

WHY FORECAST?

Jagjit S. Chadha*

“The role of prediction in economics involves a fundamental tension. On the one hand, much of the economics is concerned with prediction. On the other, economic predictions are notoriously unreliable. It is, in fact, tempting to see the economist as the trapeze-performer who tends to miss the cross-bar, or as the jockey who keeps falling off his horse... However, it cannot be doubted that prediction is one of the central pre-occupations of economics. Policy prescriptions will inter alia involve relating alternative courses of actions to predicted outcomes. Even the description of observed trends of unemployment, poverty, living standards, etc., would tend to lead one to ask questions about the future. Not all of economics is concerned with predicting, but the central role of prediction in economics can scarcely be denied.”

A. K. Sen, *Prediction and Economic Theory*, 1986

There has been an intense debate about the rationale behind economic prediction or forecasting, triggered by a sequence of forecast errors before and after the financial crisis and more recently by a ‘surprisingly’ buoyant economy after last year’s referendum on the UK’s membership of the European Union. Some economists argue that the value of a forecast is strictly related to its forecast accuracy. Others argue that what matters is less the forecast errors but the stories that are revealed by such errors. The former might be thought to relate the value of economic forecasting solely in terms of a statistical criterion and the latter to the need to concentrate on structural relationships between economic variables that will be subject to errors (or shocks) but which can be treated as stable. I argue that the forecast process is inherently subject to large errors, and so is a hazardous exercise, but that does not by itself invalidate the exercise because both the producers and consumers of forecasts understand that errors will occur. And this knowledge throws up a clear obligation

for producers to explain errors before the fact by use of uncertainty or scenario plots and for consumers to treat the forecasts with caution.

A mug’s game

Economic forecasters ought to be thankful for pollsters otherwise they might look very bad indeed. The story that has frequently been repeated is that a recession was forecast in the event of a vote to leave the European Union and because there has been no recession, economic forecasters have let us down. This story is not quite the truth. For example, in May 2016 the *National Institute Economic Review* simply argued that growth would be broadly unaffected in 2016 by a vote to leave the European Union and projected to be almost 1 percentage point lower than the baseline in 2017; a baseline which assumed that the UK would stay in the European Union. The accuracy of the central forecast for 2016 was reasonable and we will all watch carefully what happens this year.

*NIESR. E-mail: j.chadha@niesr.ac.uk. This commentary responds to a current debate about the value of economic forecasting. It is not written as an academic paper and I apologise both for the absence of references, as it reflects many conversations and papers I have read, and for any errors, which are entirely mine. In particular though I would like to thank colleagues at the National Institute of Economic and Social Research: Angus Armstrong, Paola Buonadonna, Roger Farmer, Simon Kirby, Rebecca Piggott and James Warren for helpful conversations as well as David Chambers, Anthony Garratt, Andrew Harvey, James Mitchell, Hamid Sabourian, Richard Smith, Shaun Vahey, Ken Wallis and Mike Wickens.

That said, the substantive part of the economic impact of an exit from the EU single market story was on the long run, what we might think of as the move from one pattern of trading relationships to a view about the new pattern of those relationships. This question formed the focus of most analysis because the empirical relationship between growth and trade is well established, such that a reduction in overall trade seems likely to imply a reduction in economic growth relative to the case when there is no change in the overall level of trade. Note that this is an example of a *conditional* forecast. No-one is saying that growth will be negative in the future without qualification. It is simply that, compared to any other view we may hold about the future state of the economy, one with significantly less trade is likely to have lower economic growth, at least for a while. To be clear the *unconditional* forecast predicts the date at which you will die, the *conditional* one says that if you smoke you will die a number of years earlier.

But more important than any point forecast is the need to provide distributions that capture the measured extent of our uncertainty about forecasts, which we tend to term risk. The unmeasured uncertainty is usually called Knightian uncertainty and we will return to this a little later. The way that forecast risk is presented then is that the set of errors from previous forecasts can be used to create a measure of the likely errors from the current most likely forecast path. Even a projection of normal times ahead would have a corridor of uncertainty around it. And so when we combine a lower expected rate of growth with the distribution of measurable outcomes, which is what we can calculate from previous episodes of economic news, a larger fraction of that distribution of likely outcomes will be below zero. If we think of recessions as periods of negative growth, then it is simply the increased possibility of growth outcomes below zero that represents any heightened possibility of recession.

Playing with dice

If you play with dice, think of forecasting as a game in the following way. Roll two standard dice that have one to six on their faces and note that you will get paid the amount rolled in pounds and now forecast what number will come up. The most likely number for two fair dice is 7, for which we can think of six combinations, which would then be your rational point forecast and so you might expect a £7 payoff from a dice throw and that is what you would be willing to pay to play the game. We can all agree that the correct forecast is 7 but we know even while making the forecast that we will be ‘wrong’ 5/6th of the time. The producer of the forecast knows that and the consumer must act accordingly. Two

points follow. We cannot get at the underlying risk in our forecasts by asking even a thousand people for their point forecast because, if rational, they will all say 7. Secondly, the set of point forecasts superficially will look like people are herding around a particular view, therefore the producers of forecasts need to ask supplementary questions to evaluate the ‘true’ risk to our forecasts. And as consumers we should not act on the point forecasts because we would tend to underestimate the extent of both risk and uncertainty.

Back to our game. Think of this return of £7 as akin to a normal times scenario. As already explained, the forecast is made in the full knowledge that it is much more likely to be wrong than right but the question then facing us is the extent to which such a forecast may be useful. Yes: because first it pins down the most likely number that may be useful for planning. We call this a measure of central tendency. If we want to plan more broadly we can allow for some small errors, which will be clearer if we say our forecast says that the dice will equal 7 plus or minus two, which will be the case two thirds of the time. Indeed this fraction broadly corresponds to the amount of time the economy spends in normal times. The producer of the forecast has articulated central likelihood and a notion of how much of the time we can expect normality rather than feast or famine. This statement may well be useful to consumers of forecasts.

But the randomness implied by any one event means we cannot be at all sure what number will be returned following one draw (or shock) but under repeated trials we can form distributions that equal the likely probabilities of various events. So what you will also value is the distribution or what we also call the risks to the forecast, as they will correspond to particular draws on the dice or equivalently economic shocks or better still scenarios, by which I mean an articulated story about economics events such as oil prices, lending conditions, confidence or monetary conditions. Even then the impact of the change in a particular variable on the economic outlook depends on the shock. An unanticipated increase in oil prices may result from growth to world demand or the imposition of supply constraints on production. In the former case, the oil price change is endogenous to higher levels of growth in the world and may not signal an economic slowdown. But in the latter case, firms and households may have to reduce demand in order to budget for the higher costs of oil.

Now if we were to take the view that an economic cost (or tax) were to come along and this meant that it would cost you £3 of whatever you rolled, you would

then forecast that the return from the game would only be £4 (i.e. £7–£3) and so the probability of a negative return has increased from zero to 1/12. The forecasts of the impact of exiting the European Union, or indeed any view about a deteriorating economic outlook, reflected exactly that thought experiment and although we might expect worse returns overall, and as a result there was now a possibility of a negative return (or recession), we still might get lucky and the dice might fall favourably.

Another way we can play this game is to increase the variance without changing the expectation. This can be achieved by having one or more dice with zeros, positive and negative numbers. For example, two zeros, –2, –1, +1 and +2 would have an expected value of zero but also raise the variance of possible returns. Clearly in this case, forecast errors will tend to be higher but we would still not change our model expectation. What we would need to do though is to think as producers how we can explain that forecasts are likely to have larger errors without necessarily being treated as failures and as consumers how we should react to these more uncertain times.

It is quite obvious that we cannot know the future. But equally it is also quite obvious that we cannot afford not to think and plan for the future. Projections about future states of the world depend on a combination of information and models, which are essentially devices for turning that unstructured information into a view. Even if we make the most extreme assumptions and assume that all relevant information is free, we would still not say that this structural view is anything other than a false depiction of the world that will transpire. So they need to be treated with care, particularly when they are used to inform policy. Forecasts can be used or abused but they need to be made. How we should treat the falsity of forecasts is the subject to which I now turn.

Are forecast errors bad?

The implicit yet incorrect assumption made by many is that forecast accuracy is the overwhelming metric by which to judge the forecasting process. That is, in my view, an erroneous assumption. There are a number of value functions we might employ to assess the worth of a single or series of forecasts and only in some specific circumstances are minimised forecast errors, by which we mean the difference between projection and outcome, necessarily preferred.

Let us first imagine that all forecast errors whether small or large lead to a loss in value, or utility. Let us further suppose that larger errors lead to a greater loss than smaller errors, whilst it is true in this formulation that

no forecast errors will lead to no loss in utility. When there is a lack of perfect foresight, one might also then be indifferent between a large number of small errors and a small number of large errors.

But the recent discussion about economic forecasts seems to me to be about large errors, with the implication that small errors are not terribly significant. This kind of reasoning is reflected in the Bank of England's inflation target, which penalises large errors in outcomes (rather than forecasts) of 1 percentage point or more with the need for the Governor of the Bank to write a public letter rather than small errors, which are not punished much at all. If we want models that concentrate on predicting large errors, which are rare, then we have to think of models that may not deal very well with the day-to-day or mundane. Ultimately the choice depends on the social welfare function. But one can imagine a world in which policymakers are told to avoid extreme events in all circumstances, which is quite different from the current imperative.

If we are simply worried about extreme outcomes, one is quickly into the world of min-max, which is not so much about consistently high forecast accuracy but about minimising the losses from the maximum forecast error. So the errors we care about are only the large ones in this world. And policymakers will be acting on indicators of elevated levels of risk alone. This approach may not mitigate risk very much and may even induce policy inspired fluctuations. Paul Samuelson wrote, “to prove that Wall Street is an early omen of movements still to come in GNP, commentators quote economic studies alleging that market downturns predicted four out of the last five recessions. That is an understatement. Wall Street indexes predicted nine out of the last five recessions! And its mistakes were beauties”. If it is hard to predict with any degree of certainty the most likely throw of a dice at 7, it seems to me that it is several orders of magnitude harder to predict when we will throw a double 6 or double 1.

Another possibility is that the direction of the forecast error might matter. When scrambling to get to Kings Cross for a train, the loss inflicted from being 5 minutes early is not the same as being 5 minutes late. The forecast error direction might well matter more than the strict size of the error. In this case one would prefer 10 minutes early to 1 minute late. That one might prefer being more wrong to being less wrong can seem odd but is perfectly sensible given the nature of our preferences. The direction of error might also matter from the perspective of a trader on the financial markets who may

take a position based on whether a particular economic indicator, such as non-farm payrolls, comes out above or below market expectations. Note that these market expectations themselves are produced from surveys of forecasts from professional economists. A trader though will make money if she has placed herself on the right side of the surprise. It matters little how close her view is to the outcome, only whether they are on the same side of the surprise as the market release.

And so many studies concentrate on whether forecasts make money for a portfolio. These studies are moving away from statistical assessment to financial returns, which may be more concerned with the classification for forecasts in a non-parametric manner such as a contingency table of whether the forecasts and outcomes tend to be the same side of surprise more often than we might expect by chance. Furthermore in this case, financial returns will depend on the costs of carrying a portfolio and the execution costs of a trade, which will encompass market liquidity. In other words the attitudes to forecast errors are state dependent. A small loss may not matter unless it takes the fund into bankruptcy and a large loss might be absorbed if there is sufficient capital. Even Keynes, writing in 1938, realised that trading on economic forecasts and information about firms may not always lead to high returns: “we have not proved able to take much advantage of a general systematic movement out of and into ordinary shares as a whole at different phases of the trade cycle”.

This kind of problem might be the case in the policy world as well from time to time. And if forecasts represent a game between the producer of forecasts and their consumers, in which convincing the latter may be an important part of the former’s value function, things can get complicated very quickly. As we shall now consider.

Forecasting and giving advice

We may be fortunate to have a consensus view on the most likely outcome and this may come from the dice story I have outlined. But how should a rational agent ‘consume’ this advice? I shall suggest that simple theory tends to say that we should be wary of the motivation of those who forecast at the extreme and that we should still put weight on the central case. Although this essay is about forecasts and their errors, we are also quite accustomed to the entertaining spectator sport of economists disagreeing, often in very strong terms. But the lead up to the Referendum vote introduced to the public the rather strange sight of economists coming together nearly as one. The consensus was rather clear. The long-run central case was that re-orienting trade

and capital flows away from our largest trading partners in the EU would tend to reduce output in the long and short run. I do not want to go into the mechanics of these forecasts here but want to assess the problem from that of the consumer.

Actually like many contrary-minded people, I am often wary of the consensus as it may turn out to be misleading. Experts may tend to offer advice that favours the *status quo* because there is little substantive evidence in favour of an altered state of nature. They may also herd around one opinion because they do not want to be seen as outliers with their professional colleagues or because they have not been able to develop a truly independent view; either of these possibilities may lead to what economists call an informational cascade in which information that may be valuable is lost. Ideally, we would like forecasters’ private information to be decanted into the public space so that we can benefit from the available views and act accordingly.

Some may think that all or some of the experts may also all be lying to us in order to gain some advantage later or are motivated by their own private returns from the choices of forecast consumers. We shall concentrate shortly on the incentives to influence consumers. But first we might dismiss the economic consensus on at least one of three grounds: (i) that experts are too risk averse, or (ii) short-sighted or (iii) have formed a mendacious conspiracy? In the face of such possible biases, can yet more economic analysis help us again? Even if the signal from the experts is clear, the problem then is how we consume it. As members of the public, we may have pretty much one answer or signal but before acting on it we have to decide on how to treat any biases.

Assuming that all the economists who have provided various analyses want their advice to be heeded, they may also have tailored their advice so as to maximise the probability that the advice will be used by amplifying the signal: thinking that it is better to exaggerate the impact on the up or downside so that people will understand the qualitative stance on ‘good’ or ‘bad’ more easily. For example, to influence a trader, rather than forecasting a small surprise in the non-farm payrolls we might have to forecast a big surprise so as to outweigh her natural suspicions or uncertainties in the forecasts produced. An extreme view may be a cleaner signal.

If economists wished to influence opinion in this manner, the set of published views on either side would then actually become more dispersed as economists would exaggerate their claims in order to get more attention.

Those making forecasts in the tail of the distribution are thus taking a very strong position, as they are putting a lot of weight on an unlikely scenario. And if this happened there is further, more subtle impact. The increased polarisation of opinion on the impact of a particular event will then lead to increased uncertainty over the future as rational agents will today attach some possibility to either extreme states. Accordingly we would (almost) certainly observe an increase in uncertainty as these two regimes tussle out the consequences of a given policy for the UK. So uncertainty itself then may lead to some delay in both consumption and investment, which may not be completely resolved until the episode and the policy shock have worked their way through the system.

The disinterested yet rational economist then has another choice. Let us suppose that whatever forecast or advice she wishes to publish, she wants the output consequences to be minimised. She will know that her advice may not be followed and that the other side may win the day. So she may have to change her advice in order to militate against the costs of her advice being ignored. How does she do this? She then forecasts away from the extreme and so reduces the range of possible outcomes. Our sensible economist wants to reduce uncertainty because she does not want to impart a shock on the economy as rational agents respond first to uncertainty and then to the expectation of a large change in circumstance. The public then might be best choosing the consensus when accepting advice: if the economists were truly disinterested in the results of an episode one way or other they would not wish to exaggerate the consequences because that would by itself negatively impact on the economy. And so we as the public should, in the absence of an ability to referee or replicate the analysis produced, be most wary of the outliers or extremes on either side, so perhaps the truly contrary thing in terms of the economics and the referendum is to agree with the remaining 'trimmed' consensus. Discard the extremes and place weight on the central view.

The model as time machine

The Institute has been developing its analysis of economic prospects and the causes of change since its establishment in 1938. But this process was heightened with the publication of this *Review* from 1959. By November 1963 forecasts of GDP started to be published in the *Review*. And since then there has been an intense quarterly effort from economists running models, assessing data, understanding deviations of outcomes from expectations and applying dollops of judgement. When the output is brought together, Institute staff start to make some sense of what has happened and think

about what might happen. Even though many of our thoughts may not come to pass, like nightmares or dreams, it is quite necessary to think through the future using models.

An economic model is a parsimonious and, by definition, imperfect reflection of reality. It might be derived from first principles and respect economic theory. Or it might be a set of relationships that capture observations and derive from empirical observation. The former and the latter are sometimes barely on speaking terms. But however we derive those relationships in a model, they collectively describe our view of how the world works. A forecaster, and perhaps with her judgement, would then crank that model into an unknowable future and would trace a number of possible futures: some more likely than others.

The economic forecast is thus ultimately an experiment in time travel not much different from those outlined by authors such as H.G. Wells, Ray Bradbury or Douglas Adams. The forecast allows the economist to articulate a future state of the world where each macroeconomic variable is consistent with every other macroeconomic variable. And a good forecaster will articulate a large number of possible states of the world but where each set of macroeconomic variables is consistent with each other. All measureable scenarios will give us our fan chart of possible outcomes. So as well as an artificial universe, the modeller is inventing parallel universes. Think of a set of statements about output, inflation, exchange rates, productivity, unemployment and asset prices which are all consistent with each other in each possible state of the world. Even better if we can incentivise different groups of modellers to articulate their model consistent views of the world, so the genuine uncertainty we have about models and data can be reduced somewhat by more information.

The problem for the evaluation of forecasts arises because from the perspective of today many possible states might obtain tomorrow but when we get to that tomorrow only one state will have obtained. And that will mean that a forecast comprising many states will tend to look as though it is 'wrong'. As already outlined we know this very well: forecast accuracy does not imply the absence of forecast errors. One can go further because we not only expect less than perfect forecast accuracy when forecasting, we might actually welcome that. And there are three broad reasons. First if we collectively use information efficiently all that is left to explain the future is what we do not now know and because we do not now know it, the future will be unknown and a surprise,

or what economists call ‘news’. Secondly, if we use the forecast to plan and set policy in order to minimise the worst expectations that will arise from our forecast, we will change the future. And the forecast will turn out, perhaps thankfully, to have a large error induced by our own policy actions.

Finally, and most importantly, we want to use the forecast errors to understand the news that has accrued since we made our original forecast. Without the forecast, which is what we anticipate, we cannot decompose future outcomes into what was anticipated and what was news. The anticipated part reflects the projection of key inter-relationships in future time. The error from that anticipation or news ought to allow us to understand the economic story behind the forecast error but with a set of stringent side constraints. So if consumption is higher than we expected given our path for income, wages and the supply of funds, we have to construct a story that explains higher consumption but also then does not then fail to explain the subsequent path of income, wages and the supply of loanable funds. The model does not allow completely free thinking, like a crossword the answer must fit the letters of previously identified clues. An economic model does not admit anarchy.

None of Wells, Bradbury or Adams quite got our present, as their future, quite right. Equally forecasters prior to the Great Recession did not either. And they would not have expected to be quite right. But elements of truth are there and those elements are useful. Wells’ vision of a society dominated by the young, Bradbury’s point about small events in the past having large effects in future

and Adams’ guide for hitchhikers is really a smartphone. And in some cases it is too early to tell how inaccurate they are and so it is the same with any recent economic forecast.

Concluding remarks

Academic macroeconomic forecasting transformed following the financial crisis of the previous decade. Research produced a greater emphasis on model and data uncertainties, a push for more and better quality macroeconomic data, and consideration of new models better suited to gauging risk. Risk can be captured by the measurable. But by using surveys, digital web-scraping and measures of confidence researchers are also trying to get at what was previously unmeasured, the Knightian world. The production of forecasts now stresses risk and uncertainty. And within that risk and uncertainty we need to explain better the various scenarios, stories or states of nature, that add up to our measures of risk or uncertainty. The consumption of forecasts should now do so as well.

It might ultimately be time to start thinking of economic forecasting as akin, at least in the first pass, to projection. The forecast is essentially projecting a snapshot, perhaps one that is slightly out of focus, onto another plane. The further we project the image, the less well defined it will tend to be. We might then ask why do we project? The simple answer is so that more can see. The projection gives us a scenario or set of scenarios to evaluate to think about and discuss. And that might be just enough. Why forecast? So we can think about possible futures and plan accordingly.