North American Journal of Economics and Finance 22 (2011) 43–60

Contents lists available at ScienceDirect

North American Journal of Economics and Finance

Decision-making in hard times: What is a recession, why do we care and how do we know when we are in one?☆

Kevin Lee[a], Kalvinder K. Shields[b,*]

[a] University of Leicester, UK
[b] Department of Economics, University of Melbourne, Victoria 3010, Australia

ARTICLE INFO

JEL classification:
E52
E58

Keywords:
Recession
Probability forecasts
Real time

ABSTRACT

Defining a recessionary event as one which impacts adversely on individuals' economic well-being, the paper argues that recession is a multi-faceted phenomenon whose meaning differs from person to person as it impacts on their decision-making in real time. It argues that recession is best represented through the calculation of the nowcast of recession event probabilities. A variety of such probabilities are produced using a real-time data set for the US for the period, focusing on the likelihood of various recessionary events through 1986q1-2008q4 and on prospects beyond the end of the sample.

1. Introduction

In December 2008, the Wall Street Journal carried the front-page news that "the US entered recession in December 2007" based on the NBER's announcement that the previous peak of activity had been in the fourth quarter of 2007. The fact that this was the lead article in a journal with a daily circulation of more than 2 million readers shows that there is considerable interest in the business cycle and the timing of business cycle events. There has also accumulated a voluminous academic literature

☆ The authors are grateful to the Editors and referees and to participants of the RBNZ's Nowcasting and Model Uncertainty Workshop in December 2008, for useful comments. This research was funded by the Australian Research Council under the Discovery-Projects scheme grant number DP0988112.

* Corresponding author. Tel.: +61 3 83443549; fax: +61 3 83446899.
E-mail address: k.shields@unimelb.edu.au (K.K. Shields).

doi:10.1016/j.najef.2010.12.001

concerned with the same issues (see, for example, van Dijk, van Dijk, & Frances's, 2005 special issue of *Journal of Applied Econometrics* for an overview).

Despite this interest, it is difficult to find a straightforward explanation for why business cycle pronouncements of this sort generate such interest. For example, the definition of recession used by the NBER for the US is only vaguely expressed (as a "significant decline in activity spread across the economy lasting more than a few months") and the process by which its Business Cycle Dating Committee forms its subjective qualitative judgements are obscure. The news that recession started twelve months earlier also seems a little out-of-date. It is not entirely clear, then, what the Wall Street Journal's readership thought it was reading about or why it cared.

Many academic commentators have sought to clarify matters by suggesting algorithms that define recession explicitly in terms of specified economic events and which are judged according to the extent to which their assessments of the cycle match that of the NBER. Harding and Pagan (2006) and Leamer (2008) provide good examples of this approach based on data-analytic methods while Chauvet and Hamilton (2006) provide a good illustration of the approach based on econometric modelling methods. The implication of this work is that, despite what the NBER says, there is a single definition of recession and a fixed rule that the NBER could employ to capture this definition in making their judgements. It is this single rule that this part of the academic literature has attempted to reveal.

A second strand to the literature on business cycle dating has focused on the extent to which the dating is vulnerable to data revisions. As illustrated above, NBER statements on business cycle dates are typically made with a considerable delay specifically to avoid making announcements that turn out to be misjudged subsequently simply because of inaccuracies in the available data. But the recent literature has noted that, if business cycle information is to be used in real-time decision-making rather than as an after-the-event characterisation of historical events, then the delays in publication of the information are extremely unhelpful (see Aruoba, Diebold, & Scotti, 2009, or Chauvet & Piger, 2003). It is argued that more straightforward algorithms for business cycle dating are useful, then, because they allow a more timely statement on business cycle conditions for real-time decision-making.

This paper agrees that the primary purpose of studying the timing of recession is for real-time decision-making but it starts from the viewpoint that the public experience recessions not as binary statements on business cycle dates but as events that span several dimensions. The paper aims to place the definition of recession in the context of individuals' decision-making to demonstrate why different individuals will be concerned about different recessionary events and hence explain why a recession is multi-dimensional and why there is no single event or rule that can be used to definitively define one. It concludes that the best way to characterise recession is through the production of event probability forecasts that convey the likelihood of experiencing the various events of interest to different individuals in a systematic and comprehensive way. If these are to be helpful in decision-making, then they have to be available in real time, not after a period of reflection. The paper shows that this is possible with a reasonably high degree of precision.

The NBER's Dating Committee does not use a fixed definition of economic activity but its statement on the determination of dates of turning points refers to many of the events that we consider and for which we present probability forecasts. There is clearly a link between our approach and NBER statements on recession therefore. But our approach provides an explanation for the NBER's choice of events and why its Dating Committee is reluctant to provide a straightforward algorithmic definition of recession. It also provides a possible interpretation of the NBER's pronouncements on recession as a summary statement of the committee members' views while emphasising that we believe the production of a range of recessionary-event probability forecasts is a more productive approach to characterising recession than the NBER's dichotomous announcements on recession/expansion.

The remainder of the paper is organised as follows. Section 2 elaborates on the link between decision-making in real-time and the definition of recession. Formalising the link through a discussion of loss functions and the density forecasts that describe the range of possible macroeconomic outcomes and their likelihoods, the section discusses the use of event probability forecasting in characterising recession events and briefly comments on how these can be obtained for use in real-time decision-making. Section 3 applies the methods to a real-time dataset for the US using a VAR model that accommodates information on variables as they are first-released, on their subsequent revision and on their expected future value (as observed directly from surveys and other sources). Probability forecasts are produced for a variety of alternative recessionary events over the period 1986q1-2008q4 based only on data that was available at the time to provide a sophisticated picture of recession as experienced in real-time. These capture both the broad characteristics of the macroeconomy and the uncertainties associated with the definition of recessionary events and their likelihood of occurring. A more detailed analysis is also provided of the prospects for 2009, viewed from the perspective of an individual making decisions in the turbulent circumstances of 2008q4. Section 4 concludes.

2. Recession as a decision-based phenomenon

The typical characteristics of a recession, considering the popular usage of the term, are well-understood. A recession is a period associated with reduced activity and economic hardship for a substantial number of people. Many firms' order books dry up as diverse types of consumption fall, investment opportunities are generally reduced and some workers lose their jobs. Large numbers of households' incomes fall as unemployment rises and/or real wages moderate and/or wealth is eroded. Firms in different sectors and different households feel the effects of the recession more or less strongly and at different times. But a defining feature of recession is that the reduced activity impacts on virtually everybody's decision-making in some way for a protracted period. The recurrent nature of these phases of recession is what most people think of when they describe a business cycle.

This non-specific popularist view of what constitutes a recession has been translated into the study of various economic magnitudes in the academic literature, usually involving output. The most regularly-used definition of recession is two consecutive quarters of negative output growth, basing the recessionary event on a zero output growth threshold therefore. Leamer (2008) shows that the NBER-identified periods of recession have typically been observed when growth in industrial production measured over a six-month period falls below −3% and when growth in payroll employment measured over a six-month period falls below −0.5%. This illustrates, then, that the variable of interest does not have to be GDP, that the timing of the event does not have to be two consecutive quarters and that the threshold for identifying recession does not need to be zero. But the growth threshold is nevertheless at the heart of the event. Turning point cycles, of the type described in Harding and Pagan (2002), also rely on a (zero) output growth threshold. Here, a peak is dated when quarterly output growth is greater than zero and then less than zero in consecutive periods (and a trough is similarly defined). Attention is paid here to the second derivative of the output series too but it is the growth threshold which plays the primary role with recession usually defined as the period from peak to trough.[1]

Recessionary events based on growth thresholds focus attention on the deterioration of opportunities faced by individuals as the level of activity falls from its previous peak over a protracted period. Previous decisions based on an assumption of continued positive growth will have to be revised and there will be a direct impact on individuals' utility to the extent that loss functions accommodate 'ratchet effects' to reflect individuals' positions relative to recent experience. Recessionary events defined as such are less relevant to individuals whose loss function is based on absolute income levels though. A period of recession defined as peak to trough ends as output begins to rise. This is good news for those who take pleasure from the improving opportunities that this implies but it is little consolation for those who care that their output is lower than the previous peak. For the latter group, Beaudry and Koop's (1993) measure of the "current depth of recession" is a more relevant magnitude, defined by the level of output relative to the previous peak and identifying recession as those periods when the variable falls below zero.

These examples of recessionary events suggested in the literature emphasise the fact that individuals' experience of recession is dependent on their individual objectives and subjective preferences. These are not usually made explicit by agents so the definition of recession is ambiguous by its nature.

[1] The Markov-Switching approach to business cycle analysis introduced by Hamilton (1989), also focuses on growth, allowing for two distinct growth states in an underlying econometric model of output. Recession is identified with the case where the estimated probability of being in the low growth state exceeds a critical threshold (possibly subject to some smoothing criteria to avoid abrupt changes in status; see Chauvet & Hamilton, 2006, for recent discussion). This is a clearly defined event although, obviously, it is model-dependent in the sense that it can only be defined with reference to the estimated econometric model.

One counter-example to this point is provided in the recent literature describing Dynamic Stochastic General Equilibrium models of the macro-economy; see Woodford (2003) for a detailed textbook exposition of the approach. Here, a model of the macroeconomy is derived with explicit micro-foundations fully describing individual household and firm decision-making in the face of imperfectly-competitive labour and product markets and in the presence of nominal rigidities. In this context, Woodford (2001) derives an explicit welfare-theoretic loss function depending on inflation and the deviation of output from a 'natural' output level, defined as the output level that would be obtained if there were perfectly-flexible prices. This loss function reflects the deadweight loss experienced by a representative household as the average level of output across goods deviates from its efficient level (the gap term) and as the output of each individual good deviates from the average level (a term proportional to inflation). Inflation is, in turn, influenced by the gap and by expected future values of the gap through a New Keynesian Phillips curve so that this gap measure is at the heart of decision-making, being the key determinant in the explicitly-derived loss function. While there might be difficulties in measuring the natural level of output, particularly in real time (cf. Orphanides, Porter, Refschneider, Tetlow, & Finan, 2000), it is clear that the only concept relevant to defining recession in this modelling framework is a negative gap measure.

2.1. Loss functions and event probabilities

While there are obviously common themes running through the recessionary events considered in the literature discussed above, the variety and range of events considered reflect the idea at the heart of the popularist view that recession impacts on different firms and different households in different ways and there is no single event that adequately reflects a recession. Rather, all of these events reflect the various aspects of recession that are important to some agents at some times.

This idea can be formalised a little in a standard decision-theoretic framework using the loss function $\lambda_i(q_T, \mathbf{Z}_{T+1,T+H})$. This function characterises the costs and benefits to individual i at time T when the variables in the m-vector $\mathbf{z}_t = \{z_{1t}, z_{2t}, \ldots, z_{mt}\}$ take specified values over the forecast horizon $T+1, \ldots, T+H$, using the notation $\mathbf{Z}_{T+1,T+h} = (\mathbf{z}_{T+1}, \ldots, \mathbf{z}_{T+H})'$ to denote all these future values. The loss also depends on the value of q_T, a decision variable to be chosen by the individual. In order to explain how the analysis of recessionary events can be related to real-time decision-making, in much of what follows, we assume that the vector $\mathbf{z}_t$ contains three measures relating to output, although the discussion could be readily extended to consider many other economic variables. The three measures focus on the first release of output data, revisions in the data, and direct measures of expectations that are available (including survey-based expectations data or market-based financial data, say). This is important in real-time analysis if we are to properly take into account the information that was available to agents at the time decisions are made (including direct measures of expectations) and if we are to take into account the fact that data revisions have a systematic content that agents recognise when making their decisions (which are typically concerned with post-revision magnitudes). Denoting the (log) of output at time t by y_t, the three measures that we consider are: ${}_ty_{t-1}$ which denotes the measure of output at time $t-1$ as published in the official first-release publication at time t (assuming a one period publication delay); ${}_ty_t^e$ which denotes a direct measure of the nowcast of output at time t as published in a survey, say, at time t (prior to the official estimate); and ${}_ty_{t-2}$ which, assuming that data is revised just once, provides the post-revision measure of output at time $t-2$ as published in time t. For the purpose of exposition, then, take $\mathbf{z}_t = \{{}_ty_{t-1}, {}_ty_t^e, {}_ty_{t-2}\}$.

The density function $f_T(\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_T, \mathbf{z}_{T+1}, \ldots, \mathbf{z}_{T+H}; \boldsymbol{\theta})$ describes the probability of obtaining specified values of the observed and forecast data in $\mathbf{z}_t$ over the estimation and forecast horizons, $t = 1, \ldots, T$ and $t = T+1, \ldots, T+H$ respectively, based on a given model indexed by the $k \times 1$ vector of parameters $\boldsymbol{\theta}$.[2] This probability density function (pdf) can be decomposed into the product of the conditional distributions of the successive observations on $\mathbf{z}_t$, to write

[2] We assume that the outcome of the variables in $\mathbf{Z}_{T,T+H}$ are outside the individual's control and independent of q_T.

$$f_T(\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_T, \boldsymbol{z}_{T+1}, \ldots, \boldsymbol{z}_{T+H};\, \boldsymbol{\theta})$$

$$= f_T(\boldsymbol{z}_{T+1}, \ldots, \boldsymbol{z}_{T+H}|\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_T;\, \theta) \prod_{s=1}^{T} f_T(\boldsymbol{z}_s|\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_{s-1};\, \boldsymbol{\theta}) \tag{2.1}$$

$$= f_T(\boldsymbol{Z}_{T+1,T+H}|\Omega_T;\, \theta) \prod_{s=1}^{T} f_T(\boldsymbol{z}_s|\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_{s-1};\, \boldsymbol{\theta}) \tag{2.2}$$

denoting the information available at time T by $\Omega_T = \{\boldsymbol{z}_1, \boldsymbol{z}_2, \ldots, \boldsymbol{z}_T\}$.

A general event relating to the variables in $\boldsymbol{z}_t$ at $T+1$ and over the forecast horizon can be defined by $R:\{\phi_l(\boldsymbol{Z}_{T+1,T+H}) < a_l$ for $l = 1, 2, \ldots, L\}$, or, equivalently, $R:\{\phi(\boldsymbol{Z}_{T+1,T+H}) < \boldsymbol{a}\}$ where $\boldsymbol{\phi}(.) = (\phi_1(.), \phi_2(.), \ldots, \phi_l(.))'$ and $\boldsymbol{a} = (a_1, \ldots, a_l)'$ are $l \times 1$ vectors. The event R is defined by the simultaneous occurrence of l (possibly interdependent) individual events then. In the context of real-time analysis, it is the post-revision data that is usually of interest in decision-making and, in this case, events of interest at time T might include

$A:\quad \left\{({}_{T+1}y_T - {}_T y_{T-1} < 0) \cap ({}_T y_{T-1} - {}_{T-1}y_{T-2} < 0)\right\};$
i.e. a nowcast of two consecutive periods of negative growth observed in T;

$B:\quad \{{}_{T+1}y_T < \max({}_T y_{T-1}, {}_{T-1}y_{T-2}, {}_{T-2}y_{T-3}, \ldots)\};$
i.e. period T output lies below its previous peak level;

$C:\quad \left\{{}_{T+1}y_T < {}_{T+1}\tilde{y}_T,\ \text{where}\ {}_{T+1}\tilde{y}_T = \frac{1}{5}({}_{T-1}y_{T-2} + {}_T y_{T-1} + {}_{T+1}y_T + {}_{T+2}y_{T+1} + {}_{T+3}y_{T+2})\right\};$
i.e. output lies below trend, defined as the centred five period moving-average of output;

and so on. These examples show that events can involve complicated non-linear functions of variables and can involve variables dated at a variety of different forecast horizons.

Using the expressions defined above, the probability forecast associated with event *R:* conditional on information in Ω_T is given by

$$\pi_T(a, H, \boldsymbol{\phi}(.), \boldsymbol{\theta}) = Pr[\boldsymbol{\phi}(\boldsymbol{Z}_{T+1,T+H}) < \boldsymbol{a}|\Omega_T;\, \boldsymbol{\theta}] = \int_R f_T(\boldsymbol{Z}_{T+1,T+H}|\Omega_T;\, \boldsymbol{\theta})\, d\boldsymbol{Z}_{T+1,T+H} \tag{2.3}$$

assuming that the model parameters $\boldsymbol{\theta}$ are known.

The expected loss associated with any outcome for an individual with loss function $\lambda_i(q_T, \boldsymbol{Z}_{T+1,T+H})$ depends on the mitigating actions that the individual takes. Standard expected utility maximisation means the individual sets her decision variable at time T at its optimal value, q^*_{iT}, taking into account the range of possible outcomes for $\boldsymbol{Z}_{T+1,T+H}$ so that

$$q^*_{iT} = \min_{q_T} \left\{ \int \lambda_i(q_T, \boldsymbol{Z}_{T+1,T+H}) f_T(\boldsymbol{Z}_{T+1,T+H}|\Omega_T;\, \boldsymbol{\theta}) d\boldsymbol{Z}_{T+1,T+H} \right\}. \tag{2.4}$$

The expected loss associated with the specific event R is then

$$L^R_{iT}(a, H, \boldsymbol{\phi}(.), \boldsymbol{\theta}) = \int_R \lambda_i(q^*_{iT}, \boldsymbol{Z}_{T+1,T+H}) f_T(\boldsymbol{Z}_{T+1,T+H}|\Omega_T;\, \boldsymbol{\theta}) d\boldsymbol{Z}_{T+1,T+H}.$$

While the event of interest here is defined with respect to a common set of variables $\boldsymbol{z}_{T+1}, \ldots, \boldsymbol{z}_{T+H}$, typically including current and forecast outputs, the loss function and therefore the expected loss can be quite different from individual to individual. An event might be defined as being significantly damaging for an individual if, for some threshold d_i, the expected loss when this event occurs is substantially worse than the typical outcome, so that

$$\frac{L^R_{iT}(a, H, \boldsymbol{\phi}(.), \boldsymbol{\theta})}{L^R_{iT}(\infty, H, \boldsymbol{\phi}(.), \boldsymbol{\theta})} > d_i, \text{ say.} \tag{2.5}$$

The discussion above suggests that, where the variables in $\boldsymbol{z}_T$ are associated with economic activity, (2.5) might reasonably define a "recessionary event" for individual *i*. The sort of events emphasised in the literature suggest that the loss functions most frequently associated with recession are those accommodating non-linearities and/or asymmetries and which reflect particular vulnerability to growth below some threshold over a protracted period, or simply a drop below a previous peak or below a trend. But the point to emphasise is that events *A*, *B*, *C* or any number of other events involving output might constitute a recessionary event for someone depending on their decision-making circumstances.

2.2. Characterising recession

Focusing on recession as a decision-based phenomenon in this way has a number of implications for the way that recession should be characterised and reported. Most obviously, the discussion suggests that it is generally unhelpful, if not misleading, to suggest that recession can be defined with respect to a single fixed rule. Rather, a description of recession should attempt to describe the widest possible range of recessionary events as any one of these could be of interest to someone. Further, in the absence of detailed information on individuals' loss functions and without providing overwhelming detail on the density forecasts, the most useful way of presenting information on the variables of interest is in the form of a set of recessionary-event probability forecasts described in (2.3) above since this is the form in which agents can interpret the information in their own individual decision-making. Of course, in practice, there is a limit to the number of recessionary-event probability forecasts that can be published. But the emphasis of the NBER on the simple dichotomous statement that there is or is not a recession at any point cannot satisfy the public's need for business cycle dating for use in their decision-making. The production of probability forecasts for a small number of frequently-cited events is certainly possible and can convey some of the required detail.

An important aspect of the discussion of the previous section is that the definition of recession is typically quite independent of the data generating process underlying the measures of economic activity. Hence, a simple model of the important macroeconomic aggregates can be estimated and used to generate density forecasts of output and other activity-related variables no matter how complex or ambiguously-defined the recessionary events of interest become. This means that it is very straightforward to generate density forecasts and the associated recessionary-event probabilities in real time.

This point is worth elaborating. Macroeconomic modelling can, of course, be based on models with a large number of variables or few variables and can incorporate more or less structural content. Simple Vector Autoregressive (VAR) models are widely used because of their simplicity and their ability to capture the complicated macroeconomic dynamics present in the data. They are also able to accommodate a wide variety of structural models as special cases, allowing the underlying theory to be tested; see the discussion in Garratt, Lee, Pesaran, and Shin (2006) text. In the case where we are interested in modelling real-time datasets, the VAR framework also allows direct measures of future expectations and data on revisions to be included in the model in a straightforward way. So, for example, concentrating once more on output data only, with $\boldsymbol{z}_t = ({}_t y_{t-1}, {}_t y_t^e, {}_t y_{t-2})'$, and using the data vintages released up to *T*, it is straightforward to estimate the reduced form VAR(1):

$$\boldsymbol{z}_T = \boldsymbol{C}\boldsymbol{z}_{t-1} + \boldsymbol{u}_t, \text{ for } t = 1, \ldots, T \tag{2.6}$$

with estimated parameters $\widehat{\boldsymbol{C}}$ and estimated covariance matrix $\widehat{\sum}$, say.[3] Having modelled this data generating process, the methods for the calculation of probability forecasts and pdf's are relatively straightforward to implement using simulation methods.[4] For example, abstracting from parameter uncertainty, one can use the estimated parameters of (2.6) to generate *S* replications of the future

[3] See Lee, Olekalns, and Shields (2009) for a detailed discussion of the relationship between this reduced form model and a structural model incorporating economically-meaningful shocks to the first-release, expectations and post-revision series.

[4] The methods, including those that accommodate model uncertainty and parameter uncertainty as well as the stochastic uncertainty considered here, are described in detail in Garratt, Lee, Pesaran, and Shin (2003).

vintages of data, denoted $\widehat{\mathbf{z}}_{T+h}^{(s)}$ for $h=0, 1, \ldots, H$ and $s=1, \ldots, S$. These simulations give directly the forecast pdf's of the first-release, expected and post-revision output series over the relevant forecast horizon; i.e. the estimated density forecasts $\widehat{f_T}(\mathbf{Z}_{T+1,T+H}|\Omega_T;\, \boldsymbol{\theta})$. Simply counting the number of times an event occurs in these simulations also provides a forecast of the probability that the event will occur, $\widehat{\pi_T}(a, H, \boldsymbol{\phi}(.), \boldsymbol{\theta})$; for example, the fraction of the simulations in which $\{(\widehat{_{T+1}y_T}^{(s)} - \widehat{_Ty_{T-1}}^{(s)} < 0) \cap (\widehat{_Ty_{T-1}}^{(s)} - \widehat{_{T-1}y_{T-2}}^{(s)} < 0)\}$ provides an estimate of the forecast probability of event A. The definition of the event of interest and characterisation of recession is entirely separate from the model used to characterise the data generating process.[5]

The separation between the model estimated to characterise the data-generating process and the recession-defining event is also important in judging the reliability of any statements on recession. The density forecasts for output obtained from the model are obviously central to the construction of the recession-probability forecasts and there are now a variety of tests available for the evaluation of the density forecast performance of a model; see, for example, Diebold, Gunther, and Tay (1998), Diebold, Hahn, and Tay (1999) and Clements and Smith (2000). These tests are based on the probability integral transforms of the observed data over an evaluation period obtained using the estimated density forecasts and make no reference to the decision-making context in which the model will be used. They provide a general statistical test of the adequacy of the forecasts and the underlying model judged from the perspective of *the producer of the forecast*. However, forecast evaluation is very different in a decision-based context when the perspective of *the user of the forecast* is paramount; see, for example, Pesaran and Skouras (2002) and references within for further discussion. Here the primary concern is the extent to which forecasts help to improve decision-making and a more appropriate criterion for evaluating the density forecasts is the realised mean loss function for the individual

$$\overline{\lambda_{iT}} = \frac{1}{N}\sum_{i=0}^{N-1}\lambda(q_{iT+i}^{*}, \mathbf{Z}_{T+1+i,T+H+i})$$

reflecting the losses incurred over the evaluation period $T, \ldots, T+N$ when using the estimated model to make decisions. This can be compared to the losses achieved using an alternative 'benchmark' model (e.g. a random walk) to provide more economically-meaningful decision-based forecast comparisons. As emphasised in the literature, there is no reason to believe the forecast evaluation of a model based on statistical criteria will match that based on the decision-based criterion (indeed, if the decision variable has relatively little impact on the loss, then all models will be equally unhelpful). But in those cases where recessionary events can occur (i.e. events where substantial losses, as defined in (2.5), can be incurred), the decision-based approach seems most appropriate given that it is able to take into account the individual's particular perspective. Of course, the problem is that this approach requires an explicit statement on the form of the loss function and this is not usually available.[6]

An intermediate line on forecast evaluation is to focus on the recession-probability forecasts. For example, if we choose the value of 0.5 as a threshold probability above which we say we believe the event will happen, then we can apply a contingency table-based test to judge the "hit rate" (where the model correctly predicts the event will or will not happen) against the "false alarm rate"; see for example the test proposed in Pesaran and Timmermann (1992). This sort of test may lack power because it wastes some of the information contained in the recession probabilities and it does not properly take into account the losses incurred by individuals. But it has the advantage that it can be calculated using only the density forecasts and it is also able to reflect the broad nature of the recessionary events of concern to individuals in the absence of detailed information on loss functions. Tests relating to different events could also arrive at different conclusions reflecting the fact that a particular model might be more or less relevant for different individuals depending on their perspective on what constitutes a recession.

[5] This is not the case in the Markov switching models noted earlier where it is assumed the recession event of interest to individuals itself influences the data generating process.

[6] See Garratt and Lee (2006) for an example in which the decision context is articulated explicitly in an analysis of portfolio decisions involving domestic and foreign investment.

Finally here, it is worth noting that the approach to characterising recession outlined above provides an interpretation of the NBER's dichotomous statements on recession/expansion that is arguably closer to the way in which the NBER functions than can be captured by any single fixed rule. This is because the NBER announcements are the outcome of discussion among the economists of the Dating Committee all of whom might have slightly different views of what constitutes a recession. The committee announces that the economy is in recession if a consensus is formed among the members. A consensus can be seen as a complicated but nevertheless identifiable event the likelihood of which can be readily obtained from an underlying macroeconomic model. For example, if a consensus means the majority of the committee members believe there is a recession and if there were two members each defining recession with regard to events *A*, *B* and *C* above, then the NBER will announce a recession if at least two of the three events is expected to occur. The forecast of the probability of this joint event will summarise the committee's interpretation of the underlying events and this could be converted to the simple dichotomous indicator by announcing recession when this probability is greater than 0.5, say, or using some other threshold or sequence of probabilities to introduce smoothness. The discussion above should make it clear that this sort of summary statistic loses considerable information, however, and is not our preferred approach.

3. Characterising US recessions in real-time

In this section, we apply the methods described above to provide a picture of US recessions since the mid-1980s as they would have been experienced in real time. Our intention is to describe the various dimensions of recession by looking at a number of events that might be of interest to different individuals. To emphasise the idea that the recessionary events will be important in real-time decision-making, the focus of the discussion is on estimated nowcasts of the probabilities that the recessionary events occurred in each quarter based only on information that was available at the time. The analysis is particularly pertinent at the time of writing at the end of 2008 given the turmoil in the world's financial markets and the widespread anxieties about recession in the US economy.[7] We also provide a more detailed description of recessionary event probabilities for the current period therefore.

The real time dataset we use is obtained from the Federal Reserve Bank of Philadelphia at www.phil.frb.org/econ/forecast/ and consists of 161 quarterly vintages of data; the first was released in 1965q1 and the final vintage is dated 2008q4. For US aggregate output, data on real GDP in quarter t is released for the first time at the end of the first month of quarter $t+1$. This figure is reported in the Federal Reserve Bank of Philadelphia's real time data set as the mid-point of the $(t+1)$th quarter and it is denoted by ${}_{t+1}y_t$, where y_t is the logarithm of real GDP. In contrast to the illustrative model of the previous section, the empirical model accommodates the possibility of up to four revisions in the output data. Revisions that take place in output measures in the months up to the mid-point of the $(t+2)$th quarter are given by ${}_{t+2}y_t$. Likewise, ${}_{t+3}y_t$ incorporates any revisions that are then made up to the mid-point of the $(t+3)$th quarter, and so on.

In order to capture US macro-dynamics as accurately as possible, our empirical analysis considers interest rates, money and price measures in addition to output data. In this analysis, p_{t-1} refers to the *average* value of the (logarithm of) the consumer price index (CPI) over the three months of quarter $t-1$. The observation for prices in the third month of quarter $t-1$ is not released until the end of the first month of quarter t and so, matching the timing of the release of the output data, we take each quarter's price observation to be released at the mid-point of the succeeding quarter, denoted ${}_tp_{t-1}$ The timing of the release of data on the M1 measure of the money supply is exactly the same and so ${}_tm_{t-1}$ also refers to the average of the data relating to the three months of quarter $t-1$ released for the first time at the mid-point of quarter t. Our measure of the quarterly interest rate, ${}_tr_t$, is the Federal Funds rate as observed at the beginning of January, April, July, and October, i.e. the interest rate holding on the first day of the relevant quarter.

[7] The first version of the paper was written at the end of 2008 and presented at the Reserve Bank of New Zealand's Conference on Nowcasting and Model Uncertainty in December 2008. While subsequent revisions to the paper have been made during the journal review process, the bulk of the empirical work remains as it was undertaken in 2008q4 and the associated commentary is also dated at that time.

To investigate the informational content of 'forward-looking' variables, we make use of the interest rate spreads (to reflect market expectations of future rates) and experts' forecasts on output and prices as provided in the Federal Reserve Bank of Philadelphia's *Survey of Professional Forecasters* (SPF). The spread is denoted ${}_tsp_t$ and is defined as the difference between the three-month Treasury Bill Secondary Market Rate and the market yield on US Treasury securities. Forecasts taken from the SPF are made around the mid-point of quarter t. The nowcasts relating to quarter t's output and price level are denoted by ${}_ty_t^f$ and ${}_tp_t^f$, and the forecasts of quarter $t+s$'s output and price level, $s>0$, are denoted by ${}_ty_{t+s}^f$ and ${}_tp_{t+s}^f$, respectively.

Our empirical model specification for producing forecasts is a simple VAR comparable to (2.6), using

$$\boldsymbol{z}_t = ({}_tr_t, ({}_ty_{t-1}-{}_{t-1}y_{t-2}), ({}_tp_{t-1}-{}_tp_{t-2}), ({}_tm_{t-1}-{}_tm_{t-2}), ({}_tp_t^f-{}_tp_{t-1}), ({}_ty_t^f-{}_ty_{t-1}), ({}_tp_{t+1}^f-{}_tp_t^f),$$

$$({}_ty_{t+1}^f-{}_ty_t^f), ({}_ty_{t+2}^f-{}_ty_{t+1}^f), ({}_ty_{t+3}^f-{}_ty_{t+2}^f), ({}_ty_{t+4}^f-{}_ty_{t+3}^f), {}_tsp_t, ({}_ty_{t-2}-{}_{t-1}y_{t-2}), ({}_ty_{t-3}-{}_{t-1}y_{t-3}),)',$$

for $t=1, \ldots, T$, although the model used in the empirical work was of order two rather than one as in (2.6). The model therefore explains, simultaneously, the growth in first-release output data, the nowcast and expected one-period-ahead output growth, and two revisions in output data. In addition, it incorporates first-release data on interest rates, inflation and money growth plus nowcasts and expected one-, two-, three- and four-period ahead inflation and expected future interest rates.

Details of the estimated model are provided in Lee et al. (2009), including a thorough analysis of the statistical importance of the revision data and of the forward-looking variables included in this, our preferred, specification. Among other findings, this analysis shows that there is systematic content in output revisions for up to two quarters, so that ${}_{t+3}y_t$ represents the 'final', post-revision measure of output at time t. Probabilistic statements on the likelihood of events of interest that might occur today typically revolve around forecasts of the revised output measures that will be released in three periods time therefore. Along with the other variables in the model, previous revisions data and expectations of future output movements are shown to have considerable power to explain and forecast the post-revision output series. To use the model in a real-time analysis, it was estimated first using data for the period $t=1969\text{q}1, \ldots, 1986\text{q}1$ and this was used to produce nowcasts and forecasts relating to events in 1986q1. The model was then reestimated using data for the period $t=1969\text{q}1, \ldots, 1986\text{q}2$ and the nowcasts and forecasts for 1986q2 produced, and so on. The event probabilities only make use of data and models that were available at the time therefore.

3.1. US recessionary event probabilities, 1986q1–2008q4

3.1.1. Recession defined using growth thresholds

Fig. 1a provides the nowcast probabilities of two periods of consecutive negative output growth having occurred in T. Bearing in mind that output data is released with a one quarter lag and is then subject to systematic revision for a further two periods, the empirical counter-part to event A described above is actually $pr\{[(\widehat{{}_{T+3}y_T - {}_{T+2}y_{T-1}}) < 0] \cap [(\widehat{{}_{T+2}y_{T-1} - {}_{T+1}y_{T-2}}) < 0]\}$ based on data available at time T. The figure plots these probabilities over the period 1986q1–2008q4 based on the relevant recursively estimated version of our VAR model. The figure also plots (using shading) where we now know that two periods of consecutive negative growth actually occurred based on final, post-revision data. This event occurs just twice in the period for which we have final post-revision data, namely 1991q1 and 2001q4, and these dates coincide with the only two occasions on which the nowcast probability of the event is greater then 50%. This very strong coincidence of prediction and realisation is reflected by the corresponding Pesaran and Timmermann (1992) [PT] test statistic which takes a value of 9.49 when using this 50% figure as the threshold for predicting recession to occur. The PT statistic reflects the correspondence between prediction and realisation and has a standard normal distribution when there is no association. The observed value provides very strong evidence to support the use of the model if it is evaluated on the basis of predictions using this particular recession event probability therefore. However, Fig. 1a also shows that there were a number of other periods in which there

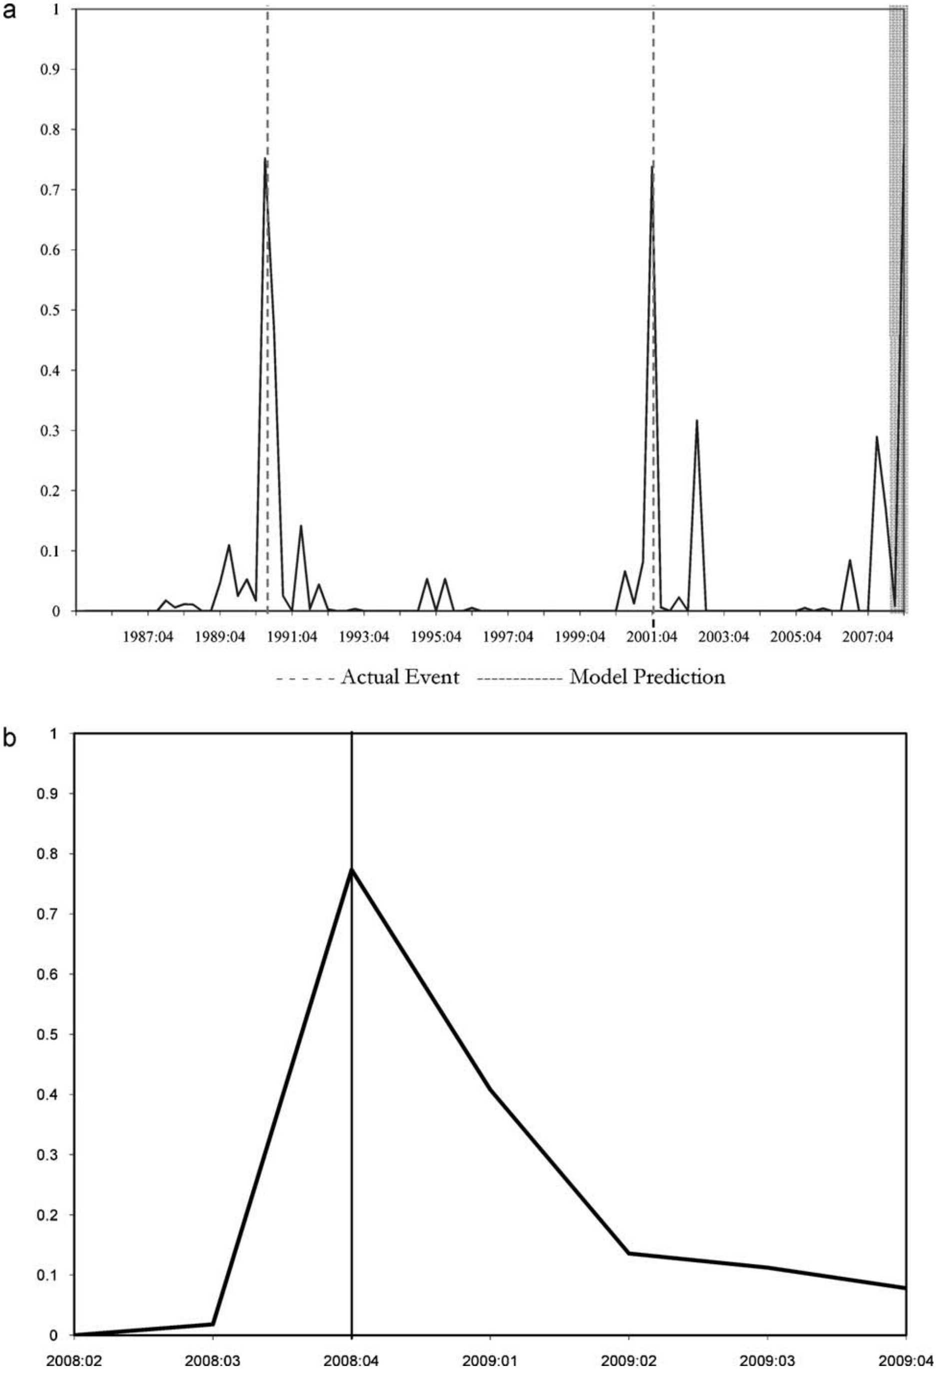

Fig. 1. (a) “Nowcast” probabilities of two periods of consecutive negative growth. (b) Forecast probabilities of two periods of consecutive negative growth; based on information available in 2008q4.

was reasonable possibility (>20%) of the event occurring and individual decision-making, reflecting expected loss minimisation of the form in (2.4), would have sensibly taken these probability forecasts into account as well.

Fig. 1a includes a shaded area covering the interval 2008q2–2008q4 indicating that, at the time of writing at the end of 2008, the final post-revision data for this period is not yet available. However, on the basis of the available data, the nowcast probability of the event occurring in 2008q4 is 78% so it seems very likely that, by this definition, the US is in recession. Fig. 1b elaborates on the current position by plotting the forecast probability of two consecutive periods of negative growth over the coming two years based on the data available at 2008q4. This indicates that, by this definition at least, the recession will be reasonably short-lived. The recession probability remains high in 2009q1 but falls quite rapidly thereafter to around 10% for the remainder of 2009. Of course, this analysis is based on the properties of the underlying estimated model and if the current position is unprecedented and renders past experience uninformative on the future, as some commentators believe, then these probability forecasts are unreliable.[8]

But the model's diagnostics suggest that the model performs well in capturing the US macroeconomic dynamics over the last forty years and it incorporates expert opinion and market information on what is likely to happen to output, prices and interest rates. Based on the data available at the end of 2008, the model suggests that recession defined in this way is unlikely to last beyond the end of 2009 and this is potentially useful information for those for whom the experience of two periods of negative growth would impact on decision-making.

Fig. 2a illustrates the likelihood of another recession event based on growth thresholds but elaborated to match with turning point analysis. Here we note, following Harding and Pagan (2005), that a peak in output at time T is nowcast to occur when

$$(\widehat{{}_{T+3}y_T} - \widehat{{}_{T+2}y_{T-1}}) > 0; \quad (\widehat{{}_{T+3}y_T} - \widehat{{}_{T+2}y_{T-1}}) - (\widehat{{}_{T+2}y_{T-1}} - \widehat{{}_{T+1}y_{T-2}}) > 0$$
$$(\widehat{{}_{T+4}y_{T+1}} - \widehat{{}_{T+3}y_T}) < 0; \quad (\widehat{{}_{T+5}y_{T+2}} - \widehat{{}_{T+4}y_{T+1}}) - (\widehat{{}_{T+4}y_{T+1}} - \widehat{{}_{T+3}y_T}) < 0$$

and a corresponding definition holds for a trough. A period of recession can be defined as the interval starting one period after a peak and ending in the period of a trough; i.e. there is recession in period T if there is a peak at time $T-s$, for some $s = 1, 2, \ldots$, and no trough has occurred subsequently. This definition of a recession is based on a complicated function of output outcomes over various periods but the likelihood of it happening can be readily computed using the simulation methods described above and the recursively estimated nowcast probabilities are shown in Fig. 2a. The years 1991 and 2001 are reasonably unambiguously identified as periods of recession by this definition, reflected by real-time probabilities in excess of 60%, although the very high nowcast probabilities of the former recession pre-dated the actual occurrence (as identified by application of the algorithm to the post-revision data) by one or two quarters. The PT statistic takes the value of 1.68 in this case, again using the 50% threshold to predict the event will occur, showing a positive relation but one that is only significant at the 10% significance level. There are also reasonably high probabilities (in excess of 20%) for much of the sample indicating that a concern over this aspect of recession could impact on decision-making in nearly all periods (not just the eight quarters ultimately identified by the dating algorithm applied to the final series) if the individual's loss function is sensitive to this event. Of course, in 2008q4, it is not possible to determine whether recession has actually occurred beyond 2007q2, given the 3-quarter delay in observing post-revision data and the need for observations on output levels two periods into the future to implement the dating algorithm. But the nowcast probabilities for 2008q3 and 2008q4 in Fig. 2a and the forecasts for 2009 in Fig. 2b indicate that recession is extremely likely at the end of 2008 but fall off quite quickly through 2009, matching the pattern shown in Fig. 1.

[8] It is relatively straightforward to extend the analysis presented here to accommodate model uncertainty which would allow for an alternative data generating process to be considered too. But this could not accommodate models with no precedent at all.

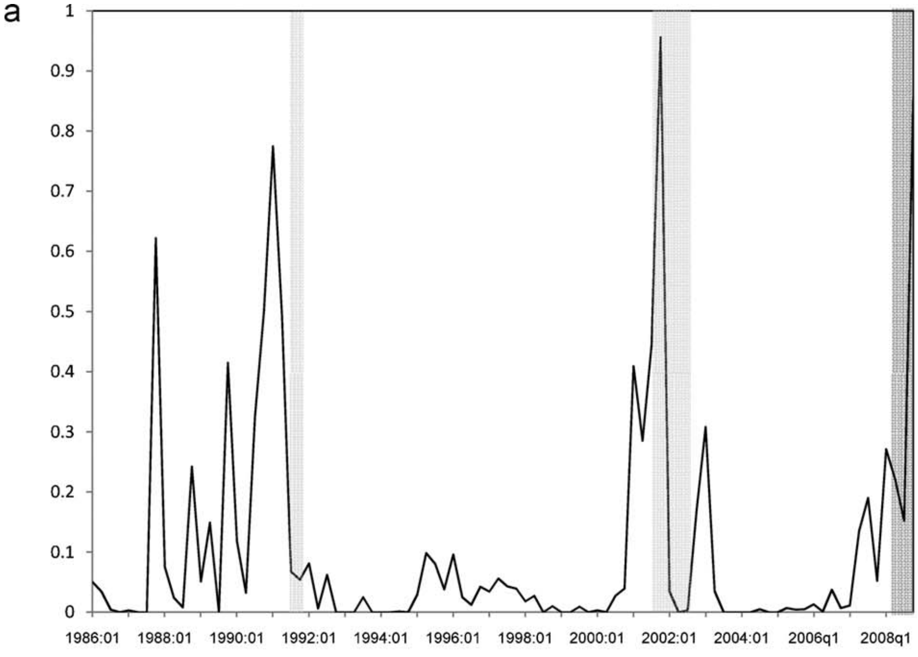

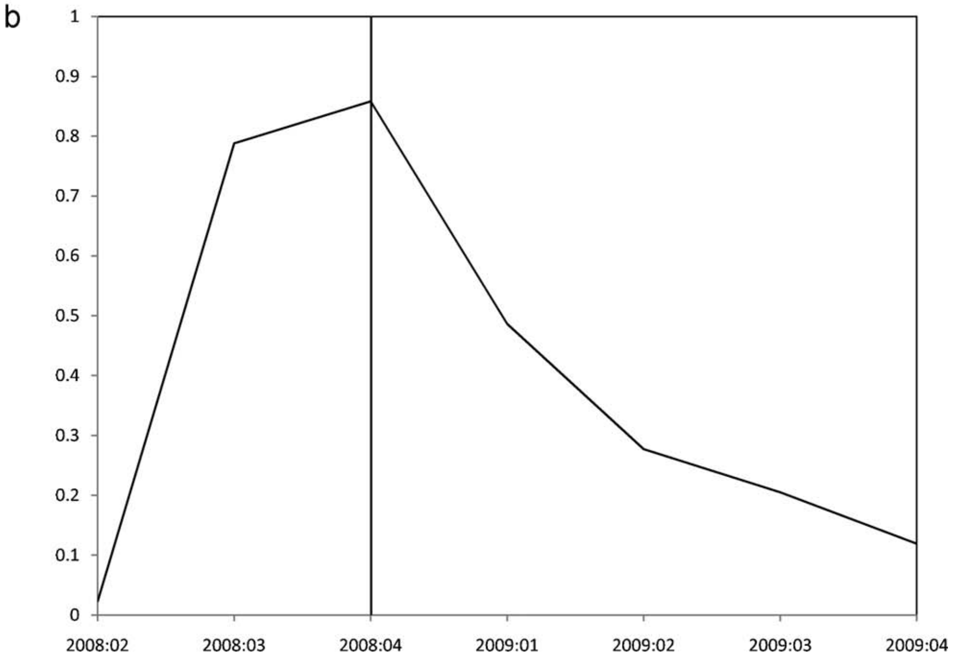

Fig. 2. (a) Nowcast probabilities of a recession based on turning points. (b) Forecast probabilities of a recession based on turning points based on information available in 2008q4.

3.1.2. Recession defined using output levels

As noted earlier, a definition of recession likely to be of more interest to those concerned with absolute levels of income might be one based on whether the output level is lower than its previous peak. A third definition of recession in time T considered in this empirical exercise, corresponding to event B discussed above, is where $\{\widehat{_{T+3}y_T} < \max(\widehat{_{T+2}y_{T-1}}, \widehat{_{T+1}y_{T-2}}, {_T}y_{T-3}, \ldots)\}$ therefore. Fig. 3a plots these probabilities, calculated in real-time, showing that, taking the period as a whole, this recessionary event generally is considered much more likely to occur than either of the events described above. The highest probabilities (in excess of 80%) broadly coincide with the periods of high probability in Figs. 1a and 2a and the PT statistic takes a value of 4.05 in this case, but reasonably-sized probabilities (in excess of 20%) are also found over a substantial part of the sample. This means that, for those who care about absolute income levels, the possibility of this type of recession will again impact on decision-making for much of the time. This observation carries over to the plots of Fig. 3b too which conveys a more pessimistic view on the likelihood of recession over the coming year. The nowcast probability of recession is virtually one in 2008q4, but the probability remains greater than 50% throughout 2009 and there is a 20% chance that the economy has not returned to its previous peak by the end of 2009.

There may be some agents who are concerned with both growth rates and absolute income levels and a fourth definition of recession in time T might therefore be defined by the event $\{N_T < 0\}$ where

$$N_T = -0.65 - 158.37(\widehat{_{T+3}y_T} - \widehat{_{T+2}y_{T-1}}) - 58.86(\widehat{_{T+2}y_{T-1}} - \widehat{_{T+1}y_{T-2}}) + 0.03\,ICDR_T + 1.15 ICDR_{T-1} + e_t, \tag{3.7}$$

and $ICDR_t$ is an indicator variable taking the value one when the event $\{\widehat{_{T+3}y_T} < \max(\widehat{_{T+2}y_{T-1}}, \widehat{_{T+1}y_{T-2}}, {_T}y_{T-3}, \ldots)\}$ occurs and zero otherwise. This appears to be a relatively arbitrary event at first sight but, as explained in detail in Lee et al. (2009), it actually represents the outcome of a Probit analysis of the NBER announcements of recession. While we have been keen in this paper to stress that recession is best seen as a multifaceted phenomenon, it is interesting to look for an event that corresponds to the NBER announcements for the purpose of straight comparison and also to try to establish the uncertainty associated with the NBER announcements if they had been made in real time. Fig. 4a plots the probabilities of $\{N_T < 0\}$ along with shadings to show the NBER's actual judgements. The heavy shading over the period 2008q1–2008q4 once more acknowledges that, at the time of writing, we do not know what the NBER will say about most of 2008.

The figure shows that the highest nowcast probabilities (in excess of 80%) do coincide with the subsequent NBER announcements of recession in 1991, 2001 and the announcement at the start of 2008 is matched with a nowcast probability of nearly 60%. This alignment of prediction and realisation is reflected in a PT statistic value of 5.82. But the plots show too that the event $\{N_T < 0\}$ had probabilities in excess of 20% for times in 1990, 1995 and 2007, showing that the dichotomous NBER statements abstract (in retrospect) from anxieties that might be important in decision-making for some individuals. These reasonably high probabilities are also reflected in Fig. 4b which show probabilities in excess of 50% to mid-2009 but remain above 20% throughout the year.

3.1.3. Recession defined using the output gap

This section concludes with a final set of probability forecasts relating to a recessionary event defined using the output gap. Fig. 5a describes the probability of the event $\{\widehat{_{T+3}y_T} < \tilde{y}_T\}$ where the trend $\tilde{y}_T$ is measured by fitting the Hodrick-Prescott (HP) filter to the series comprising the final post-revision data on output to $T-4$ augmented by the model-based forecasts of $\widehat{_{T+h}y_{T+h-3}}$ for $h = 1, 2, \ldots$. There is widespread use of HP (and other) filters in defining gaps and Orphanides and van Norden (2002) showed that these are typically very vulnerable to real-time analysis. However, Garratt, Lee, Mise, and Shields (2008) showed that the method described above, applying smoothing techniques to forecast-augmented series, can considerably improve the precision of the estimates of an output gap at the end-of-sample. This is, of course, extremely important in real-time decision-making. The HP filter underlying the probabilities shown in Fig. 5a is again a complicated function of information

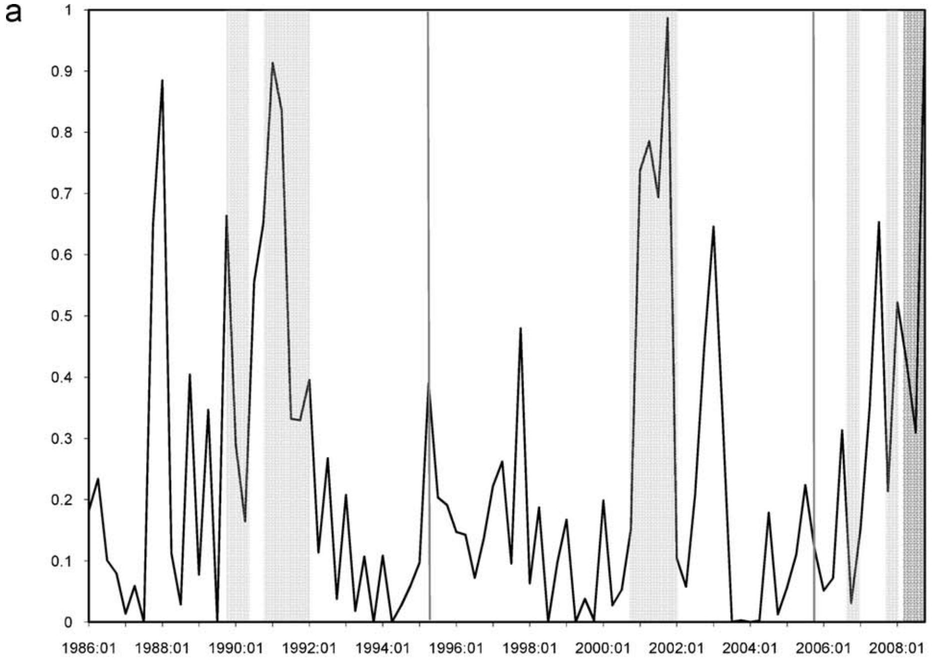

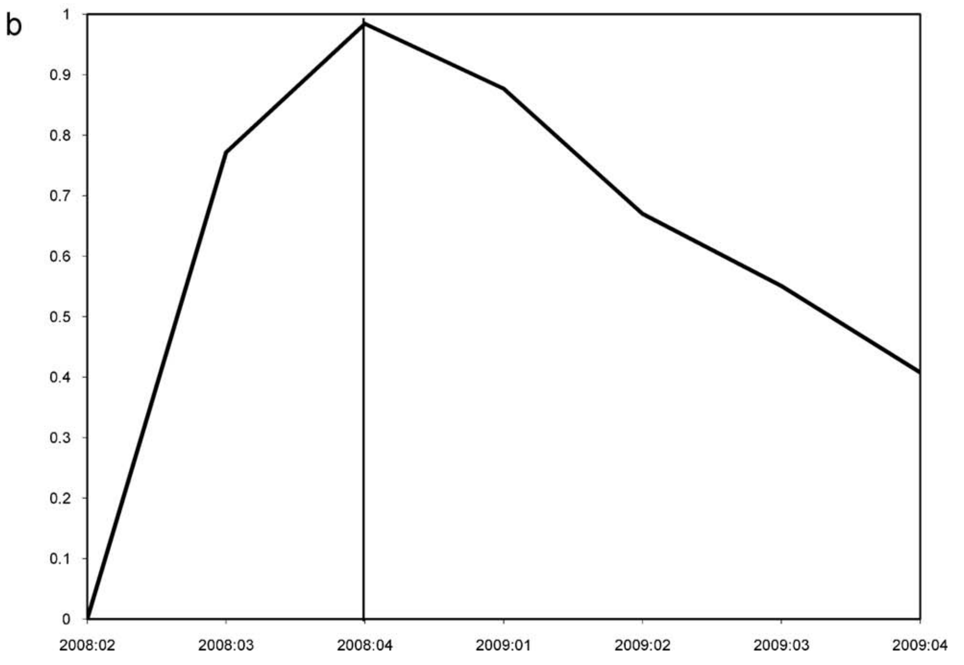

Fig. 3. (a) Nowcast probabilities of "Current Depth of Recession". (b) Forecast probabilities of "Current Depth of Recession" based on information available in 2008q4.

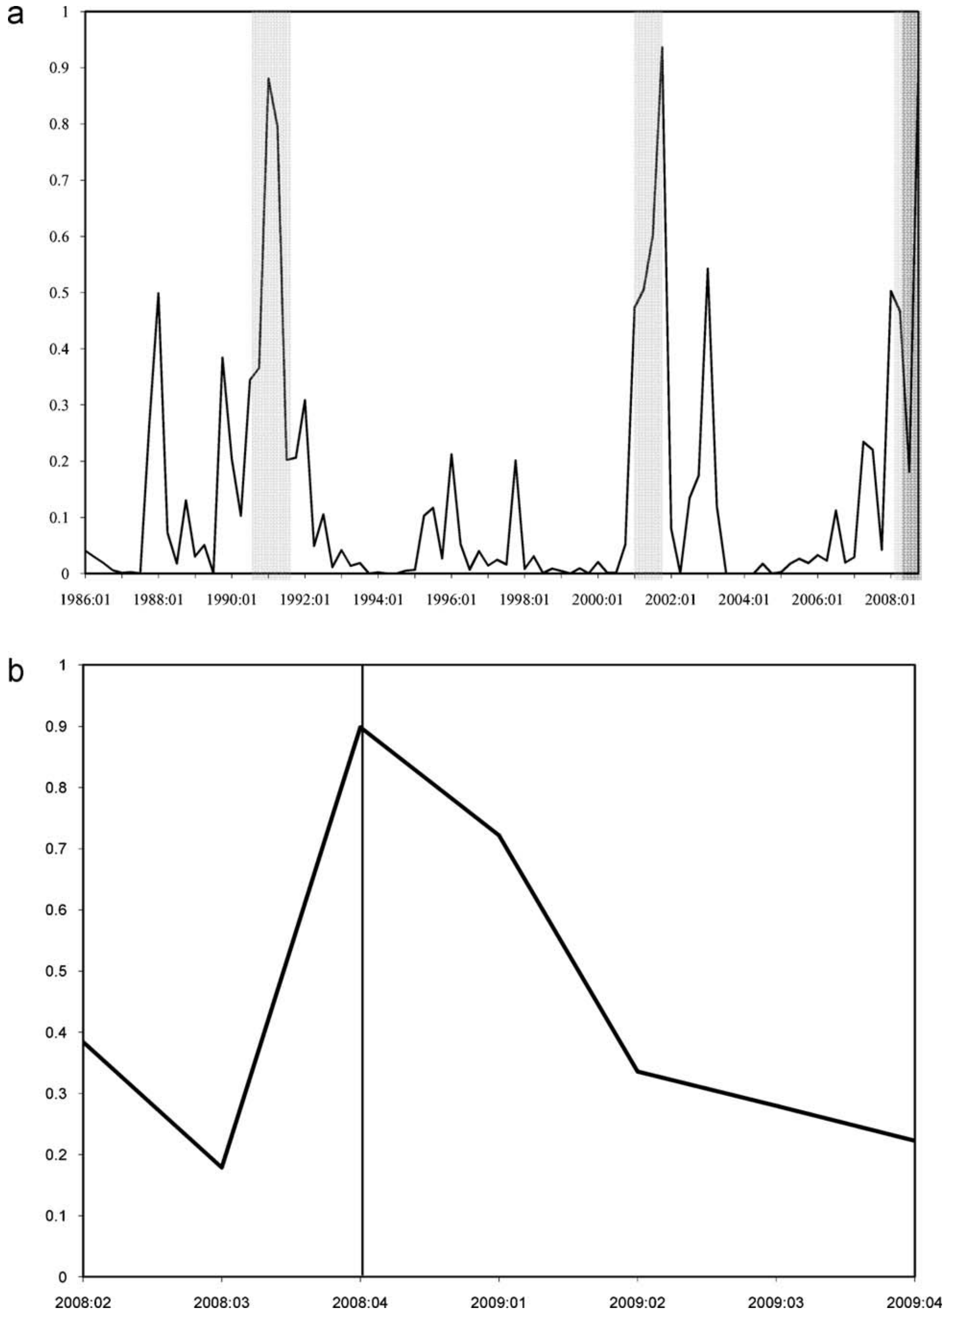

Fig. 4. (a) "Nowcast" probabilities of NBER periods of contraction. (b) Forecast probabilities of NBER periods of contraction based on information available in 2008q4.

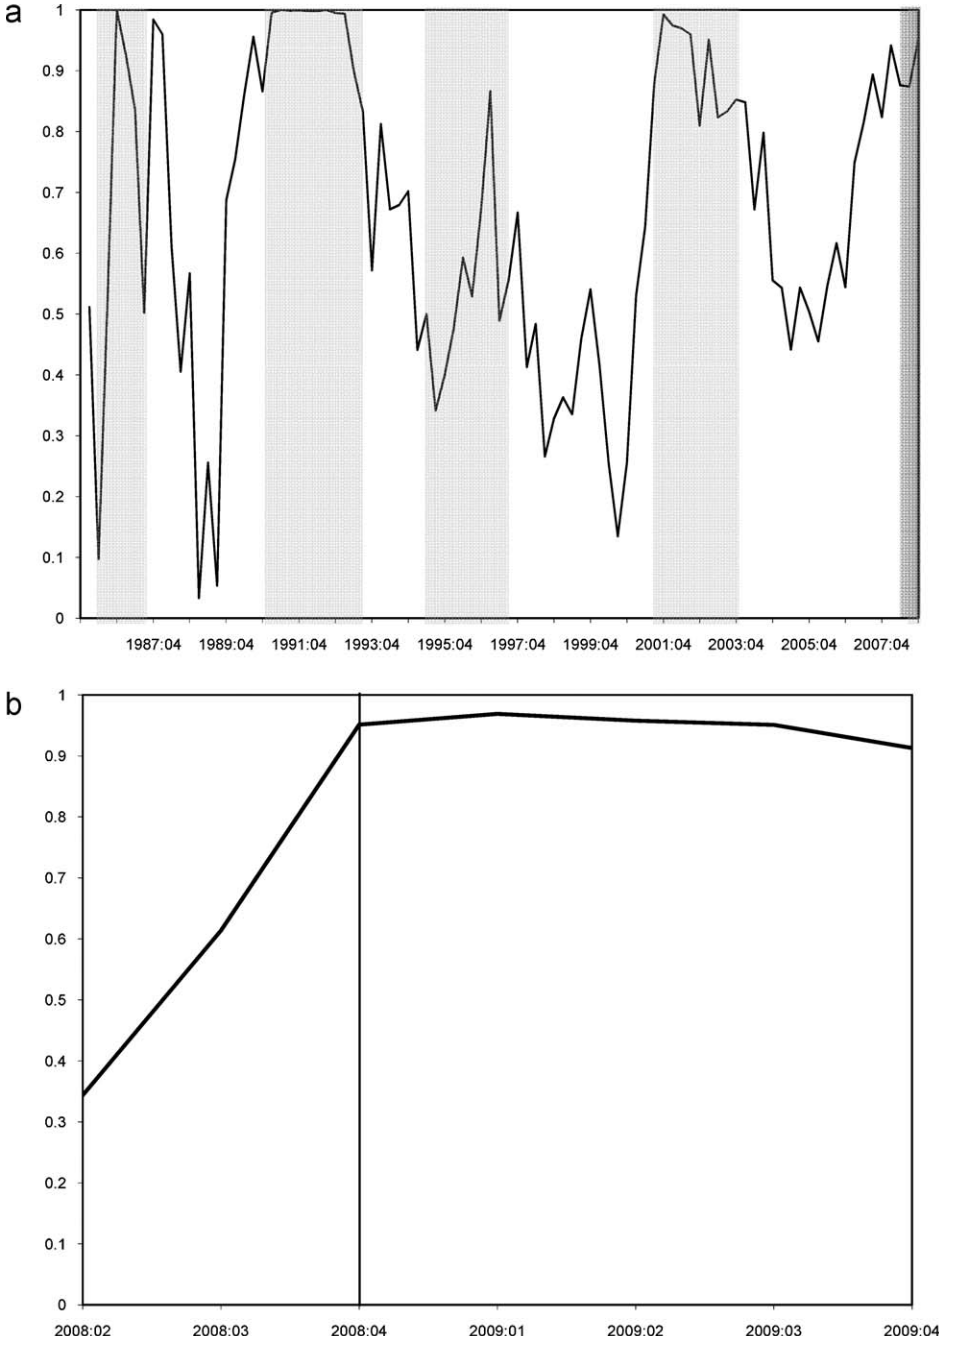

Fig. 5. (a) Nowcast probabilities of a negative output gap: forecast-augmented Hodrick–Prescott filter. (b) Forecast probabilities of a negative output gap based on information available in 2008q4: Hodrick–Prescott filter.

known at time *T* therefore but one for which probabilistic statements can be calculated readily using our proposed simulation methods.[9]

The figure itself provides quite a different perspective on recession than that captured by the previous plots. Particularly high probabilities (in excess of 80%) are observed around 1991/92 and 2001/02 and most recently, broadly matching the occurrence of the recessionary events discussed in relation to previous figures. But output was observed to be below trend for considerable intervals through the sample, and this is reflected by probabilities in excess of 50% for large part of the sample outside the three intervals highlighted in the earlier plots. The periods of gap-based recession that were actually observed come in relatively distinct intervals of two years or so, reflecting the smooth nature of the underlying trend, and the nowcast probabilities also display this sort of pattern. The correspondence between the periods of high probabilities and the actual occurrence of recession is not as clear-cut as in some of the previous plots, although the PT test statistic takes a value of 4.02 showing a high degree of consensus between prediction and realisation. Further, based on the data available at the end of 2008 and looking to the outlook over the coming year, it is very clear from Fig. 5b that negative gaps are likely to continue for some time, with forecast probabilities in excess of 90% throughout 2009.

4. Concluding comments

The figures of the previous section describe various aspects of recession. They build a sophisticated picture of the decision-making context faced by individuals in real-time conveying both the general macroeconomic prospects and the extent to which these might translate to events of specific interest to different individuals. The general macroeconomic prospects are reflected by common patterns in the nowcast probabilities, with all figures showing high probabilities of the various recessionary events occurring around 1991, 2001 and at the end of the sample in 2008. This reflects the general deterioration in macroeconomic activity at these times although the precise timing of recession even at these times is certainly not the same for all of the events. Moreover, the probabilities also describe the degree of conviction with which the recessionary events are nowcast to occur, showing that for most events, there remains a reasonable possibility (probability greater than 20%) of one or more recessionary events occurring at almost all times. Dichotomous statements on the likely occurrence of events based on thresholds are not affected by the continuum in the probability measures. But the decision-making of optimising agents should, and would, take this into account and the strength of the tests of the forecast performance of the model provided in the empirical section shows that the recessionary-event probability forecasts would be very useful in real-time decision-making.

The results of the paper show that, far from being a straightforward dichotomous event, recession is a complicated multifaceted phenomenon which will impact on the decision-making of different individuals in different ways in most time periods. The use of nowcast probabilities of the various recessionary events provides a useful means of characterising these various facets and demonstrates both the sophistication necessary to answer the question "when do we know we are in recession?" and the means of providing an answer.

References

Aruoba, S. B., Diebold, F. X., & Scotti, C. (2009). Real time measurement of business conditions. *Journal of Business and Economic Statistics*, *27*(4), 417–427.

Beaudry, P., & Koop, G. (1993). Do recessions permanently change output? *Journal of Monetary Economics*, *31*, 149–163.

Chauvet, M., & Hamilton, J. D. (2006). Dating business cycle turning points. In D. van Dijk, C. Milas, P. A. Rothman, & D. E. Wildasin (Eds.), *Nonlinear time series analysis of business cycles* (pp. 1–54). Elsevier's Contributions to Economic Analysis Series.

Chauvet, M., & Piger, J. (2003). Identifying business cycle turning points in real time. In *Federal Reserve Bank of Atlanta Working Paper* (pp. 47–62).

[9] The gap based on the forecast-augmented HP filter is, of course, only one of many alternative measures. Probabilistic statements on recession can be made for any of the available alternatives and indeed an overarching recession probability forecast can be readily obtained by aggregating across these alternative probabilities; see Garratt et al. (2008), Garratt, Lee, Mise, and Shields (2009) for discussion.

Clements, M. P., & Smith, J. (2000). Evaluating the forecast densities of linear and nonlinear models. *Journal of Forecasting*, *19*, 255–276.
Diebold, F. X., Gunther, T. A., & Tay, A. S. (1998). Evaluating density forecasts with applications to financial risk management. *International Economic Review*, *39*, 863–884.
Diebold, F. X., Hahn, J., & Tay, A. S. (1999). Multivariate density forecast evaluation and calibration in financial risk management: High-frequency returns on foreign exchange. *Review of Economics and Statistics*, *81*, 661–673.
Garratt, A., and Lee, K. (2006). Investment Under Model Uncertainty: Decision Based Evaluation of Exchange Rate Forecasts in the US, UK and Japan. Birkbeck Working Papers in Economics and Finance, 0616.
Garratt, A., Lee, K., Mise, E., & Shields, K. (2008). Real time representations of the output gap. *Review of Economics and Statistics*, *90*(4), 792–804.
Garratt, A., Lee, K., Mise, E., & Shields, K. (2009). Real time representations of the UK output gap in the presence of model uncertainty. *International Journal of Forecasting*, *25*, 81–102.
Garratt, A., Lee, K., Pesaran, M. H., & Shin, Y. (2003). Forecast uncertainty in macroeconometric modelling: An application to the UK economy. *Journal of the American Statistical Association*, *98*(464), 829–838.
Garratt, A., Lee, K., Pesaran, M. H., & Shin, Y. (2006). *National and global macro-econometric modelling: A long-run structural modelling approach*. Oxford University Press.
Hamilton, J. (1989). A new approach to the economic analysis of non-stationary time series and the business cycle. *Econometrica*, *57*, 357–384.
Harding, D., & Pagan, A. (2002). Dissecting the cycle: a methodological investigation. *Journal of Monetary Economics*, *49*(2), 365–381.
Harding, D., & Pagan, A. (2005). A suggested framework for classifying the modes of business cycle research. *Journal of Applied Econometrics*, *20*(2), 151–161.
Harding, D., & Pagan, A. (2006). Synchronisation of cycles. *Journal of Econometrics*, *132*(1), 59–79.
Leamer, E. E. (2008). *What's a recession, anyway?* NBER Working Paper Series, No. 14221.
Lee, K. C., Olekalns, N., Shields, K. (2009), *Nowcasting, business cycle dating and the interpretation of new information when real time data are available*. University of Leicester Working Paper, July 2009.
Orphanides, A., Porter, R. D., Refschneider, D., Tetlow, R., & Finan, F. (2000). Errors in the measurement of the output gap and the design of monetary policy. *Journal of Economics and Business*, *52*, 117–141.
Orphanides, A., & van Norden, S. (2002). The unreliability of output gap estimates in real time. *Review of Economics and Statistics*, *84*(4), 569–583.
Pesaran, M. H., & Skouras, S. (2002). Decision-based methods for forecast evaluation. In M. P. Clements, & D. F. Hendry (Eds.), *A companion to economic forecasting* (pp. 241–267). Oxford: Basil Blackwell.
Pesaran, M. H., & Timmermann, A. (1992). A simple nonparametric test of predictive performance. *Journal of Business and Economic Statistics*, *10*, 461–465.
van Dijk, D., van Dijk, H., & Frances, P. H. (2005). On the dynamics of business cycle analysis: Editors' introduction. *Journal of Applied Econometrics*, *20*(2), 147–150.
Woodford, M. (2001). The Taylor rule and optimal monetary policy. *American Economic Review*, *91*(2), 232–237.
Woodford, M. (2003). *Interest and prices: foundations of a theory of monetary policy*. Princeton University Press.