EGRIE Keynote Address

Risk and the Endogenous Economist: Some Comparisons of the Treatment of Risk in Physics and Economics*

Neil Doherty

Wharton University of Pennsylvania, 255 S. 38th St. Philadelphia, PA 19104.6359, USA.

This is a talk, rather than a research or survey paper. Very little of what I say will be original, but I wish to stimulate discussion on a set of issues that arise from the nature of risk and that I consider problematic to our profession. The paper is not exhaustive of references and many of my arguments have been treated elsewhere. However, I suspect few will have approached the issues from the same starting point and assembled them in the same way.

The Geneva Risk and Insurance Review (2013) **38**, 1–22. doi:10.1057/grir.2012.6

Keywords: risk; economics; physics; endogenous economist

Introduction

Let me start with definitions. I would venture that, at root, all the risk we contemplate in economics is epistemic. But it depends on definition.¹ Often the term “epistemic risk” is used to refer to model risk, as opposed to “aleatory risk”, which is risk that arises from incomplete data. However, if we define epistemic risk more broadly as that which arises from incomplete knowledge, then one could argue that the risk inherent in the inputs to our models (data) also is epistemic. Consider the uncertain commodity prices we put into our trade models, the earthquake or hurricane probability distributions we put into calculation of insurance premiums, the distributions of stock or bond returns we use in asset market equilibrium models, or indeed practically any other uncertain input we use in our economic models. Each of these model inputs is estimated in turn from prior models that themselves depend on their own data inputs. There is backward regression. And if only these earlier models were better specified, or

*This paper is based on the 24th Geneva Risk Economics Lecture.

¹ For example, the following is extracted from a risk assessment primer disseminated by the U.S. Department of Transportation, Federal Highway Administration: “Aleatory (data) risks refer to uncertainty associated with the data used in risk calculations. An example of an aleatory risk is the uncertainty surrounding the cost of a material (i.e., steel or asphalt). Epistemic (model) risks refer to risks that arise from the inability to accurately calculate a value”.

had better data, we would reduce (in principle eliminate all) uncertainty. Of course, we should not hold our breath. As a practical matter, we are not going to eliminate all risk.

Consistent with this broader definition of epistemic risk, we can define aleatory risk alternatively as risk that is inherent in a process. This definition implies that aleatory risk cannot be reduced or removed by more knowledge; risk is a fundamental property of the process itself. Thus, epistemic risk now refers to our lack of knowledge about the process. And this definition would seem to be more consistent with the verbal root, epistemology.

So here is where I am. Using this broader classification, all risk we address in economics is, at heart, epistemic. Not so in physics. When Einstein objected that “God doesn’t play dice with the universe”, he was objecting to the view in the (then) new quantum mechanics where risk was held to be a property of nature. Einstein was, of course, wrong.

Risk and uncertainty are fundamental properties of the smallest things in nature, sub-atomic particles. Even if we had perfect information, we still could not state with certainty the properties of these particles, such as their position, or spin, or velocity. This notion is, of course, the “uncertainty principle” and is encapsulated in Erwin Schrodinger’s famous equation (and even more famously in the metaphor of Schrodinger’s cat). It is not that we know nothing about properties of particles, such as the location, velocity, spin, etc. Indeed, we can specify a probability distribution. Furthermore, pairs of properties are linked in an interesting stochastic manner. For example, the more we know about the position, the less we know about the velocity. Thus, the strange message of quantum mechanics is that nature is inherently uncertain. And in case we do not take this to heart, the empirical evidence in support of this model is unbelievably strong and it is often described as the most successful theory in all of science.²

Although I am arguing that the nature of risk is different in economics and quantum science, I want to argue that there are some important lessons we economists can learn from physicists and the manner in which they have dealt

² Early evidence came in the form of shooting electrons through two slits in a shield and plotting their impact on a screen. The resulting pattern of impacts on the screen shows an interference pattern that suggests a wave function (a wave of chance) for the electrons’ paths. The only coherent explanation for such an interference pattern is that the path of the electron is a wave function, and the mainstream explanation is that this wave is a probability wave (distribution). Spectacular confirmation more recently came from the background microwave mapping of the observable universe. If indeed the universe started with a big bang and unimaginably rapid inflation of space, then the quantum fluctuations on the smallest of scales in the first instance of the universe would be stretched out to determine the large-scale structure of the entire universe. The correspondence between the very subtle observed deviations from an almost uniform structure of the observable universe and those predicted by quantum theory is truly amazing.

with risk. Indeed, the very fact that risk is different confronts physicists with issues in scientific methodology that have been largely unattended to in our profession. And on the more positive side, the alternative way in which physicists have characterised risk presents ideas for reframing some basic issues in economics that might (we will have to see) be fruitful.

The basic issue that I want to address is the relationship between the scientist and the subject of any experiment. Any observation is a joint statement about the thing being studied and the properties of the observer. In a sense, the observer must always contaminate any experiment. And the secondary issue I want to address is whether reality is not just a realisation from some probability distribution, but rather the distribution itself. In transferring this idea to economics, I do not want to get bogged down in philosophical mumbo jumbo—so I will not take this second question too seriously. But I do want to ask whether we can get any interesting insights from contemplating the prospect that some economic realities may be usefully considered as distributions (rather than as realisations from those distributions).

The observer and the experiment—Endogenising the observer

I have mentioned that quantum mechanics is possibly the best tested theory ever. However, at first sight, this may seem surprising. Indeed, it seems on its face to be an untestable theory.

Here is the rub. You can refer to footnote 2 to see that the first tests took the form of setting up experiments to see where electrons landed on a sensitive collecting screen when they were fired from a “gun”. We should see that they all end up in different positions on the screen according to the probability distribution that governed their positions. The experimenters looked at the histogram of final resting places on the screen for all the electrons and compared this to the predicted probability distribution. But wait a minute—the theory does not say that the electron will end up at some actual position selected randomly from the distribution. The theory says the position CANNOT be known, the position itself is stochastic. The fact that we seemed to observe the final positions seems to CONTRADICT the theory. And since we cannot (as far as I know) observe distributions, it would appear that the theory is inherently untestable.

There are two ways in which this dilemma has been approached by empirical physicists: the “Copenhagen View” and the “Many Worlds” theory.

1. *Copenhagen view*: This view is named after the Copenhagen school and is associated notably with Niels Bohr. There is a “don’t ask, don’t tell” quality to this view. Somehow, the act of observation itself forces the uncertain particle to resolve its uncertainty and to congeal to a fixed and observable position (the usual lingo is “collapse of the wave function”).

Location uncertainty resolves on observation, but is restored when one looks away. In other words, the theory (properties of particles are uncertain) is somehow suspended on observation. But as soon as the observer looks away, the particle resumes its normal stochastic location.³

2. “*Many worlds*”⁴ and similar views: Quantum theory states that we have a density function defined over location (velocity, spin, etc.) space. In the “many worlds” theory, every possible location, velocity, spin, etc., actually occurs. Thus, if the particle *can* appear in many different locations, each with a specified probability, it *does* appear in *all* of those locations. Thus, it is no puzzle that we happen to see that electron at one location—because simultaneously there are many more realities in which an observer saw it at another location. In other words, reality is continuously multiplying itself.⁵ And before you dismiss this as science fiction, this is probably becoming the mainstream view among physicists.

In both views, interaction between the observer and the experiment is important. In the *Copenhagen view*, somehow (unspecified) the act of observation forces the particle to give up its property of uncertainty (at least while it is being observed). It is also acknowledged by this school that observation itself is a physical act that interferes with the observed particle. For example, photons from the equipment used to make the observation can interact with the particle.

In the *many worlds* view, each possible location of the particle is a separate reality, which implies that if one observer sees the particle *here*, then there is another reality in which the same observer sees it *there*. Thus, it would seem that, conditional on the particle having some revealed property, the observer will actually observe that property for certain; that is, the conditional “me” that exists conditional on the particle being at position “x” will for certain

³ The famous metaphor is Schrodinger’s cat. Imagine a cat in a sealed container that is subject to a “wave function” spanning two states, “the cat is dead” and “the cat is alive”. The cat’s condition is as described by the function, alive *and* dead. However, if the container is opened, one will observe one of the two conditions, either “alive” or “dead”. The wave function collapses on observation and a realization is observed.

Oh, curséd cat, how can one be
alive and dead, just like a zombie?
alive and dead, it makes one weep,
alive and dead—*unless I peep*.

⁴ The “many worlds” idea was derived by the American physicist Hugh Everett in his 1957 Princeton PhD thesis.

⁵ This leaves a puzzle as to what happened to probability. One view could be that, if all possible realities are realised, then somehow the force of those different realities must be weighted by their prior probabilities. So if one location has a higher probability than another, then the realisation at the former location is somehow “more real” than the latter.

observe the article at “ x ”. This implies that there is another conditional “me” that sees the particle at position “ y ”. This would seem to resolve the tension between a theory that says all properties are stochastic and experiments in which the observed realisations of the particle follow a histogram echoing the density function. But it is more complex, since all the “stuff” making up the observer is also subject to the same quantum uncertainty and each basic particle making up the observer also can have many realisations. Again, the interaction between the observer and the experiment is central.

I want to take away two thoughts from this discussion. The first is to reflect on the way in which physicists have grappled with the problem that they themselves are part of the natural order that they seek to model and test. The second, how are we to deal with a theory that says that risk is an inherent property of the phenomena we seek to observe. Let me now ask similar questions of economics.

The economist as a strategic agent

In quantum mechanics there has been a very serious effort to come to terms with the relationship between the observer and the experiment. The strategy seems to have two parts. The first is to address the contaminating effect of the presence of an observer on the experiment under way. And the second is to deal with this contamination by including the presence of the observer in the physical model of reality. I would suggest that, while the relationship between the economist and economics is equally important, it has largely been ignored. The economist is a strategic player in the creation of models and in their empirical verification. We have paid little attention to either role.

The economist as a strategic empiricist⁶

Physics appears to have a clearer division of labour between theory and empiricism than does economics. Perhaps this has much to do with the inherent complexity of the subject matter, which creates a need for specialisation. If so, it certainly has the fortunate side effect that those conducting empirical work are not really swayed by the incentives associated with ownership of ideas. Of course, there is some specialisation in economics, but less so. Many economists either test their own models or co-author with others to do so. This would not be an issue if there were strict protocols on empirical work, but this is not the case.

⁶ Most of the issues raised here were discussed many years ago by Ed Leamer in his famous article “Let’s take the ‘con’ out of econometrics”. See also his recent talk: http://www.econtalk.org/archives/2010/05/leamer_on_the_s.html.

Indeed, testing of models calls for many judgements to be made, which are rarely open to inspection. For example:

- choosing which of multiple predictions of a model are to be tested;
- choosing which proxies to use for unobservable variables;
- choosing a data set;⁷
- deciding how to censor the data;
- deciding which statistical or econometric procedure to be used given the properties of the data; and
- choosing a stopping rule when sequentially addressing the previous issues.⁸

Rarely are the answers to these questions clear-cut. And this leaves enormous flexibility for any economist who has a vested interest in the outcome of the test to search for the “desired” result. Of course, we are all quite familiar with the ways in which this engineering occurs. We start with an “honest” attempt of finding the best answers to the above questions and run our first test accordingly. We then look at the results, giving much more scrutiny to the “bad” results (i.e. those that are not confirming of our priors). Then we start the “real” empirical work.

- If our model is tightly specified, the predictions will usually be quite clear. However, with a loose “model” (i.e. a collection of more-or-less related statements, each graced with the title “hypothesis”) we can start selectively reworking those hypotheses that are not helped by the data, while smugly accepting hypotheses that are confirmed by our first results.
- We can ask if our data-censoring procedure is too lax, or too severe.
- We can add or delete control variables.
- We can try other statistical procedures to accommodate the vagaries of the data and the model.
- And we can keep repeating these experiments until we find a set of results that provide the strongest confirmation of our model.

Thus, the implied stopping rule is when we feel we have the most positive results. Sounds familiar?

After the fact, we will have convinced ourselves, and will be armed to persuade others, that all this trial and error was a learning process, by which we

⁷ In his Nobel Lecture many years ago, Hayek addressed agency issues in the choices of data sets.

⁸ Of course, there is a well-known theory of stopping rules and there are some well-known solutions for well-defined problems such as the “secretary problem”. For a more recent example, which will appeal to this audience, see the recent paper by Carvagnaro *et al.* (2012). They examine a search process to find the optimal decision stimuli to test (and discriminate between) various utility-based decision models.

have truly converged on the most impartial way of testing our model. However, if we are really honest with ourselves, it is all pretty bogus.

There is nothing original in my thoughts here. Edward Leamer pointed out all these issues many years ago, as did his namesake Edward Glaeser in a fairly recent NBER paper. And more recently, my colleagues Joseph Simmons and Uri Simonsohn, together with Leif Nelson, have demonstrated how easy it is to generate false-positive results by this progressive tinkering. Yet despite the obvious issues, we blithely continue along our way. We know it is wrong and we jokingly talk of “torturing the data”, but there is no will to do anything different.

Of course, this would not be an issue if all results were cross-checked. And, in a sense, they potentially are. In the refereeing process, there is an opportunity for the reviewer and the journal editor to comment on the empirical procedure. But how often will they really check the results. And, even if they do, the reviewer is always responding to the selected empirical results that have survived the author’s search. How many of us, as reviewers, have seen a complete file of ALL empirical runs conducted by the author? Never, I suspect, will we reach the level of cross-checking that was seen when many thousands of scientists were scouring data generated at the Large Hadron Collider in CERN in the now apparently successful search for the Higgs boson. Or I doubt whether any paper in economics will receive the scrutiny plied on Andrew Wiles’ first incorrect, and second successful, solution to Fermat’s Last Theorem.

It is not difficult to derive protocols to address this moral hazard. However, we must be mindful of the scarcity of reviewing resources and the enforceability of any protocol. Here are some thoughts.⁹

- (a) With unlimited resources, we could set up an empirical adversary for each paper submitted for publication, but this is not very practical. Or we could require that each time we test a theory, the authors are required to pre-commit to a single protocol for handling data and choice of statistical method—then conduct their empirical run (A SINGLE RUN) accordingly. This would probably be unenforceable. So let us get real.
- (b) A simple rule would be to establish a differential hurdle between pure theory papers and those papers in which the author(s) both present a theory and also offer an empirical test. We should discriminate against the latter. The discrimination could take the form of either a higher burden of proof (e.g. differential confidence levels), or perhaps asking those who test

⁹ See Simmons *et al.* (2011) for similar suggestions.

their own theory to pay a higher fee and provide all their data, so that the editor can employ someone to test the results.¹⁰

- (c) We could require that each submitted paper contain the entire file, together with a summary, of ALL empirical runs, showing how the strength of the results varied as the authors experimented in their selection of data and in their choices of empirical method. The referees would then see how the empirical results have trended with the search process.
- (d) Given the scarcity of resources, it may be that only premier journals adopt such a protocol. This would establish a two-tier process, by which we would see that results published in premium journals would be more credible and results in lower-tier journals would be treated as empirical conjecture.

Of course, you may object that all this is too much. Currently, we do have a reviewing process and if some “statistically tortured” results may get through the net, then so be it. Any important papers will eventually be challenged and tested by others. Indeed, careers can be made by toppling paradigms that are based on false empirics. We can all come up with our own favourite examples of economic wisdoms that have overturned with more careful econometric analysis. However the short-run costs of misinformation can be high.

Finally, as a segue to the following section on the economist as a strategic theorist, I will float a question. Do we have a principal-agent model that allows a theorist-cum-empiricist to secure a benefit from having his model validated, and which can maximise utility by choice of empirical protocol?

The economist as a strategic theorist

In quantum theory, as argued, considerable intellectual effort has been devoted to understanding and modelling the interaction between the observer and the experiment. This cannot be said for economics. It is true that we have a voluminous literature on principal-agency issues. Indeed, few relationships have escaped this principal-agent scrutiny. And of course, we can consider the economist as an agent in such a relationship, with the principal being a bank, investor, government or university that employs him or her. But when was the last time we saw an economic model of some external function or institution that contained an explanation of why the economist constructed that particular model. For example, suppose we wished to explain how, in a competitive market, certain contractual features might be adopted in insurance contracts to address adverse selection. Does that model also explain why the economist

¹⁰ Although one can question whether an impartial testing is itself optimal, it may be better to have advocates for rival theories who are more clearly rewarded and motivated (see Dewatripont and Tirole, 1999).

(who is indeed an economic agent who buys insurance) might (adversely) select that adverse selection model? Or, if an economist is building a model of bank regulation, can we ignore the fact that the economist's career prospects will be affected by the regulatory solution? Indeed, I would offer that I cannot recall ever seeing an economic paper that modelled some institution or behaviour and, at the same time, subjected the economist who is writing the model to the same behavioural assumptions and decision sequences as all other actors in the model. Likewise, in general equilibrium models of economic systems, is the economist that produces the model an integral part of the system that is being modelled? Do we ever really endogenise the economist? We are somewhat like the cobbler who fails to discern that her own children have no shoes.

Why is it important to endogenise the economist? One can argue that our models are incomplete and inelegant if they fail to do so. For example, take a behavioural model in which a highly rational economist chooses to model the behaviour of other (non-economist) agents as boundedly rational. One should be obviously sceptical of such a model that suggests that economists are smarter than everyone else. On the other hand, take a rational expectations model. Is it reasonable for the economist who publishes a model to assume that all agents in her model implicitly favour the assumptions of her model over any competing economic models that are out there? In which case, is the economist lacking in rational foresight? I will return to these points in a while. But perhaps a more compelling reason why we should endogenise the economist is simply that the economist might have a vested interest in any policy choices that may flow from her work.

A literature on the issue on "the endogenous economist" does exist, though, as far as I can find, it is very thin and largely neglected. A paper of this title was published by Bernard Gauci and Thomas Baumgartner in 1992.¹¹ And there are other works that sidestep into this issue; for example, a recent book by MacKenzie *et al.*¹² that examines the notion that an economic theory possibly can be self-realising, in the sense that it changes reality in its own image. And of course economists sometimes do take stock of their own profession. But that is not the same as including themselves as agents in their own models. Indeed, given the ubiquity of principal-agent models in our field it is amazing that we have applied them to practically everyone but ourselves.

The Gauci–Baumgartner (G&B) paper posits two aspects of endogeneity. First, in construction of economic models, is the economist viewed as an object of study under the same paradigm as other agents? Second, do economists

¹¹ Despite the intriguing title of this paper "The Endogenous Economist", the paper seems to have been mostly ignored having just a couple of citations in some 20 years.

¹² MacKenzie *et al.* (2009).

examine whether they influence the objects of their study? The G&B paper focuses on the second issue and specifically asks whether other actors (non-economists) are influenced by the statements made by economists in their models. And it is this question that perhaps has its closest parallel to the debate in the physical sciences. The very act of observing natural phenomena is a physical interaction between the observer and the subject of the experiment, and the ongoing concern is to allow for this interaction in the interpretation of experimental results.

Now there are, as the G&B paper points out, weak examples of endogenisation. For example, consider models in the rational expectations class, such as a financial pricing model. In one sense, these models do endogenise the economist because *all* agents (including economists) must behave in the same rational fashion. Thus, one can claim that economists are endogenised. But this is trivial and usually there are no specific statements coming from such models about the behaviour of the economist *per se* and of why the economist chooses to construct such a model in the first place. But on G&B's second aspect of endogeneity (the influence of the economist on other actors), the rational expectations models only partially endogenise the economist. This class of models is heavy handed, for they posit that all actors accept the unique truth of the model and behave accordingly, that is rationally.¹³ However, economic discourse is not limited to rational expectations models; there are other models and there is an ongoing empirical competition between them. And we, as economists, rarely ask how agents are affected by this lack of a clear message from the economic community. It would be silly to think that all actors simply ignore all economic debate and choose to act selectively according to the postulates of the model being tested at any particular time.

I am not sure whether the failure to endogenise the economist in our models reflects an economy of effort in the face of a difficult problem, or simple arrogance. It cannot be easy to add another principal-agent relationship into every model we build. But it also creates the self-serving illusion that we economists are set apart from the incentives that confound the frail relationships between all other agents. And this can lead to absurdities. For example, it is a curious spectacle for a perfectly rational economist to produce a beautiful and internally consistent model of the boundedly rational behaviour of other agents. Yet we see this all the time. Are economists the noble and rational gods in a society of self-serving and irrational mortals?

Let me now make some more formalised illustrations of the issue. How might we endogenise the economist? I stress that I am not trying to construct an economics model, so do not get too hung up on the specific notation—rather

¹³ The book by MacKenzie *et al.* (2009) examines this issue.

I am simply trying to indicate what issues might be addressed in such models. I merely ask some questions.

Model 1: A minimal structure with rational players

Consider the most basic way to capture the idea of an endogenous economist. Suppose that we wish to construct a model of the production–trading relationships of two parties, X and Y . However, we build a model with three parties, X , Y and an economist E . All parties have a utility function of the form:

$$U_i = U(C_i - c_i),$$

where $i = X, Y, E$; C = consumption; c_i = cost of an action A_i .

The production function for social wealth is

$$P = P(A_E, A_X, A_Y).$$

Social production is apportioned over the three parties who consume all income¹⁴

$$C_X = (1 - \alpha)P - T_X^E - c_X; \quad C_Y = \alpha P - T_Y^E - c_Y; \quad C_E = T_X^E(P; \alpha) + T_Y^E(P; \alpha) - c_E,$$

where T_i^E = transfer from “ i ” to the economist. The economist is the strategic player

$$A_E = f(T_X^E, T_Y^E; c_E).$$

Now, in the normal paper in an economics journal, the economist is concerned about the relationship between X and Y —how they jointly or competitively engage in productive activities, how the product is allocated, and so forth. But I have included an intermediary who I have called an economist. So far, this minimal structure could apply to any type of intermediary (and there are countless intermediary models out there). What makes it a story about an economist is the contextual structure on which this is mounted. That is, the structure put around actions, A , and the manner in which they impact the production P and transfers α and T . Thus, we can think of A_E as the promotion

¹⁴ Note that I have treated the costs of the three parties as a deadweight loss in this three-person economy. They could of course be income to external parties.

of an economics model, either by construction of a new model or by bolstering the credibility of a pre-existing model through empirical tests.

- *Example 1a: The economist's contribution is to enhance efficiency.* If the economics model led to a better understanding between X and Y of the social production function, and of the actions that X and Y could take separately (or jointly) to increase P , then we would have $\partial P / \partial A_i = f(A_E)$, $i = X, Y$. If we stop here, it seems plausible that the economist would always wish to build a model that maximises net surplus. And the economist would be compensated for so doing with a fee not exceeding, and scaled to, the increase in social production. Thus, the economist would be motivated to offer the "best" advice. However, I am having some discomfort with this structure. If indeed X and Y do not really understand the social production function, it supposes some limit on their economic literacy. So perhaps we should start delving deeper into economic literacy and the economist's contribution to this. We do so below in model 2.
- *Example 1b: The economist's contribution affects allocation (and possibly efficiency).* Suppose that one party can "capture" the economist to secure trading advantage. Thus $\alpha = f(A_E)$. Also, suppose that this wealth transfer entails an efficiency loss (reduction in P). Now various possibilities arise depending on the enforceability of contracts. If enforceable contracts can be written between X and Y , it would seem safe to conjecture that they would commit not to employ the economist. If such enforceable contracts could not be written, then perhaps one party will employ the economist. If agent X were able to write an exclusive contract, then he/she would be willing to pay up to a maximum of $(1 - \alpha^{wx})P^{wx} - (1 - \alpha^{wy})P^{wy}$; and Y would be willing to pay a maximum of $\alpha^{wy}P^{wy} - \alpha^{wx}P^{wx}$ (where wx denotes that the advice is proffered to X and wy denotes advice to given to Y). Thus, it seems to be a straight contest as to who will employ the economist. But what makes this scenario interesting is to identify the conditions under which contracts may be enforceable or not. Consider:
 - To what extent does the ability to contract depend on the particular relationship between X and Y and the behaviour that is being modelled? Try to think of examples. Are X and Y already in a contractual relationship? If so, it would seem quite likely that they would extend the terms of that contract to commit not to employ the economist for unilateral capture. If no pre-existing contractual relationship exists between X and Y , then such commitment would be more difficult. Can we think of illustrations under which the qualitative and quantitative nature of the relationship between X and Y would make commitment unlikely? One answer is that, if there are many parties in our model then a collective agreement in which each party commits not to employ the

economist involves complex coordination even though such an agreement would avoid a negative sum game. And this coordination problem grows with the number of players.

- Can the economist act strategically in order to influence the ability for X and Y to write enforceable contracts? For example, if the parties themselves do not recognise their joint ability to commit not to employ the economist (thus avoiding efficiency loss), then the economist is hardly likely to tell them. Or if the economist is being relied on to provide the counterfactual measure of P (i.e. the level of P that would prevail if she was not being employed to influence the value of α to the benefit of her client), then the economist can defuse the incentive for the parties to commit not to employ her, by manipulating the counterfactual.
- Quite another set of issues arises if the economic activity can both influence the allocation parameter, α , and *increase* P . But I will not explore this here.

I have also implied a rational setting. Can this now be extended to bounded rationality, which applies to all parties? Either all parties must have the same utility functions, or differences in utilities must be endogenous to the model.

Model 2: Bounded, but endogenous, rationality

Again consider three parties: two producing/trading parties and an economist. The context I wish to consider is that parties might not choose the action that would maximise their utility, A^* , but an action that is sub-optimal, A . Of course, this type of bounded rationality problem is well explored by others and I just want the economist included. Consider a simple case where A is a random displacement from A^* . However, parties can become more “rational” by investing in economic literacy L . The expected value of the displacement, $A^* - A$, is a function of economic education L . For example, we could have $A^* - A = \tilde{a}$, where $\tilde{a}$ is distributed in non-negative space and the expected displacement, $\bar{a}$, is a function of L . That is,

$$E(A^* - A) = \bar{a}(L).$$

It may be that an infinite investment in literacy, L^∞ , is required to produce perfectly rational decisions, A^* . And, of course, A^* is unknown if this infinite investment is not made. The potential maximal utility of any party, “ i ” (conditional on perfect literacy), can be stated as:

$$U_i^* = U_i(C_i\{A_i(L_i^\infty); A_j\} - c_{Ai}(A_i(L_i^\infty)) - c_{Li}(L_i^\infty)),$$

where C_i =consumption generated *inter alia* by own optimal action “ $A_i^*=A_i(L_i^\infty)$ ” and by the actions of others, and A_j and $c_{Ai}(A_i(L_i)) - c_{Li}(L_i)$

are, respectively, the cost of action and the cost of the investment in literacy. We can now state the boundedly maximised utility as:¹⁵

$$\hat{U}_i = U_i(C_i(A_i^* - \tilde{a}; A_j) - c_{Ai}(A_i^* - \tilde{a}) - c_{Li}(L_i)); \quad \text{where } E(\tilde{a}) = \bar{a}$$

So far, nothing unusual (except clumsy specification). So let us address the endogenous economist idea. As before, the production function for social wealth $P = P(A_E, A_X, A_Y)$ depends on the actions of all three parties:

$$C_X = (1 - \alpha)(P - T_X^E) - c_{AX}(A_X(L_X) - c_{LX}(L_X));$$

$$C_Y = \alpha(P - T_Y^E) - c_{AY}(A_Y(L_Y) - c_{LY}(L_Y));$$

$$C_E = T_X^E(P; \alpha) + T_Y^E(P; \alpha) - c_{AE}(A_E(L_E) - c_{LE}(L_E)).$$

T_i^E = transfer from “i” to the economist.

The economist is a strategic player as before,

$$A_E = f(T_X^E, T_Y^E; c_{AE}; c_{LE}; L_E(c_{LE})).$$

However, the economist must not only determine the action, for example, the choice of economic model, but also choose how much of an investment to make in his/her own literacy.

Non-economist agents are maximising with respect to their choice of economic literacy and their action conditional on the chosen level of literacy. One way to capture the impact of the economist is to allow the literacy levels of X and Y to be augmented by the actions of the economist.

$$L_i = f(c_{Li}, A_E), \quad i = X, Y.$$

One can now see a backward induction problem in a form such as this:

<i>Agent</i>	<i>Decision variable</i>	<i>Time</i>
E	LE
E	$A_E(L_E)$	
X, Y	$L_X(A_E); L_Y(A_E)$	
X, Y	$A_X(L_X(A_E); A_E); A_Y(L_Y(A_E); A_E)$	

¹⁵ Note that the only source of risk here is the risk inherent in a sub-optimal decision $A \neq A^*$.

We could now start asking some interesting questions such as those for model 1—that is, questions about allocative efficiency and capture of the economist by one agent. But there is another set of questions that go begging.

- We might smugly assume that economists are more economically literate than other agents. What determines whether there is a balance of literacy in this direction? Is it that some of us have a natural proclivity—that is, our costs of acquiring literacy are lower than for other agents? Or is it that the cost functions for all agents encourage specialisation?
- Can we have an equilibrium in which economists are perfectly rational, but other agents are not?
- What are the mechanisms by which the economist is active? Three suggest themselves:
 - direct actions (A_E) in the form of creating knowledge;
 - direct actions (A_E) in the form of consulting to a particular party;
 - teaching, that is, enhancing the literacy of others, L_X and L_Y .

Model 3: The economist chooses strategically the testing procedure (e.g. the stopping rule) for testing her own economic theory

Suppose that it is costly to monitor and verify any empirical procedure adopted by the economist to test her theoretical model. Consider a principal-agent model in which the economist tests her own theory and chooses a stopping rule in which she undertakes “ n ” searches for a test confirming the theory. Of the “ n ” attempts, she chooses to present only that test that provides maximal confirmation. Moreover, assume that the higher the level of confirmation, the higher the probability that non-economist agents, X and Y , will adopt the theory (because it is more likely to be published in a journal). However, a high level of confirmation as revealed by this selection protocol does not mean that the theory is more likely to be correct (as a positive theory) or useful (as a normative theory). Thus, the consumer of the theory will discount its veracity (probability of being true) according to the observed (or imputed) number of trials “ n ”—assuming of course the consumer knows (or can deduce) how many trials “ n ” have been undertaken.

Now what would be a useful way of pursuing this idea? First, we could set up a familiar principal-agent model in which the former decides on a compensation structure; for example, taxpayers (as consumers of economic ideas) decide how academic economists should be rewarded. And to be quite routine, we could evoke the revelation principle and include a truth-telling constraint. By doing this, we are forcing all actors to behave as if they understood the theory that is attempting to explain their behaviour and that there were no competing

theories to influence their behaviour. In this case, the consumer/principal could infer the probable veracity of the theory by anticipating the economist's privately optimal choice of stopping rule. We would conclude that there is an "optimal" level of searching that trades-off cost and efficiency much like any other principal-agent problem.¹⁶

But this lets economists off the hook too easily for their self-serving empirical protocols. One will readily conclude that it is OK for economists to choose an empirical protocol that IS NOT an unbiased test of the underlying theory. Rather, economists (and everyone else) recognise that the monitoring of the economist is costly and the socially optimal test is one that trades-off the costs of monitoring against the veracity of the test.

The point that is raised by Gauci and Baumgartner¹⁷ is that the principal's information set may include knowledge of the model under test, as well as information about competing economic models. While the economist may be promoting her chosen model as the sole true model of reality, other actors do not necessarily buy into this propaganda or act as if it were true. Indeed, the public is exposed to several competing economic models and will ascribe different levels of credibility to each. Except in the extreme case where principals provide full credibility to the model under test (and zero credibility to all other models), the rationale for the truth-telling constraint is violated. Therefore, if we were to try to build a model with these wider information assumptions, one could no longer argue that the economist's privately optimal search could be inferred from observable variables. Thus, we would have little idea about how far economists were playing the stopping rule game.

And here lies the challenge in constructing such a model. If we really wish to investigate whether the adoption of privately optimal stopping rules by economists leads to truly inefficient outcomes (i.e. levels of bias that do not reflect an appropriate trade-off between the costs of monitoring and the costs of bias), then one cannot use the revelation principle. Indeed, the revelation principle assumes away the problem we wish to investigate. I suspect that what one would need to do is structure a model in which non-economist agents are exposed to competing economic models and ascribe differing levels of credibility to each. Then, if the economist wishes to test model "A", she must anticipate that some consumers might act as though competing model "B" were valid. Thank goodness I am now retired and I do not have to add more detail.

¹⁶ See Glaeser (2006).

¹⁷ Gauci and Baumgartner (1992).

“In-the-moment” utility vs. “contemplative” utility: Are the fundamental measures of economic performance realisations or distributions?

At least at the quantum level in physics, reality itself may be the evolution of a probability distribution or wave function rather than a temporal sequence of realisations from that distribution.¹⁸ Are there any useful perspectives here that we can transfer to economics? We can quickly get bogged down into a weird philosophical quagmire here—so let us be careful not to venture too far. Of course, I believe that we do see realisations of economic variables, and what I am trying to do is see whether there is any *additional* insight in viewing (at least part of) our changing economic reality in the form of an evolving distribution.

We assume that the end purpose of economic activity is to produce some measure of benefit to agents. How do we identify and measure this benefit? We normally use observable proxies such as consumption, income, etc., or self-assessed (e.g. by questionnaire) measures of welfare such as happiness. Now most economic models adopt an “in-the-moment” convention, whereby welfare is a function of the period-by-period realisations of consumption or income. For example, consider a lifetime consumption model where the objective is in the form:

$$W = \sum_t \int U(C_t) f(C_t) dC. \quad (1)$$

Despite the inter-temporal nature of this model, the point is that it is contemporaneous consumption that directly determines utility.

Since we do not have a clearly defined (let alone clearly measured) index of welfare, are we really sure that we should use the sequentially (or continuously) realised values of the proxy as the measure of welfare? People do not simply live “in the moment”. At an intuitive level, I think we recognise that our well-being is both shaped by the past and contemplates the future.

However let me be a little clearer about what I mean by contemplative. In lifetime models as argued, we can *optimise over anticipations of the utility of future contemporaneous consumption*. But it is consumption itself that directly impacts utility. This is very different from saying that *utility may be derived directly from the contemplation of future consumption or income*.¹⁹ For example,

¹⁸ This version of reality certainly appears to be the theoretical prediction of quantum physics as I have indicated earlier. However, the two competing interpretations, the Copenhagen approach and the Many Worlds approach, both allow for the observer to perceive a unique realisation and to interpret this as (at least one) reality.

¹⁹ Note the difference between “anticipation of utility” vs. “utility of anticipation”. This has a Jensen’s inequality ring to it.

if utility in this contemplative form depends on the first moment, then we might have the following:

$$W = V\left(\sum_i \int C f(C_i) dC\right). \quad (2)$$

If we draw an analogy between people and firms, we can think of the value of the person as the present value of future incomes (i.e. net worth including human capital). Can the individual draw welfare directly from this value as opposed to deriving utility from the stream of consumption that the net worth is expected to yield?²⁰

Notice in Eq. (2) that it is the distribution over future consumption that yields value—not the anticipated realisations from the distribution. And herein lies the analogy with quantum physics. But we should not let a mere analogy with a more fundamental science dictate the form of our economics. The important issue is whether the analogy is useful. This is clearly a subject for the behavioural economists. And the subject of such contemplative utility has indeed been around for a long time, though perhaps not with quite the motivation I am providing here.

It is quite plausible to think of the welfare of sophisticated, social, sentient subjects (humans) as measured by the evolving path of their net worth. I suspect that a major component of our (non-pathological) welfare lies in the contemplation of our future—it is about planning and conducting our journey as much as it is about the stops on the way. And our feedback in measuring our performance in that journey may be largely summarised in the trajectory of the distribution over our future wealth or consumption. The idea is neatly captured in the title of Leonard Woolf's (husband of Virginia) final, 1969, autobiographical book, *The Journey Not the Arrival Matters*.

Let me give intuitive illustrations of what I mean.²¹ We plan to go out tonight for dinner and we anticipate some utility from the dinner. However, suppose instead we delay our dinner—we will go next week instead. This allows

²⁰ I mentioned a moment ago that the utility functions normally used in personal and household economics focus on the anticipated realisations of consumption (etc.) from some probability distribution. By contrast, in corporate finance, the objective function of choice is the value of the firm which is the capitalised value of expected future earnings. And so value changes for the firm are nothing more than changes in the distribution of future earnings expectations. Thus, at first sight, we have a comparison with quantum physics. Of course, it is quite reasonable to value *any* liquid financial asset as the present value of its future cash flows, because this tracks the value that can be realised by the owner at any time. Moreover, this analogy is stressed a little since firm value evolves as only the first moment of the distribution—whereas the “Many Worlds” approach in quantum mechanics relates to the evolution of the whole distribution.

²¹ I would like to thank Howard Kunreuther here for helpful discussions of these ideas.

us to contemplate the future dinner and derive enjoyment from the contemplation itself—looking forward to the dinner itself gives us utility. Or, working in the opposite direction, take two people who currently have similar jobs, incomes, homes, etc. but who have had very different backgrounds. One had a privileged childhood and had expected all along to attain his current socio-economic status; the other came from a working-class background and has struggled hard to achieve the current status. I suspect we would agree that the latter enjoys some measure of utility from the evolutionary path that has brought him to his current position.

Here are some popular wisdoms that would give further support to such an approach (at least on a limited basis). Rich people are often said to derive satisfaction for their wealth *per se* (or relative to that of others). This is captured in sayings such as “money is simply how we keep score”. Status-conscious people might assess their welfare in terms of relative wealth. Pathologically risk-averse or insecure people might live entirely in the future, with current consumption never playing a direct welfare role.

As another insight into the differences between *anticipations of the utility of future consumption* and *the utility of anticipations of future consumption*, consider the following contrasting notions of risk aversion. The notion of risk aversion that is more familiar is the aversion to variability in period-by-period consumption. Thus, if income and consumption varies randomly year-by-year, the individual takes a welfare hit. An alternative notion of risk aversion would focus on random variation in the person’s value or net worth. If we recognize time diversification, a person who was value risk averse (as implied by the form of “ V ” in Eq. (2)), could tolerate large potential variation in future consumption on the expectation that this would diversify out in period-by-period measures of net worth.

In measuring utility relative to the evolving distribution in Eq. (2), I illustrated a simple form where only the first moment matters. Of course, one can imagine more complex forms where utility is derived from higher moments.²² But whatever the properties of the value function “ V ”, the consumer’s challenge can be represented as a control problem in which she optimises over the temporal path of the value distribution.

²² The properties of function “ V ” in the value-based control problem are up for grabs. For example, value-based risk aversion, as described above, would suggest that deviations from a smooth path be penalised, though this might not impose any penalty on variation in period-by-period consumption. We might also wish to benchmark anticipative/value utility, perhaps against the contemporaneous value of the societal mean wealth distribution. In which case, we might ask whether our benchmarking against this mean is cardinal or ordinal.

In addressing such anticipative/value utility problems, it would be important to identify the conditions under which this problem collapses into a more normal “in-the-moment” utility form.

Conclusion

If you have reached this point in the paper, a thought has undoubtedly occurred to you. It is all very well to endogenise the economist; it has a pleasing circularity to it. But why would an economist argue that one should endogenise the economist? What is in it for me to be writing this paper? We can soon slip into an infinite regression. So at this stage, I will conveniently develop a headache and thank you for your time.

Acknowledgements

I would like to thank Howard Kunreuther and Harris Schlesinger for helpful discussions.

References

- Carvagnaro, D.R., Gonzalez, R., Myung, J.I. and Pitt, M.A. (2012) 'Optimal decision stimuli for risky choice experiments: An adaptive approach', *Management Science*, forthcoming.
- Dewatripont, M. and Tirole, J. (1999) 'Advocates', *The Journal of Political Economy* **107**(1): 1–39.
- Gauci, B. and Baumgartner, T. (1992) 'The endogenous economist', *American Journal of Economics and Sociology* **51**(1): 71–85.
- Glaeser, E.L. (2006) *Researcher incentives and empirical methods*, Discussion Paper 2122, Harvard Institute of Economic Research.
- Leamer, E. (1983) 'Let's take the 'con' out of econometrics', *The American Economic Review* **73**(1): 31–43.
- MacKenzie, D., Muiesi, F. and Sui, L. (2009) *On the Performativity of Economics*, Princeton, NJ: Princeton University Press.
- Simmons, J.P., Nelson, L.D. and Simonsohn, U. (2011) 'False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant', *Psychological Science* **22**(11): 1359–1366.

About the Author

Neil Doherty is an Emeritus Professor of Business Economics and Public Policy at the Wharton School of the University of Pennsylvania.

Appendix²³

The irrational economist

When we think of evolution, we recite the case with ease.
How all species are related, in continuous family trees.
There a systematic order in natural selection
So imagine our surprise when we see a rare defection
Each species sires another and says it was “begat of us”
But who can claim the credit for the wondrous duck billed platypus?
For those who do not know it, the question that it begs.
Is how a warm blood mammal, can stoop to lay some eggs?
If you’re puzzled over platypus, the next one’s quite audacious.
A fish that went extinct, right back in the Cretaceous
Spent sixty million years playing “catch me if you can’t”
Then it lost its little game, the unlucky cealacanth.

Now Darwin’s latest puzzle, creates an astronomic fuss
That most elevated of species, *homo economicus*
Now evolution served us with a firm, if hidden, hand
And bred shrewd and cunning schemers to populate our land.
Along came Adam Smith from out the Scottish mist
And with rational self-interest, the markets boiled and hissed.
And as for Jerry Bentham, philosopher and contrarian,
He showed that economic man was quite utilitarian.
Now *hombre economico* ascends unto his throne
When Von Neumann and Morgenstern offer up their tome
And showed (as Wordsworth showed before)
“the lore, of the nicely calculated less or more”

There he stands that pinnacle of Darwinian integration
A confident and optimizing, and *boring*, automaton
There he stands in splendor, the most rational of species
Till evolution hiccups — and the fan has met the feces!

Ah! the feces hit the fan, - enter Ellsberg and Allais
And economic man withers in their gaze
Forget about that dribble over minimums and maximums

²³ The poem is an extract from a poem read at the festschrift of Howard Kunreuther in 2008.

He stumbles and he stalls on simple little axioms
He likes A to B and B to C, but really what's the point
'cos he goes and chooses C when his nose is out of joint
His framing is atrocious, his causes are Quixotic
Expected futility! Dean of the myopic
L'homme Économiques, that object of buffoonery
Is hardly fit to organize a piss-up in a brewery
He can't obey his axioms, his decisions are perverse. He
needs a little help from Kahnemann and Tversky

And now that we've established that *man* is quite irrational
(and this wanting of the species is really international)
But I know that you are thinking that the point is surely missed
The point of this poem is the irrational *economist*
But tarry a while and I'll clear away the fog
Then we'll see that the master grows to look much like the dog

Now Philosophers and Christians, love to join the scrimmage
Did God make man, or Man make god, in the other's image?
So did the heirs of Adam Smith, irrational and weary
craft consumers and producers to populate their theory?²⁴
Or was it just the other way, did agents fat and jolly,
choose sadly flawed economists to document their folly?²⁵
Whichever way causation flows, and really here's the rub
Economists are quite irrational, we can't avoid the snub.

²⁴ Ask Schroedinger; tis very clear what Erwin meant

The (irrational) observer contaminates the experiment.

²⁵ T'were absurd were man irrational in every task but one.

But, when he turns to economics, this flaw is strangely gone?