

Information in Economics: A Fictionalist Account*

By Philip Mirowski**

Abstract

We propose that certain classes of economic models best be understood as ‘fictions,’ in the sense promoted by Roman Frigg and others. The structure of the argument parallels that made by Arnon Levy for information in biology. The lesson is that economists are not really all that concerned over the sorts of things, such as the nature of knowledge, that philosophers deem central to epistemology.

JEL Codes: B21, B41, D80

1. The Fictionalist Account of Models

Much of the last few decades of the philosophy of economics has been taken up with the rejection of the pretense of mathematical models in economics as (in)adequate representations of the economy. There have been a parade of critics who repetitively moan that “economic agents don’t behave like that” or “the assumptions are preposterous” for over a century now; and yet, the empirical fact of their utter irrelevance to the evolution of neoclassical economics should have put paid to simplistic notions that the models could ever be vulnerable to naive blanket accusations of lack of faithfulness of representation. History reveals time and again that neoclassical economists had neither consulted contemporary psychology nor actual market forms in the course of their development of their mathematical models.¹

With the assistance of some philosophers, perhaps this situation could be surmounted if not entirely banished. There has recently appeared in the philosophy of science literature an entirely different approach to models in science,

* This is a revision of a draft of a Keynote address to the Models and Simulations conference at Notre Dame, May 9–11, 2014.

** John J. Reilly Center, University of Notre Dame, 452 Geddes Hall, Notre Dame, Indiana 46556. The author can be reached at Philip.E.Mirowski.1@nd.edu.

¹ For psychology, see (Giocoli 2005); for market structures, see (Mirowski 2007). A broader conception of economists’ practices is provided in (Walliser 2011).

which seems to hold greater promise for understanding what has been going on in the history of economics. In contemporary philosophy, this program has been called the ‘fictionalist approach to models,’ associated with such figures as Roman Frigg, Peter Godfrey-Smith, and Arnon Levy, among others.² Although their accounts differ in various respects, they all treat models as hypothetical systems which exist as autonomous entities subject to their own rules of manipulation and critique, explicitly treated as akin to places and characters in literary fiction. While this proposition might initially strike the reader as a deliciously transgressive, or perhaps a deliberately debunking exercise, these philosophers have made a persuasive case that “[t]reating model systems to be intrinsically representational is a fundamental mistake” (Frigg 2010, 99). From this perspective, models cannot depict actual states of affairs because they are acts of imagination which then conform to local disciplinary rules of elaboration and critique, whose relationships to nature are postponed during intervals when their internal implications are worked out to the satisfaction of the relevant thought collective. Philosophy of science has long renounced the quest for a single set of unique rules of scientific model-making; the fictionalist view has the further salutary implication that different disciplines would tend to follow different rules for producing and critiquing their own science fictions (Frigg 2010, 124). Moreover, a fictionalist account of models does not preclude commitment to a realist ontology (Levy 2012); rather, one would observe a stratification and division of labor between those most concerned to develop ‘intrafictional’ propositions and those dedicated to establishment of ‘transfictional’ propositions (Frigg 2010, 117–119). This may or may not map into conventional roles of ‘theoretical’ and ‘applied’ scientists, since a body of theory may consist of something more than a collection of currently sanctioned models.

While much of the philosophical literature is concerned to elucidate what it means precisely for a model to qualify as ‘fictional,’ I should rather like to highlight how the fictionalist account of models can tell us new and interesting things about economics that were beyond the pale in previous discussions of the nature of models. For instance, there is something very special about the practice of construction of mathematical models which the new approach foregrounds. Mathematical models constitute a unique genre of fiction, one which tends to obliterate all narrative context by pretending the model in question was constructed *de novo* from first principles or axioms. The beauty of this narrative convention is not only that it is false, but more significantly, it fosters plausibility by pretending that no other rival fiction concurrently exists in the domain of that particular field of speculative endeavor. The mathematical model is thus enshrined in a mythical space empty of spectators with their rival models and their troublesome preoccupations, as though there were no prior history of theo-

² See (Godfrey-Smith 2009; Frigg 2010). For a critical overview, see (Weisberg 2013, chap. 4).

retical speculation, no other models in play at that moment, nor indeed any active dispute over the narrative whatsoever. This widespread modeling convention is a major technique in construction of the “view from nowhere” so central to the conjuring of objectivity in science. The fictionalist approach to models thus compares this standard mathematical trope with modernist literary fiction, which by contrast revels in the co-existence of multiple simultaneous narratives, unreliable narrators, paradoxes of self-reference, and the entire panoply of narrative techniques of the post-Joyce, post-Woolf, post-Faulkner generations. The fictionalist approach to models permits contemplation of indirect and implicit functions of models such as these narrative consequences of mathematical conventions, above and beyond the standard attributes of parsimony, predictive accuracy, logical completeness, and the like.

What might be particularly salutary about this body of philosophical literature is that it automatically avoids the unthinking reification of physics as the science *par excellence* as a template for the philosophy of science; partly the rethinking of models arose because biology had been looming large in recent philosophical discussions. A number of the philosophical advocates of models as fictions make explicit reference to various models in economics (as well as physics and biology) in their roster of exemplars.³ One should not find this odd or troublesome, since it seems there have been few disciplines more besotted with the image of themselves as tool-makers and model-builders, a ‘selfie’ that promotes their self-concept as scientists. The recent work of historian of economics and philosopher Mary Morgan, although she avoids explicit endorsement of the fictionalist approach, does seem to have been highly influenced by its proponents, writing for instance that, “[m]odel making ... is an activity of creating small worlds expressed in another medium” (Morgan 2012, 30). “From a naturalist philosophy of science viewpoint, the way that economists work with models suggests they are regarded, and may be understood, as autonomous working objects” (Morgan and Knuuttila 2012, 67). This has been seconded by a few other methodologists (e.g., Sugden 2000). Finally, it should be noted that some famous economists have been giving fictionalist accounts of their activities for a number of years now. Nowhere has this trope of models as fictions been more dominant than at the University of Chicago in the 20th century, from Milton Friedman to Freakonomics. I don’t think the fictionalist account comes out more clearly than in the methodological comments of Bank of Sweden Prize winner Robert Lucas:

Economists have an image of practicality and worldliness not shared by physicists and poets. Some economists have earned this image. Others – myself and many of my colleagues here at Chicago – have not. I’m not sure whether you will take this as a confession or a boast, but we are basically story-tellers, creators of make-believe economic systems ...

³ See, for instance, (Frigg 2010; Frigg and Hartmann 2012).

And, as we all know, the analogy that one person finds persuasive, his neighbor may well find ridiculous. Well, that is why honest people can disagree. I don't know what one can do about it, except keep trying to tell better and better stories, to provide the raw material for better and more instructive analogies. How else can we free ourselves from the limits of historical experience so as to discover ways in which our society can operate better than it has in the past?

... that is what economists do. We are storytellers, operating much of the time in worlds of make believe. We do not find the realm of imagination and ideas is an alternative to, or a retreat from practical reality. On the contrary, it is the only way we have found to think seriously about reality. In a way, there is nothing more to this method than maintaining the conviction (which I know you have after four years at Chicago) that imagination and ideas matter ...

It is fun and interesting and, really, there is no alternative (Lucas 1988).⁴

So there you have it: some recent philosophical accounts of the nature of modeling and the self-accounts of (some) economists concerning their modeling activities have tended to converge into one harmonious shared narrative, proposing that models be explicitly approached as fictions. Whilst some philosophers might find this conjunction of characterization of models reassuring, I am interested in taking it yet a bit further, to ask whether the fictionalist account of models might serve to usefully structure as-yet unwritten histories of models in economics. Here, I will follow the lead of one of the eminent proponents of the fictionalist approach, Arnon Levy.

2. Information as Exemplary Fiction

If we were to entertain an approach to economic models as fictions, it would help to focus upon something beyond an isolated equation or two here or there in order to give the philosophical notion a serious run for its money. One should explore whole classes of models, or alternatively, a conceptual protagonist that keeps showing up in one model after another. Moreover, everything that is colloquially called a 'model' may not qualify as a candidate in the fictionalist approach.⁵

Philosophers of science have not always been known to provide extensive exemplars as illustrations of their theoretical meditations, so it is all the more auspicious that one contemporary philosopher has gone to the trouble to spell out the fictionalist approach to models in one very specific scientific instance, namely, the use of the information concept in postwar biology. Arnon Levy

⁴ See also (Buchanan 2014; Sugden 2013). A sampling of the stories economists have told themselves about their models can be viewed in (Institute for New Economic Thinking 2011).

⁵ See, in the case of biology, (Levy and Currie 2014).

(2011) has noticed that there has been a fair amount of controversy over the use of the information concept in biology, and yet, nevertheless it has served as a major organizing principle in numerous attempts to model all manner of biological phenomena, extending well outside the expected domain of the ‘DNA code.’ One could approach this curious entity ‘information’ as resembling a fictional protagonist who keeps turning up like a bad penny, but nevertheless provides necessary narrative continuity across a whole series of models, rather like an Aristotelian tragic hero.

The bill of attainder concerning appeals to information in contemporary biology are familiar in the philosophy of science literature, so much so that the issues are treated as fairly pedestrian; there is even a separate entry in the *Stanford Encyclopedia of Philosophy* that conveniently summarizes them in a succinct manner.⁶ The indictment begins by parsing two different rival meanings of information. There is a weak sense of the terminology of information, which deploys it as a synonym for mere correlation between variables and states. That usage is relatively harmless, although there one could just as easily dispense with the information concept as superfluous: people do not generally cite an ‘informational’ connection between smoke and fire, or dust and a sneeze. Furthermore, it would seem otiose to consecrate ‘information’ as just another way to talk about causality. The other sense, which is often freighted in under the cover of some other metaphorical reference, is to *semantic* or *intentional* information. These locutions often arise in contexts where assertions are made that genes ‘carry a message,’ telling an embryo how it must develop. These metaphors are nearly pervasive in our culture, with information treated as the primary causal agent in a context-independent sense, even though no such high-level physical property can be said to actually exist. Some philosophers have been scathing about the widespread use of the second concept of information in biology: “there is no clear technical notion of ‘information’ in molecular biology. It is little more than a metaphor that masquerades as a theoretical concept” (Sarkar 1996, 187). “The history of information talk in molecular developmental biology is one consistent retreat from the literal to the metaphoric” (Griffiths 2001, 408).

I am certainly neither capable nor willing to pronounce with confidence upon the validity of this indictment in the sphere of biology; the pertinent task for current purposes is rather to try and understand the sources of perennial popularity of information talk in biology, given its lack of solid referent.⁷ Levy (2011) takes this as the starting point of his own article, which suggests that “information language in biology should not be treated literally” (647) as do the critics; rather, it is a prime example of recourse to a model as an explicit

⁶ See (Godfrey-Smith and Sterelny 2008), itself based on a literature which includes (Sarkar 1996; Griffiths 2001).

⁷ Some exemplary histories of this problem are (Kay 2000; Kline 2004).

fiction. The scientists engaged in building biological models regularly acknowledge that appeals to information do not allow full-blown ascription of intentional properties to the gene or the gland secreting an enzyme; they treat information as a shorthand “where one wishes to understand how the activity of one cell or structure is regulated by a distinct cell or structure, typically in a complex and adaptive manner” (649). Something is being imported from another science – be it a rough model of a Shannon communication device, a simulacrum of a computer, or perhaps even some bits of mathematics from a prior model, such as a game-theoretic setup – but this frequently is hedged about with qualifications that the author in question does not really think there are actual programs run, signals transmitted, or games being played. The last thing the biologist wants to become embroiled in is the actual metaphysics of the information entity posited, and the reason is because everyone involved understands it as a fiction. Neither is it some manner of ‘idealization,’ since in many cases Levy insists it is a liminal metaphor that escapes explicit commentary. Nevertheless, the fiction in question is neither gratuitous nor ornamental: much of the model is shaped by the previous history of information concepts in other sciences. For instance, the intentional connotations are often accessed when natural selection itself is modeled as a force bearing down upon the successful operation of the organism. Quotes like the following are nearly ubiquitous: “In biology, the statement that A causes information about B implies A has the form it does because it carries that information ... The element of intentionality comes from natural selection” (Smith 2010, 139).

The information concept brings weight and heft to the model: it induces some things to be highlighted, and others to be provisionally ignored. Moreover, it traces its legacy from a seemingly unstoppable spread of information talk throughout the other sciences: “The big bang was a bit bang. Starting from its very earliest moments, every piece of the universe was processing information” (Davies and Gregersen 2010, 96). The prescription proposed by Levy is not to marvel at such impetuous use of fictions such as these, but rather to seek to unpack the myriad ways that recourse to the fiction has shaped the details of the model in the context of the science of interest, perhaps as it gets appropriated from a different science already immersed in a different set of narrative tropes. Imagination does not run riot, but travels in some well-worn grooves. Levy, of course, is not himself an historian, but does explicitly regard his project “to provide a general view of informational notions in biology, one that takes into account their full range of application” (2011, 640).

I believe that Levy’s “information in biology” might serve to set the tone and template for how one might begin to deploy the fictionalist theory of models to actually organize the history of certain episodes in the history of the sciences in general. In the remainder of this article, I shall follow Levy’s lead, and will briefly demonstrate how one might go about this task in the history of modern

economics. Conveniently, I will focus on the very same model complex which was the designated topic of Levy's paper: the brace of models surrounding information, but only now situated squarely in the realm of economics.

3. The Three-Fold Path of Information Models in 20th Century Economic Theory

An attempt to survey the uses of 'information' in economics immediately encounters a daunting obstacle: there is currently no consensus history of information economics, nor even a standard history of the orthodoxy of information economics, to be found within the contemporary discipline. Frequently, economists bandy about the term as if they could readily presume everyone else has been using it in the same manner; but even textbooks of microeconomics admit that pretense is unfounded: twenty-five years on, David Kreps' (1990, 578 fn) warning is still good advice: "The terms of information economics, such as moral hazard, adverse selection, hidden action, hidden information, signaling, screening and so on are used somewhat differently by different authors, so you must keep your eyes open when you see any of these terms in a book or an article ... As a consumer of the literature, you should pay less attention to these labels and more to the 'rules of the game' – who knows what when, who does what when." This constitutes the first clue that the practices surrounding these models signals something unstable within the scientific community: a general uneasiness about a perceived ontological promiscuity, leavened with a blanket acceptance of a fair amount of indeterminacy. One survey in 1995 concluded: "we encounter an overabundance of results and/or equilibria; almost anything can happen" (Levine and Lippman 1995, xii). This is the first symptom of what we might reasonably expect from the vantage point of a fictionalist account of models, especially in their early stages.

The second clue to the significance of information in economics is the demonstrable fact that the supposed 'natural field of application' in the social sciences vis-à-vis the natural sciences has not been borne out by its historical development. Although it would seem obvious that economics as a field must deal in communication, which would of necessity encompass information, much of the history of neoclassical theory had denied that communication was required for markets to function properly. Hence economists were not in the forefront of development of information concepts. Indeed, the historical vectors of influence have run in reverse, from the natural sciences to economics. When searching for ways to model information, economists did not intrepidly invent their own formalisms, but rather reached out to the sciences of inanimate nature for their inspiration not once, but at least on three different occasions. Thus, the problem of intentionality cited above in biology actually festered in economics as well, although in a somewhat different manner.

Just as in the case of biology, reference to ‘information’ in postwar economic theory has been pervasive and ubiquitous; but the problem bedeviling model practice has not been a precautionary distance from intentionality when engaging with the metaphor, in a manner parallel with biology. Instead, the looming problem with information in economics might be dubbed *empty intentionality*: modeling practice has assumed the format of attempting to reify something outside of any serious commitments to mind or epistemology, avoiding both will and ratiocination, such that it might be effortlessly integrated into the prior neo-classical orthodoxy. In other words, time and again in the pageant of neoclassical economic models of information, one observes a putatively intentional object lacking an intentional agent. Partly, this could be an artifact of the appropriation of formal models from the natural sciences; but this is also a hallmark of the fictionalist nature of the models in economics.

Third, even though many economists date the inception of the economics of information from the third quarter of the 20th century, the literature has grown so massive in such short order that a literal intellectual history is unachievable in this context.⁸ In its place, we briefly characterize three major classes of economic models of information, each of which has given rise to a rather large subset of the literature, and all of which are best suited to illustrate the fictionalist approach to constructing a more elaborate history. They are: information as a commodity; information as a partition over a comprehensive state space; and information as the product of computation. They are summarized below in Table 1.

Table 1

<i>Information is:</i>		
a thing (Shannon)	an inductive index (Blackwell)	symbolic computation (Turing)
<i>Cognition is:</i>		
irrelevant	intuitive statistics & epistemic formal logic	symbol manipulation
<i>Learning is:</i>		
purchase of a commodity	statistical inference	algorithm augmentation
<i>Communication is:</i>		
same as exchange	‘signaling’	information transmission

⁸ See, however, (Mirowski and Nik-Khah, forthcoming).

3.1 Information as a Thing/Commodity

If one dominant research heuristic of postwar economists was to ‘do as little as possible to revise or alter the neoclassical theory handed down from our forebears’ when discussing the mathematics of information, then one can appreciate the path of least resistance was simply to append a subscript to the existing model. If one could get away with information as the stuff of knowledge, then cognition might remain quarantined away from economics. If information really existed as a thing-like object in nature, then it could just be subtended to the commodity space as just one more good, and apparently ‘nothing’ need be changed about the standard maximization model whatsoever. The received textbook model would be safe. Moreover, such a thing-like information concept would conveniently absolve the theorists in question of having to confront whatever model of mind which was supposedly inherent in the utility function. Although various versions were proposed beginning in the 1950s, the options tended to gel in the 1960s, with Kenneth Arrow (1962) portraying scientific knowledge as a ‘public good;’ Gary Becker (1964) lumping together knowledge, information, ideas, skills, and health of individuals, all under the rubric of ‘human capital;’ and Fritz Machlup (1962) busily taxonomizing the information commodity into different types of ‘goods’ – investment, intermediate, and consumption. Later still, information deliquesced, ‘spilling over’ into all sorts of positive externalities.

The problem immediately arose as to how to ‘measure’ or ‘quantify’ this kind of information, and that is where Claude Shannon’s “information theory” served a critical fictional function.⁹ Shannon had developed an argument which suggested information could be treated just like entropy in physics, comparing it to an enumeration of the number of ways a stochastic combination of symbols could make up a measurable macrostate of messages. A concept fashioned to discuss capacities of communication channels, devoid of all semantic content, may turn out to be utter nonsense when used to discuss the semantics of communication in trade, as many soon came to suspect. But that did not restrain its significance for economics. To the untutored, Shannon seemed to proffer a unit of measurement for information in the ‘bit.’ The Shannon enthusiasm of the first two postwar decades had the unintended consequence of bolstering the general impression that scientists could and should treat information as a quantifiable thing, and even as a *commodity*. In practice, it became quite common to conflate the embodiments and encapsulations of knowledge in objects and artifacts as mere epiphenomenal manifestations of a generic ‘thing’ called information. Some of it was cast as inherently fungible; but other portions might prove a bit refractory, and therefore were written off as ‘tacit.’ The entire parable was a reification based largely upon a misapprehension – but that

⁹ See (Shannon and Weaver 1949; Geoghegan 2008).

didn't mean that it still wouldn't have untold consequences down the line. One suppressed implication, as we would expect from the fictionalist approach to models, was the bogey of self-reflexivity: "I do not suppose that the information content of this essay could ever be quantified" (Dorfman 1960, 585) wrote one nervous economist.¹⁰

Nevertheless, once knowledge was broadly construed as a 'good,' then arguments could begin over just what special sort of good it might be. A riot of metaphorical invention ensued in economics. Perhaps it resembled a capital good, but one capable of metempsychosis, like 'human capital.' Or, perhaps its special conditions of production dictated its status as a 'public good.' The trick was to decide how some particular bit of information might ascend to this status. If you could get people to accept that knowledge was an eminent instance of such a 'good,' it helped if you then began to endow it with all sorts of peculiar qualities. Starting with Paul Samuelson (1954) and Kenneth Arrow (1962), knowledge was claimed to be a slippery elusive sort of thing whose use by one person did not restrict or prevent its use by another (in the jargon: 'non-rivalrous'); but also something from which it was intrinsically difficult to prevent another from enjoying the benefits once you bought it (in the jargon: 'non-excludable'). This created rather more problems for mathematical modeling, but more to the point, was used in the 1960s-80s to justify state subsidy and provision of this marvelous commodity.

The fictional character of these models seems fairly straightforward. No advocate of these models proceeded to resolve the 'information' involved into measurable 'bits;' neither did anyone go about modeling a 'channel' with the normal Shannon characteristics of a fixed capacity, or a noise source. Moreover, no real-life market sold anything like commodity units of 'information;' every real-life application involved sale of some other derivative object (a book, a lecture, an experience) or a set of legal property rights. As fictional stylists, economists betrayed a weakness for synecdoche, misrepresenting the part for the whole. Economists tended to conflate intellectual property with information, even though that exhibited a severe misunderstanding of the nature of patents. Any such objections were treated as mere quibbles; 'information' was pronounced the lifeblood of the New Economy, and nothing would withstand the drumbeat of reification of information into a commodity.

¹⁰ The notion that those devious relativists who thrive in science studies are the only cadre who are susceptible to the perils of reflexivity is one of the sillier arguments made by modern philosophers.

3.2 Information as an Inductive Index and/or the Stochastic Object of an Epistemic Logic

With the development of mathematical statistics, there had been efforts early in the 20th century to link intuitions of a ‘good sample’ to the amount of ‘information’ it contained, often in the tradition of R. A. Fisher. None of these proposals had amounted to much outside the ambit of a small coterie of statisticians in the 1930s. However, in the postwar period, an interesting phenomenon occurred where the statistical tools of inductive inference (having just spread throughout the social sciences) began to get conflated with more capacious models of mind. Since the story of psychology in the early 20th century consisted of a series of frontal assaults on the conscious mind as executive in charge of rationality, a revanchist movement resorted to the theory of probability to stem the tide.¹¹

The situation escalated when mathematical statisticians were brought together with operations researchers, philosophers and game theorists at the RAND Corporation in the early 1950s. There, especially in the work of David Blackwell, a practice took hold of equating ‘information’ with measures defined over partitions imposed upon an exhaustive total enumeration of states of the world, both actual and virtual.¹² Crudely, how much a procedure (it was harder to phrase this in terms of real people) ‘knew’ about a world was a function of how finely or coarsely it could divide up the possibilities, distinguish the class of outcome, and thus assign probabilities to eventual outcomes, and the sensitivity with which its detectors could discriminate which of the possibilities had actually obtained. The necessity for game theory to divide and discriminate strategies according to states of the world was an immediate inspiration, but quickly the formalism was developed in two relatively separate directions: one, as the framework for modern definitions of one version of inductive inference, and the other, as the scaffolding used to assign semantic relations to a modal logic. In an alliance with artificial intelligence, it became the basis for formal models of an important class of machine logic.

As in the case of information as a commodity, in this instance the inspiration had also come from outside the economics profession; but our limited task here is to briefly document the fictional character it assumed when it became incorporated into orthodox economic models. The class of models is usually desig-


¹¹ “History has witnessed the attempt to make probability theory coherent with what was believed to be rational thought, and it has seen efforts to reduce rational thought to probability theory. For instance, what was believed to be rational judicial and economic thought actually determined the way in which probability theory developed mathematically” (Gigerenzer and Murray 1987, 137).

¹² The historical background to this development is covered in (Mirowski 2002, 380–386). A nice introductory analytical treatment from the standpoint of epistemic logic is (Fagin et al. 1995).

nated as ‘Bayes-Nash’ models of imperfect information within the profession. When students are told that the modern economics of information has ‘solved’ the problem of asymmetric information, more often than not the speaker is making reference to a complex of modeling practices, which we will call in honor of its progenitor the “Blackwell program,” which was rapidly taken up by the economics profession after 1965. The model complex consists of four basic components, with others added to adapt the model to the specific circumstances (but not covered here).

3.2.1 Reality is a Pre-Existent State Space of Possible Worlds

‘Knowledge’ in this approach is said to consist of partitions over a totally exhaustive state space, which distinguishes possible worlds by the inclusion or exclusion of a preset menu of variables. Greater knowledge is said to be represented by finer and finer partitions over the invariant state space, allowing greater precision concerning location in the space. Notice, the ontology of this world is given and fixed prior to the analysis, and cannot be altered by any activities of the knower. Time, by construction, merely is one of the state variables. Thus, this rather resembles the ‘block universe’ of relativistic physics, where past, present and future all coexist simultaneously in this state space. A being observing the situation from the outside would notice that nothing really new or unprecedented can ever appear in this space. This is the first indication of the fictional character of the model.



3.2.2 *All Information is Inductive*

Blackwell proposed that ‘experiments’ consisted of information extracted from the state space by inductive inference of probabilities attributed to partitions of the state space. In the narrative, ‘information’ consists of signals emitted from experimental interventions, which serve to alter epistemic probabilities defined over the space. In theoretical statistics, Blackwell (1953) compared this to a ‘game against nature,’ which reveals the formal inspiration from postwar work done on game theory, particularly in the formal analogies with ‘strategies’ and payoffs.¹³ The fictional character of nature playing games should be obvious. Blackwell himself was partial to the statistical school of Bayesian inductive inference, with the scientific researcher beginning with ‘prior probabilities’ inherited from the past.

Although early game theory had originally come equipped with no particular psychological or epistemic capacities whatsoever, the mathematical similarities between the Blackwell formalism and game theory induced game theorists to explore this approach to knowledge and information in their quest to generalize its ambitions. The other major application of the theory came under the rubric of ‘machine logic’ or ‘epistemic logic,’ primarily for use in computer programming (Fagin et al. 1995). When developed in computing, the Blackwell structure was used as a convenient model of induction, and not as some general approach to human, or even artificial intelligence.

The same could not be said for the economists. It was at this stage that neo-classical economists began to conflate any problem of ‘information’ with some version of generic ‘choice under uncertainty’ (Diamond and Rothschild 1978). As Thomas Schelling once said, “[t]here is a tendency in our planning [models] to confuse the unfamiliar with the improbable” (1962, vii). Any epistemic problem of any stripe, in pursuit of analytical tractability, was to be reduced to a set of probabilities induced over utility payoffs. As economists sought to incorporate this structure within their prior models of given preferences and von Neumann-Morgenstern expected utility, they discovered that the absence of cognitive content wreaked havoc with the result. Little toy models would set out to posit as an assumption that some agents purportedly ‘knew’ something that others did not, denoting this as an example of ‘asymmetric information,’ but the mathematics proved so arbitrary as to verge on emptiness. Once one opens up this scheme to doubt, it can rapidly rot the entire model enterprise. For example, should the relevant state space formalism include not only states of nature, but also the states of mind of rival players? (Recall from 3.2.1 this would imply subjective states of mind should be included in the changeless block universe.) What would validate the truth of things believed to be ‘known’ by respective payers? Nothing in existing utility theory permitted the formalization

¹³ See (Mirowski 2002, 380–385).

of the infinite regress of “I think that you think that I think that you think that ...” (unless all players were effectively identical, so knowledge is still ‘perfect,’ as in early Nash non-cooperative games). Other subtleties popped up in the machine learning literature: “Although you may have false beliefs, you cannot know something that is false” (Fagin et al. 1995, 32). A parallel false universe was banished by construction.

3.2.3 *The Harsanyi-Nash Program*

A number of alternate responses to the conceptual problems presented by the Blackwell setup could have been explored; some economists admitted this in the early stages of the inductive approach to information: “What equilibrium is in a particular market depends on what individuals in that market know. That the converse is true – that is, that what people know (or believe) is a function of the equilibria of the markets in which they participate – is an observation which surely must precede Marx” (Rothschild 1973, cited in Diamond and Rothschild 1978, 479). However, very few economists had the stomach to explore the ways knowledge and markets interactively shape one another in an alternative approach to microeconomics.¹⁴ Instead, by the 1970s, the preferred exit out of this conundrum for game theorists involved the uncritical preservation of standard formalism of von Neumann-Morgenstern expected utility, combined with uncritical adherence to the Nash solution concept in game theory. The way this is often put is that all imperfections of information about the world in the guise of uncertainty in games must be reduced to uncertainty over parametric payoff functions. The person who supplied the escape route was John Harsanyi.

Harsanyi developed his approach to information in conjunction with other game theorists such as Robert Aumann while employed by military agencies to apply game theory to problems of nuclear war and disarmament. As he saw it, the existence of ‘incomplete information’ would lead game theorists to have to incorporate an infinite hierarchy of beliefs into their models of agency. These would consist of a pyramid of beliefs for player *i* over parameter vector *X* in the state space consisting of:

First order beliefs: Player *i*’s probability distribution over vector *X*

Second order beliefs: Player *j*’s probability distribution over player *i*’s first order beliefs

Third order beliefs: Player *i*’s probability distribution over Player *j*’s second order beliefs

... and so on ad infinitum

¹⁴ “Models of what is usually called disequilibrium behavior do not make sense and cannot serve as reliable guides to further theorizing or policy unless they are consistent and coherent” (ibid., 461). Here we observe what is usually considered the ‘Lucas critique’ was well established in orthodox microeconomics at that time.

Clearly this infinity appears daunting (not to mention the cognitive demands it imposes upon strategic choice). Harsanyi proposed to recast the model to ‘define away’ these infinities (perhaps similar to the way renormalization dispensed with infinities in the case of quantum electrodynamics) through reducing all problems in epistemology to problems in the definition of the agent. He argued for this (1967/1968) by suggesting anything a player knows *privately* that could affect the payoffs in the game about other players could be summarized in a single vector, called indifferently the player’s *information vector* or *type*. This ‘simplification’ in effect collapses everything that a player knows privately at the beginning of a game that could potentially affect his beliefs about payoffs plus the ‘types’ of all his opponents into a single roster of typology of players. ‘Information’ consequently disappears as a discrete analytical entity, only to be replaced by an artificial zoo of player ‘types,’ or as textbooks often put it, any game of incomplete or missing information becomes a game of imperfect information – the only uncertainty is over which automaton nature has bequeathed you as opponent. The artificiality of this model strategy has been acknowledged repeatedly by its second-generation of advocates: “we use type structures solely as a modeling device. Types are not real world objects” (Dekel and Siniscalchi 2014, 4). Again we find knowledgeable theorists conceding the fictional character of their models.

Harsanyi’s model device renders formal the vernacular maxim: it is not what you know, it is *who* you know. To preserve the Nash solution concept, it is the ‘type’ of the actual players that is the primary unknown in the analysis; the roster of types, along with the structure of the game itself, and the rationality of the players is given *a priori* and presumed known to all players. Since the theorist supposedly places himself on the same epistemic plane as the players, the only way to learn anything further about the game is through Bayesian inductive inference. This is the current meaning of Bayes-Nash game theory. The standard game setup ends up inverted: realized payoffs tell you who your opponents really are. Of course, once mixed strategies are allowed over types, then all meaning of player identity dissolves into thin air. How you are supposed to know who you really are under such circumstances may appear a mystery.

3.2.4 Common Knowledge

It was left to Harsanyi and Robert Aumann to draw out the final implications for knowledge of this marriage of orthodox game theory and the Blackwell formalism. The irony is that a modeling approach which sought to place inductive approaches to information on a sound theoretical footing ended up banishing ignorance altogether, at least for the Bayes-Nash agent. The involution began with knowledge of the block universe, presumed known as true by all participants, incapable of expressing falsehood. It intensified with the Harsanyi procedure of presuming that players also come equipped with full identical

knowledge of the game they are playing, which includes the roster of all player types, and shared recourse to Bayesian inference. Given the extensive shared knowledge presumed on the part of all players, Harsanyi realized that it would be arbitrary to allow that different ‘players’ start the game equipped with different Bayesian prior probabilities, since that would constitute the only thing that would be ‘unknown’ to opponents in any deep sense. The entire thrust of the Harsanyi program was to banish from the standard model any parameters that are not already ‘common knowledge’ amongst all players; so he propounded the doctrine that player types also came equipped with identical Bayesian priors.

At this point, the Bayes-Nash approach to information disappeared up its own insistence upon rigorous consistency with its own prior postulates. Robert Aumann (1976) originally used his definition of common knowledge to prove a notorious result that says that in a certain sense, agents cannot “agree to disagree” about their beliefs, formalized as probability distributions, if they start with common prior beliefs. Since agents are often portrayed as holding different opinions plus standard capacities for statistical inference, one might attribute such differences to the agents having different private information. Aumann’s incongruous result is that even if agents condition their beliefs on private information in a Blackwell setup, mere common knowledge of their conditioned beliefs and a common prior probability distribution implies that their beliefs cannot be different after all.

This seeming travesty of the economics of information has given rise to a cottage industry exploring the meaning of such locutions. One important student of Harsanyi has admitted that “[t]here is something fundamentally counterintuitive about the art of modeling [information] with Bayesian games” (Myerson 2004, 1823), and reported that Harsanyi himself sometimes seemed uncomfortable with the implications of common knowledge in his approach. Other economists, less concerned with epistemic niceties, regarded the Harsanyi setup as license to apply the model to all manner of ‘information asymmetries’ in the phenomenal world along these lines. Their version of ‘common knowledge’ was interpreted to mean everyone who was rational had to agree with their model. One can appreciate the rather self-congratulatory effect this had on the economics profession.

3.3 Information as Computation

This third version of information owes the greatest debt to the postwar development of the computer and the theory of computation, but curiously enough, has proven over time to be the least palatable to many neoclassical economists, of the three manifestations described here. It predominantly travels under the banner of “computationalism,” which tends to identify mental states with the

computational states found in (either abstract or tangible) computers. Computationalism today is comprised of many competing visions, ranging from formal symbol manipulation to ‘connectionism’ to ‘machine cognition;’ but economists have rarely been very sensitive to these controversies within artificial intelligence and cognitive science. For instance, few economists realize that the ‘connectionists’ and the proponents of genetic algorithms often praise Friedrich Hayek as an early progenitor, whereas the first generation of classical artificial intelligence theorists and complexity mavens instead celebrated Herbert Simon as their inspiration.

The computationalist turn has assumed two different formats in the history of orthodox economic theory: the first attempts to subject the standard rational choice model to be subsumed under a computationalist model of mind; while the second tends to fall under the rubric of ‘market design.’ The former cadre are a rather diverse lot, ranging from those (such as Alain Lewis) seeking to model individual rational choice as an explicit computational proposition to what has been called ‘agent-based computational models,’¹⁵ but might be better thought of as simulations of agent swarms, after the model of cellular automata. Actual experience with computers has provided all manner of heuristic suggestions as how to meld cognitive science with neoclassical economics, perhaps taken to an extreme at certain locations. Indeed, as one Clark Medal recipient has admitted, “if you try and do psychology at MIT, you study computers, not humans” (Rabin, cited in Colander et al. 2004, 141). The latter market design wing combines certain sectors of experimental economics with what might be best described as ‘engineers of automated markets,’ where both claim to have superior insight into the informational properties of markets with large numbers of participants. This latter group has ambitions to be engineers of the human soul, arguing that their purpose-built machines can force people to tell the truth even when their every intention is to be mendacious, or provide them with information that they would otherwise find inaccessible through any conventional recourse to research channels.

As with the previous cases, economists themselves did not pioneer computational theory (if you leave Herbert Simon out of the picture). It was mathematicians, computer scientists and linguists who forged the theory of symbol manipulation by automata of various computational capacities, with the Turing Machine occupying the highest rung on the computational hierarchy. In a ranking of the power of various abstract ‘machines,’ the class of Turing Machines are generally conceded to be the most powerful. The importance of the computational hierarchy is that it facilitates the proof of impossibility theorems concerning what can and cannot be computed upon machines falling within a particular computational class. Computational approaches have had the prophylactic virtue of ruling out all sorts of physically and mathematically impossible

¹⁵ See, for instance, (Tesfatsion 2016).

procedures from falling within the purview of an algorithmic conception of rationality. Treatment of infinities assumes much heightened significance; implementable algorithms are more highly regarded than in-principle proofs.

Early on, the computational metaphor of mind proved a mixed blessing for economists. If one were to seriously entertain the notion of a marketplace of ideas, the problem became where in the economy one would situate the computer. Was each agent a Turing Machine, or perhaps an automaton of less exalted capacity? The von Neumann architecture built into every laptop did seem a bit removed from human cognition, and then there were the interminable disputes of the 1960s-90s over what it was that humans could supposedly do that computers could not. Most would admit computers could readily store and manipulate information, but could a computer be seriously thought to be knowledgeable? The development of the Internet did seem to present templates for the formalization of the communication of information, however. Or perhaps edging closer to Hayek's vision, should the marketplace itself be treated as one vast Turing Machine, with agents simply plug-compatible peripherals of rather diminished personal cognitive capacities? The history of this research program reveals that certain key aspects of the neoclassical model were shown to be Turing non-computable;¹⁶ yet this did not foreclose machine models of mind for the economists. The temptation then arose to shift the location of the computer to another ontological level – namely, the level of the market itself – in order to evade the unsavory implications. Market designers took their cue, and ensconced their models inside machines that would act as markets.

The purpose of this whirlwind summary is not to do justice to the various reincarnations of 'the model' of economic man once it encountered the computer, but rather to suggest that computationalism ratcheted up the fictional quotient of economic models of information to ever newer heights. One can understand the history of this tradition as a fevered quest to find a substantial role for this new interloper (the Turing Machine, classifier systems, networks inspired by the internet) within the standard narrative of neoclassical economics. The task seemed to first imagine a computer, and then recast all the other players in a simulacrum of the market as one big information processor. Slowly, imperceptibly, people stopped talking about markets as ideal allocation devices for things. Instead, the market itself (or perhaps the boutique markets cobbled together by market designers) started to look like an all-purpose arbiter of Truth. This was a fiction far more powerful and far more consequential than any single isolated mathematical model.

¹⁶ See, in particular, (Mirowski 2002, 422–436).

4. The Lessons of the Fictionalist Account of Information in Economics

Arnon Levy has summarized the lessons he draws from his examination of the history of ‘information’ in biology:

My fictionalist proposal is motivated by the idea that even if we treat information non-literally we may still take it seriously and assign it a real role in biological understanding. But we shouldn’t take it too seriously ... The fiction-based view explains what is wrong with a metaphysics of information, and obviates the inference from the theoretical role of information descriptions to the existence of informational ‘things’ ... information enthusiasts are illicitly taking the cognitive success of information as the basis of ontological commitment” (Levy 2011, 652–3).

I would like to make a similar argument. There is no doubt that the modern economics profession believes with justification that their models are driven by a conviction that information is the central economic question which needs to be contemplated in the 21st century, and that it informs many of the mathematical innovations which undergird their cherished research programs. But, echoing Levy, the error arises when they take their limited success as a metaphysical warrant for the very existence of (one of) their (three) adopted versions of information: human capital as impersonal ‘thing’ which gives our lives meaning, the Bayes-Nash system as an airtight guarantee that rational agents must accept the economists’ own Blackwell model as the truth and never be ignorant, and most crucially, that the market which they theorize as a computer is the final arbiter of Truth in all things. Each of these propositions are unselfconsciously converted into ontological claims, and as such, are rendered completely impervious to empirical disproof, or even the observation that they are logically inconsistent with one another, since they are based upon pairwise incompatible versions of ‘information.’ The overt fictional stance towards models renders economists impervious to any such critiques.

I think the bigger issue raised by these ‘fictions’ is that the spread of models of information in economics has served to elide or suppress key questions of who is doing the thinking in the economy, and how this undermines notions of agency, ‘welfare’ and most of the other concepts economists hold dear. As it stands, stylized resort to existing models of ‘information’ suppress most of the considerations that come as second nature to someone trained in philosophy. Indeed, taken at face value, they jointly constitute an imperious rejection of philosophy itself as irrelevant to questions of knowledge in the economy.

References

- Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for invention." In *The Rate and Direction of Inventive Activity*, edited by Richard Nelson. Princeton: Princeton University Press.
- Aumann, Robert. 1976. "Agreeing to Disagree." *Annals of Statistics* 4: 1236–1239.
- Becker, Gary. 1964. *Human Capital*. New York: NBER.
- Buchanan, Mark. 2014. "Just so stories of modern economics." Accessed March 6, 2015. <https://medium.com/the-physics-of-finance/8b45018f432a>.
- Colander, David; Ric Holt and J.B. Rosser, eds. 2004. *The Changing Face of Economics*. Ann Arbor: Univ. of Michigan Press.
- Davies, Paul and Niels Gregersen, eds. 2010. *Information and the Nature of Reality*. New York: Cambridge University Press.
- Dekel, Eddie and Marciano Siniscalchi. 2014. *Epistemic Game theory*. Accessed March 6, 2015. <https://sites.google.com/site/eddiedekelsite/papers/game-theory>.
- deMarchi, Neil, ed. 1993. *Non-Natural Social Science: reflecting on the enterprise of More Heat than Light*. Durham: Duke University Press.
- Diamond, Peter and Michael Rothschild, eds. 1978. *Uncertainty in Economics*. New York: Academic Press.
- Dorfman, Robert. 1960. "Operations Research." *American Economic Review* (50): 575–623.
- Fagin, Ronald, Joseph Halpern and Yoram Moses. 1995. *Reasoning about Knowledge*. Cambridge: MIT Press.
- Frigg, Roman. 2010. "Fiction and Scientific Representation." in *Beyond Mimesis and Convention*, edited by Roman Frigg and Matthew Hunter. New York: Springer.
- Frigg, Roman and Stephan Hartmann. 2012. "Models in Science." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Accessed March 7, 2016. <http://plato.stanford.edu/archives/fall2012/entries/models-science/>.
- Geoghegan, Bernard. 2008. "The Historiographic Conception of Information: a critical survey." *History of Computing* (**): 66–81.
- Gigerenzer, Gerd and David Murray. 1987. *Cognition as Intuitive Statistics*. Hillsdale, NJ: Erlbaum.
- Giocoli, Nicola. 2003. *Modeling Rational Agents*. Cheltenham: Elgar.
- Godfrey-Smith, Peter. 2009. "Models and Fictions in Science." *Philosophical Studies* 143: 101–116.
- Godfrey-Smith, Peter and Kim Sterelny. 2008. "Biological Information." Accessed December 2014. <http://plato.stanford.edu/entries/information-biological/>.
- Griffiths, Paul. 2001. "Genetic Information: a metaphor in search of a theory." *Philosophy of Science* 68: 394–412.

- Institute for New Economic Thinking. 2011. "Bretton Woods, Past and Present: 3. Models in Economics." Accessed April 5 2016. https://www.youtube.com/watch?v=7eYnPxJpm_E#t=361.
- Harsanyi, John. 1967. "Games with Incomplete Information played by Bayesian players." *Management Science* 14: 159–182; 320–334; 486–502.
- Hellman, D., ed. 1988. *Analogical Reasoning*. Dordrecht: Kluwer.
- Hesse, Mary. 1963. *Models and Analogies in Science*. London: Sheed & Ward.
- Kay, Lily. 2000. *Who Wrote the Book of Life?* Stanford: Stanford University Press.
- Kline, Ronald. 2004. "What is Information Theory a Theory of?" in *The History and Heritage of Scientific and Technical Information Systems: Proceedings of the 2002 Conference, Chemical Heritage Foundation*, edited by W. Boyd Rayward and Mary Ellen Bowden, 15–28. Medford, NJ: Information Today.
- Kreps, David. 1990. *A Course in Microeconomic Theory*. Princeton: Princeton University Press.
- Levine, David and Steven Lippman, eds. 1995. *The Economics of Information*. Cheltenham: Edward Elgar.
- Levy, Arnon. 2011. "Information in Biology: a fictionalist account." *Nous* 45: 640–657.
- 2012. "Models, Fictions and Realism: two packages." *Philosophy of Science* 79: 738–748.
- Levy, Arnon and Adrian Currie. 2014. "Model Organisms are not (Theoretical) Models." *British Journal for the Philosophy of Science* 0: 1–22.
- Lucas, Robert. 1988. "What Economists Do." Unpublished manuscript of talk delivered to graduate students at the University of Chicago.
- Machlup, Fritz. 1962. *The Production and Distribution of Knowledge in the US*. Princeton: Princeton University Press.
- Mirowski, Philip. 1989. *More Heat than Light: economics as social physics*. New York: Cambridge University Press.
- 2002. *Machine Dreams*. New York: Cambridge University Press.
- 2007. "Markets Come to Bits." *Journal of Economic Behavior and Organization* 63: 209–42.
- Mirowski, Philip and Edward Nik-Khah. Forthcoming. *The Knowledge we have Lost in Information*. New York: Oxford University Press.
- Morgan, Mary. 2012. *The World in the Model*. Cambridge: Cambridge University Press.
- Morgan, Mary and Tarja Knuttila. 2012. "Models and Modeling in Economics." In *Philosophy of Economics*, edited by Uskali Maki. Amsterdam: North Holland.
- Myerson, Roger. 2004. "Comments on 'Games with Incomplete information Played by Bayesian Players.'" *Management Science* 50: 1818–1824.
- Samuelson, Paul. 1954. "The Pure Theory of Public Expenditure." *Review of Economics and Statistics* 36: 387–389.

- Schelling, Thomas C. 1962. Foreword to *Pearl Harbor: Warning and Decision*. Stanford, CA: Stanford University Press.
- Shannon, Claude and Warren Weaver. 1949. *The Mathematical Theory of Communication*. Urbana: University of Illinois Press.
- Smith, John Maynard. 2010. "The Concept of Information in Biology." In *Information and the Nature of Reality*, edited by Paul Davies and Niels Gregersen. New York: Cambridge University Press.
- Sugden, Robert. 2000. "Credible Worlds: the status of theoretical models in economics." *Journal of Economic Methodology* 7: 1–31.
- Sugden, Robert. 2013. "How Fictional accounts can explain." *Journal of Economic Methodology* 20: 237–243.
- Tesfatsion, Leigh. 2016. "Agent-Based Computational Economics." Accessed March 22. <http://www2.econ.iastate.edu/tesfatsi/ace.htm>
- Walliser, Bernard. 2011. *Comment raisonnent les économistes: les fonctions des modèles*. Paris: Odile Jacob.
- Weisberg, Michael. 2013. *Simulation and Similarity*. New York: Oxford University Press.