

Improving U.S. GDP Measurement: A Forecast Combination Perspective

S. Borağan Aruoba, Francis X. Diebold, Jeremy Nalewaik,
Frank Schorfheide and Dongho Song

“A growing number of economists say that the government should shift its approach to measuring growth. The current system emphasizes data on spending, but the bureau also collects data on income. In theory the two should match perfectly—a penny spent is a penny earned by someone else. But estimates of the two measures can diverge widely, particularly in the short term...”

[Binyamin Appelbaum, New York Times, August 16, 2011]

Abstract Two often-divergent U.S. GDP estimates are available, a widely-used expenditure-side version GDP_E , and a much less widely-used income-side version GDP_I . We propose and explore a “forecast combination” approach to combining them. We then put the theory to work, producing a superior combined estimate of GDP growth for the U.S., GDP_C . We compare GDP_C to GDP_E and GDP_I , with particular attention to behavior over the business cycle. We discuss several variations and extensions.

S. Borağan Aruoba (✉)
University of Maryland, College Park, MD, USA
e-mail: aruoba@econ.umd.edu

F. X. Diebold · F. Schorfheide · D. Song
University of Pennsylvania, Philadelphia, PA, USA
e-mail: fdiebold@sas.upenn.edu

F. Schorfheide
e-mail: schorf@ssc.upenn.edu

D. Song
e-mail: donghos@sas.upenn.edu

J. Nalewaik
Federal Reserve Board, Washington, DC, USA
e-mail: jeremy.j.nalewaik@frb.gov

1 Introduction

GDP growth is surely the most fundamental and important concept in empirical/applied macroeconomics and business cycle monitoring, yet significant uncertainty still surrounds its estimation. Two often-divergent estimates exist for the U.S., a widely-used expenditure-side version, GDP_E , and a much less widely-used income-side version, GDP_I . Nalewaik (2010) makes clear that, at the very least, GDP_I deserves serious attention and may even have properties in certain respects superior to those of GDP_E . That is, if forced to choose between GDP_E and GDP_I , a surprisingly strong case exists for GDP_I .

But of course one is *not* forced to choose between GDP_E and GDP_I , and a combined estimate that pools information in the two indicators GDP_E and GDP_I may improve on both. In this chapter, we propose and explore a method for constructing such a combined estimate, and we compare our new GDP_C (“combined”) series to GDP_E and GDP_I over many decades, with particular attention to behavior over the business cycle, emphasizing comparative behavior during turning points.

Our work is motivated by, and builds on, five key literatures. First, and most pleasing to us, our work is very much related to Hal White’s in its focus on dynamic modeling while acknowledging misspecification throughout.

Second, we obviously build on the literature examining GDP_I and its properties, notably Fixeler and Nalewaik (2009) and Nalewaik (2010). GDP_I turns out to have intriguingly good properties, suggesting that it might be usefully combined with GDP_E .

Third, our work is related to the literature distinguishing between “forecast error” and “measurement error” data revisions, as for example in Mankiw et al. (1984), Mankiw and Shapiro (1986), Faust et al. (2005), and Aruoba (2008). In this chapter we work largely in the forecast error tradition.

Fourth, and related, we work in the tradition of the forecast combination literature begun by Bates and Granger (1969), viewing GDP_E and GDP_I as forecasts of GDP [actually a mix of “backcasts” and “nowcasts” in the parlance of Aruoba and Diebold (2010)]. We combine those forecasts by forming optimally weighted averages.¹

Finally, we build on the literature on “balancing” the national income accounts, which extends back almost as far as national income accounting itself, as for example in Stone et al. (1942), who use a quadratic loss criterion to propose weighting different GDP estimates by the inverse of their squared “margins of error.” Stone refined those ideas in his subsequent national income accounting work, and Byron (1978) and Weale (1985) formalized and refined Stone’s approach. Indeed a number of papers by Weale and coauthors use subjective evaluations of the quality of different U.K. GDP estimates to produce combined estimates; see Barker et al. (1984), Weale

¹ For surveys of the forecast combination literature, see Diebold and Lopez (1996) and Timmermann (2006).

(1988), Solomou and Weale (1991), and Solomou and Weale (1993).² For example, Barker et al. (1984) and Weale (1988) incorporate data quality assessments from the U.K. Central Statistical Office. Weale also disaggregate some of their GDP estimates to incorporate information regarding differential quality of underlying source data. In that tradition, Beaulieu and Bartelsman (2004) use input–output tables to disaggregate GDP_E and GDP_I , using what they call “tuning” parameters to balance the accounts. We take a similar approach here, weighting competing GDP estimates in ways that reflect our assessment of their quality, but we employ more of a top-down, macro perspective.

We proceed as follows. In Sect. 2 we consider GDP combination under quadratic loss. This involves taking a stand on the values of certain unobservable parameters (or at least reasonable ranges for those parameters), but we argue that a “quasi-Bayesian” calibration procedure based on informed judgment is feasible, credible, and robust. In Sect. 3 we consider GDP combination under minimax loss. Interestingly, as we show, it does not require calibration. In Sect. 4 we apply our methods to provide improved GDP estimates for the U.S. In Sect. 5 we sketch several extensions, and we conclude in Sect. 6.

2 Combination Under Quadratic Loss

Optimal forecast combination typically requires knowledge (or, in practice, estimates) of forecast error properties such as variances and covariances. In the present context, we have two “forecasts,” of true GDP, namely GDP_E and GDP_I , but true GDP is never observed, even after the fact. Hence we never see the “forecast errors,” which complicates matters significantly but not hopelessly. In particular, in this section we work under quadratic loss and show that a quasi-Bayesian calibration based on informed judgment is feasible and credible, and simultaneously, that the efficacy of GDP combination is robust to the precise weights used.

2.1 Basic Results and Calibration

First assume that the errors in GDP_E and GDP_I growth are uncorrelated. Consider the convex combination³

$$GDP_C = \lambda GDP_E + (1 - \lambda) GDP_I,$$

² Weale also consider serial correlation and time-varying volatility in GDP measurement errors, as well as time-varying correlation between expenditure- and income-side GDP measurement errors.

³ Throughout this chapter, the variables GDP , GDP_E , and GDP_I that appear in the equations refer to growth rates.

where $\lambda \in [0, 1]$.⁴ Then the associated errors follow the same weighting,

$$e_C = \lambda e_E + (1 - \lambda)e_I,$$

where $e_C = \text{GDP} - \text{GDP}_C$, $e_E = \text{GDP} - \text{GDP}_E$ and $e_I = \text{GDP} - \text{GDP}_I$. Assume that both GDP_E and GDP_I are unbiased for GDP, in which case GDP_C is also unbiased, because the combining weights sum to unity.

Given the unbiasedness assumption, the minimum-MSE combining weights are just the minimum-variance weights. Immediately, using the assumed zero correlation between the errors,

$$\sigma_C^2 = \lambda^2 \sigma_E^2 + (1 - \lambda)^2 \sigma_I^2, \quad (1)$$

where $\sigma_C^2 = \text{var}(e_C)$, $\sigma_E^2 = \text{var}(e_E)$ and $\sigma_I^2 = \text{var}(e_I)$. Minimization with respect to λ yields the optimal combining weight,

$$\lambda^* = \frac{\sigma_I^2}{\sigma_I^2 + \sigma_E^2} = \frac{1}{1 + \phi^2}, \quad (2)$$

where $\phi = \sigma_E / \sigma_I$.

It is interesting and important to note that in the present context of zero correlation between the errors,

$$\text{var}(e_E) + \text{var}(e_I) = \text{var}(\text{GDP}_E - \text{GDP}_I). \quad (3)$$


The standard deviation of GDP_E minus GDP_I can be trivially estimated. Thus, an expression of a view about ϕ is in fact implicitly an expression of a view about not only the ratio of $\text{var}(e_E)$ and $\text{var}(e_I)$, but about their actual values. We will use this fact (and its generalization in the case of correlated errors) in several places in what follows.

Based on our judgment regarding U.S. GDP_E and GDP_I data, which we will subsequently discuss in detail in Sect. 2.2, we believe that a reasonable range for ϕ is $\phi \in [0.75, 1.45]$, with midpoint 1.10.⁵ One could think of this as a quasi-Bayesian statement that prior beliefs regarding ϕ are centered at 1.10, with a 90% prior credible interval of $[0.75, 1.45]$. In Fig. 1 we graph λ^* as a function of ϕ , for $\phi \in [0.75, 1.45]$. λ^* is of course decreasing in ϕ , but interestingly, it is only mildly sensitive to ϕ . Indeed, for our range of ϕ values, the optimal combining weight remains close to 0.5, varying from roughly 0.65 to 0.30. At the midpoint $\phi = 1.10$, we have $\lambda^* = 0.45$.

⁴ Strictly speaking, we need not even impose $\lambda \in [0, 1]$, but $\lambda \notin [0, 1]$ would be highly nonstandard for two valuable and sophisticated GDP estimates such as GDP_E and GDP_I . Moreover, as we shall see subsequently, multiple perspectives suggest that for our application the interesting range of λ is well in the interior of the unit interval.

⁵ Invoking Eq. (3), we see that the midpoint 1.10 corresponds to $\sigma_I = 1.30$ and $\sigma_E = 1.43$, given our estimate of $\text{std}(\text{GDP}_E - \text{GDP}_I) = 1.93\%$ using data 1947Q2-2009Q3.

Fig. 1 λ^* versus ϕ . λ^* constructed assuming uncorrelated errors. The *horizontal line* for visual reference is at $\lambda^* = 0.5$. See text for details



It is instructive to compare the error variance of combined GDP, σ_C^2 , to σ_E^2 for a range of λ values (including $\lambda = \lambda^*$, $\lambda = 0$, and $\lambda = 1$).⁶ From (1) we have:

$$\frac{\sigma_C^2}{\sigma_E^2} = \lambda^2 + \frac{(1 - \lambda)^2}{\phi^2}.$$

In Fig. 2 we graph σ_C^2/σ_E^2 for $\lambda \in [0, 1]$ with $\phi = 1.1$. Obviously the maximum variance reduction is obtained using $\lambda^* = 0.45$, but even for nonoptimal λ , such as simple equal-weight combination ($\lambda = 0.5$), we achieve substantial variance reduction relative to using GDP_E alone. Indeed, a key result is that *for all* λ (except those very close to 1, of course) we achieve substantial variance reduction.


Now consider the more general and empirically-relevant case of correlated errors. Under the same conditions as earlier,

$$\sigma_C^2 = \lambda^2\sigma_E^2 + (1 - \lambda)^2\sigma_I^2 + 2\lambda(1 - \lambda)\sigma_{EI}, \tag{4}$$

so

⁶ We choose to examine σ_C^2 relative to σ_E^2 , rather than to σ_I^2 , because GDP_E is the “standard” GDP estimate used in practice almost universally. A graph of σ_C^2/σ_I^2 would be qualitatively identical, but the drop below 1.0 would be less extreme.

Fig. 2 σ_C^2/σ_E^2 for $\lambda \in [0, 1]$. We assume $\phi = 1.1$ and uncorrelated errors. See text for details



$$\begin{aligned} \lambda^* &= \frac{\sigma_I^2 - \sigma_{EI}}{\sigma_I^2 + \sigma_E^2 - 2\sigma_{EI}} \\ &= \frac{1 - \phi\rho}{1 + \phi^2 - 2\phi\rho}, \end{aligned}$$

where $\sigma_{EI} = \text{cov}(e_E, e_I)$ and $\rho = \text{corr}(e_E, e_I)$.

It is noteworthy that—in parallel to the uncorrelated-error case in which beliefs about ϕ map one-for-one into beliefs about σ_E and σ_I —beliefs about ϕ and ρ now similarly map one-for-one into beliefs about σ_E and σ_I . Our definitions of σ_E^2 and σ_I^2 imply that

$$\sigma_j^2 = \text{var}[\text{GDP}_j] - 2\text{cov}[\text{GDP}_j, \text{GDP}] + \text{var}[\text{GDP}], \quad j \in \{E, I\}. \quad (5)$$

Moreover, the covariance between the GDP_E and GDP_I errors can be expressed as

$$\sigma_{EI} = \text{cov}[\text{GDP}_E, \text{GDP}_I] - \text{cov}[\text{GDP}_E, \text{GDP}] - \text{cov}[\text{GDP}_I, \text{GDP}] + \text{var}[\text{GDP}]. \quad (6)$$

Solving (5) for $\text{cov}[\text{GDP}_j, \text{GDP}]$ and inserting the resulting expressions for $j \in \{E, I\}$ into (6) yields

$$\sigma_{EI} = \text{cov}[\text{GDP}_I, \text{GDP}_E] - \frac{1}{2} \left(\text{var}[\text{GDP}_I] + \text{var}[\text{GDP}_E] - \sigma_I^2 - \sigma_E^2 \right). \quad (7)$$

Finally, let $\sigma_{EI} = \rho\sigma_E\sigma_I$ and $\sigma_E^2 = \phi^2\sigma_I^2$. Then we can solve (7) for σ_I^2 :




Fig. 3 λ^* versus ϕ for various ρ values. The horizontal line for visual reference is at $\lambda^* = 0.5$. See text for details

$$\sigma_I^2 = \frac{\text{cov}[\text{GDP}_I, \text{GDP}_E] - \frac{1}{2}(\text{var}[\text{GDP}_I] + \text{var}[\text{GDP}_E])}{\rho\phi - \frac{1}{2}(1 + \phi^2)} = \frac{N}{D}. \quad (8)$$

For given values of ϕ and ρ we can immediately evaluate the denominator D in (8), and using data-based estimates of $\text{cov}[\text{GDP}_I, \text{GDP}_E]$, $\text{var}[\text{GDP}_I]$, and $\text{var}[\text{GDP}_E]$ we can evaluate the numerator N .

Based on our judgment regarding U.S. GDP_E and GDP_I data (and again, we will discuss that judgment in detail in Sect. 2.2), we believe that a reasonable range for ρ is $\rho \in [0.30, 0.60]$, with midpoint 0.45. One could think of this as a quasi-Bayesian statement that prior beliefs regarding ρ are centered at 0.45, with a 90 % prior credible interval of $[0.30, 0.60]$.⁷

⁷ Again using GDP_E and GDP_I data 1947Q2-2009Q3, we obtain for the numerator $N = -1.87$ in Eq. (7) above. And using the benchmark values of $\phi = 1.1$ and $\rho = 0.45$, we obtain for the denominator $D = -0.61$. This implies $\sigma_I = 1.75$ and $\sigma_E = 1.92$. For comparison, the standard deviation of GDP_E and GDP_I growth rates is about 4.2. Hence our benchmark calibration implies that the error in measuring true GDP by the reported GDP_E and GDP_I growth rates is potentially quite large.




Fig. 4 λ^* versus ρ for various ϕ values. The horizontal line for visual reference is at $\lambda^* = 0.5$. See text for details

In Fig. 3 we show λ^* as a function of ϕ for $\rho = 0, 0.3, 0.45$, and 0.6 ; in Fig. 4 we show λ^* as a function of ρ for $\phi = 0.95, 1.05, 1.15$, and 1.25 ; and in Fig. 5 we show λ^* as a bivariate function of ϕ and ρ . For $\phi = 1$ the optimal weight is 0.5 for all ρ , but for $\phi \neq 1$ the optimal weight differs from 0.5 and is more sensitive to ϕ as ρ grows. The crucial observation remains, however, that under a wide range of conditions it is optimal to put significant weight on both GDP_E and GDP_I , with the optimal weights not differing radically from equality. Moreover, for all ϕ values greater than one, so that less weight is optimally placed on GDP_E under a zero-correlation assumption, allowance for positive correlation further decreases the optimal weight placed on GDP_E . For a benchmark calibration of $\phi = 1.1$ and $\rho = 0.45$, $\lambda^* \approx 0.41$.

Let us again compare σ_C^2 to σ_E^2 for a range of λ values (including $\lambda = \lambda^*$, $\lambda = 0$, and $\lambda = 1$). From (4) we have:

$$\frac{\sigma_C^2}{\sigma_E^2} = \lambda^2 + \frac{(1-\lambda)^2}{\phi^2} + 2\lambda(1-\lambda)\frac{\rho}{\phi}.$$

In Fig. 6 we graph σ_C^2/σ_E^2 for $\lambda \in [0, 1]$ with $\phi = 1.1$ and $\rho = 0.45$. Obviously the maximum variance reduction is obtained using $\lambda^* = 0.41$, but even for nonoptimal λ ,

Fig. 5 λ^* versus ρ and ϕ . See text for details





Fig. 6 σ_C^2/σ_E^2 for $\lambda \in [0, 1]$. We assume $\phi = 1.1$ and $\rho = 0.45$. See text for details



such as simple equal-weight combination ($\lambda = 0.5$), we achieve substantial variance reduction relative to using GDP_E alone.

2.2 On the Rationale for our Calibration

We have thus far implicitly asked the reader to defer to our judgment regarding calibration, focusing on $\phi \in [0.75, 1.45]$ and $\rho \in [0.30, 0.60]$ with benchmark midpoint values of $\phi = 1.10$ and $\rho = 0.45$. Here we explain the experience, reasoning, and research that supports that judgment.

2.2.1 Calibrating ϕ

The key prior view embedded in our choice of $\phi \in [0.75, 1.45]$, with midpoint 1.10, is that GDP_I is likely a somewhat more accurate estimate than GDP_E . This accords with the results of Nalewaik (2010), who examines the relative accuracy of the GDP_E and GDP_I in several ways, with results favorable to GDP_I , suggesting $\phi > 1$.

Let us elaborate. The first source of information on likely values of ϕ is from detailed examination of the source data underlying GDP_E and GDP_I . The largest component of GDP_I , wage, and salary income, is computed using quarterly data from tax records that are essentially universe counts, contaminated by neither sampling nor nonsampling errors. Two other very important components of GDP_I , corporate profits, and proprietors' income, are also computed using annual data from tax records.⁸ Underreporting and nonreporting of income on tax forms (especially by proprietors) is an issue with these data, but the statistical agencies make adjustments for misreporting, and in any event the same misreporting issues plague GDP_E as well as GDP_I , as we discuss below.

In contrast to GDP_I , very little of the quarterly or annual data used to compute GDP_E is based on universe counts.⁹ Rather, most of the quarterly GDP_E source data are from business surveys where response is voluntary. Nonresponse rates can be high, potentially introducing important sample-selection effects that may, moreover, vary with the state of the business cycle. Many annual GDP_E source data are from business surveys with mandatory response, but some businesses still do not respond to the surveys, and surely the auditing of these nonrespondents is less rigorous than the auditing of tax nonfilers. In addition, even the annual surveys do not attempt to collect data on some types of small businesses, particularly nonemployer businesses (i.e., businesses with no employees). The statistical agencies attempt to correct some of these omissions by incorporating data from tax records (making underreporting and nonreporting of income on tax forms an issue for GDP_E as well as GDP_I), but it is not entirely clear whether they adequately plug all the holes in the survey data.

⁸ The tax authorities do not release the universe counts for corporate profits and proprietors' income; rather, they release results from a random sample of tax returns. But the sample they employ is enormous, so the variance of the sampling error is tiny for the top-line estimates. Moreover, the tax authorities obviously know the universe count, so it seems unlikely that they would release tabulations that are very different from the universe counts.

⁹ Motor vehicle sales are a notable exception.

Although these problems plague most categories of GDP_E , some categories appear more severely plagued. In particular, over most of history, government statistical agencies have collected annual source data on less than half of personal consumption expenditures (PCE) for services, a very large category comprising between a quarter and a half of the nominal value of GDP_E over our sample. At the quarterly frequency, statistical agencies have collected even less source data on services PCE.¹⁰ For this reason, statistical agencies have been forced to cobble together less-reliable data from numerous nongovernmental sources to estimate services PCE.

A second source of information on the relative reliability of GDP_E and GDP_I is the correlation of the two measures with other variables that should be correlated with output growth, as examined in Nalewaik (2010). Nalewaik (2010) is careful to pick variables that are not used in the construction of either GDP_E or GDP_I , to avoid spurious correlation resulting from correlated measurement errors.¹¹ The results are uniformly favorable to GDP_I and suggest that it is a more accurate measure of output growth than GDP_E . In particular, from the mid-1980s to the mid-2000s, the period of maximum divergence between GDP_E and GDP_I , Nalewaik (2010) finds that GDP_I growth has higher correlation with lagged stock price changes, the lagged slope of the yield curve, the lagged spread between high-yield corporate bonds and Treasury bonds, short and long differences of the unemployment rate (both contemporaneously and at leads and lags), a measure of employment growth computed from the same household survey, the manufacturing ISM PMI (Institute for Supply Management, Purchasing Managers Index), the nonmanufacturing ISM PMI, and dummies for NBER recessions. In addition, lags of GDP_I growth also predict GDP_E growth (and GDP_I growth) better than lags of GDP_E growth itself.

It is worth noting that, as regards our benchmark midpoint calibration of $\phi = 1.10$, we have deviated only slightly from an “ignorance prior” midpoint of 1.00. Hence our choice of midpoint reflects a conservative interpretation of the evidence discussed above. Similarly, regarding the width of the credible interval as opposed to its midpoint, we considered employing intervals such as $\phi \in [0.95, 1.25]$, for which $\phi > 1$ over most of the mass of the interval. The evidence discussed above, if interpreted aggressively, might justify such a tight interval in favor of GDP_I , but again we opted for a more conservative approach with $\phi < 1$ over more than a third of the mass of the interval.

2.2.2 Calibrating ρ

The key prior view embedded in our choice of $\rho \in [0.30, 0.60]$, with midpoint 0.45, is that the errors in GDP_E and GDP_I are likely positively correlated, with a

¹⁰ This has begun to change recently, as the Census Bureau has expanded its surveys, but ϕ is meant to represent the average relative reliability over the sample we employ, so these facts are highly relevant.

¹¹ For example, the survey of households used to compute the unemployment rate is used in the construction of neither GDP_E nor GDP_I , so use of variables from that survey is fine.

moderately but not extremely large correlation value. This again accords with the results in Nalewaik (2010), who shows that 26% of the nominal value of GDP_E and GDP_I is identical. Any measurement errors in that 26% will be perfectly correlated across the two estimates. Furthermore, GDP_E and GDP_I are both likely to miss fluctuations in output occurring in the underground or “gray” economy, transactions that do not appear on tax forms or government surveys. In addition, the same price deflator is used to convert GDP_E and GDP_I from nominal to real values, so any measurement errors in that price deflator will be perfectly correlated across the two estimates.

These considerations suggest the lower bound for ρ should be well above zero, as reflected in our chosen interval. However, the evidence favoring an upper bound well below one is also quite strong, as also reflected in our chosen interval. First, and most obviously, the standard deviation of the difference between GDP_E and GDP_I is 1.9%, far from the 0.0% that would be the case if $\rho = 1.0$. Second, as discussed in the previous section, the source data used to construct GDP_E is quite different from the source data used to construct GDP_I , implying the measurement errors are likely to be far from perfectly correlated.

Of course, ρ could still be quite high if GDP_E and GDP_I were contaminated with enormous common measurement errors, as well as smaller, uncorrelated measurement errors. But if that were the case, GDP_E and GDP_I would fail to be correlated with other cyclically-sensitive variables, such as the unemployment rate, as they both are. The R^2 values from regressions of the output growth measures on the change in the unemployment rate are each around 0.50 over our sample, suggesting that at least half of the variance of GDP_E and GDP_I is true variation in output growth, rather than measurement error. The standard deviation of the residual from these regressions is 2.81% using GDP_I and 2.95% using GDP_E . For comparison, taking our benchmark value $\phi = 1.1$ and our upper bound $\rho = 0.6$ produces $\sigma_I = 2.05$ and $\sigma_E = 2.25$. Increasing ρ to 0.7 produces $\sigma_I = 2.36$ and $\sigma_E = 2.60$, approaching the residual standard error from our regression. This seems like an unreasonably high amount of measurement error, since the explained variation from such a simple regression is probably not measurement error, and indeed some of the unexplained variation from the regression is probably also not measurement error. Hence the upper bound of 0.6 for ρ seems about right.

3 Combination Under Minimax Loss

Here we take a more conservative perspective on forecast combination, solving a different but potentially important optimization problem. We utilize the minimax framework of Wald (1950), which is the main decision-theoretic approach for imposing conservatism and therefore of intrinsic interest. We solve a game between a benevolent scholar (the Econometrician) and a malevolent opponent (Nature). In that game the Econometrician chooses the combining weights, and Nature selects the stochastic properties of the forecast errors. The minimax solution yields the

combining weights that deliver the smallest chance of the worst outcome for the Econometrician. Under the minimax approach knowledge or calibration of objects like ϕ and ρ is unnecessary, enabling us to dispense with judgment, for better or worse.

We obtain the minimax weights by solving for the Nash equilibrium in a two-player zero-sum game. Nature chooses the properties of the forecast errors and the Econometrician chooses the combining weights λ . For expositional purposes, we begin with the case of uncorrelated errors, constraining Nature to choose $\rho = 0$. To impose some constraints on the magnitude of forecast errors that Nature can choose, it is useful to re-parameterize the vector $(\sigma_I, \sigma_E)'$ in terms of polar coordinates; that is, we let $\sigma_I = \psi \cos \varphi$ and $\sigma_E = \psi \sin \varphi$. We restrict ψ to the interval $[0, \bar{\psi}]$ and let $\varphi \in [0, \pi/2]$. Because $\cos^2 \varphi + \sin^2 \varphi = 1$, the sum of the forecast error variances associated with GDP_E and GDP_I is constrained to be less than or equal to $\bar{\psi}^2$. The error associated with the combined forecast is given by

$$\sigma_C^2(\psi, \varphi, \lambda) = \psi^2 \left[\lambda^2 \sin^2 \varphi + (1 - \lambda)^2 \cos^2 \varphi \right], \quad (9)$$

so that the minimax problem is

$$\max_{\psi \in [0, \bar{\psi}], \varphi \in [0, \pi/2]} \min_{\lambda \in [0, 1]} \sigma_C^2(\psi, \varphi, \lambda), \quad (10)$$

The best response of the Econometrician was derived in (2) and can be expressed in terms of polar coordinates as $\lambda^* = \cos^2 \varphi$. In turn, Nature's problem simplifies to

$$\max_{\psi \in [0, \bar{\psi}], \varphi \in [0, \pi/2]} \psi^2 (1 - \sin^2 \varphi) \sin^2 \varphi,$$

which leads to the solution

$$\varphi^* = \arcsin \sqrt{1/2}, \quad \psi^* = \bar{\psi}, \quad \lambda^* = 1/2. \quad (11)$$

Nature's optimal choice implies a unit forecast error variance ratio, $\phi = \sigma_E/\sigma_I = 1$, and hence that the optimal combining weight is $1/2$. If, instead, Nature set $\varphi = 0$ or $\varphi = \pi/2$, that is $\phi = 0$ or $\phi = \infty$, then either GDP_E or GDP_I is perfect and the Econometrician could choose $\lambda = 0$ or $\lambda = 1$ to achieve a perfect forecast leading to a suboptimal outcome for Nature.

Now we consider the case in which Nature can choose a nonzero correlation between the forecast errors of GDP_E and GDP_I . The loss of the combined forecast can be expressed as

$$\sigma_C^2(\psi, \rho, \varphi, \lambda) = \psi^2 \left[\lambda^2 \sin^2 \varphi + (1 - \lambda)^2 \cos^2 \varphi + 2\lambda(1 - \lambda)\rho \sin \varphi \cos \varphi \right]. \quad (12)$$

It is apparent from (12) that as long as λ lies in the unit interval the most devious choice of ρ is $\rho^* = 1$. We will now verify that conditional on $\rho^* = 1$ the solution in

(11) remains a Nash equilibrium. Suppose that the Econometrician chooses equal weights, $\lambda^* = 1/2$. In this case

$$\sigma_C^2(\psi, \rho^*, \varphi, \lambda^*) = \psi^2 \left[\frac{1}{4} + \frac{1}{2} \sin \varphi \cos \varphi \right].$$

We can deduce immediately that $\psi^* = \bar{\psi}$. Moreover, first-order conditions for the maximization with respect to φ imply that $\cos^2 \varphi^* = \sin^2 \varphi^*$ which in turn leads to $\varphi^* = \arcsin \sqrt{1/2}$. Conditional on Nature choosing ρ^* , ψ^* , and φ^* , the Econometrician has no incentive to deviate from the equal-weights combination $\lambda^* = 1/2$, because

$$\sigma_C^2(\psi^*, \rho^*, \varphi^*, \lambda) = \frac{\bar{\psi}}{2} \left[\lambda^2 + (1 - \lambda)^2 + 2\lambda(1 - \lambda) \right] = \frac{\bar{\psi}}{2}.$$

In sum, the minimax analysis provides a rational for combining GDP_E and GDP_I with equal weights of $\lambda = 1/2$.

To the best of our knowledge, this section's demonstration of the optimality of equal forecast combination weights under minimax loss is original and novel. There does of course exist some related literature, but ultimately our approach and results are very different. For example, a branch of the machine-learning literature (e.g., Vovk 1998; Sancetta 2007) considers games between a malevolent Nature and a benevolent "Learner." The learner sequentially chooses weights to combine expert forecasts, and Nature chooses realized outcomes to maximize the Learner's forecast error relative to the best expert forecast. The Learner wins the game if his forecast loss is only slightly worse than the loss attained by the best expert in the pool, even under Nature's least favorable choice of outcomes. This game is quite different and much more complicated than ours, requiring different equilibrium concepts with different resultant combining weights.

4 Empirics

We have shown that combining using a quasi-Bayesian calibration under quadratic loss produces λ close to but less than 0.5, given our prior means for ϕ and ρ . Moreover, we showed that combining with λ near 0.5 is likely better—often much better—than simply using GDP_E or GDP_I alone, for wide ranges of ϕ and ρ . We also showed that combining under minimax loss always implies an optimal λ of exactly 0.5.

Here we put the theory to work for the U.S., providing arguably-superior combined estimates of GDP growth. We focus on quasi-Bayesian calibration under quadratic loss. Because the resulting combining weights are near 0.50, however, one could also view our combinations as approximately minimax. The point is that a variety of perspectives lead to combinations with weights near 0.50, and they suggest that




Fig. 7 U.S. GDP_C and GDP_E growth rates. GDP_C constructed assuming $\phi = 1.1$ and $\rho = 0.45$. GDP_C is solid and GDP_E is dashed. In the top panel we show a long sample, 1947Q2–2009Q3. In the bottom panel, we show a recent sample, 2006Q1–2009Q3. See text for details

such combinations are likely superior to using either of GDP_E or GDP_I alone, so that empirical examination of GDP_C is of maximal interest.

4.1 A Combined U.S. GDP Series

In the top panel of Fig. 7 we plot GDP_C constructed using $\lambda = 0.41$, which is optimal for our benchmark calibration of $\phi = 1.1$ and $\rho = 0.45$, together with the “conventional” GDP_E . The two appear to move closely together, and indeed they do, at least at the low frequencies emphasized by the long time-series plot. Hence for low-frequency analyses, such as studies of long-term economic growth, use of GDP_E , GDP_I or GDP_C is not likely to make a major difference.

At higher frequencies, however, important divergences can occur. In the bottom panel of Fig. 7, for example, we emphasize business cycle frequencies by focusing on a short sample 2006–2010, which contains the severe U.S. recession of 2007–2009. There are two important points to notice. First, the bottom panel of Fig. 7 makes clear that growth-rate assessments on particular dates can differ in important ways depending on whether GDP_C or GDP_E is used. For example, GDP_E is strongly positive for 2007Q3, whereas GDP_C for that quarter is close to zero, as GDP_I was strongly negative. Second, the bottom panel of Fig. 7 also makes clear that differing assessments can persist over several quarters, as for example during the financial crisis episode of 2007Q1–2007Q3, when GDP_E growth was consistently larger than GDP_C growth. One might naturally conjecture that such persistent and cumulative

data distortions might similarly distort inferences, based on those data, about whether and when the U.S. economy was in recession. We now consider recession dating in some detail.

4.2 U.S. Recession and Volatility Regime Probabilities

Thus far we have assessed how combining produces changes in measured GDP. Now we assess whether and how it changes a certain important *transformation* of GDP, namely measured probabilities of recession regimes or high-volatility regimes based on measured GDP. We proceed by fitting a regime-switching model in the tradition of Hamilton (1989), generalized to allow for switching in both means and variances, as in Kim and Nelson (1999a),

$$\begin{aligned} (\text{GDP}_t - \mu_{s_{\mu t}}) &= \beta(\text{GDP}_{t-1} - \mu_{s_{\mu t-1}}) + \sigma_{s_{\sigma t}} \varepsilon_t & (13) \\ \varepsilon_t &\sim iidN(0, 1) \\ s_{\mu t} &\sim \text{Markov}(P_{\mu}), \quad s_{\sigma t} \sim \text{Markov}(P_{\sigma}). \end{aligned}$$

Then, conditional on observed data, we infer the sequences of recession probabilities [$(P(s_{\mu t} = L))$, where L (“low”) denotes the recession regime] and high-volatility regime probabilities [$(P(s_{\sigma t} = H))$, where H (“high”) denotes the high-volatility regime]. We perform this exercise using both GDP_E and GDP_C , and we compare the results.

We implement Bayesian estimation and state extraction using data 1947Q2-2009Q3.¹² In Fig. 8 we show posterior median smoothed recession probabilities. We show those calculated using GDP_C as solid lines with 90% posterior intervals, we show those calculated using GDP_E as dashed lines, and we also show shaded NBER recession episodes to help provide context. Similarly, in Fig. 9 we show posterior median smoothed volatility regime probabilities.

Numerous interesting substantive results emerge. For example, posterior median smoothed recession regime probabilities calculated using GDP_C tend to be greater than those calculated using GDP_E , sometimes significantly so, as for example during the financial crisis of 2007. Indeed, using GDP_C one might date the start of the recent recession significantly earlier than did the NBER. As regards volatilities, posterior median smoothed high-volatility regime probabilities calculated by either GDP_E or GDP_C tend to show the post-1984 “great moderation” effect asserted by McConnell and Perez-Quiros (2000) and Stock and Watson (2002). Interestingly, however, those calculated using GDP_E also show the “higher recession volatility” effect in recent decades documented by Bloom et al. (2009) (using GDP_E data), whereas those calculated using GDP_C do not.

¹² We provide a detailed description in Appendix.

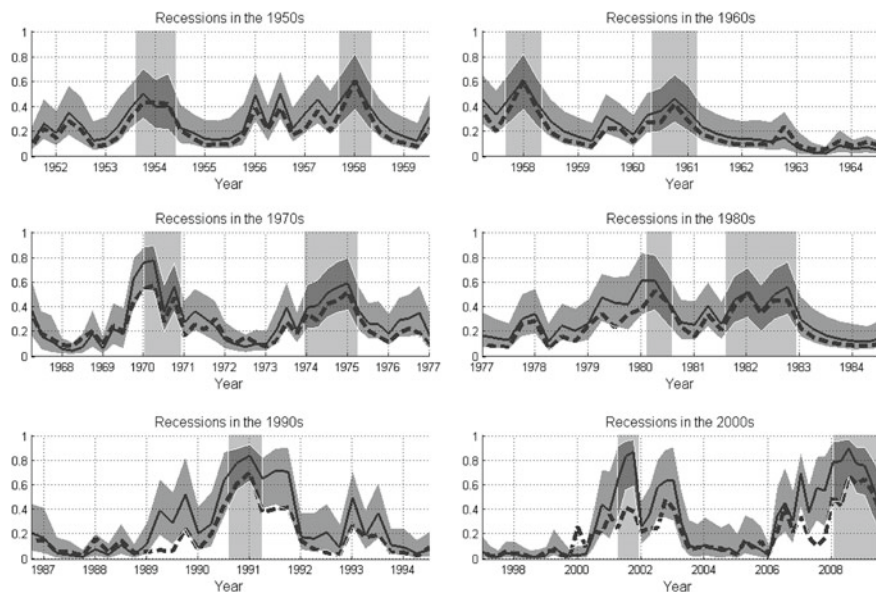


Fig. 8 Inferred U.S. Recession Regime Probabilities, calculated using GDP_C versus GDP_E . *Solid lines* are posterior median smoothed recession regime probabilities calculated using GDP_C , which we show with 90% posterior intervals. *Dashed lines* are posterior median smoothed recession regime probabilities calculated using GDP_E . The sample period is 1947Q2–2009Q3. *Dark shaded bars* denote NBER recessions. See text and appendix for details

For our present purposes, however, none of those substantive results are of first-order importance, as the present chapter is not about business cycle dating, low-frequency versus high-frequency volatility regime dating, or revisionist history, *per se*. Indeed, thorough explorations of each would require separate and lengthy papers. Rather, our point here is simply that one’s assessment and characterization of macroeconomic behavior can, and often does, depend significantly on use of GDP_C versus GDP_E . That is, choice of GDP_C versus GDP_E can *matter* for important tasks, whether based on direct observation of measured GDP, or on transformations of measured GDP such as extracted regime chronologies.

5 Extensions

Before concluding, we offer sketches of what we see as two important avenues for future research. The first involves real-time analysis and nonconstant combining weights, and the second involves combining from a measurement error as opposed to efficient forecast error perspective.




Fig. 9 Inferred U.S. high-volatility regime probabilities, calculated using GDP_C versus GDP_E . *Solid lines* are posterior median smoothed high-volatility regime probabilities calculated using GDP_C , which we show with 90% posterior intervals. *Dashed lines* are posterior median smoothed high-volatility regime probabilities calculated using GDP_E . The sample period is 1947Q2-2009Q3. *Dark shaded bars* denote NBER recessions. See text and appendices for details.

5.1 Vintage Data, Time-Varying Combining Weights, and Real-Time Analysis

It is important to note that everything that we have done in this chapter has a retrospective, or “off-line,” character. We work with a single vintage of GDP_E and GDP_I data and combine them, estimating objects of interest (combining weights, regime probabilities, etc.) for any period t using *all* data $t = 1, \dots, T$. In all of our analyses, moreover, we have used time-invariant combining weights. Those two characteristics of our work thus far are not unrelated, and one may want to relax them eventually, allowing for time-varying weights, and ultimately, a truly real-time-analysis.

One may want to consider time-varying combining weights for several reasons. One reason is of near-universal and hence great interest, at least under quadratic loss. For any given vintage of data, error variances and covariances may naturally change, as we pass backward from preliminary data for the recent past, all the way through to “final revised” data for the more distant past.¹³ More precisely, let t index time measured in quarters, and consider moving backward from “the present” quarter $t = T$.

¹³ This is the so-called “apples and oranges” problem. To the best of our knowledge, the usage in our context traces to Kishor and Koenig (2011).

At instant $v \in T$ (with apologies for the slightly abusive notation), we have vintage- v data. Consider moving backward, constructing combined GDP estimates $GDP_{C,T-k}^v$, $k = 1, \dots, \infty$. For small k , the optimal calibrations might be quite far from benchmark values. As k grows, however, ρ and ϕ should approach benchmark values as the final revision is approached. The obvious question is how quickly and with what pattern should an optimal calibration move toward benchmark values as $k \rightarrow \infty$. We can offer a few speculative observations.

First consider ρ . GDP_I and GDP_E share a considerable amount of source data in their early releases, before common source data are swapped out of GDP_I (e.g., when tax returns eventually become available and can be used). Indeed Fixler and Nalewaik (2009) show that the correlation between the earlier estimates of GDP_I and GDP_E growth is higher than the correlation between the later estimates. Hence ρ is likely higher for dates near the present (small k). This suggests calibrations with ρ dropping monotonically toward the benchmark value of 0.45 as k grows.

Now consider ϕ . How ϕ should deviate from its benchmark calibration value of 1.1 is less clear. On the one hand, early releases of GDP_I are missing some of their most informative source data (tax returns), which suggests a lower-than-benchmark ϕ for small k . On the other hand, early releases of GDP_E growth appear to be noisier than the early releases of GDP_I growth (see below), which suggests a higher-than-benchmark ϕ for small k . All told, we feel that a reasonable small- k calibration of ϕ is less than 1.1 but still above 1.

Note that our conjectured small- k effects work in different directions. Other things equal, bigger ρ pushes the optimal combining weight downward, away from 0.5, and smaller ϕ pushes the optimal combining weight upward, toward 0.5. In any particular data set the effects could conceivably offset more-or-less exactly, so that combination using constant weights for all dates would be fully optimal, but there is of course no guarantee.

Several approaches are possible to implement the time-varying weights sketched in the preceding paragraphs. One is a quasi-Bayesian calibration, elaborating on the approach we have taken in this chapter. However, such an approach would be more difficult in the more challenging environment of time-varying parameters. Another is to construct a real-time data set, one that records a snapshot of the data available at each point in time, such as the one maintained by the Federal Reserve Bank of Philadelphia. The key is to recognize that each quarter we get not simply one new observation on GDP_E and GDP_I , but rather an entire new vintage of data, all the elements of which could (in principle) change. One might be able to use the different data vintages, and related objects like revision histories, to infer properties of “forecast errors” of relevance for construction of optimal combining weights across various k .

One could go even further in principle, progressing to a truly real-time analysis, which is of intrinsic interest quite apart from addressing the issue of time-varying combining weights in the above “apples and oranges” environments. Tracking vintages, modeling the associated dynamics of revisions, and putting it all together to

produce superior combined forecasts remains an outstanding challenge.¹⁴ We look forward to its solution in future work, potentially in the state-space framework that we describe next.

5.2 A Model of Measurement Error

In parallel work in progress (Aruoba et al. 2011), we pursue a complementary approach based on a state-space model of measurement error. The basic model is

$$\begin{bmatrix} \text{GDP}_{E,t} \\ \text{GDP}_{I,t} \end{bmatrix} = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \text{GDP}_t + \begin{bmatrix} \varepsilon_{E,t} \\ \varepsilon_{I,t} \end{bmatrix}$$

$$\text{GDP}_t = \beta_0 + \beta_1 \text{GDP}_{t-1} + \eta_t, \quad (14)$$

where $\varepsilon_t = (\varepsilon_{E,t}, \varepsilon_{I,t})' \sim WN(0, \Sigma_\varepsilon)$, $\eta_t \sim WN(0, \sigma_\eta^2)$, and ε_t and η_t are uncorrelated at all leads and lags. In this model, both GDP_E and GDP_I are noisy measures of the latent true GDP process, which evolves dynamically. The expectation of true GDP conditional upon observed measurements may be extracted using optimal filtering techniques such as the Kalman filter.

The basic state-space model can be extended in various directions, for example to incorporate richer dynamics, and to account for data revisions and missing advance and preliminary releases of GDP_I .¹⁵ Perhaps most important, the measurement errors ε may be allowed to be correlated with GDP, or more precisely, correlated with GDP innovations, η_t . Fixler and Nalewaik (2009) and Nalewaik (2010) document cyclicalities in the “statistical discrepancy” ($\text{GDP}_E - \text{GDP}_I$), which implies failure of the assumption that ε_t and η_t are uncorrelated at all leads and lags. Of particular concern is contemporaneous correlation between η_t and ε_t . The standard Kalman filter cannot handle this, but appropriate modifications are available.

6 Conclusions

GDP growth is a central concept in macroeconomics and business cycle monitoring, so its accurate measurement is crucial. Unfortunately, however, the two available expenditure-side and income-side U.S. GDP estimates often diverge. In this chapter, we proposed a technology for optimally combining the competing GDP estimates,

¹⁴ Nalewaik (2011) makes some progress toward real-time analysis in a Markov-switching environment.

¹⁵ The first official estimate of GDP_I is released a month or two after the first official estimate of GDP_E , so for vintage v the available GDP_E^v data might be $\{\text{GDP}_{E,t}^v\}_{t=1}^{T-1}$, whereas the available GDP_I^v vintage might be $\{\text{GDP}_{I,t}^v\}_{t=1}^{T-2}$. Note that for any vintage v , the available GDP_I data differ by at most one quarter from the available GDP_E data.

we examined several variations on the basic theme, and we constructed and examined combined estimates for the U.S.

Our results strongly suggest the desirability of separate and careful calculation of both GDP_E and GDP_I , followed by combination, which may lead to different and more accurate insights than those obtained by simply using expenditure-side or estimates alone. This prescription differs fundamentally from U.S. practice, where both are calculated but the income-side estimate is routinely ignored.

Our call for a combined U.S. GDP measure is hardly radical, particularly given current best-practice “balancing” procedures used at various non-U.S. statistical agencies to harmonize GDP estimates from different sources. We discussed U.K. GDP balancing at some length in the introduction, and some other countries also use various similar balancing procedures.¹⁶ All such procedures recognize the potential inaccuracies of source data and have a similar effect to our forecast combination approach: the final GDP number lies between the alternative estimates.

Other countries use other approaches to combination. Indeed Australia uses an approach reminiscent of the one that we advocate in this chapter, albeit not on the grounds of our formal analysis.¹⁷ In addition to GDP_E and GDP_I , the Australian Bureau of Statistics produces a production-side estimate, GDP_P , defined as total gross value added plus taxes and less subsidies, and its headline GDP number is the simple average of the three GDP estimates. We look forward to the U.S. producing a similarly-combined headline GDP estimate, potentially using the methods introduced in this chapter.

Acknowledgments We dedicate this chapter to Hal White, on whose broad shoulders we stand, on the occasion of his sixtieth birthday. For helpful comments we thank the editors and referees, as well as John Geweke, Greg Mankiw, Matt Shapiro, Chris Sims, and Justin Wolfers. For research support we thank the National Science Foundation and the Real-Time Data Research Center at the Federal Reserve Bank of Philadelphia. For research assistance we thank Ross Kelley and Matthew Klein.

Appendix: Estimation of U.S. Recession Probabilities

Here we provide details of Bayesian analysis of our regime-switching model.

A.1 Baseline Model

We work with a simple model with Markov regime-switching in mean and variance:

¹⁶ Germany’s procedures, for example, are described in Statistisches Bundesamt (2009).

¹⁷ See <http://www.abs.gov.au>, under Australian National Accounts, Explanatory Notes for Australia.

$$(\text{GDP}_t - \mu_{s_{\mu t}}) = \beta(\text{GDP}_{t-1} - \mu_{s_{\mu t-1}}) + \sigma_{s_{\sigma t}} \varepsilon_t \quad (\text{A.1})$$

$$\varepsilon_t \sim iidN(0, 1)$$

$$s_{\mu t} \sim \text{Markov}(P_{\mu}), \quad s_{\sigma t} \sim \text{Markov}(P_{\sigma}), \quad (\text{A.2})$$

where P_{μ} and P_{σ} denote transition matrices for high and low mean and variance regimes,

$$P_{\mu} = \begin{bmatrix} p_{\mu H} & 1 - p_{\mu H} \\ 1 - p_{\mu L} & p_{\mu L} \end{bmatrix}$$

$$P_{\sigma} = \begin{bmatrix} p_{\sigma H} & 1 - p_{\sigma H} \\ 1 - p_{\sigma L} & p_{\sigma L} \end{bmatrix}.$$

Overall, then, there are four regimes:

$$S_t = 1 \text{ if } s_{\mu t} = H, \quad s_{\sigma t} = H \quad (\text{A.3})$$

$$S_t = 2 \text{ if } s_{\mu t} = H, \quad s_{\sigma t} = L$$

$$S_t = 3 \text{ if } s_{\mu t} = L, \quad s_{\sigma t} = H$$

$$S_t = 4 \text{ if } s_{\mu t} = L, \quad s_{\sigma t} = L.$$

For $t = 0$ the hidden Markov states are governed by the ergodic distribution associated with P_{μ} and P_{σ} .

A.2 Bayesian Inference

Priors. Bayesian inference combines a prior distribution with a likelihood function to obtain a posterior distribution of the model parameters and states. We summarize our benchmark priors in Table A.1. We employ a normal prior for μ_L , a gamma prior for $\mu_H - \mu_L$, inverted gamma priors for σ_H and σ_L , beta priors for the transition probabilities, and finally, a normal prior for β . Our prior ensures that $\mu_H \geq \mu_L$ and thereby deals with the “label switching” identification problem.

For μ_L , the average growth rate in the low-growth state, we use a prior distribution that is centered at 0, with standard deviation 0.70%. Note that a priori we do not restrict the average growth rate to be negative. We also allow for (mildly) positive values. We choose the prior for $\mu_H - \mu_L$ such that the mean difference between the average growth rates in the two regimes is 2.00%, with standard deviation 1.00%. Our priors for the transition probabilities p_{μ} and p_{σ} are symmetric and imply a mean regime duration between three and 14 quarters. Finally, our choice for the prior of the autoregressive parameter β is normal with zero mean and unit variance, allowing a priori for both stable and unstable dynamics of output growth rates.

Table A.1 Prior choices and posterior distributions

	Prior Choice	GDP _E			GDP _C		
		Median	5 %	95 %	Median	5 %	95 %
$\mu_H - \mu_L$	Gamma(2, 1)	–	–	–	–	–	–
μ_H	–	3.50	[3.03	4.12]	3.76	[2.97	4.28]
μ_L	Normal(0, 0.5)	1.25	[0.34	2.29]	0.82	[0.17	1.64]
σ_H	InvGamma(2, 2)	4.82	[4.35	5.43]	4.64	[4.21	5.13]
σ_L	InvGamma(1, 2)	1.92	[1.55	2.34]	1.71	[1.74	2.05]
β	Normal(0, 1)	0.31	[0.17	0.45]	0.37	[0.27	0.53]
p_{μ_H}	Beta(25, 5)	0.91	[0.82	0.96]	0.92	[0.85	0.96]
p_{μ_L}	Beta(25, 5)	0.79	[0.64	0.87]	0.80	[0.67	0.88]
p_{σ_H}	Beta(25, 5)	0.91	[0.83	0.96]	0.91	[0.83	0.96]
p_{σ_L}	Beta(25, 5)	0.89	[0.81	0.95]	0.91	[0.85	0.95]

Implementation of Posterior Inference. Posterior inference is implemented with a Metropolis-within-Gibbs sampler, building on work by Carter and Kohn (1994) and Kim and Nelson (1999b). We denote the sequence of observations by $\text{GDP}_{1:T}$. Moreover, let $S_{1:T}$ be the sequence of hidden states, and let

$$\theta = (\mu_H, \mu_L, \sigma_H, \sigma_L, \beta)', \quad \text{and} \quad \phi = (p_{\mu_H}, p_{\mu_L}, p_{\sigma_L}, p_{\sigma_H})'.$$

Our Metropolis-within-Gibbs algorithm involves sampling iteratively from three conditional posterior distributions. To initialize the sampler we start from (θ^0, ϕ^0) .

Algorithm: Metropolis-within-Gibbs Sampler

For $i = 1, \dots, N$:

1. Draw $S_{1:T}^{i+1}$ conditional on $\theta^i, \phi^i, \text{GDP}_{1:T}$. This step is implemented using the multi-move simulation smoother described in Sect. 9.1.1 of Kim and Nelson (1999b).
2. Draw ϕ^{i+1} conditional on $\theta^i, S_{1:T}^{i+1}, \text{GDP}_{1:T}$. If the dependence of the distribution of the initial state S_1 on ϕ is ignored, then it can be shown that the conditional posterior of ϕ is of the Beta form (see Sect. 9.1.2 of Kim and Nelson 1999b). We use the resulting Beta distribution as a proposal distribution in a Metropolis–Hastings step.
3. Draw θ^{i+1} , conditional on $\phi^{i+1}, S_{1:T}^{i+1}, \text{GDP}_{1:T}$. Since our prior distribution is nonconjugate, we are using a random-walk Metropolis step to generate a draw from the conditional posterior of θ . The proposal distribution is $N(\theta^i, c\Omega)$.

We obtain the covariance matrix Ω of the proposal distribution in Step 3 as follows. Following Schorfheide (2005) we maximize the posterior density,

$$p(\theta, \phi | \text{GDP}_{1:T}) \propto p(\text{GDP}_{1:T} | \theta, \phi) p(\theta, \phi),$$

to obtain the posterior mode $(\tilde{\theta}, \tilde{\phi})$. We then construct the negative inverse of the Hessian at the mode and let Ω be the submatrix that corresponds to the parameter subvector θ . We choose the scaling factor c to obtain an acceptance rate of approximately 40%. We initialize our algorithm choosing (θ^0, ϕ^0) in the neighborhood of $(\tilde{\theta}, \tilde{\phi})$ and use it to generate $N = 100,000$ draws from the posterior distribution.¹⁸

Posterior Estimates. Table A.1 also contains percentiles of posterior parameter distributions. The posterior estimates for the volatility parameters and the transition probabilities are similar across GDP_E and GDP_C . However, the posterior estimate for μ_L is higher using GDP_E than using GDP_C , while the opposite is true for β . Moreover, the differential between high and low mean regimes is bigger in the case of GDP_C , all of which can influence the time-series plot of the recession probabilities.

The Markov-switching means capture low-frequency shifts while the autoregressive coefficient captures high-frequency dynamics. Thus, the presence of the autoregressive term may complicate our analysis, because we are trying to decompose the GDP measurement discrepancy into both low- and high-frequency components. As a robustness check, we remove the autoregressive term in (A.1) and estimate an *iid* model specification. Although the posterior estimates for μ_L change, the remaining parameters are essentially identical to Table A.1. The smoothed recession probabilities remain nearly identical to Fig. 8.

References

- Aruoba, B. (2008), “Data Revisions are not Well-Behaved”, *Journal of Money, Credit and Banking*, 40, 319–340.
- Aruoba, S.B. and F.X. Diebold (2010), “Real-Time Macroeconomic Monitoring: Real Activity, Inflation, and Interactions”, *American Economic Review*, 100, 20–24.
- Aruoba, S.B., F.X. Diebold, J. Nalewaik, F. Schorfheide, and D. Song (2011), “Improving GDP Measurement: A Measurement Error Perspective”, Manuscript in progress, University of Maryland, University of Pennsylvania and Federal Reserve Board.
- Barker, T., F. van der Ploeg, and M. Weale (1984), “A Balanced System of National Accounts for the United Kingdom”, *Review of Income and Wealth*, 461–485.
- Bates, J.M. and C.W.J. Granger (1969), “The Combination of Forecasts”, *Operations Research Quarterly*, 20, 451–468.
- Beaulieu, J. and E.J. Bartelsman (2004), “Integrating Expenditure and Income Data: What To Do With the Statistical Discrepancy?” FEDS Working Paper 2004, 39.
- Bloom, N., M. Floetotto, and N. Jaimovich (2009), “Really Uncertain Business Cycles”, Manuscript, Stanford University.
- Byron, R. (1978), “The Estimation of Large Social Accounts Matrices”, *Journal of the Royal Statistical Society Series A*, 141, Part 3, 359–367.
- Carter, C.K. and R. Kohn (1994), “On Gibbs Sampling for State Space Models”, *Biometrika*, 81, 541–553.
- Diebold, F.X. and J.A. Lopez (1996), “Forecast Evaluation and Combination”, In G.S. Maddala and C.R. Rao (eds.) *Handbook of Statistics (Statistical Methods in Finance)*, North-Holland, 241–268.

¹⁸ We performed several tests confirming that our choice of N yields an accurate posterior approximation.

- Faust, J., J.H. Rogers, and J.H. Wright (2005), "News and Noise in G-7 GDP Announcements", *Journal of Money, Credit and Banking*, 37, 403–417.
- Fixler, D.J. and J.J. Nalewaik (2009), "News, Noise, and Estimates of the "True" Unobserved State of the Economy", Manuscript, Bureau of Labor Statistics and Federal Reserve Board.
- Hamilton, J.D. (1989), "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle", *Econometrica*, 57, 357–384.
- Kim, C.-J. and C.R. Nelson (1999a), "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle", *Review of Economics and Statistics*, 81, 608–616.
- Kim, C.-J. and C.R. Nelson (1999b), *State Space Models with Regime Switching*, MIT Press.
- Kishor, N.K. and E.F. Koenig (2011), "VAR Estimation and Forecasting When Data Are Subject to Revision", *Journal of Business and Economic Statistics*, in press.
- Mankiw, N.G., D.E. Runkle, and M.D. Shapiro (1984), "Are Preliminary Announcements of the Money Stock Rational Forecasts?" *Journal of Monetary Economics*, 14, 15–27.
- Mankiw, N.G. and M.D. Shapiro (1986), "News or Noise: An Analysis of GNP Revisions", *Survey of Current Business*, May, 20–25.
- McConnell, M. and G. Perez-Quiros (2000), "Output Fluctuations in the United States: What Has Changed Since the Early 1980s?" *American Economic Review*, 90, 1464–1476.
- Nalewaik, J.J. (2010), "The Income- and Expenditure-Side Estimates of U.S. Output Growth", *Brookings Papers on Economic Activity*, 1, 71–127 (with discussion).
- Nalewaik, J.J. (2011), "Estimating Probabilities of Recession in Real Time Using GDP and GDI", *Journal of Money, Credit and Banking*, in press.
- Sancetta, A. (2007), "Online Forecast Combinations of Distributions: Worst Case Bounds", *Journal of Econometrics*, 141, 621–651.
- Schorfheide, F. (2005), "Learning and Monetary Policy Shifts", *Review of Economic Dynamics*, 8, 392–419.
- Solomou, S. and M. Weale (1991), "Balanced Estimates of U.K. GDP 1870–1913", *Explorations in Economic History*, 28, 54–63.
- Solomou, S. and M. Weale (1993), "Balanced Estimates of National Accounts When Measurement Errors Are Autocorrelated: The U.K., 1920–1938", *Journal of the Royal Statistical Society Series A*, 156 Part 1, 89–105.
- Statistisches Bundesamt, Wiesbaden (2009), "National Accounts: Gross Domestic Product in Germany in Accordance with ESA 1995 - Methods and Sources", *Subject Matter Series*, 18.
- Stock, J.H. and M.W. Watson (2002), "Has the Business Cycle Changed and Why?" In M. Gertler and K. Rogoff (eds.), *NBER Macroeconomics Annual*, Cambridge, Mass.: MIT Press, 159–218.
- Stone, R., D.G. Champernowne, and J.E. Meade (1942), "The Precision of National Income Estimates", *Review of Economic Studies*, 9, 111–125.
- Timmermann, A. (2006), "Forecast Combinations", In G. Elliot, C.W.J. Granger and A. Timmermann (eds.), *Handbook of Economic Forecasting*, North-Holland, 136–196.
- Vovk, V. (1998), "A Game of Prediction with Expert Advice", *Journal of Computer and System Sciences*, 56, 153–173.
- Wald, A. (1950), *Statistical Decision Functions*, John Wiley, New York.
- Weale, M. (1985), "Testing Linear Hypotheses on National Accounts Data", *Review of Economics and Statistics*, 90, 685–689.
- Weale, M. (1988), "The Reconciliation of Values, Volumes, and Prices in the National Accounts", *Journal of the Royal Statistical Society Series A*, 151 Part 1, 211–221.