

Federal Reserve Bank of New York
Staff Reports

Macroeconomic Nowcasting and Forecasting with Big Data

Brandyn Bok
Daniele Caratelli
Domenico Giannone
Argia Sbordone
Andrea Tambalotti

Staff Report No. 830
November 2017



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

Macroeconomic Nowcasting and Forecasting with Big Data

Brandyn Bok, Daniele Caratelli, Domenico Giannone, Argia Sbordone, and Andrea Tambalotti

Federal Reserve Bank of New York Staff Reports, no. 830

November 2017

JEL classification: C32, C53, C55, E32

Abstract

Data, data, data . . . Economists know it well, especially when it comes to monitoring macroeconomic conditions—the basis for making informed economic and policy decisions. Handling large and complex data sets was a challenge that macroeconomists engaged in real-time analysis faced long before “big data” became pervasive in other disciplines. We review how methods for tracking economic conditions using big data have evolved over time and explain how econometric techniques have advanced to mimic and automate the best practices of forecasters on trading desks, at central banks, and in other market-monitoring roles. We present in detail the methodology underlying the New York Fed Staff Nowcast, which employs these innovative techniques to produce early estimates of GDP growth, synthesizing a wide range of macroeconomic data as they become available.

Key words: monitoring economic conditions, business cycle, macroeconomic data, large data sets, high-dimensional data, real-time data flow, factor model, state space models, Kalman filter

Bok, Giannone, Sbordone, Tambalotti: Federal Reserve Bank of New York. Corresponding author: Giannone (email: domenico.giannone@ny.frb.org). Caratelli contributed to this paper while working at the Federal Reserve Bank of New York. The authors thank Grant Aarons, Patrick Adams, Matthew Cocci, Jessica Iannuzzi, Theresa Izzillo, Shahrouz Khalili, Eric Qian, Camilla Schneier, and Anna Snider, who contributed at various stages to the development of the New York Fed Staff Nowcast, as well as Dick Peach and many colleagues at the New York Fed for helpful discussion. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

“[O]nly by analyzing numerous time series, each of restricted significance, can business cycles be made to reveal themselves definitely enough to permit close observation.”

Burns and Mitchell, eds (1946), *Measuring Business Cycles* (p.11)

I Introduction

For more than a century, government agencies and private institutions have been collecting and organizing information on many facets of the economy, and over time the scope of data collection has grown and the quality of data has improved. Today, data are released to the public on a regular schedule: almost every day new data become available and are analyzed, commented and interpreted. Real-time monitoring of macroeconomic conditions has become the full-time job of dedicated economists at central banks, government agencies and the corporate world, who sift through big and complex data to distill all relevant information. Releases that come out as surprises move markets, sometimes significantly, as investors re-assess their expectations about the state of the economy.

Although the term “big data” typically conjures the image of data collected via the Internet about individual habits on consumption and social media, big data was a challenge to macroeconomists well before the collection of more granular data became pervasive in other disciplines. From the pioneering search for patterns and regularities in the data that led Burns and Mitchell to identify the business cycle and the parallel effort of Kuznets to build the National Income and Product Accounts, to the vast array of expert data collection and analysis done today, macroeconomists have embraced the big data challenge, pushing the frontier of statistical methods and refined measurement.

New methodologies in time-series econometrics developed over the past two decades have made possible the construction of automated platforms for monitoring macroeconomic conditions in real time. Giannone, Reichlin and Small (2008) built the first formal and internally consistent statistical framework of this kind by combining models for big data and filtering techniques. Because of the emphasis on the present, they dubbed it “nowcasting,” a term originally used in meteorology for forecasting the weather in the present and in the next few hours. The model that we describe in this paper, developed at the Federal Reserve Bank of New York, builds on that contribution. This nowcasting model extracts the latent factors that drive the movements in the data and produces a forecast of each economic series

that it tracks: when the actual release for that series differs from the model’s forecast, this ‘news’ impacts the nowcast of GDP growth. This approach formalizes key features of how market participants and policymakers have traditionally produced forecasts, a process that involves monitoring many data releases, forming expectations about them, and then revising the assessment of the state of the economy whenever facts differ from those expectations. The model combines in a unified framework a variety of approaches developed over time for monitoring economic conditions.

Figure 1.1, whose detailed explanation we defer to Section 5, illustrates the evolution of the nowcast of annualized real GDP growth for 2016:Q4.

[Insert Figure 1.1 here.]

The black diamonds represent the weekly updates of the nowcast, i.e., the predictions of the model based on the information available at the dates indicated on the horizontal axis. Their progression reflects how the news in the data released each week change the nowcast for that week. The impact on the nowcast of news from a week’s data releases is visualized by the colored bar of that week, where the colors identify the categories of the data releases as indicated in the legend. For example, on November 11, the nowcast of real GDP growth for the fourth quarter of 2016 was 1.6 percent; a string of positive surprises during the following week, primarily from consumption data (represented by the green segment of that week’s bar) and housing market data (the red segment), only partially offset by negative surprises from manufacturing data (the orange segment), increased the nowcast to 2.4 percent.

Before moving to a detailed description of the New York Fed model, in the next section we review the variety of methods developed over time to monitor macroeconomic conditions. We then discuss issues of data collection and measurement, with an emphasis on the nature of macroeconomic time series data and their real-time flow (Section 3). In Section 4 we present the theoretical framework for nowcasting with a large dataset focusing on the parsimonious aspect of the dynamic factor model methodology. In Section 5 we dig into the specifics of the New York Fed Staff Nowcast. Section 6 concludes.

II Monitoring economic conditions

Every day economists parse the trove of economic data released by statistical agencies, private and public surveys, and other sources to assess the health of the economy. Separating

meaningful signals from noise is not an easy task, and several approaches have been developed and applied over time to tackle it. These range from detecting business cycle turning points and constructing indexes of economic activity to forecasting comprehensive macroeconomic measures of the state of the economy with formal models and judgment.

The first systematic analysis of economic fluctuations dates back to Burns and Mitchell, the economists that pioneered business cycle analysis at the National Bureau of Economic Research (NBER) in the late 1930s. Faced with the complexity of the economic system, Burns and Mitchell attacked their investigation as a big data problem: they scrutinized hundreds of data series in search for patterns and regularities.¹ What they uncovered was a systematic co-movement among the series and pervasiveness of fluctuations across different sectors and different kinds of economic activities. This led them to identify the broad recurrence of two states in the economy, expansions and recessions. And so they defined the business cycle as the “type of fluctuation found in the aggregate economic activity of nations that organize their work mainly in business enterprises: a cycle consists of expansions occurring at about the same time in many economic activities, followed by similarly general recessions,” (p. 3).² What makes Burns and Mitchell’s work so important and innovative is the fact that the pattern they were looking for was unknown—in modern language we would call it “unsupervised classification.” We could argue that their careful screening for pattern recognition is what many decades later became machine learning.

Pervasiveness in the movement of various indicators (across sectors and activities) remains central to the definition of business cycles currently used by the NBER Business Cycle Dating Committee: “During a recession, a significant decline in economic activity spreads across the economy and can last from a few months to more than a year. Similarly, during an expansion, economic activity rises substantially, spreads across the economy, and usually lasts for several years.”³

In a modern context, as formalized by Harding and Pagan (2016), determining business cycle turning points from a variety of series can be seen as a two-step process: first identifying

¹Burns and Mitchell classified 71 out of the original 487 economic time series as the most trustworthy indicators of business cycle revivals: “[...] we have drawn up a list of statistical series differing widely in other respects but alike in that each has in the past proved to be a fairly consistent indicator of cyclical movements in general business. We regard this list not as a ‘forecasting’ machine, but rather as a registering device that may be useful to those who are trying to interpret the general drift of current fluctuations in different types of business activity,” *Statistical Indicators of Cyclical Revivals* NBER, 1938 (p. 1).

²The most well-known research of Burns and Mitchell is collected in the NBER 1946 volume.

³See <http://www.nber.org/cycles/recessions.html>

turning points in each of the large variety of data and then constructing reference turning points based on the distribution of the individual series' turning points. The first step was initially based on judgment, and was later automated by Bry and Boschan (1971). The identification of clusters of turning points to determine aggregate turning points has been formalized by Harding and Pagan (2002) and Stock and Watson (2010, 2014).

The 'dating' of the business cycle represents one of the most common and robust summaries of the economy, widely understood not only by experts, but by the public at large. The NBER dating committee today continues the work of Burns and Mitchell, determining the official turning points in the U.S. economy. It examines and compares the behavior of a variety of broad and comprehensive economic activity measures, primarily real GDP and most recently real gross domestic income (GDI), labor market indicators and industrial production, together with other less broad but highly informative indicators, to determine when a recession starts and when it ends.

Monitoring a large number of different variables can enhance timeliness and accuracy in assessing the health of the economy: it not only enables one to exploit different sampling frequencies and different timing of macroeconomic data releases but also mitigates the risk of overweighting idiosyncratic fluctuations as well as measurement errors. However, as many indicators move together over the cycle, the behavior of multiple series providing a similar signal can be well summarized with low-dimensional indexes, which can be broadly considered as indexes of business cycles.

Indexes of economic activity such as the leading, coincident and lagging indicator indexes for the U.S. economy constructed by the Conference Board are in this tradition. The OECD also publishes a composite leading indicator (CLI) index for 21 member countries and three zone aggregates (OECD, 2012), and for the euro area the Conference Board publishes coincident and leading indicators.⁴ More recently, indexes of economic indicators have been constructed using dynamic factor models, which, as we will argue at length in Section 4, amount essentially to using model-based aggregation schemes. The use of factor models for the construction of business cycle indexes was pioneered by Stock and Watson (1989).

From an econometric perspective, the use of factor models to monitor macroeconomic conditions stems from the basic insight that information about different aspects and sectors of the economy can be considered as imperfect measures of a latent common business cycle factor. A robust finding of this literature is that a few common factors can capture the

⁴For a survey of indexes of economic indicators, see Marcellino (2006).

salient features of business cycle fluctuations. First documented by Sargent and Sims (1977), this result has more recently been confirmed with high-dimensional macroeconomic data, as shown by Giannone et al. (2005) and Watson (2005).

Economists also focus on some key and comprehensive indicators of economic activity, such as the growth rate of real gross domestic product (GDP). Indeed, the business cycle turns out, ex post, to be very close to the peaks and troughs of this single comprehensive measure of economic activity.⁵ Moreover, the journalistic definition of a “technical recession,” one when there are two consecutive quarters of negative real GDP growth, is a popularized version of algorithms derived to identify business cycle turning points. This bridges with the careful work dedicated to the construction of GDP data in the National Income and Product Accounts, as we will discuss in Section 3. However, since comprehensive measures are available only with a delay, it is customary to make predictions for the official figures while waiting for their release, pooling information from a variety of economic series.

Forecasting is essential to central banks in informing their policy decisions and communicating their economic outlook to the public. Central bank staff typically use a suite of models and a fair amount of expert judgment to arrive at their forecasts.⁶ Businesses and consumers, lacking individual expertise, also rely on forecasts by professional economists to inform their spending and investment decisions.

The collection of expert forecasts has a long tradition. The oldest quarterly survey of macroeconomic forecasts is the Survey of Professional Forecasters (SPF), which began in 1968 and is currently conducted by the Federal Reserve Bank of Philadelphia.⁷ Forecasters provide quarterly projections of U.S. GDP growth, measures of inflation, unemployment and payroll employment for the current quarter and subsequent four quarters as well as annual projections for the current and following year.⁸

⁵See for example Hamilton (1989) and Harding and Pagan (2002).

⁶Sims (2002) provides an insightful review and assessment of the forecasting activity at several major central banks.

⁷The survey started in 1968:Q4 and was initially conducted by the American Statistical Association (ASA) and the NBER, known as the ASA-NBER Survey. See Victor Zarnowitz, 1968, [The New ASA-NBER Survey of Forecast by Economic Statisticians](#). The survey was taken over by the Federal Reserve Bank of Philadelphia in 1990:Q2.

⁸Another old survey is the Livingston Survey, started in 1946 and currently run by the Federal Reserve Bank of Philadelphia. This survey is semi-annual and consists of forecasts of 18 different quarterly and monthly variables describing national output, prices, unemployment, and other macroeconomic data. Also, since 1999 the ECB runs a Survey of Professional Forecasters similar to the U.S. survey that collects expectations for the rates of inflation, real GDP growth and unemployment in the euro area for several horizons,

Professional forecasters typically use a combination of approaches for forecasting. A special survey conducted by the Real-Time Data Research Center at the Philadelphia Fed in 2009 revealed that the majority of the SPF panelists use mathematical models to form their projections, but also apply subjective adjustments to their pure model-generated forecasts. Interestingly, the use of models is predominant for short-horizon forecasts, less so for long-horizon projections. However, not all forecasters monitor economic conditions at high frequency: only 5 out of 25 respondents seem to update their forecasts at higher than monthly frequency.

Alongside professional forecasters, market analysts also strive to understand where the economy currently is and to forecast in which direction it is going. They track major data releases to detect early signals: news in the data, relative to their expectations, lead them to update their projections. Market forecasts are also collected in surveys, of which the most widely followed is the one conducted by Bloomberg. When releases come out as surprises they move markets (Gürkaynak et al., 2005; Bartolini et al., 2008; Gürkaynak and Wright, 2013). In fact, macroeconomic surprises explain a large part of asset prices fluctuations, up to one third of the quarter-to-quarter fluctuations in government bond yields (Altavilla et al., 2014). This evidence suggests that investors continuously update and reassess their expectations about the future path of the economy and the policy reaction.

How successful are professional forecasts? Apparently, there is little predictability of real GDP growth beyond the current and next quarter, as shown in Table 2.1, which reports the SPF forecast error statistics alongside those of a naïve statistical model: the big gain of SPF forecasts is at horizon 0 (forecast of current quarter). For reference, the table also includes the root-mean-square error of the BEA advance GDP release assessed relative to its most recent, revised value.

[Insert Table 2.1 here.]

Forecasts appear most helpful when one wants to understand where the economy is now, but predicting the present requires tracking a large and complex set of data as it becomes available continuously in real time. Traditionally, this was achieved using a combination of data scrutiny, a variety of simple models, and expert judgment. As discussed in the

together with a quantitative assessment of the uncertainty surrounding them. In addition to other private surveys, such as Blue Chip and Consensus Economics, several institutions such as the OECD and the IMF communicate and summarize their economic outlook by publishing forecasts.

introduction, in the last two decades new methodologies in time-series econometrics have made possible the development of platforms for real-time forecasting that combine formal models for big data and filtering into nowcasting. The model we describe in this paper is one such platform, which we would argue unifies several analytical approaches for monitoring current economic conditions that are typically used independently. As indexes of coincident and leading indicators do, our model characterizes current economic activity by condensing the information into a few factors that summarize business cycle conditions. The model mimics the behavior of market participants and professional forecasters, by tracking all relevant measures of economic activity, making predictions that are constantly updated in response to unexpected developments in economic releases. The general finding is that these automated forecasts are as accurate as, and highly correlated with, the forecasts produced by institutions and experts.

Unlike professional forecasters that combine a variety of unrelated models and apply some form of judgment, using a single formal model allows for a transparent and internally coherent analysis of the real-time data flow. The model, in essence, codifies within an econometric framework the best practice and expert knowledge in business cycle analysis. This is a significant change in paradigm that is well summarized by Stock and Watson (2017).⁹

III Big data and the real-time information flow

Parallel to the development of various ways of monitoring economic conditions just described is the advancement of measurement. The effort to collect a very large and complex set of measurements on the economy and to organize and synthesize them in a system of coherent aggregates—the national accounting—was first undertaken during the Great Depression, at the same time as macroeconomics emerged as an independent discipline. If macroeconomics

⁹“Twenty years ago, economists who monitored the economy in real time used indexes of economic indicators and regression models for updating expectations of individual releases [...], combined with a large dose of judgment based on a narrative of where the economy was headed. While this approach uses data, it is not scientific in the sense of being replicable, using well-understood methods, quantifying uncertainty, or being amenable to later evaluation. Moreover, this method runs the risk of putting too much weight on the most recent but noisy data releases, putting too little weight on other data, and being internally inconsistent because each series is handled separately. [...] The current suite of tools for handling large series and complicated data flows [...] using a single model to evaluate these releases—rather than a suite of small models or judgment—provides a scientific way to use the real-time data flow” (p. 71).

was the answer to the challenges posed to economics by the events of the Great Depression, national accounting was its counterpart in the realm of economic measurement. In the United States, Simon Kuznets developed the National Income and Product Accounts (NIPAs) to provide a comprehensive picture of what was happening in the economy during that time of crisis, as well as to monitor the effect of the many policies put in place by President Roosevelt to fight the Depression. Subsequently, the NIPAs became a crucial tool in the efforts to transform the economy in support of the war effort.¹⁰

Nowadays, academic analysis largely focuses on a few macroeconomic aggregates that comprise the national accounts, such as GDP, consumption or investment. These time series result from a complex and systematic effort to measure all the economic activity taking place in the U.S. economy within a formalized and coherent framework. Conceptually, this framework is based on accounting principles, rather than on statistical or economic models, but it too is a formalized answer to a big data challenge: how to describe and track over time the evolution of a complex and continually evolving system like the U.S. economy.

For GDP, the most representative and cited of all macroeconomic variables, the Bureau of Economic Analysis (BEA) produces every five years benchmark estimates based on an economic census that covers virtually all the roughly seven million businesses with paid employees in the United States and over 95 percent of the expenditures included in GDP. Economic data hardly gets any bigger than this! Between these benchmark estimates, which provide an accurate, comprehensive and detailed snapshot of the U.S. economy, annual and quarterly estimates are based on surveys also conducted by the Census Bureau, with about 150,000 and 35,000 reporting units respectively, as well as on administrative data (for instance from the Internal Revenue Service) and by extrapolations based on past patterns or other source data (for instance on employment, hours, and earnings from the Bureau of Labor Statistics). This process, whereby very detailed microeconomic information is aggregated into a coherent set of national accounts, produces a regular stream of GDP estimates and subsequent revisions. After the first, or advance, quarterly print, released about one month after the end of the quarter in question, a second and a third quarterly estimate are released in subsequent months, and comprehensive revisions follow, which incorporate methodological advances that update the accounts to reflect changes in the economy.

One of the primary considerations in the design of this release schedule, and of the data collection efforts that underpin it, is the trade-off between accuracy on the one hand and

¹⁰See Landefeld et al. (2008) for more details.

timeliness and frequency of the estimates on the other. Maximum accuracy is achieved with the benchmark releases, since they are based on a census, but this is only carried out every five years. At the other extreme, the advance release is available every quarter and with less than a month delay, but only about half of the included expenditures data reflect survey-based information for all three months of the quarter. The rest is based on information for two months and on extrapolations. As a result of these statistical “shortcuts,” which are the inevitable cost of a timely release, the initial estimates are subject to potentially sizable revisions as more comprehensive and reliable information is folded into the accounts.

The statistical imprecision inherent in the quarterly GDP estimates, together with the fact that even the first estimate is available with a delay of nearly a month, pose a significant challenge to policymakers and other observers with an interest in monitoring the state of the economy in real time. As a result, as we discussed in the previous section, most of these observers rely on alternative indicators of the health of the economy that become available over the course of the quarter to form a real-time view of economic developments. Table 3.1 contains a list of the releases by both government agencies and the private sector that contain the most widely followed of these indicators. These releases are followed closely not just by economists, but by market participants, people in business and the media. The bars in the first column of the table provide a measure of the relevance of each release based on the percentage of Bloomberg users that subscribe to related alerts.

[Insert Table 3.1 here.]

Perhaps the most prominent among these releases is the Bureau of Labor Statistics’ Employment Situation report, which is issued on the first Friday of every month, as described in the third column of Table 3.1. This report, which includes data on payroll employment, unemployment, earnings, and many other aspects of the labor market, is of independent interest because it provides an in depth picture of a particular segment of the economy, which is not covered in as much detail in the national accounts. Yet, the nature of business cycles, in which most sectors of the economy tend to move together, implies that good news for the labor market—or for manufacturing, construction, retail trade, and so on—usually reflects good news for the economy as a whole. Therefore, the information in the Employment Report, along with that contained in all the other releases listed in Table 3.1, can be used to extract a signal on the current overall level of economic activity well before the first GDP estimate is available.

Of course, this exercise is subject to a trade-off between accuracy and timeliness similar to the one we discussed above in relation to the successive GDP releases. None of the releases listed in Table 3.1 is quite as comprehensive in its coverage of economic activity as the NIPAs. Moreover, the surveys underlying the releases vary widely in size and hence in their statistical reliability. In general, indicators released closer to their reference period are bound to be less accurate. Therefore, no one indicator can be a silver bullet that solves the problem of accurately tracking the evolution of the economy in real time. A more promising approach is instead combining the information contained in the many available releases. Given the number of these releases, and the hundreds of statistics that they often include, designing such an approach is once again a big data challenge, essentially the same one faced by Kuznets in developing the national accounts: how to synthesize the complexity of the U.S. economy through one summary statistic. GDP provides an answer to this question based on accounting principles. Nowcasting addresses the same challenge through statistical modeling, as we will discuss in detail in the next two sections.

In principle, the nowcasting solution to this problem is straightforward. The data is summarized through a few common factors whose evolution is tracked in real time through filtering techniques. In practice, however, the implementation of this idea is complicated by the intricate nature of the information being tracked. As shown in Table 3.1, economic indicators are released on a nearly continuous basis over the course of a quarter. This trickling of information over time is often referred to as the data flow, but it is actually less smooth than the term might suggest, although it does follow an entirely predictable calendar. The earliest available information for the national economy on any given quarter is provided by the ISM manufacturing report for the first month of that quarter, which is released on the first business day of the following month (second row in Table 3.1). On that same day, the Census Bureau's Construction Spending report is also made available (first row of Table 3.1) but, unlike the ISM report, it refers to two months prior. Therefore, it does not carry relevant information for the current quarter until the third month of the quarter. Similar considerations hold for the international trade and manufacturers' shipments report. Next are two closely followed reports on the labor market. The most important is the already mentioned Employment Situation Report by the BLS, which is released on the first Friday of every month. Since 2006, this report has been preceded on Wednesday by the ADP National Employment Report. ADP is a large private payroll processing company that has assembled a nationally representative sample of firms among its clients, which allows it to estimate total payroll employment. From its relatively short track record, the ADP payroll estimate

appears to be noisier than that produced by the BLS. However, the fact that it is available two days earlier makes it a potentially useful input in any effort to track the economy in real time; it offers a nice illustration of the trade-off between accuracy and timeliness that we discussed above.

Table 3.1 provides a qualitative snapshot of the richness and complexity of the regularly available data tracked by economists and other observers, and that are accordingly used as an input in the New York Fed model. Figure 3.1 provides a complementary, more quantitative perspective on these same data, by plotting their joint evolution over time. The heat map on the horizontal plane highlights the degree of co-movement in the series, with more intense red denoting realizations further below the mean and brighter yellow realizations further above the mean. The red “ridges” in the early 1990s, early 2000s, and most notably in 2008 and 2009 emphasize quite clearly the three recessions in the sample. This visualization captures in a simple, unstructured way the co-movement that drives the data flow over the business cycle, but also the extent to which idiosyncratic factors drive each series in different directions at any given point in time. Nowcasting is essentially a structured, formal, and efficient way of extracting the business cycle signal and discarding the idiosyncratic “noise” from this rich and complex set of data. The next section introduces the econometric theory behind this exercise, while Section 5 presents more details on how these ideas are implemented in the New York Fed Staff Nowcast.

[Insert Figure 3.1 here.]

IV Dealing with big data: econometric models

As discussed extensively in the previous sections, monitoring macroeconomic conditions in real time is inherently a big data problem. It crucially relies on the availability and the exploitation of a large amount of complex data. Increasing complexity of the data leads to increasing complexity in the models, with a growing number of parameters to estimate. Indeed, dealing with large data sets using overly simplified models may lead to misspecification as important features are omitted. On the other hand, modeling the interaction among a large number of variables leads to a proliferation of parameters: that implies large estimation uncertainty which makes the results from traditional tools unreliable and unstable. This fact is often referred to as the “curse of dimensionality:” the modeler faces a trade-off

between excessive simplicity (leading to misspecification) and excessive complexity (leading to instabilities). The econometrics of big data aims at turning the curse of dimensionality into a blessing by capturing in a parsimonious manner the salient features of the interactions among many series.

From an econometric perspective, estimation is challenging whenever the number of parameters is large relative to the number of observations. This is known in statistics as the “large p , small n ” paradigm, where p stands for the number of variables and n indicates the number of observations. Given their long tradition in handling a large amount of heterogeneous and complex data with a short time span, it is not surprising that macroeconomists have pioneered the statistical analysis of big data. At the root of the recent statistical developments is the key insight of Burns and Mitchell that we discussed: the pervasiveness of common fluctuations across different sectors of the economy implies strong cross-sectional correlations, suggesting that the bulk of fluctuations are essentially driven by a few common sources. Dynamic factor models (DFMs) build on this basic fact to provide a parsimonious and yet suitable representation for the macroeconomic series; they are one of the main tools that macroeconomists today use to handle big data.

A dynamic factor model assumes that many observed variables $(y_{1,t}, \dots, y_{n,t})$ are driven by a few unobserved dynamic factors $(f_{1,t}, \dots, f_{r,t})$, while the features that are specific to individual series, such as measurement errors, are captured by idiosyncratic errors $(e_{1,t}, \dots, e_{n,t})$. The empirical model can be summarized in the following equation:

$$y_{i,t} = \lambda_{i,1}f_{1,t} + \dots + \lambda_{i,r}f_{r,t} + e_{i,t}, \text{ for } i = 1, \dots, n \quad (1)$$

which relates the data $y_{i,t}$ to the r latent common factors $f_{1,t}, \dots, f_{r,t}$ through the factor loadings $\lambda_{i,1}, \dots, \lambda_{i,r}$. All together, they give the common component $\sum_{j=1}^r \lambda_{i,j}f_{j,t}$, while the idiosyncratic component $e_{i,t}$ captures the movements specific to each variable i .

As discussed in Section 2, factor models have a long tradition in the statistical and econometric literature. However, the application to big data is relatively recent. The earliest contributions are Forni et al. (2000) and Stock and Watson (2002a,b) who introduced principal components estimators for large dynamic factor models in economics. Associated with these developments were also the earliest occurrences of the term big data in the academic context.¹¹ Simultaneously, West (2003) pioneered Bayesian inference with large factor

¹¹See Diebold (2012) for an interesting discussion on the origin(s) of the term ‘big data.’

models in statistics and introduced the “large p , small n ” paradigm, which translates into “large n , small T ” in our context, where T is the sample size. These pioneering papers have led to many further advances in this research field, which have been recently surveyed by Stock and Watson (2016).

As initially pointed out by Giannone et al. (2008), dynamic factor models are particularly suitable for nowcasting and monitoring macroeconomic conditions in real time. This is because these models are naturally cast in a state-space form and hence inference can be performed using Kalman filtering techniques, which in turn provide a convenient and natural framework for handling the irregularities of the data in real time (i.e., mixed frequencies and non synchronicity of the data releases) and updating the predictions. Indeed, the Kalman filter digests incoming data in a coherent and intuitive way: it updates the predictions of the model recursively by weighting the innovation components of incoming data on the basis of their timeliness and their quality. Moreover, as the model produces forecasts for all variables simultaneously, the analysis of the flow of data does not require piecing together many separate, unrelated models.

In order to conduct inference in DFMs using likelihood based methods and Kalman filtering techniques, the common factors and the idiosyncratic components are modeled as Gaussian autoregressive processes, which account for their serial correlation and persistence.

$$f_{j,t} = a_j f_{j,t-1} + u_{j,t}, \quad u_{j,t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}\left(0, \sigma_{u_j}^2\right) \quad \text{for } j = 1, \dots, r \quad (2)$$

$$e_{i,t} = \rho_i e_{i,t-1} + \varepsilon_{i,t}, \quad \varepsilon_{i,t} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}\left(0, \sigma_{\varepsilon_i}^2\right) \quad \text{for } i = 1, \dots, n \quad (3)$$

Equations 1, 2, and 3 form a state space model where the common factors and the idiosyncratic components are unobserved states. Equation 1 is known as measurement equation and links the data to the unobserved states. Equations 2 and 3, known as the transition equations, describe the dynamics of the system. In order to avoid the proliferation of parameters it is typical to maintain a parsimonious empirical specification for the idiosyncratic components and assume that they are cross-sectionally orthogonal.

A model of this kind was first used by Stock and Watson (1989) to extract a single common factor ($r = 1$) from a small set of monthly core indicators including employment, industrial production, sales, and income ($n = 4$). Mariano and Murasawa (2003) extended it to include GDP growth ($n = 5$) and Aruoba et al. (2009) to include weekly unemployment

claims ($n = 6$). This framework also encompasses simpler approaches to the construction of business cycle indexes. In particular, the extracted common factor corresponds to the principal components if the empirical model is constrained to be static (i.e., assuming serially uncorrelated factors ($a_1 = \dots = a_r = 0$) and idiosyncratic components ($\rho_1 = \dots = \rho_n = 0$) and with homogenous signal-to-noise ratio ($\sigma_{\varepsilon_1}^2 = \dots = \sigma_{\varepsilon_n}^2$). If there is only one factor ($r = 1$) loaded homogeneously by all variables ($\lambda_{1,1} = \dots = \lambda_{n,1}$), the extracted common factor corresponds to cross-sectional averages.

The use of likelihood-based methods for factor models with big data was advocated by Doz et al. (2012), who established the viability of the approach. It was shown that estimates of the common factors are not only consistent when the sample size T and the cross-sectional dimension n are large, but also robust to cross-sectional misspecification, time-series correlation of the idiosyncratic components, and non-Gaussianity. In this respect, this is a quasi-maximum likelihood estimator in the sense of White (1982). Importantly, no constraint is required on the number of series n that can be handled for a given sample size T , which assures that the approach is suitable for the “large n , small T ” paradigm.

In practice, the estimates can be conveniently computed iteratively using the Kalman smoother and the EM algorithm. The algorithm is initialized by computing principal components, and the model parameters are estimated by OLS regression, treating the principal components as if they were the true common factors. This is a good initialization especially with big data given that principal components are reliable estimates of the common factors. In the second step, given the estimated parameters, an updated estimate of the common factors is obtained using the Kalman smoother. Stopping at the second step gives the two-step estimate of the common factors used in Giannone et al. (2008) and studied by Doz et al. (2011). Maximum likelihood estimation is obtained by iterating the two steps until convergence, taking into account at each step the uncertainty related to the fact that factors are estimated.¹²

The model we described is the one that we use for the New York Fed Staff Nowcast, which we illustrate in the next section. We kept it as simple as possible in order to minimize pre-testing and specification choices, but it can be easily extended to include more lags in the autoregressive process and allow for dynamic interactions among common factors. Other extensions can be useful to further improve accuracy. D’Agostino et al. (2016) extended the model to accommodate heterogeneity in the lead-lag relationships of different indicators

¹²Bańbura and Modugno (2014) show how to perform the parameter estimation step in presence of arbitrary patterns of missing data.

along the business cycles; Kim and Nelson (2001) introduced time-varying parameters in the form of Markov Switching; Antolin-Diaz et al. (2017) introduced time variation in the intercept; and Marcellino et al. (2016) introduced time varying volatility.

It is worth emphasizing that in spite of, and indeed thanks to, its simplicity, this basic model has been successfully applied to nowcasting in many economies with very different characteristics. These include large developed economies and small open economies, as well as emerging-market and developing economies.¹³

Bayesian vector autoregression (BVAR) offers an alternative modeling framework for nowcasting and monitoring macroeconomic conditions with big data in real time. The basic idea consists in addressing the curse of dimensionality by using a parsimonious naïve prior to discipline the estimation of a very flexible, densely parameterized and complex model. Vector autoregressions (VARs) are the most general linear model and are widely used in macroeconomics: every variable depends on its own past and on the past of each other variable, and the pattern of correlation of the forecast errors in different variables is left unconstrained. In BVARs, this high level of complexity is combined with a naïve prior model that assumes that all the variables are independent white noise or random walks. These models have been advocated by the earliest proponents of VAR models in economics (Sims, 1980; Doan et al., 1984). Recent research has shown that they are strictly connected with factor models and are suitable for the analysis of big data (De Mol et al., 2008; Banbura et al., 2010). For a recent survey of Bayesian vector autoregressive models, see Karlsson (2013) and Koop (2017). BVARs are also suitable for nowcasting since they can be cast in a state-space form allowing for conveniently handling data in real time using filtering techniques, in the same way described for the factor model. This is an important line of research since Bayesian inference provides a coherent probabilistic framework that can be exploited to greatly reduce the number and importance of subjective choices, such as data transformations or the informativeness of the prior (see Giannone et al. (2015) and Carriero

¹³For instance, see Giannone et al. (2008), Banbura et al. (2012), Lahiri and Monokroussos (2013), and Liebermann (2014) for nowcasting the United States; Angelini et al. (2011) and Banbura et al. (2011) for the aggregate euro area economy alongside Runstler et al. (2009) for individual member countries; and Bragoli (2017) for Japan. See Chernis and Sekkel (2017) for Canada; Aastveit and Trovik (2012) and Luciani and Ricci (2014) for Norway; Matheson (2010) for New Zealand; and Anesti et al. (2017) for the UK. See Dahlhaus et al. (2015) for BRIC countries and Mexico; Bragoli et al. (2015) for Brazil; Bragoli and Fosten (2016) for India; Yiu and Chow (2010) for China; Caruso (2015) for Mexico; Luciani et al. (2017) for Indonesia; and Kabundi et al. (2016) for South Africa. For an extensive survey, see Banbura et al. (2011, 2012).

et al. (2015)).¹⁴

V Nowcasting in practice

In this section, we detail the platform used at the Federal Reserve Bank of New York to track U.S. real GDP growth (FRBNY, 2016).¹⁵

The New York Fed Staff Nowcast is based on the dynamic factor model described in Section 4 and incorporates all the releases listed in Table 3.1. As described in Section 3, these are the releases most widely followed by market participants. While our nowcasting model is parsimonious and can accommodate many data series, we include only the headlines of each release, those that ‘move’ markets and make front page news. We do not include disaggregated data since there are no substantial gains in prediction from including them, although they can be useful for interpretation.¹⁶ Therefore, for example, from the Employment Situation report, the model incorporates the unemployment rate and non-farm payrolls, but does not include information on employment by age or industry; similarly, the model includes only total indexes for industrial production and capacity utilization, disregarding sectoral disaggregation. We should note that financial variables are not included in the model: although they are timely information, they tend to be quite volatile and have a limited role in nowcasting GDP growth once a rich set of macroeconomic variables have been included (see Banbura et al., 2012; Knotek and Zaman, 2017). However, going beyond the prediction of the central tendency, financial conditions can provide valuable information, especially for monitoring downside risks (Adrian et al., 2016).

All variables are listed in the first column of Table 5.1. The colored box next to the series’ name indicates the category to which the series is grouped. Each series enters the model in stationary form: in most cases this requires simply including the series in the same

¹⁴These models are increasingly used for nowcasting, as in Schorfheide and Song (2015), McCracken et al. (2015), Brave et al. (2016), and Knotek and Zaman (2017).

¹⁵A similar platform is used at the Atlanta Fed for nowcasting U.S. GDP growth, called GDPNow (Higgins, 2014). That model also builds on Giannone et al. (2008); however, it combines the components’ forecasts with a sort of ‘bean counting,’ trying to mimic the key elements of the data construction machinery used by the BEA, that we discussed briefly in Section 3. Moreover, the GDPNow model construction does not target the systematic component of GDP growth, as the New York Fed model does, but rather aims at an accurate prediction of the BEA’s advance estimate of GDP growth.

¹⁶For a detailed discussion see Banbura et al. (2011, 2012).

units as it is tracked by markets.¹⁷ The last column of Table 5.1 indicates the units in which each series enters the model. In the end, the factor loading resolves into a single factor, the global factor (G) which affects all variables. In addition, a few local blocks are included to control for idiosyncrasies in particular subgroups of series; this can improve inference even though the model is robust to the presence of local correlations. Specifically, to model the local correlations in survey data, we include the soft block (S), on which only variables representing economic agents' perceptions and sentiments load. Two additional local blocks are included for real (R) and for labor (L) variables; the structure of factor loadings is given in the second column of Table 5.1.

[Insert Table 5.1 here.]

Figure 5.1 below reports the standardized data (dotted colored lines) along with the global factor (solid black line) estimated from the model specification of Table 5.1. It is clear that the common (global) factor captures the bulk of the covariation between the variables.

[Insert Figure 5.1 here.]

The New York Fed Staff Nowcast framework specified here is run and updated daily at 10 a.m. whenever new data releases are issued, and updates to the nowcast are published weekly every Friday at 11:15 a.m. on the New York Fed's public website.¹⁸ This daily model updating allows to read the real-time data flow and quantify how each data release contributes to updates in the forecasts of all other variables. The factors, and hence the nowcast, change if either the data changes or the model parameters are re-estimated. The data change constantly, not only because new data are released, but also because all data are revised. Model parameters are re-estimated at the beginning of every quarter using the most recent 15 years of data. In the nowcast detail table, 'parameter revisions' are changes to the nowcast due to the re-estimation while 'data revisions' reflect changes in already-released data; both are classified within the Other (gray) category.

Figure 5.2 reports the evolution of the nowcast of real GDP growth in 2016:Q4. This is the same Figure reported in the introduction, to which we add shades to provide information

¹⁷There are a few exceptions, though. Housing starts, for instance, is typically tracked in levels by market participants, but because of evident trends this series enters the model in monthly changes.

¹⁸<https://www.newyorkfed.org/research/policy/nowcast>. Note: The nowcast can be made available to the public because it is model-based and does not reflect the outlook of the New York Fed.

about forecasting uncertainty. Precisely, the shades represent the 68th percentile of the empirical distribution of the forecast errors evaluated in real time for the sample 2000 to 2016. We will discuss forecasting performance in more detail later, but it should be noticed here that the bands narrow as the quarter progresses and information accumulates. This suggests that the data contain useful information that the model is able to exploit in real time. Notice also the substantial uncertainty also present in the official release of GDP, as illustrated in the figure by the error bar around the release, which reflect data revisions. This uncertainty is similar in magnitude to that of the last model forecast. In a sense, the model predictions are as accurate as the first release in predicting the latest estimates of GDP growth.

The solid black line in the figure is the progression of the nowcast throughout the period of update, which starts one month before the reference quarter and comprises the three months of the reference quarter and one month after the reference quarter, until the BEA publishes the advance GDP release; after this release the nowcast for that quarter is no longer updated. Each black diamond on the line is the nowcast for the particular week indicated on the x-axis; the colored bar corresponding to that diamond represents the contribution of the surprises in the releases of that week to the change in the nowcast. Each bar has potentially many segments of different colors, which represent the net contributions of different categories of data. Accompanying this plot is a table which breaks down the colored bars into the contribution of the news in individual series. In the figure we report the details pertaining to the week ending on September 16, 2016. For each series, the table reports the day and time of the release (first column), reference period (third column), as well as the units in which the series is reported (fourth column). In the next four columns the table reports, respectively, the model prediction for the series [a], the actual value of the series in the release [b], the weight [c] that the model assigns to the surprise (or news) [$(b - a)$] in the nowcast of GDP growth, and the impact [$c(b - a)$]: how this surprise changes the nowcast.

[Insert Figure 5.2 here.]

Reading Figure 5.2, one sees that the initial model prediction for 2016:Q4 GDP growth on August 19, 2016 was around 2.0 percent. After an initial fall to a low of 1.2 percent due largely to negative surprises from survey, manufacturing, and retail and consumption data, the nowcast steadily increased until, in the middle of the quarter, on the week of November 18th, it jumped up almost one full percentage point to 2.6 percent, due to positive surprises

from housing data and retail and consumption data. This increase was partially reverted just a few weeks later, on December 16, due to negative news from manufacturing and housing data. The nowcast moved slightly upwards in the following six weeks and was last recorded at 2.0 percent on January 27, just before the advance GDP release. By comparison, the BEA advance estimate of real GDP growth was 1.9 percent (circle in Figure 5.2), and the latest official estimate was 2.1 percent (square).

We performed a comprehensive backtesting to evaluate the real-time performance of the model. We computed the nowcast recursively on real-time vintages of data reconstructed to replicate the data exactly as they were available at the time.¹⁹ The nowcast for any given week is therefore precisely what would have been obtained by a forecaster running the model at that time. As we do currently, the parameters were estimated recursively at the beginning of every quarter and the incoming data was automatically incorporated at the end of any given week.

The dots in Figure 5.3 report the errors of the real-time nowcast, relative to the most recent GDP data, for all quarters in the evaluation sample. The dots at horizon zero report the revision error of the advance GDP release. Considering the universe of errors (dots), we uncovered a preponderance of negative errors, i.e., an upward bias in the nowcast. This pattern is related to the much discussed issue of residual seasonality in first quarter real GDP growth.²⁰ Therefore, we partitioned the error distribution into first quarter errors (blue dots) and errors for all quarters except first quarters (black dots) and report in the figure 68th percentile error bands separately for first quarters (blue shaded area) and for all others quarters (gray shaded area). These bands are superimposed on the nowcast progression plot for 2016:Q4 in Figure 5.2, in order to illustrate a measure of forecast uncertainty. We also provide a progression plot for the 2017:Q1 nowcast in Figure 5.4. In the figures, we report 68 percentile error bars for the revision errors of advance GDP, separately for first quarters (blue error bar) and for all others quarters (black error bar).

Inspecting the shaded bands reveals that much of the upward bias comes from first quarter nowcasts, in line with the issue of residual seasonality just discussed. The asymmetry is evident not only in the nowcasts but also in the revision errors of advance GDP estimates.

¹⁹The weekly vintages go from the first week of January 2000 through the last week of January 2017.

²⁰The BEA reported its findings on residual seasonality and its plan to publish a not seasonally adjusted (NSA) GDP series; see Moulton and Cowan (2016). See also Charles E. Gilbert and Sahm (2015), Stark (2015), Glenn D. Rudebusch and Mahedy (2015), Groen and Russo (2015), Phillips and Wang (2016), Kliesen (2017), and Lunsford (2017) for reports from the Federal Reserve Board as well as from the district Federal Reserve Banks of Philadelphia, San Francisco, New York, Dallas, St. Louis, and Cleveland, respectively.

Beyond issues pertaining to residual seasonality, the error distributions clearly show the attributes of a good forecast. Contrary to the first quarters, in which the errors exhibit significant downward bias, the errors for the other quarters are generally distributed symmetrically around zero. Further, the bands get narrower as time goes on, indicating on average a more accurate prediction of GDP growth over the nowcasting period as more information about the economy is released. Finally, at the end of the nowcast updating period the bands are similar to the error bars for the advance GDP release, indicating that the uncertainty surrounding the final nowcast made for each quarter is similar to that of the BEA's first estimate in predicting the true value of aggregate output growth in the economy.

[Insert Figure 5.3 here.]

[Insert Figure 5.4 here.]

We conclude this analysis by asking what are the most important variables driving the nowcast, when and why. Figure 5.5 reports the average (absolute) weekly impact of each series, grouped by category, computed in real-time over the sample 2000 to 2016. From this figure three features are evident. First, the most prominent colors are blue, green, orange and red, indicating that survey, consumption, manufacturing, and housing data are the main contributors to changes in the nowcast. Second, the bell shape of the plot indicates that the most useful information for the nowcast arrives in the middle of the nowcasting period, when data for the reference period first become available. Conversely, surprises move the nowcast less both at the beginning and at the end of the period: at the beginning, because at that time signals for GDP growth are still too weak, and at the end because by then most useful information has become available and there is little room for improvement. Third, we see from the contribution of surveys that soft data have a large impact at the beginning of each nowcasting timeframe. Later, as more information accumulates, their impact diminishes and hard data become more important. This confirms that timeliness is just as important as quality of the data.

Overall, each of the series tracked by the nowcast provides relevant information at various updating horizons. The combined information allows us to track the economy during each quarter and interpret the changes in nowcast with increasing precision, highlighting the importance of closely monitoring and continuously updating the outlook of the economy via the real-time data flow.

[Insert Figure 5.5 here.]

VI Conclusion

In this paper we illustrate the application of recent statistical techniques for the construction of an automated platform to process the real-time data flow—nowcasting, which we place in the context of various approaches developed over time to monitor and measure economic conditions.

Nowcasting is a relatively new field in time-series econometrics, and it is likely to continue to be developed on many fronts.

First of all, jointly modeling macroeconomic and financial conditions would provide an interface between finance and the macroeconomy. This would present a coherent framework to study the mechanisms through which macroeconomic news are transmitted to financial markets. Furthermore, it would allow to go beyond the prediction of the central tendency toward nowcasting vulnerabilities and risks to the outlook, along the line of Adrian et al. (2016).

Second, nowcasting can be developed in a structural environment. Monti et al (2016) proposed a nowcasting framework for a dynamic stochastic general equilibrium model. A benefit of this analysis is that it would allow to compute real-time estimates of model-based variables that are not directly observable, such as the output gap (which captures the difference between actual GDP and its potential value) and the natural rate of interest. Reading the data flow through the lens of a structural model would also make it possible to identify meaningful shocks to the economy in real time.

Finally, it is important to continue to refine the communication and sharing of nowcasting, a step we have taken by publishing the interactive version of this model on the New York Fed’s public website. This will foster interaction with other analysts and forecasters and help maintain the development of the model attuned to changes in market practices.

References

Aastveit, Knut and Torres Trovik, “Nowcasting norwegian GDP: the role of asset prices in a small open economy,” *Empirical Economics*, February 2012, 42 (1), 95–119.

- Adrian, Tobias, Nina Boyarchenko, and Domenico Giannone**, “Vulnerable growth,” Staff Reports 794, Federal Reserve Bank of New York September 2016.
- Altavilla, Carlo, Domenico Giannone, and Michele Modugno**, “Low Frequency Effects of Macroeconomic News on Government Bond Yields,” Finance and Economics Discussion Series 2014-52, Board of Governors of the Federal Reserve System (U.S.) June 2014.
- Anesti, Nikoleta, Ana Galvao, and Silvia Miranda-Agrippino**, “Uncertain Kingdom: A New Framework to Nowcast GDP and its Revisions,” 2017. Bank of England.
- Angelini, Elena, Gonzalo Camba-Mendez, Domenico Giannone, Lucrezia Reichlin, and Gerhard Rünstler**, “Short-term forecasts of euro area GDP growth,” *The Econometrics Journal*, 2011, 14 (1).
- Antolin-Diaz, Juan, Thomas Drechsel, and Ivan Petrella**, “Tracking the Slowdown in Long-Run GDP Growth,” *The Review of Economics and Statistics*, May 2017, 99 (2), 343–356.
- Aruoba, S. Boragan, Francis X. Diebold, and Chiara Scotti**, “Real-Time Measurement of Business Conditions,” *Journal of Business & Economic Statistics*, 2009, 27 (4), 417–427.
- Banbura, Marta, Domenico Giannone, and Lucrezia Reichlin**, “Large Bayesian vector auto regressions,” *Journal of Applied Econometrics*, 2010, 25 (1), 71–92.
- , – , and – , “Nowcasting,” in M. Clemens and D. Hendry, eds., *Oxford Handbook on Economic Forecasting*, 2011, pp. 193–224.
- , – , **Michele Modugno, and Lucrezia Reichlin**, “Now-Casting and the Real-Time Data Flow,” 2012. Handbook of Economic Forecasting, G. Elliott and A. Timmermann, eds., Volume 2.
- Bartolini, Leonardo, Linda S Goldberg, and Adam Sacarny**, “How economic news moves markets,” 2008.
- Bañbura, Marta and Michele Modugno**, “Maximum Likelihood Estimation Of Factor Models On Datasets With Arbitrary Pattern Of Missing Data,” *Journal of Applied Econometrics*, 01 2014, 29 (1), 133–160.
- Bragoli, Daniela**, “Now-casting the Japanese economy,” *International Journal of Forecasting*, 2017, 33 (2), 390–402.
- and **Jack Fosten**, “Nowcasting Indian GDP,” University of East Anglia School of Economics Working Paper Series 2016-06, School of Economics, University of East Anglia, Norwich, UK. June 2016.
- , **Luca Metelli, and Michele Modugno**, “The importance of updating: Evidence from a Brazilian nowcasting model,” *OECD Journal: Journal of Business Cycle Measurement and Analysis*, 2015, 2015 (1), 5–22.
- Brave, Scott, R. Andrew Butters, and Alejandro Justiniano**, “Forecasting Economic Activity with Mixed Frequency Bayesian VARs,” Working Paper Series WP-2016-5, Federal Reserve Bank of Chicago May 2016.
- Bry, Gerhard and Charlotte Boschan**, *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs* NBER Books, National Bureau of Economic Research, Inc, 1971.

- Burns, Arthur F. and Wesley C. Mitchell**, eds, *Measuring Business Cycles*, NBER Book Series Studies in Business Cycles, 1946.
- Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino**, “Bayesian VARs: Specification Choices and Forecast Accuracy,” *Journal of Applied Econometrics*, 01 2015, 30 (1), 46–73.
- Caruso, Alberto**, “Nowcasting Mexican GDP,” ECARES Working Papers, ECARES October 2015.
- Chernis, Tony and Rodrigo Sekkel**, “A dynamic factor model for nowcasting Canadian GDP growth,” *Empirical Economics*, Aug 2017, 53 (1), 217–234.
- D’Agostino, Antonello, Domenico Giannone, Michele Lenza, and Michele Modugno**, “Nowcasting Business Cycles: a Bayesian Approach to Dynamic Heterogeneous Factor Models,” in “Dynamic Factor Models,” Emerald Group Publishing Limited, 2016, pp. 569–594.
- Dahlhaus, Tatjana, Justin-Damien Guénette, and Garima Vasishtha**, “Nowcasting BRIC+M in Real Time,” Staff Working Papers 15-38, Bank of Canada 2015.
- Diebold, Francis X**, “On the Origin (s) and Development of the Term ‘Big Data’,” 2012. University of Pennsylvania.
- Doan, Thomas, Robert Litterman, and Christopher A. Sims**, “Forecasting and conditional projection using realistic prior distributions,” *Econometric Reviews*, 1984, 3, 1–100.
- Doz, Catherine, Domenico Giannone, and Lucrezia Reichlin**, “A two-step estimator for large approximate dynamic factor models based on Kalman filtering,” *Journal of Econometrics*, September 2011, 164 (1), 188–205.
- , – , and – , “A quasi–maximum likelihood approach for large, approximate dynamic factor models,” *Review of economics and statistics*, 2012, 94 (4), 1014–1024.
- Forni, Mario, Marc Hallin, Marco Lippi, and Lucrezia Reichlin**, “The Generalized Dynamic-Factor Model: Identification And Estimation,” *The Review of Economics and Statistics*, November 2000, 82 (4), 540–554.
- FRBNY**, “Nowcasting Report,” 2016.
- Giannone, Domenico, Lucrezia Reichlin, and David Small**, “Nowcasting: The real-time informational content of macroeconomic data,” *Journal of Monetary Economics*, 2008, 55 (4), 665–676.
- , – , and **Luca Sala**, “Monetary Policy in Real Time,” in “NBER Macroeconomics Annual 2004, Volume 19” NBER Chapters, National Bureau of Economic Research, Inc, Jan-Jun 2005, pp. 161–224.
- , **Michele Lenza, and Giorgio E. Primiceri**, “Prior Selection for Vector Autoregressions,” *The Review of Economics and Statistics*, May 2015, 97 (2), 436–451.
- Groen, Jan and Patrick Russo**, “The Myth of First-Quarter Residual Seasonality,” *New York Fed Liberty Street Economics*, 2015.
- Gürkaynak, Refet S. and Jonathan H. Wright**, “Identification and Inference Using Event Studies,” *Manchester School*, 09 2013, 81, 48–65.

- , **Brian Sack**, and **Eric Swanson**, “The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models,” *American Economic Review*, March 2005, *95* (1), 425–436.
- Hamilton, James D.**, “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle,” *Econometrica*, March 1989, *57* (2), 357–384.
- Harding, D. and A. Pagan**, *The Econometric Analysis of Recurrent Events in Macroeconomics and Finance* The Econometric and Tinbergen Institutes Lectures, Princeton University Press, 2016.
- Harding, Don and Adrian Pagan**, “Dissecting the cycle: a methodological investigation,” *Journal of Monetary Economics*, March 2002, *49* (2), 365–381.
- Higgins, Patrick C.**, “GDPNow: A Model for GDP “Nowcasting”,” FRB Atlanta Working Paper 2014-7, Federal Reserve Bank of Atlanta July 2014.
- Kabundi, Alain, Elmarie Nel, and Franz Ruch**, “Nowcasting Real GDP growth in South Africa,” Working Paper Series WP/16/01, South African Reserve Bank February 2016.
- Karlsson, Sune**, *Forecasting with Bayesian Vector Autoregression*, Vol. 2 of *Handbook of Economic Forecasting*, Elsevier, January
- Kim, Chang-Jin and Charles R Nelson**, “A Bayesian Approach to Testing for Markov-Switching in Univariate and Dynamic Factor Models,” *International Economic Review*, November 2001, *42* (4), 989–1013.
- Kliesen, Kevin L.**, “Residual Seasonality: The Return of an Old First-Quarter Friend?,” *St. Louis Fed On the Economy*, 2017.
- Knotek, Edward S. and Saeed Zaman**, “Financial Nowcasts and Their Usefulness in Macroeconomic Forecasting,” Working Paper 1702, Federal Reserve Bank of Cleveland March 2017.
- Koop, Gary**, “Bayesian methods for empirical macroeconomics with big data,” *Review of Economic Analysis*, March 2017.
- Lahiri, Kajal and George Monokroussos**, “Nowcasting US GDP: The role of ISM business surveys,” *International Journal of Forecasting*, 2013, *29* (4), 644–658.
- Landefeld, J. Steven, Eugene P. Seskin, and Barbara M. Fraumeni**, “Taking the Pulse of the Economy: Measuring GDP,” *Journal of Economic Perspectives*, June 2008, *22* (2), 193–216.
- Liebermann, Joelle**, “Real-Time Nowcasting of GDP: A Factor Model vs. Professional Forecasters,” *Oxford Bulletin of Economics and Statistics*, December 2014, *76* (6), 783–811.
- Luciani, Matteo and Lorenzo Ricci**, “Nowcasting Norway,” *International Journal of Central Banking*, December 2014, *10* (4), 215–248.
- , **Madhavi Pundit, Arief Ramayandi, and Giovanni Veronese**, “Nowcasting Indonesia,” *Empirical Economics*, Jun 2017.
- Lunsford, Kurt G.**, “Lingering Residual Seasonality in GDP Growth,” *Cleveland Fed Economic Commentary*, March 2017, *2017* (6).
- Marcellino, Massimiliano**, “Chapter 16 Leading Indicators,” *Handbook of Economic Forecasting*, 2006, *1*, 879 – 960.

- , **Mario Porqueddu, and Fabrizio Venditti**, “Short-Term GDP Forecasting With a Mixed-Frequency Dynamic Factor Model With Stochastic Volatility,” *Journal of Business & Economic Statistics*, January 2016, *34* (1), 118–127.
- Mariano, Roberto S. and Yasutomo Murasawa**, “A new coincident index of business cycles based on monthly and quarterly series,” *Journal of Applied Econometrics*, 2003, *18* (4), 427–443.
- Matheson, Troy D.**, “An analysis of the informational content of New Zealand data releases: The importance of business opinion surveys,” *Economic Modelling*, January 2010, *27* (1), 304–314.
- McCracken, Michael W., Michael T. Owyang, and Tatevik Sekhposyan**, “Real-Time Forecasting with a Large, Mixed Frequency, Bayesian VAR,” Working Papers 2015-30, Federal Reserve Bank of St. Louis October 2015.
- Mol, Christine De, Domenico Giannone, and Lucrezia Reichlin**, “Forecasting using a large number of predictors: Is bayesian shrinkage a valid alternative to principal components?,” *Journal of Econometrics*, October 2008, *146* (2), 318–328.
- Morin, Andrew D. Paciorek Charles E. Gilbert Norman J. and Claudia R. Sahm**, “Residual Seasonality in GDP,” *FEDS Notes*, 2015.
- Moulton, Brent R. and Benjamin D. Cowan**, “Residual Seasonality in GDP and GDI: Findings and Next Steps,” *Survey of Current Business*, July 2016, *96* (7).
- OECD**, *Business cycle indicators handbook* 2012.
- Phillips, Keith R. and Jack Wang**, “Residual Seasonality in U.S. GDP Data,” *Dallas Fed Working Papers*, 2016.
- Rudebusch, Daniel Wilson Glenn D. and Tim Mahedy**, “The Puzzle of Weak First-Quarter GDP Growth,” *FRBSF Economic Letter*, 2015, (16).
- Runstler, G., K. Barhoumi, S. Benk, R. Cristadoro, A. Den Reijer, A. Jakaitiene, P. Jelonek, A. Rua, K. Ruth, and C. Van Nieuwenhuyze**, “Short-term forecasting of GDP using large datasets: a pseudo real-time forecast evaluation exercise,” *Journal of Forecasting*, 2009, *28* (7), 595–611.
- Sargent, Thomas J. and Christopher A. Sims**, “Business Cycle Modeling Without Pretending to Have Too Much A-Priori Economic Theory,” in C. Sims et al., ed., *New Methods in Business Cycle Research*, Federal Reserve Bank of Minneapolis, 1977.
- Schorfheide, Frank and Dongho Song**, “Real-Time Forecasting With a Mixed-Frequency VAR,” *Journal of Business & Economic Statistics*, July 2015, *33* (3), 366–380.
- Sims, Christopher A.**, “Macroeconomics and Reality,” *Econometrica*, 1980, *48* (1), 1–48.
- , “The Role of Models and Probabilities in the Monetary Policy Process,” *Brookings Papers on Economic Activity*, 2002, *2002* (2), 1–40.
- Stark, Tom**, “The Myth of First-Quarter Residual Seasonality,” *Philadelphia Fed Research Rap*, 2015.
- Stock, James H. and Mark W. Watson**, “New Indexes of Coincident and Leading Economic Indicators,” in “NBER Macroeconomics Annual 1989, Volume 4” NBER Chapters, National Bureau of Economic Research, Inc, 1989, pp. 351–409.

- **and** –, “Forecasting Using Principal Components from a Large Number of Predictors,” *Journal of the American Statistical Association*, 2002, *97*, 147–162.
 - **and** –, “Macroeconomic Forecasting Using Diffusion Indexes,” *Journal of Business and Economics Statistics*, 2002, *20*, 147–162.
 - **and** –, “Indicators for Dating Business Cycles: Cross-History Selection and Comparisons,” *American Economic Review*, May 2010, *100* (2), 16–19.
 - **and** –, “Estimating turning points using large data sets,” *Journal of Econometrics*, 2014, *178* (P2), 368–381.
 - **and** –, “Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics,” *Handbook of Macroeconomics*, 2016, *2*, 415 – 525.
 - **and** –, “Twenty Years of Time Series Econometrics in Ten Pictures,” *Journal of Economic Perspectives*, May 2017, *31* (2), 59–86.
- Watson, Mark W**, “Comment on “Monetary Policy in Real Time,” by Domenico Giannone, Lucrezia Reichlin, and Luca Sala,” in “NBER Macroeconomics Annual 2004, Volume 19” NBER Chapters, National Bureau of Economic Research, Inc, Jan-Jun 2005, pp. 216–221.
- West, Mike**, “Bayesian factor regression models in the “large p, small n” paradigm,” in JM Bernardo, MJ Bayarri, JO Berger, AP Dawid, D Heckerman, AFM Smith, and M West, eds., *Bayesian statistics*, Vol. 7, Oxford University Press, 2003, pp. 733–742.
- White, Halbert**, “Maximum Likelihood Estimation of Misspecified Models,” *Econometrica*, January 1982, *50* (1), 1–25.
- Yiu, Matthew S. and Kenneth K. Chow**, “Nowcasting Chinese GDP: information content of economic and financial data,” *China Economic Journal*, 2010, *3* (3), 223–240.

A Appendix

List of Figures




Figure 1.1 New York Fed Staff Nowcast, 2016:Q4.

Source. Federal Reserve Bank of New York “Nowcasting Report” March 3, 2017. Retrieved at <https://www.newyorkfed.org/research/policy/nowcast.html>.




Figure 3.1 Big data in macroeconomics.

The three-dimensional surface plot displays the standardized time series for the major economic indicators since 1985, which are colored by category as indicated in the legend. The heat map on the horizontal plane shows positive (yellow) and negative (red) readings of the data.

Source. Authors' calculations, Haver Analytics, Federal Reserve Economic Database (FRED).





Figure 5.1 Estimated global factor.

The solid black line illustrates the global factor estimated from the dynamic factor model. The colored dotted lines report all data series that enter the model, plotted in standard deviations from their mean. Gray shaded areas indicate NBER recessions.

Source. Authors' calculations, Haver Analytics, Federal Reserve Economic Database (FRED), National Bureau of Economic Research.



		Data Releases	
◆	New York Fed Staff Nowcast	■	Housing and construction
○	Advance GDP estimate	■	Manufacturing
□	Latest GDP estimate	■	Surveys
		■	Income
		■	Labor
		■	International trade
		■	Retail and consumption
		■	Others

Update	Release Date	Data Series	Reference Period	Units	Forecast	Actual	Weight	Impact	Nowcast GDP Growth
					[a]	[b]	[c]	[c(b - a)]	
Sep 09	8:30AM Sep 14	Import price index	Aug	MoM % chg.	-0.119	-0.248	0.029	-0.004	1.69
	8:30AM Sep 14	Export price index	Aug	MoM % chg.	0.058	-0.826	0.053	-0.047	
	8:30AM Sep 15	Retail sales and food services	Aug	MoM % chg.	0.489	-0.295	0.206	-0.161	
	8:30AM Sep 15	PPI: Final demand	Aug	MoM % chg.	-0.021	0.000	0.040	0.001	
	8:30AM Sep 15	Empire State Mfg. Survey	Sep	Index	-0.583	-1.99	0.014	-0.019	
	8:30AM Sep 15	Philadelphia Fed Mfg. Business Outlook	Sep	Index	-1.50	12.8	0.016	0.228	
	9:20AM Sep 15	Industrial production index	Aug	MoM % chg.	-0.051	-0.433	0.415	-0.158	
	9:20AM Sep 15	Capacity utilization	Aug	Ppt. chg.	-0.100	-0.357	0.520	-0.134	
	8:30AM Sep 16	CPI-U: All items	Aug	MoM % chg.	0.048	0.199	0.062	0.009	
	8:30AM Sep 16	CPI-U: All items less food and energy	Aug	MoM % chg.	0.142	0.252	0.021	0.002	
			Data revisions					-0.037	
Sep 16								1.37	

Figure 5.2 New York Fed Staff Nowcast, 2016:Q4 – detail table and empirical uncertainty.

The solid black line indicates the nowcast of annualized real GDP growth for 2016:Q4. For comparison, the circle and square at the right of the chart show respectively the advance and latest official estimates of real GDP growth. Colored bars reflect the weekly contribution of data releases to changes in the nowcast, aggregated by categories given in the key. The uncertainty bands around the solid black line represent the the 68th percentile of the empirical error distribution, and the error bar around the circle represents the 68th percentile of revision errors from the advance estimate of GDP to the latest, both of which are computed over the 2000-2016 sample, as illustrated in Figure 5.3. The table details the real-time data flow for a given week, where the weighted average of the news during the week gives the forecast revision from one week to another.

Source. Authors' calculations, Federal Reserve Bank of New York "Nowcasting Report."

- nowcast errors for first quarters ■ 68% bands for first quarters
- nowcast errors for all other quarters ■ 68% bands for all other quarters




Figure 5.3 Empirical nowcast error distribution.

Errors are computed as the difference between the nowcast computed in real time and observed real GDP growth as of the latest release. The dots indicate errors for individual quarters in the evaluation sample (2000:Q1 to 2016:Q4). The x-axis refers to the point in the quarter when the nowcasts were made, measured in the number of weeks before the first official GDP release. Motivated by residual seasonality in first-quarter GDP growth, errors are partitioned in two groups: errors for first quarters and errors for all other quarters; shaded areas show the 16th and 84th percentiles of the empirical distribution of the errors in the respective partitions. At week 0 the errors refer to the difference between the advance release of GDP and the latest release, with bars indicating the 16th and 84th percentiles of their empirical distribution.

Source. Authors' calculations, Federal Reserve Bank of New York "Nowcasting Report."




Figure 5.4 New York Fed Staff Nowcast, 2017:Q1 – empirical uncertainty.

The solid black line indicates the nowcast of annualized real GDP growth for 2017:Q1. For comparison, the circle and square at the right of the chart show respectively the advance and latest official estimates of real GDP growth. Colored bars reflect the weekly contribution of data releases to changes in the nowcast, aggregated by categories given in the key. The blue uncertainty bands around the solid black line represent the the 68th percentile of the empirical error distribution, and the error bar around the circle represents the the 68th percentile of revision errors from the advance estimate of GDP to the latest, both of which are computed over the 2000-2016 sample, as illustrated in Figure 5.3. The table details the real-time data flow for a given week, where the weighted average of the news during the week gives the forecast revision from one week to another.

Source. Authors' calculations, Federal Reserve Bank of New York "Nowcasting Report."

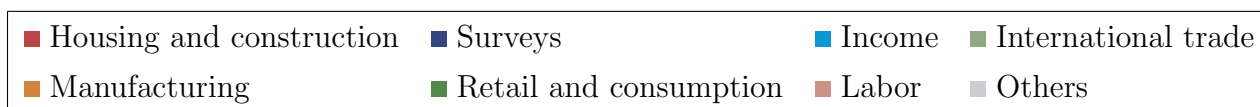



Figure 5.5 Historical impact of data releases on the nowcast.

The colored bars report the absolute impact of each data series on the nowcast, grouped by category. The x-axis refers to the point in the quarter when the nowcasts were made, measured in the number of weeks before the first official GDP release. The impacts are computed in real time and averaged over the evaluation sample (2000:Q1 to 2016:Q4).

Source. Authors' calculations.

List of Tables

	horizon (h quarters ahead)					
	-1	0	1	2	3	4
BEA	1.61					
Naïve		2.43	2.46	2.55	2.55	2.55
SPF		1.94***	2.21**	2.40	2.47	2.52

Table 2.1 Horizons of predictability.

The table reports root-mean-square errors (RMSEs) for h -quarter ahead forecasts (SPF and naïve) and estimates (BEA) of GDP growth, over the sample 1985Q1 to 2014Q4. Errors are computed as the difference between the latest GDP growth release and various projections of GDP. For BEA, the errors refer to the difference between the first official estimate, released one month after the quarter ends ($h = 1$), and the latest available estimate. SPF refers to the median forecasts from the Survey of Professional Forecasters, which is conducted in the second month of every quarter. Naïve refers to iterative forecasts from an autoregressive model calculated by the Philadelphia Federal Reserve. See “Forecast Error Statistics for the Survey of Professional Forecasters” on their public website for more details. Based on the Diebold-Mariano test, SPF forecasts that are significantly more accurate than the naïve model at the 1% (***), 5% (**), and 10% (*) level are indicated.

Source. Authors’ calculations, Bureau of Economic Analysis, Federal Reserve Bank of Philadelphia.

Release	Timing	Delay	Source
■ Construction Spending	first business day of the month, two months prior	33	Census Bureau
■ ISM Manufacturing Report on Business	first business day of the month, one month prior	3	ISM
■ ISM Non-Manufacturing Report on Business	third business day of the month, one month prior	5	ISM
■ U.S. International Trade in Goods and Services	first full week of the month, two months prior	35	BEA, Census Bureau
■ Manufacturers' Shipments, Inventories, and Orders	first week of the month, two months prior	35	Census Bureau
■ ADP National Employment Report	first Wednesday of the month, one month prior	5	ADP
■ Employment Situation	first Friday of the month, one month prior	7	BLS
■ Manufacturing and Trade Inventories	first full week of the month, two months prior	44	Census Bureau
■ Job Openings and Labor Turnover	second week of the month, two months prior	42	BLS
■ U.S. Import and Export Price Indexes	middle of the month, one month prior	13	BLS
■ Retail Trade	ninth business day of the month, one month prior	14	Census Bureau
■ Producer Price Index	middle of the month, one month prior	14	BLS
■ Wholesale Trade	middle of the month, two months prior	37	Census Bureau
■ Empire State Manufacturing Survey	15th of the month, current month	-14	New York Fed
■ Manufacturing Business Outlook Survey	third Thursday of the month, current month	-11	Philadelphia Fed
■ Industrial Production and Capacity Utilization	middle of the month, one month prior	17	Federal Reserve Board
■ Consumer Price Index	middle of the month, one month prior	18	BLS
■ New Residential Construction	12th business day of the month, one month prior	16	Census Bureau
■ Advance Economic Indicators	last week of the month, one month prior	28	Census Bureau
■ New Residential Sales	17th business day of the month, one month prior	26	Census Bureau
■ Advance Durable Goods	third week of the month, one month prior	26	Census Bureau
■ Personal Income and Outlays	last week of the month, one month prior	30	BEA
■ Gross Domestic Product	last week of the month, prior quarter	28	BEA
■ Productivity and Costs	first week of the month, prior quarter	34	BLS

Table 3.1 Macroeconomic data releases.

The table shows the major macroeconomic releases in the United States, ordered by when the release is issued over time. Each release is accompanied by the Bloomberg relevance index, a timing of the publication of the release, the delay (in days) of the release with respect to the end of the reference period and the source of the release. Delays are based on releases in 2017 but may vary over time.

Source. Bloomberg LP.

Data Series	Block	Units
	G S R L	
■ All employees: Total nonfarm	■ ■ ■ ■	Level change (thousands)
■ Real gross domestic product	■ ■ ■ ■	QoQ % change (annual rate)
■ ISM mfg.: PMI composite index	■ ■ ■ ■	Index
■ CPI-U: All items	■ ■ ■ ■	MoM % change
■ Manufacturers new orders: Durable goods	■ ■ ■ ■	MoM % change
■ Retail sales and food services	■ ■ ■ ■	MoM % change
■ New single family houses sold	■ ■ ■ ■	MoM % change
■ Housing starts	■ ■ ■ ■	MoM % change
■ Civilian unemployment rate	■ ■ ■ ■	Ppt. change
■ Industrial production index	■ ■ ■ ■	MoM % change
■ PPI: Final demand	■ ■ ■ ■	MoM % change
■ ADP nonfarm private payroll employment	■ ■ ■ ■	Level change (thousands)
■ Empire State Mfg. Survey: General business conditions	■ ■ ■ ■	Index
■ Merchant wholesalers: Inventories: Total	■ ■ ■ ■	MoM % change
■ Value of construction put in place	■ ■ ■ ■	MoM % change
■ Philly Fed Mfg. business outlook: Current activity	■ ■ ■ ■	Index
■ Import price index	■ ■ ■ ■	MoM % change
■ ISM nonmanufacturing: NMI composite index	■ ■ ■ ■	Index
■ ISM mfg.: Prices index	■ ■ ■ ■	Index
■ Building permits	■ ■ ■ ■	Level change (thousands)
■ Capacity utilization	■ ■ ■ ■	Ppt. change
■ PCE less food and energy: Chain price index	■ ■ ■ ■	MoM % change
■ CPI-U: All items less food and energy	■ ■ ■ ■	MoM % change
■ Inventories: Total business	■ ■ ■ ■	MoM % change
■ Nonfarm business sector: Unit labor cost	■ ■ ■ ■	QoQ % change (annual rate)
■ JOLTS: Job openings: Total	■ ■ ■ ■	Level change (thousands)
■ Real personal consumption expenditures	■ ■ ■ ■	MoM % change
■ PCE: Chain price index	■ ■ ■ ■	MoM % change
■ ISM mfg.: Employment index	■ ■ ■ ■	Index
■ Export price index	■ ■ ■ ■	MoM % change
■ Manufacturers shipments: Durable goods	■ ■ ■ ■	MoM % change
■ Mfrs. unfilled orders: All manufacturing industries	■ ■ ■ ■	MoM % change
■ Manufacturers inventories: Durable goods	■ ■ ■ ■	MoM % change
■ Real gross domestic income	■ ■ ■ ■	QoQ % change (annual rate)
■ Real disposable personal income	■ ■ ■ ■	MoM % change
■ Exports: Goods and services	■ ■ ■ ■	MoM % change
■ Imports: Goods and services	■ ■ ■ ■	MoM % change

■ Housing and construction	■ Manufacturing	■ Surveys	■ Retail and consumption
■ Income	■ Labor	■ International trade	■ Others

Table 5.1 New York Fed Staff Nowcast model specification.

The table reports series included in the model (first column), the blocks they load on in the dynamic factor model (second column), and the units in which they enter the model (third column). G, S, R, and L indicate the global, soft, real and labor factors, respectively.