-i}) = p(x_i \mid x_{A(i)-\{i\}})$ whereas $p_R(x_j \mid x_{-j}) = p(x_j \mid x_{A(j)-\{j\}})$. Since $A(i) \neq A(j)$ and $A(i) - \{i\} \neq \emptyset$, we can find p for which $p(x_i \mid x_{A(i)-\{i\}}) \neq p(x_i \mid x_{A(j)-\{j\}})$.

Proposition 6

Assume the contrary - i.e., that there exist non-equivalent R, R' that are not linear orderings, such that R is more rational than R' . For notational convenience, enumerate $Y = \{1, \dots, L\}$. Fix p and denote $q = p_R$, $r = p_{R'}$. Both q and r are probability vectors of length L . Define the length- L vector z as follows. For each y , $z(y) = u(a, y) - u(a', y)$.

Consider the $L \times 3$ matrix

$$A = \begin{pmatrix} r(1) & -q(1) & -p(1) \\ r(2) & -q(2) & -p(2) \\ \vdots & \vdots & \vdots \\ r(L) & -q(L) & -p(L) \end{pmatrix}$$

Let $b = (-\varepsilon, -\varepsilon, -\varepsilon)$ be a vector in \mathbb{R}^3 , where $\varepsilon > 0$ is arbitrarily small. The assumption that R is more rational than R' thus implies that there exists no z that satisfies the inequality $Az > b^T$. By Farkas' Lemma, this means that there is a vector $a > 0$ in \mathbb{R}^3 , such that there $A^T a = 0$ (and since $a > 0$, $ba^T < 0$). Thus,

$$r(y) = \frac{a^2}{a^1} q(y) + \frac{a^3}{a^1} p(y)$$

for every y . Since $\sum_y r(y) = \sum_y q(y) = \sum_y p(y) = 1$ by assumption, $a^1 = a^2 + a^3$, such that the claim holds with $\alpha = a^3/(a^2 + a^3)$.

We have thus established that for any p , we can find $\alpha \in (0, 1)$ such that $p_R = \alpha p + (1 - \alpha)p_{R'}$. In particular, for any p that is consistent with R , $p_R = p$ and so the equation reduces to $p_R = p_{R'}$. Likewise, for any p that is consistent with R' , $p_{R'} = p$ and again we obtain $p_R = p_{R'}$. It follows that the sets of distributions that are consistent with R and R' are identical, contradicting the assumption that R and R' are not equivalent.