Preliminary

NEW EVIDENCE ON THE IMPACT OF FINANCIAL CRISES IN ADVANCED COUNTRIES

Christina D. Romer

David H. Romer

University of California, Berkeley

October 12, 2014

We are grateful to the Center for Equitable Growth at the University of California, Berkeley for financial support, to Demian Pouzo and James Powell for helpful discussions, and to Marc Dordal i Carreras and Dmitri Koustas for research assistance.

NEW EVIDENCE ON THE IMPACT OF FINANCIAL CRISES IN ADVANCED COUNTRIES

ABSTRACT

This paper revisits the aftermath of financial crises in advanced countries in the decades before the Great Recession. We construct a new series on financial distress in 24 OECD countries for the period 1967–2007. The series is based on narrative assessments of the health of countries' financial systems that were made in real time; and it classifies financial distress on a relatively fine scale, rather than treating it as a 0-1 variable. We find little support for the conventional wisdom that the output declines following financial crises are uniformly large and long-lasting. Rather, the declines are highly variable, on average only moderate, and often temporary. One important driver of the variation in outcomes across crises appears to be the severity and persistence of the financial distress itself: when distress is particularly extreme or continues for an extended period, the aftermath of a crisis is worse.

Christina D. Romer
Department of Economics
University of California, Berkeley
Berkeley, CA 94720-3880
cromer@econ.berkeley.edu

David H. Romer
Department of Economics
University of California, Berkeley
Berkeley, CA 94720-3880
dromer@econ.berkeley.edu

I. INTRODUCTION

Recent work on financial crises by Reinhart and Rogoff (2009) and others has led to a new conventional wisdom: recessions following financial crises are typically severe and long-lasting. Journalists and policymakers have embraced this view as an explanation for the large downturn and the slow recovery in the United States and Europe following the 2008 financial crisis. And economists have begun to refine their models and textbooks to reflect this new consensus.

Issues. Several features of the existing research, however, suggest that the widespread acceptance of this view may be premature, particularly in the case of advanced countries. One issue is that much of the evidence of severe effects from financial crises comes from the pre-World War II period and from emerging economies. Whether the aftermaths of crises in modern, advanced economies are similarly dire has been much less thoroughly analyzed.

A second issue involves the identification of financial crises. The specification of what constitutes a crisis is often imprecise. It can sometimes have a "we know one when we see one" feel. To the degree that the specification is made concrete, it is often based on indirect indicators, such as whether there were bank failures or government bailouts of the financial system, which may be highly imperfect indicators of disruptions to credit supply. Another feature of the identification of crises is that it is typically done by scholars who know what happened to output afterward. The interaction of this ex post identification and the somewhat vague specification of a crisis makes it possible for inaccuracy to creep in. There is perhaps a natural tendency to look a little harder for a financial crisis before a known severe recession, or to identify the start of a crisis earlier than was perhaps apparent in real time. This may skew the empirical results toward finding that the effects of financial crises are particularly severe and long-lasting.

A third issue is that existing crisis chronologies almost all use a 0-1 classification: either a country experienced a crisis or it did not. A few studies differentiate between systemic and nonsystemic crises, but do not go further than that. This binary classification surely obscures

some important information about the variation in the severity of crises. It also means that errors in classification are likely very consequential. As a result, the empirical estimates of the real impact of crisis are likely to be imprecise and potentially inaccurate.

A fourth issue is that most existing studies use very simple empirical techniques. For example, Reinhart and Rogoff focus on the peak-to-trough fall in annual real GDP per capita around the start of crises. Similarly, other studies, such as Bordo et al. (2001), compare the severity of recessions accompanied by financial crises with those not experiencing financial difficulties. While simple summary statistics are often helpful and illuminating, sometimes they may lead to questionable conclusions. Looking at the peak-to-trough decline in output around a financial crisis may understate the impact of a crisis in a country with very strong trend growth. On the other hand, in cases where the crisis occurred relatively late in the downturn, it may attribute to the crisis declines in output caused by other factors. More generally, reverse causation is inherently a problem in the study of financial crises: crises may depress output, but falls in output may also cause crises. While no statistical procedure can avoid these problems entirely, using recessions as the unit of observation may tend to magnify them.

Overview. Our paper seeks to provide new evidence on the impact of financial crises in modern advanced countries that deals with some of the issues surrounding existing studies. We create a new, semiannual measure of financial distress in a sample of 24 advanced economies from 1967 to 2007. This measure is derived from real-time narrative accounts of country conditions prepared by the Organisation for Economic Co-Operation and Development (OECD).

Our measure of financial distress is designed to capture a reduction in the supply of credit resulting from a rise in the cost of credit intermediation. Using both keyword searches and a detailed reading of the relevant sections, we look for discussions in the *OECD Economic Outlook* of factors such as balance sheet problems and rising loan defaults, which could reduce the willingness of banks to lend. We also look for discussions of creditworthy firms and households having trouble getting credit.

We classify the degree of financial distress on a scale from 0 to 15. We place particular emphasis on ensuring that episodes are scaled consistently across time and countries. Compiling a continuous measure of financial distress allows us to take into account the severity and duration of financial distress, and to analyze how crises emerge and progress.

Our resulting new series on financial distress in advanced countries captures many well-known modern crises, such as those in Japan and Sweden in the 1990s. It also identifies episodes of more modest financial disruption, such as that in France in the mid-1990s following the rescue of Crédit Lyonnais. At the same time, the new series finds no financial distress in some other commonly identified crisis episodes, such as that in Spain in the late 1970s or in Turkey in the early 1980s. The new series also tends to date the periods of financial stress later than do most existing chronologies.

One concern about relying on a single narrative source to identify distress is that it might be idiosyncratic. For a set of key episodes, we therefore check the assessments based on the *Economic Outlook* against the information in two other real-time sources: the annual reports of the central bank of the country and articles in the *Wall Street Journal.* We find that although these sources do not agree with the OECD in every detail, they suggest that the assessments based on the *Economic Outlook* provide a reasonably accurate summary of what a range of observers perceived at the time. We also find that the evidence from the other sources further strengthens the case that using an all-or-nothing measure of financial problems omits a large amount of information about the evolution of distress.

We use the new series on financial distress to investigate the behavior of industrial production and real GDP following financial crises. We run straightforward panel regressions to summarize the relationship between financial distress and economic activity. More specifically, we use Jordà's (2005) local projection method to estimate the response of output at different horizons to an innovation in the financial distress variable.

These panel regressions show more complicated results than the simple summary of the

existing literature that the real effects of crises are severe and long-lasting. For industrial production, we find that a moderate crisis (a 7 on our scale from 0 to 15 of financial distress) is associated with a contemporaneous fall in output that is highly statistically significant, but quite modest in absolute size. Moreover, the effect starts to recede after six months, and is completely gone after two years.

For GDP, a moderate crisis is associated with a contemporaneous fall in output that is again only moderate, but in this case very persistent. However, this persistence is driven entirely by the experience of Japan, which had a very prolonged slowdown in GDP growth starting around the same time as its financial crisis. When Japan is excluded from the sample, the results for GDP mimic those for industrial production. Taken together, these results suggest that in the forty years before the 2008 global financial crisis, the output declines following financial crises in advanced countries were on average moderate and largely temporary.

Our source and approach do not allow us to separate financial distress arising from a decline in output from financial distress due to more exogenous factors. When we examine the behavior of the financial distress variable itself, we find that it is moderately predictable based on lagged output—suggesting that omitted variable bias may be present. Thus, our estimates are, if anything, likely to be an overestimate of the impact of financial distress. Consistent with this, we find that the results, though quite robust to most differences in specification, are much weaker when distress is not allowed to affect output contemporaneously.

We also compare the impact of financial crises estimated using our new series with that estimated using alternative chronologies for the same sample of advanced countries in the postwar era. The impacts using our same empirical approach and these alternative chronologies are slower and more persistent, but actually smaller than those using our new distress measure. This suggests that much of the conventional wisdom that the output consequences of crises are very large is due to the simple empirical approaches of previous studies, and to the inclusion of crises in developing countries and the prewar era.

Finally, the sensitivity of many of our findings to the inclusion of Japan leads us to investigate the variation in the response of output to financial distress across our sample of advanced countries. A comparison of simple autoregressive forecasts and actual outcomes following periods of significant financial distress reveals substantial differences in experiences across countries and episodes. For example, output fell little relative to its pre-crisis path following the financial crises in Norway and Sweden, but dramatically following the crises in Japan and Turkey. We also find that the severity and persistence of the financial distress itself accounts for much of the observed variation.

Related Literature. Though a starting point of our research is Reinhart and Rogoff's influential book, theirs is obviously only one of many studies of the effects of financial crises. In identifying when modern financial crises occurred, Reinhart and Rogoff draw heavily on the work of Caprio and Klingebiel (1996, 1999, 2003). Caprio and Klingebiel base their crisis chronology primarily on the retrospective assessments of experts on episodes in various countries. In recent years, scholars at the International Monetary Fund (IMF) have refined the Caprio and Klingebiel dates using more precise criteria and some quantitative indicators (see Laeven and Valencia, 2014, for the most recent description of the IMF chronology). Importantly, the IMF chronology is still based partly on retrospective assessments.[1]

Studies have investigated the impact of financial crises on the real economy in a variety of ways. As described above, Reinhart and Rogoff look at the peak-to-trough fall in output per capita around crises; Cecchetti, Kohler, and Upper (2009) use a similar approach and also conclude that the fall in output around crises is large. Bordo et al. (2001), IMF (2009a), Jordà, Schularick, and Taylor (2013), and Claessens, Kose, and Terrones (2014) not only examine recessions around financial crises, but explicitly compare recessions with and without crises. These studies find that recessions accompanied by financial crises are more severe. Similarly,

[1] Other chronologies also use retrospective judgments. For example, Schularick and Taylor (2012) make some modifications to a set of crisis dates drawn from Reinhart and Rogoff (2009), Bordo et al. (2001), and the IMF based in part on additional retrospective information.

Claessens, Kose, and Terrones (2009) compare recessions with and without "credit crunches," where credit crunches are identified based on the magnitudes in the declines in credit. Kaminsky and Reinhart (1999) look at simple averages of the behavior of output and other variables before and after the start of crises, compared with averages in "tranquil" times.

Several studies, such as Hoggarth, Reis, and Saporta (2002), IMF (2009b), and Laeven and Valencia (2014), compare the path of output following crises with projections of pre-crisis trends. Those studies, however, use asymmetric rules in making these comparisons. For example, if actual output following a crisis is on average above the pre-crisis trend, Laeven and Valencia report an "output loss" of zero. As a result, although the median output loss (as measured by the sum of the shortfalls of GDP from the pre-crisis trend in the four years following a crisis) in their sample is just 2½ percent of a year's GDP, they report an average output loss of 20 percent.

A few studies use standard regression analysis of postwar data. Cerra and Saxena (2008) look at the behavior of output following the starting dates of the banking crises identified by Caprio and Klingebiel (2003). They find large and persistent falls in output after the onset of crises. Gourinchas and Obstfeld (2012), combining dates of banking crises from a range of existing chronologies, estimate updated versions of regressions analogous to the averages reported by Kaminsky and Reinhart (1999). They note in passing that their findings indicate that output is moderately but persistently below trend following the starts of banking crises in advanced economies.[2]

Many studies consider banking crisis in samples that combine advanced and other countries. Since only a small fraction of the crises in standard chronologies are in advanced countries, the results of these studies are likely to be dominated by the experiences of developing

[2] A study that is similar to ours in approach but covers a very different time period is Jalil (2013). Jalil constructs a new series on banking panics for the United States back to the early 1800s using contemporary newspaper accounts. He is able to scale panics into major and minor crises, and to identify a handful of panics that appear to have been caused by factors other than a decline in output. Using simple time-series regressions, he finds that crises have large and persistent real effects in the period before 1929.

countries. A few studies, such as Cerra and Saxena (2009), IMF (2009b), Gourinchas and Obstfeld (2012), and Claessens, Kose, and Terrones (2009, 2014) report results for advanced or high-income countries separately. In general, these studies find that though the effects of financial crises are less severe in advanced countries, they are still quite large.

Outline. Our paper is organized as follows. Section II discusses the derivation of our new measure of financial distress. It also presents the new measure, and compares it with other chronologies of financial crises for the same sample of countries. Section III presents the statistical analysis of the relationship between financial distress and economic activity in advanced countries. In addition to the baseline regressions, it discusses numerous robustness checks and compares our results with those using other chronologies. Section IV investigates the variation in the apparent response of output to financial distress in different episodes. Finally, Section V presents our conclusions and discusses the implications of our findings.

II. NEW MEASURE OF FINANCIAL DISTRESS

As previously discussed, there are reasons to be cautious about the accuracy of existing chronologies of financial crises. Moreover, the 0-1 nature of most classifications potentially suppresses important variation in the severity of financial distress both across and within episodes. For these reasons, we create a new continuous measure of financial distress for 24 advanced countries for the period 1967–2007.

A. Definition and Approach

Conceptually, we think of financial distress as corresponding to increases in what Bernanke (1983) refers to as the "cost of credit intermediation." This cost includes both the cost of funds for financial institutions relative to a safe interest rate, and their costs of screening, monitoring, and administering loans and other types of financing. A rise in the cost of intermediation makes it more costly for financial institutions to extend loans to firms and households, and thus

reduces the supply of credit. Importantly, we do not consider reductions in lending stemming from increases in all interest rates (as a result of tighter monetary policy, for example) as representing financial distress. The question of how monetary policy and the overall level of interest rates affect the economy is different from the issue of the effects of disruptions to the financial system, and so we do not want to confound the two.[3]

Following most previous work, we do not rely on statistical indicators of financial distress. Using data on quantities, such as the growth rate of bank lending, would mix increases in the cost of credit intermediation not just with shifts in monetary policy, but also with a host of factors affecting credit demand and the creditworthiness of borrowers. Similarly, measures of government intervention, such as spending to aid failed or distressed banks, may be only remotely related to the cost of intermediation. Aggressive government intervention, rather than being a sign of a large rise in the cost of intermediation, might prevent any significant rise; or, long-delayed intervention might clean up institutions that had long since become insolvent and whose lending activities had already been superseded by healthier institutions. Likewise, measures of failures of financial institutions are at best very noisy indicators of financial distress. Most obviously, institutions' cost of credit intermediation, and hence their ability to lend, can change greatly without their outright failure—particularly in the presence of regulatory forbearance, or of just enough government intervention to prevent outright failure.

A more promising avenue to a statistical measure of financial-market problems would involve data on spreads between funding costs for financial institutions and safe interest rates. However, two considerations prevent us from taking that route. The conceptual problem is that, as many authors have emphasized, allocations in credit markets often occur through rationing rather than through changes in interest rates. As a result, spreads may not rise greatly even in

[3] Bernanke also includes influences on credit availability and interest rates resulting from changes in the creditworthiness of borrowers in his definition of the cost of credit intermediation. Again, however, because our goal is to examine the effects of financial distress and because considering the creditworthiness of borrowers blurs the line between loan supply and loan demand, we focus only on the condition of financial firms.

times of substantial distress. The practical problem involves data limitations. To be useful, we would need data on funding costs for all major classes of financial institutions. This is true because a financial-market disruption may be concentrated in a narrow subset of the financial sector, such as savings and loans institutions in the United States in the 1980s; a measure of funding costs for large banks would miss such an episode. Unfortunately, such disaggregate data are not available even for many advanced countries.

In light of these complications, we rely on more qualitative evidence about the health of the financial system to construct our index of financial distress. A key feature of our approach is the use of a consistent real-time source of this information. We study comprehensive summaries of conditions in each country that were written by the OECD at the time, and use those to identify times of financial distress. The use of contemporaneous accounts should help us avoid the bias inherent in retrospective identification of financial crises. The use of a single source that covers many countries over a long period of time helps ensure consistency in the analysis across countries and episodes.

A second important feature of our measure is that we do not treat financial crises as a 0-1 variable, or divide crises into just two groups, such as minor and major or nonsystemic and systemic. Both logic and descriptions of actual episodes of financial distress suggest that financial-market problems come much closer to falling along a continuum than to being discrete events that are all of similar severities, or that fall into just a few categories. Treating a continuous variable as discrete introduces measurement error, both because the variation across crises is omitted and because a small inaccuracy in evaluating an observation can cause a large change in the value assigned to it.

B. Source and Methods

Source. The real-time source we use is the *OECD Economic Outlook*. The *Economic Outlook* has been providing descriptions of economic conditions in each member of the OECD

twice a year (at mid-year and at year-end) since 1967. It is fairly consistent over time in terms of the amount of detail it provides and the features of economies that it discusses.[4] The semiannual publication of the *Economic Outlook* means that our new measure is semiannual as well.

To have a relatively consistent sample and to keep the focus on advanced countries, we restrict the sample to the twenty-four members of the OECD as of 1973. These are Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, Japan, Luxembourg, Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, Turkey, the United Kingdom, and the United States. And because our goal is to assess the evidence from before the recent financial crisis, we end the sample prior to the widespread outbreak of the crisis. Concretely, the last issue of the *OECD Economic Outlook* we examine is that for the first half of 2007, at which point only a few countries were facing financial-market difficulties that the OECD considered noteworthy.

***Methods*.** To derive our new scaled measure of financial distress, we read the *Economic Outlook* to see if OECD analysts described a rise in the cost of credit intermediation for individual countries. To narrow the amount of the volumes we need to study closely, we do a keyword search for terms likely to appear in accounts of financial distress. The most important are "bank" and "financial," but we also search for "crisis," "rescue," "bailout" (and "bail-out"), "crunch," and "squeeze."[5] Importantly, our assessments of the health of countries' financial systems are not based on a mechanical rule, such as the frequency with which these words are used. Instead, we read the material surrounding the places where these words occur to see if the OECD appeared to be describing financial distress. We also read all the potentially relevant information about a country in the *Economic Outlook* if we find information elsewhere in the

[4] As we discuss in the appendix, there is some variation in the length of the country write-ups over time and across countries. We take this variation into account in deciding how significant a discussion of financial distress is in different time periods and different countries.

[5] We also experimented with searching for "credit" in a few issues of the *Economic Outlook*. This did not yield any noteworthy passages that were not identified by the other search terms.

document suggesting that there were financial-market difficulties in the country, or if the previous issue of the *Economic Outlook* suggested financial distress in the country.

In interpreting the material in the *Economic Outlook*, we put the most weight on factors that are clear markers for increases in the cost of intermediation. We look for discussions of such developments as increases in financial institutions' costs of obtaining funds relative to safe interest rates; general increases in the perceived riskiness of financial institutions; financial institutions facing problems obtaining funds; reductions in financial institutions' willingness to lend; and difficulties of creditworthy borrowers obtaining funds because of problems at financial institutions. In addition to looking for descriptions of factors directly linked to the cost of intermediation, we look for references to developments likely to weaken financial institutions, and so reduce their ability to perform their normal functions. Examples include rising loan defaults, increases in nonperforming loans, balance sheet problems, and erosion of their capital.

For the accounts that suggest financial distress, we group them according to the severity of the difficulties. To scale the degree of financial distress, we look for signs of more or less change in the indicators mentioned above. Was the rise in the perceived riskiness of financial institutions relatively minor, or so large that it is described as a widespread panic? Was the effect on the willingness to lend described as minor or extreme? Was the rise in nonperforming loans thought to be small or large? In this ranking, we also consider some indirect proxies for the rise in the cost of intermediation. For example, we put some weight on descriptions of government intervention in the financial system as an indicator of the perceived severity of balance sheet and funding problems. Likewise, the OECD's description of the actual or anticipated impact of financial troubles on spending and the economy is often a useful summary statistic for the perceived severity of financial distress.[6]

We view a central aspect of this classification as comparative: we attempt to group

[6] Importantly, we see no evidence in the *Economic Outlook* that OECD analysts were deducing financial distress from declines in spending and output. Rather, they were trying to explain recent economic developments and project future performance, and they viewed distress as one influence on those outcomes.

problems that the *Economic Outlook* describes in similar terms together, and to place ones that it describes as more severe in higher categories. Thus, much of our classification involves comparing episodes to try to make our assignments as consistent as possible.

Because we did not read each issue of the *Economic Outlook* in real time, the issue of possible bias based on information about subsequent outcomes that we raised about other chronologies also potentially applies to our own. We take a number of steps to minimize such bias. First, and most important, we set out relatively clear criteria for classifying episodes based on the descriptions in the *Economic Outlook*, and we document in online Appendix A the key passages and descriptions that lead us to classify episodes as we do. This forces us to be rigorous in our assessment of the narrative accounts, and allows others to check our classification. Second, we read the documents in order, so that we base our interpretation of the health of a country's financial system only on contemporaneous descriptions. We never revisit our interpretation of the description of a situation based on later descriptions of developments in the country. Third, we did not examine any data on economic outcomes until after we completed our classification. Because we are dealing with a large number of countries, our prior knowledge of the high-frequency macroeconomic history of most of the countries in our sample was very limited.

***Criteria for the Different Categories*.** The categories to which we assign episodes have natural interpretations. Our main ones are "credit disruption," "minor crisis," "moderate crisis," "major crisis," and "extreme crisis." In keeping with the fact that the accounts suggest that financial-market problems fall along continuum, we subdivide each category into "regular," "minus," and "plus." Thus, for example, an episode of relatively minor financial distress could be classified as "credit disruption–minus," "credit disruption–regular," or "credit disruption–plus."

In our empirical work, we convert these categories into a numerical scale. Cases where there is no financial distress are assigned a zero. Positive levels of distress start at 1 for a credit

disruption–minus and go through 15 for an extreme crisis–plus. Table 1 lists the full set of categories and the values we assign to them.

As much as possible, we try to use specific criteria to classify episodes into categories. It is therefore useful to describe the characteristics common to the various groupings briefly. The hallmark of the episodes that we identify as credit disruptions is that the OECD perceived financial-market problems or increases in the cost of credit intermediation that were important enough to be mentioned, but that it did not believe were having significant macroeconomic consequences. A common form for this to take was for the OECD to describe the problems not as directly affecting its outlook for the country, but as posing a risk to the outlook. Other possibilities are that the OECD viewed the problems as affecting only a narrow part of the economy; that it mentioned them in passing or explicitly identified them as minor; or that it described the financial system as improved but not fully healed following a situation that we classify as a minor crisis. An example of a regular credit disruption occurred in Germany in 1974:2 (that is, the second half of 1974), where the OECD described "strains" in the banking system and the extension of special credit facilities to help small and medium-sized companies obtain credit (pp. 50 and 26, respectively).

A canonical case of a minor crisis has three characteristics: a perception by the OECD that there were significant problems in the financial sector; a belief that they were affecting credit supply or the overall performance of the economy in a way that was clearly nontrivial, and not confined to a minor part of the economy; and a belief that they were not so severe that they were central to recent macroeconomic developments or to the economy's prospects. An example of a regular minor crisis is France in 1996:1, where the OECD described serious problems in the banking sector, including "high refinancing ... costs and large provisions for bad debts," as well as government intervention to support some financial institutions, but did not give banking problems a central role in its discussion of the outlook (p. 78).

A moderate crisis, in our classification, involves problems in the financial sector that are

widespread and severe, that are central to the performance of the economy as a whole, and that are not so serious that they could reasonably be described as taking the form of the financial system seizing up entirely. One way that we judge whether the OECD viewed problems as being central to the performance of the economy is by whether it mentioned them prominently, for example in the opening summary of the entry on a country. Another is by whether it discussed impacts on credit supply or real activity repeatedly. We also use the descriptions of the size and scope of any government interventions in the financial system as one indicator of the severity of the disruption. Our definition of a moderate crisis appears to roughly correspond to the cutoff in other chronologies, such as Caprio et al. (2005) and Laeven and Valencia (2014), between a crisis and no crisis, or between a systemic and a nonsystemic crisis. An example of a regular moderate crisis is the United States in 1990:2, where the *Economic Outlook* referred to an "apparent credit 'crunch'," and cited difficulties in the financial system as one of the main factors, but not the only one, slowing the economy (pp. 14 and 48, respectively).

At the severe end of the spectrum are major and extreme financial crises. These are situations where there are large impediments to normal financial intermediation throughout virtually all of the financial system. We look for such markers as the unreserved use the term "crisis" in referring to the financial system, and for such terms as "dire," grave," "unsound," and "paralysis." We also look for clear-cut statements that the financial-sector disruptions were having an important effect on credit supply and macroeconomic outcomes. In addition, we view references to major government interventions as suggesting that the problems were severe. There are only two episodes in our sample that we classify as major or extreme crises, Japan in 1998:1 and 1998:2. The more significant is 1998:2, which we classify as an extreme crisis–minus. In that case, the OECD referred to the "breakdown in the credit creation mechanism," to the "the severe and prolonged crisis in the banking system," and to banks being in "dire straits" (pp. 44, 20, and 45, respectively).

Our subdivision of the various broad categories into minor, regular, and plus is based on

the specifics of the discussions within these general rubrics. In the case of credit disruptions, for example, we tend to place disruptions that the OECD described as posing major risks to the outlook in higher categories than ones that it viewed as posing minor risks. Similarly, if the OECD reported that a disruption was serious enough that it had caused authorities to make some type of intervention in credit markets to improve credit flows, we tend to classify the disruption as more serious.

***Documentation*.** Online Appendix A provides more information about our criteria for the different categories of financial distress and our procedures for classifying episodes using the accounts in the *Economic Outlook*. Table 2 lists each episode for which we find that the *Economic Outlook* was describing financial distress. The bulk of the appendix provides episode-by-episode explanations of the analysis and discussion in the *Economic Outlook* that lead to our classifications. Exhibit 1 reproduces the appendix entries for the four episodes cited above: Germany in 1974:2 (credit disruption–regular), France in 1996:1 (minor crisis–regular), the United States in 1990:2 (moderate crisis–regular), and Japan in 1998:2 (extreme crisis–minus).

C. New Series

Figure 1 shows our new measure of financial distress for the period 1967:1 to 2007:1 for the ten OECD countries that had some nonzero values of our measure. The other fourteen countries that we analyze had no times in our sample period where the OECD noted financial distress.

Several features are clear from the figure. Most obviously, there were essentially no episodes of financial distress, and certainly nothing that would count as a significant crisis, in the 1970s and 1980s. For advanced countries, these two decades were a time of financial calm, despite oil price shocks and severe moves toward disinflation in many countries. In contrast, the 1990s were a period of extensive financial distress. Our new measure captures the well-known financial troubles in a number of Nordic countries and Japan in this period. It also identifies significant distress in the United States at the turn of the decade related to the savings

and loan crisis and other disruptions.

Another thing that is clear from the figure is the tremendous variation in how crises evolve. Some, such as the crisis in Sweden in 1992–1993, became acute almost instantaneously, and then resolved just as quickly. Others, such as the distress in Japan, built slowly before eventually erupting into severe distress. Japan also stands out as a case where the financial distress lingered—not just for years, but for well over a decade. In other episodes, such as France in the mid-1990s, a country may suffer mild distress for a prolonged period, but never have it erupt into a full-fledged crisis.

D. Comparison with Other Chronologies

It is natural to ask how our new measure of financial distress compares with other crisis chronologies for the same countries over the period we consider. We focus on two alternatives: the dates of crises considered by Reinhart and Rogoff (2009), and the latest version of the dates in the IMF Systemic Banking Crises Database (Laeven and Valencia, 2014). Both chronologies use a 0-1 classification and typically date crises in years (for example, 1984–1991). To make each of these series comparable to our semiannual series, we generally put the beginning of a crisis in the first half of the year in which the chronology identifies the start of a crisis, and the end in the second half of the year that it lists as the last one of the crisis. Occasionally, a chronology gives a particular month for the start of a crisis; in this case, we place the start in the half-year corresponding to that month.[7]

The Reinhart and Rogoff chronology includes a fairly wide range of episodes, while the Laeven and Valencia chronology only identifies systemic crises. Laeven and Valencia identify eight systemic crises in advanced countries over the period 1970–2006: Spain in the late 1970s and early 1980s, the United States in the late 1980s, Sweden, Finland, and Norway in the early

[7] The Reinhart and Rogoff dates are from their Table A.4.1, pp. 348–392. The Laeven and Valencia start dates are from their Appendix Tables 2A.1 and 2A.3, pp. 94–112 and 118–135. Their end dates are from Laeven and Valencia (2013, Table A1, pp. 254–259). The Reinhart–Rogoff chronology covers the period 1800–2008; the Laeven-Valencia chronology covers 1970–2007. Our comparison considers the period 1970–2006.

1990s, Japan in the late 1990s and early 2000s, and Turkey in the early 1980s and around 2000. Reinhart and Rogoff also identify crises in these same periods. But they list an additional sixteen crises in the countries we consider over our sample period.

Our new measure derived from OECD reports identifies significant financial distress in six of the eight episodes Laeven and Valencia identify. There is no discussion of financial distress in Spain in the *OECD Economic Outlook* at any time, nor in Turkey in the 1980s.[8] Figure 2 compares the three crisis series in each of the remaining six episodes. In particular, for each episode, we show the start and end dates as identified by Reinhart and Rogoff and Laeven and Valencia, along with our continuous indicator of financial distress over the same period.[9]

The biggest differences are between Reinhart and Rogoff's chronology and both the Laeven and Valencia chronology and our new measure of financial distress. For example, for the crisis in Norway, Reinhart and Rogoff identify the start in 1987:1, while Laeven and Valencia date it in 1991:2, and the new measure does not spike up until the same half-year. Likewise, for the United States, Reinhart and Rogoff date the crisis as running from 1984 to 1991, while Laeven and Valencia date it as being limited to just 1988, and the new measure identifies distress in the period 1990:1 to 1992:1. And for Japan, the new measure shows mild to moderate distress in the years Reinhart and Rogoff identify as the crisis (1992 to 1997), but the peak in distress is well after Reinhart and Rogoff say the crisis ended. Likewise, Reinhart and Rogoff date the end of the Japanese crisis in the same half-year that Laeven and Valencia date the start.

Though the differences between the new measure and the Laeven and Valencia chronology are smaller than those with the Reinhart and Rogoff series, many are still substantial. As already described, for the United States, Laeven and Valencia identify a crisis as occurring two years before the new measure derived from OECD records shows noticeable distress. Similarly,

---

[8] This remains the case even when we do not use our keyword search as a starting point, but examine every reference to Spain in the late 1970s and early 1980s, and every reference to Turkey in the first half of the 1980s.

[9] In panel (e) (Turkey), Reinhart and Rogoff identify four separate crises, two of which have start and end dates in the same half-year (1991:1 and 1994:1).

Laeven and Valencia date the end of the Japanese crisis a half-year before the new measure shows a second acute rise in distress, and several years before it shows the distress ending. And, Laeven and Valencia date the end of the crisis in Sweden two years after distress returns to zero in the new measure.

E. Additional Evidence on the Accuracy of the New Series

Given that the new measure of financial distress derived from OECD documents differs in important ways from the two alternative chronologies considered, it is prudent to look for additional evidence with which to either corroborate or question the new series. To do this, we consider the descriptions of financial distress in two other real-time sources: the annual reports of the relevant central banks and the *Wall Street Journal*. To keep the task manageable, we consider these additional sources only around the eight systemic crises in advanced countries between 1970 and 2006 identified by Laeven and Valencia. Online Appendix B provides a detailed episode-by-episode discussion of this additional evidence. Here, we summarize the findings from this analysis.

Perhaps the most important conclusion is a methodological one. Examining the additional sources strengthens the case for a continuous measure of financial distress rather than a (0,1) crisis classification. Like the OECD, both the central bank reports and the *Wall Street Journal* described a range of financial troubles in the various countries at different times. While it is clear that financial distress was worse in some half-years than in others, it is often very hard to see where one would draw the line between a "crisis" and not. This is particularly obvious in the case of Japan. Like the OECD, the additional sources described growing financial troubles over the early and mid-1990s; much more severe problems in the late 1990s; improvement and then another round of severe distress in the early 2000s; then, finally, gradual recovery in the mid-2000s. Attempting to reduce this complex experience to a limited period of crisis would be both difficult and counterproductive.

In terms of the accuracy of the new measure of financial distress, we find that the descriptions in the additional sources are typically in close agreement with the new series. The correlation is certainly not perfect, but it is high enough to suggest that the *OECD Economic Outlook* is a reliable and accurate source.

The agreement between the new series and the additional sources is highest in the cases where the new series differs most from the alternative chronologies. For example, in the case of the United States, the *Annual Report of the Board of Governors of the Federal Reserve System* agrees with the new measure that distress was concentrated in the early 1990s, and not significant in the mid-1980s as Reinhart and Rogoff suggest or in 1988 as Laeven and Valencia identify. Likewise, for Japan, the *Bank of Japan Annual Review* and the *Wall Street Journal* agree that there was some distress over almost all of the period 1990:1 to 2005:1 when the new measure shows positive levels of distress, and that it peaked in 1998 and 2002 when the new series does. This is very different from the early crisis (1992 to 1997) shown by the Reinhart and Rogoff chronology and the relatively short crisis (1997 to 2001) shown by the Laeven and Valencia chronology. Finally, for Norway, where the new series and the Laeven and Valencia chronology identify the start of problems in late 1991 and Reinhart and Rogoff date it in 1987, the additional sources strongly support the later date.[10]

For the two cases where Laeven and Valencia and Reinhart and Rogoff identify a crisis while the new series shows no financial distress, the results are somewhat more mixed. In the case of Spain, the alternative sources suggested there was some financial distress, but it was small. Both the *Wall Street Journal* and the reports of the Bank of Spain described a number of bank failures over the period 1978–1983, and both occasionally used the term "crisis." But neither suggested more than minor increases in the cost of credit intermediation, and both had

[10] Importantly, while the additional evidence in these cases corroborates the new measure, it also often provides clues as to why the alternative chronologies date crises as they did. For example, the records of the Federal Reserve showed a small amount of concern about the banking system in 1984 and 1985 when Reinhart and Rogoff date the start of the crisis. Similarly, the Norges Bank discussed minor financial problems in 1987, the start date of the Norwegian crisis in the Reinhart and Rogoff chronology. In both cases, however, the additional sources were quite clear that the problems were not of crisis proportions.

extensive discussions of the forces influencing the economy that included little role for financial distress. Thus, these alternative sources largely agree with the OECD. In the case of Turkey, both the reports of the Turkish central bank and the *Wall Street Journal* described significant distress in Turkey in the early 1980s—thus conflicting with the new measure. The level of distress described in the additional sources is noticeably less than in other crisis episodes, such as those in the Nordic countries in the early 1990s, but still substantial.

One somewhat systematic discrepancy between the new series derived from the *OECD Economic Outlook* and the evidence from the additional sources is that the new series appears to be slightly slow to identify the start of financial distress. In the cases of Finland, Norway, Sweden, and Turkey (in the early 2000s), the new series does not depart from zero until about a half-year later than the start of descriptions of distress in the additional sources. In some cases, this is due to the fact that, because of production lags, the *OECD Economic Outlook* discussed events occurring late in the half-year in the subsequent issue.[11] But in others, it seems to represent either some inertia on the part of the OECD or a slightly high bar for mentioning distress.

On the end of financial distress, the new series also differs occasionally from the evidence in the additional sources. However, the differences are not systematic. The OECD was slightly later in seeing the end of distress than the additional sources for Norway and Turkey, but slightly earlier for Sweden. It was roughly in the middle for Finland and Japan, where the additional sources disagree with each other somewhat. Importantly, because all three real-time sources suggested at least somewhat gradual changes, even where there are disagreements about when distress reached zero, there is close accord that distress was low and falling.

[11] The mid-year *Economic Outlook* contains information through roughly May, and the end-year volume through November. As a result, financial distress that occurs very late in the half-year will be discussed in the next volume. This means that in our series, financial distress occurring late in the half-year will also be dated in the subsequent half-year. Since in the empirical analysis we pair the semiannual financial distress data with output data for the second and fourth quarters, this timing convention for distress makes the two series roughly contemporaneous. Output effects occurring very late in the half-year also show up predominantly in the subsequent half-year.

The bottom line of this consideration of additional evidence is that the new series is surely not perfect, but it is very good. And in the cases where it differs most from the alternative crisis chronologies, the new series appears to be substantially more accurate. It also has the virtue of being continuous, which the additional evidence suggests is appropriate and valuable.

III. RELATIONSHIP BETWEEN FINANCIAL DISTRESS AND ECONOMIC ACTIVITY

Having created a new, continuous measure of financial distress for a sample of advanced countries, the obvious next step is to see what it shows about the relationship between financial distress and economic activity.

Before looking at this relationship, it is important to spell out a key limitation of the analysis. Our new measure is designed to identify times of financial distress. We feel it does this in a sensible way that is relatively uncontaminated by information about subsequent economic developments. As a result, it should be largely free of some of the biases that may be present in existing chronologies. However, there is another bias that our new series does not address. Our analysis tells us nothing about the ultimate cause of the financial distress. In particular, it does not separate distress caused by relatively exogenous factors, such as managerial malfeasance, from financial problems caused or exacerbated by a decline in output or by forces that reduce output directly. As a result, a regression of output on our new distress variable may, and indeed is likely to, suffer from omitted variable bias. This means that any relationship we find may well not be causal, and may overstate the sensitivity of output to financial distress.

A. Data and Specification

Data. To analyze the impact of financial distress, we focus on two broad indicators of economic activity: industrial production and real GDP. Both series are available quarterly from

the OECD for the 24 countries in our sample.[12] Industrial production has the virtue of being a relatively straightforward series to produce, and so is likely to be quite accurate and consistent across countries. GDP is more complicated and less transparent; as a result, it may be less consistent across countries. GDP's obvious virtue is that it is the broadest measure of economic activity available.

Our measure of financial distress is semiannual. We therefore convert the output data to a semiannual frequency as well. We do this by taking the quarterly values for the second and fourth quarters of each year. Since the *OECD Economic Outlook* is issued at mid-year and at year-end, the timing of the output data roughly corresponds with the timing of the OECD's descriptions of country conditions.

In our analysis we work with data in calendar time. That is, we do not use recessions as our unit of analysis, as is more common in this literature. Moving away from analyzing peak-to-trough changes in output allows us to focus more closely on the timing of the relationship, and to limit the analysis to what happens to output following financial crises. This can help reduce the degree to which the analysis is picking up reverse causation.

Specification. To estimate the relationship between output and financial distress, we create a panel dataset including output and the new distress variable for the 24 OECD countries we consider for the period 1967:1 to 2007:1. We then use the Jordà local projection method to estimate the behavior of output in the wake of financial distress. The Jordà method runs separate regressions for output at various horizons starting at time t on the crisis variable at time t. The sequence of coefficient estimates for the various horizons provides a nonparametric estimate of the impulse response function.

---

[12] The data, which are ultimately collected and reported by the individual countries, are available on the OECD website: http://www.oecd.org/statistics/. The industrial production data are from the Production and Sales Dataset, production of total industry. The GDP data are from the Quarterly National Accounts Dataset, series VPVOBARSA. Both series were downloaded 7/6/2014. There are some minor gaps in the two series for some countries. Industrial production data are missing for Australia before 1974Q3; Denmark before 1974Q1; Iceland before 1998Q1; Ireland before 1975Q3; New Zealand before 1977Q2; and Turkey before 1985Q1. Real GDP data are missing for Iceland before 1997Q1 and for Greece after 1999Q4.

The particular specification that we estimate is:

$$(1) \qquad y_{j,t+i} = \alpha_j^i + \gamma_t^i + \beta^i F_{j,t} + \sum_{k=1}^{4} \varphi_k^i F_{j,t-k} + \sum_{k=1}^{4} \theta_k^i y_{j,t-k} + e_{j,t}^i,$$

where the *j* subscripts index countries, the *t* subscripts index time, and the *i* superscripts denote the horizon (half-years after time *t*) being considered. $y_{j,t+i}$ is the log of output (either industrial production or real GDP) for country *j* at time *t+i*. $F_{j,t}$ is the financial distress variable for country *j* at time *t*. We include four lags of both the distress variable and the output variable as controls.[13]

We also include country fixed effects (the α's) to capture the fact that normal output behavior may differ across countries. Similarly, we include time fixed effects (the γ's) to control for economic developments facing all countries in a given year. The inclusion of time fixed effects means that we are only using the cross-section variation in output and financial distress to estimate the relationship between the two series.

We estimate equation (1) for values of *i* from 0 to 10 half-years. That is, we consider horizons up to five years after time *t*. The sequence of coefficients on the financial distress variable at time *t* shows the behavior of output in response to an innovation in the distress variable of 1. To make the interpretation of the impulse response function more straightforward, we multiply the coefficients by 7, which is the value of our financial distress measure corresponding to the start of the "moderate crisis" category. This transformed impulse response function thus shows the behavior of output following a relatively large innovation in financial distress.

In the baseline estimates we include all 24 OECD countries for which we have the new measure. As discussed above, Japan is a substantial outlier in terms of the severity and duration

[13] Because we include four lags, the estimation begins in 1969:1. And because we do not want the behavior of output in the recent period to be driving the results, we do not use output data past 2007:1. The sample for horizon *i* therefore ends *i* half-years before 2007:1. For *i* = 4, for example, *t* runs from 1969:1 to 2005:1.

of its financial distress. Since this may give it an outsized weight in the estimation, we also consider the 23-country sample excluding Japan. We discuss a number of additional permutations in the specification and sample in the subsection on robustness below.

B. Behavior of Output Following Financial Crises

We estimate equation (1) for various horizons using both industrial production and real GDP. Even at quite distant horizons, the sum of the coefficients on lagged output is close to one. As a result, the country fixed effects essentially capture differences in average growth rates across countries. The hypothesis that the country fixed effects are all zero is strongly rejected for both output measures at all horizons (except the contemporaneous horizon for industrial production). Similarly, the hypothesis that the time fixed effects are all zero is overwhelmingly rejected for both output measures at all horizons.

Figure 3 shows the impulse response functions for the two output series estimated over the full sample of 24 advanced countries, together with the two-standard-error bands. Panel (a) shows the results for industrial production, and panel (b) for real GDP.

Industrial production falls noticeably at the time of the impulse to the financial distress variable. The *t*-statistic on the contemporaneous relationship is over 4 (in absolute value). The negative effect increases slightly in the half-year after the impulse. However, the absolute size of the decline following a moderate crisis is modest: the peak effect is a decline of 3.9 percent ($t = -3.2$). To put this number in perspective, Romer and Romer (1989) find that industrial production fell roughly 12 percent following relatively exogenous shifts to contractionary monetary policy in the United States in the postwar period. Another noteworthy feature of the response of industrial production is how quickly the negative impact of a financial crisis dissipates. The effect is a decline of just 1.7 percent ($t = -1.2$) within a year after the impulse and zero within two years.

Real GDP also appears to decline contemporaneously with the impulse to the financial

distress variable. The impact of a moderate crisis is a decline in GDP of 3.0 percent ($t = -6.1$). This negative effect grows slightly over the 3½ years following the impulse, peaking at 4.2 percent ($t = -2.6$). Since industrial production is normally substantially more volatile than GDP, the fact that the percentage decline in the two series following a financial crisis is roughly similar suggests that the impact on GDP is stronger.

The more striking difference between the responses of industrial production and GDP is that the effects of a crisis on GDP appear to be much more persistent. Though the standard errors increase substantially at longer horizons, the point estimate of the impact of a moderate crisis is strongly negative for the full five years that we consider.

This persistence, however, is driven almost entirely by the special case of Japan. Japanese real GDP growth slowed from over 4 percent per year to under 2 percent around the time that its financial troubles began in the early 1990s, and then remained low. As a result, Japan's experience points strongly in the direction of financial distress having large negative long-run effects. Moreover, as can be seen in Figure 1, a considerable fraction of the variation in financial distress in our sample comes from the observations for Japan. As a result, the Japanese experience gets a large weight in the estimation. Panel (b) of Figure 4 shows the impulse response function for GDP when the sample excludes Japan. The maximum fall in GDP following a moderate crisis is now 3.0 percent ($t = -5.2$), compared with a decline of 4.2 percent in the full sample. More notable is the difference in persistence. Without Japan in the sample, the estimated impact of a financial crisis on GDP begins to dissipate after six months. By two years after the crisis, the point estimate is effectively zero. After that, the point estimates turn positive, though the standard errors also become large.

Panel (a) of Figure 4 shows that excluding Japan also affects the estimated impulse response function for industrial production. The effects of a moderate crisis, which are small in the full sample, are now virtually nonexistent. The contemporaneous effect is a fall in industrial production of just 2.5 percent ($t = -2.5$). The effect, which dissipates fairly quickly in the full

sample, goes away almost instantaneously in the no-Japan sample. By a year after the crisis, the effect is positive; by two years after, the positive effect is actually significant.

The bottom line of the focal empirical results is that the relationship between output and financial crises is not as dire as the modern conventional wisdom would lead one to believe. Using our new, real-time measure of financial distress and sensible time-series estimation methods, we find that the impact of a financial crisis on output is quite small. And, with the exception of Japan, the impact does not appear to be very long-lasting. Indeed, for industrial production, one would have to say that the negative effects are remarkably transitory.

C. Behavior of the Financial Distress Variable

The local projection approach to estimating the impulse response function for output does not provide any information about the evolution of the distress variable itself. However, it is straightforward to examine this issue.

Exogeneity. One question involves the exogeneity of our new measure of financial distress. Does lagged output predict the distress variable? To analyze this, we regress the distress variable on four own lags and four lags of either log industrial production or log real GDP. Since we are working with panel data, we also include country and time dummies.

Table 3 shows the results of these Granger causality tests. The hypothesis that all of the coefficients on lagged output are zero is rejected at the 3 percent confidence level for industrial production and the 1 percent level for real GDP. This suggests that reverse causation is likely present. The coefficients on the various lags of the output variables, however, suggest that the relationship is quantitatively small, and more complicated than just that output declines tend to lead to financial distress. A fall in GDP of 1 percent raises the crisis variable six months later by just 0.006, and then the coefficients alternate sign.

Excluding Japan from the sample weakens the predictive power of the output variables somewhat. The hypothesis that all of the coefficients on industrial production are zero is

rejected at the 8 percent confidence level in the no-Japan sample. The hypothesis that all of the GDP coefficients are zero is now rejected only at the 20 percent level.

That the financial distress variable is somewhat predictable based on lagged output (and that the relationship is negative in the short run) may affect how one should interpret the impulse response functions for the output variables. Recall that the largest estimated impact of a financial crisis occurs in the contemporaneous period. Since causation appears to run both directions, it is highly unlikely that this contemporaneous effect entirely reflects an impact of financial distress on output.

***The Persistence of Financial Distress*.** Looking at the impulse response function of the financial distress variable to itself can also aid in the interpretation of the impulse response function for output to the distress variable. To do this, we estimate equation (1), replacing the left-hand-side variable with the distress variable at different horizons after time *t*. We use GDP as the output control.

Panel (a) of Figure 5 shows the impulse response function of the distress variable to itself, estimated over the full sample of 24 countries. As with the impulse response function for output, we simulate the impact of an innovation of 7 in the distress variable (a moderate crisis). The figure shows that there is important serial correlation in the financial distress variable. An impulse of 7 is followed by a value of the new measure of 5.9 a half-year later. After two periods (1 year), the rate of decay speeds up noticeably, so that by 2½ years after the impulse any effect of distress on itself is almost gone.

Panel (b) shows the impulse response function for the distress variable estimated for the sample excluding Japan. Excluding Japan, where financial distress was extraordinarily persistent, results in a somewhat more rapid dying out of the impulse. Otherwise, the results are similar.

That there is substantial serial correlation in the distress variable, particularly at near horizons, suggests that some of the near-term persistence we find in the impact of financial

distress on output is likely due to persistence in the distress itself. It is not necessarily that financial crises have long-lasting effects, but rather that crises themselves tend to last for a while. This possibility, and the role that differences in the persistence of crises across episodes play in explaining the variation in output behavior, is analyzed further in Section IV.

D. Robustness

We examine the robustness of our findings along numerous dimensions. Here we highlight some of the most important.

Alternative Specifications. We consider several variations of our baseline specification. The first is to replace the Jordà approach with a single-equation autoregressive specification. In particular, we estimate an equation of the form:

$$\Delta y_{j,t} = \alpha_j + \gamma_t + \sum_{k=0}^{6} \varphi_k F_{j,t-k} + \sum_{k=1}^{6} \theta_k \Delta y_{j,t-k} + e_{j,t}, \quad (2)$$

where the output variable (either industrial production or real GDP) is now expressed as the change in logarithms. We include the contemporaneous value and six lags of the crisis variable and six lags of the change in the output variable. We then simulate the impact of a realization of the distress variable equal to 7 (a moderate crisis–minus on our scale). This specification does not provide any information about the evolution of the crisis variable, but the simulation requires an assumption about the values of the variable after the initial realization. There are two natural baseline cases: a one-period crisis ($F_{j,t} = 7$ with $F_{j,t+k} = 0$ for $k > 0$), and a permanent crisis ($F_{j,t+k} = 7$ for all $k \geq 0$). Since most crises are relatively short, we focus on the temporary case.

The results are very similar to those of the baseline specification at near horizons. The contemporaneous effect of a crisis is a fall in industrial production of 3.5 percent ($t = -4.2$) and a fall in GDP of 2.9 percent ($t = -5.9$). The results for industrial production are similar to those of the baseline even at longer horizons; in both specifications, the effects are quite short-lived—

disappearing entirely within two years. For GDP, however, the results of the single-equation model are substantially different from the baseline at longer horizons. Even in the full sample, the effects of a crisis are quite temporary in the single-equation alternative. In response to a one-time realization equal to 7 on our scale, real GDP is down only 1.0 percent after six months, and 0.2 percent after a year.

A sensible way to improve this autoregressive specification and bring it closer to the Jordà approach is to run vector autoregressions (VARs). Adding a second equation for the financial distress variable allows the data to determine how the distress variable evolves following an innovation. The two equations then jointly imply how output and financial distress respond to an innovation in distress. As a result, the impulse response functions are analogous to those from the Jordà approach, but with the responses at all horizons computed from the two equations rather than estimated separately at each horizon.

Specifically, we consider VARs with two variables, output (either industrial production or GDP), entered in log levels, and our new distress measure. Paralleling the timing assumption in our baseline specification, we order the distress variable first. The VARs include six lags. The VAR results are very similar to our baseline results. That both the one-equation autoregressive approach and VARs yield results similar to the baseline approach suggests that our findings are not being driven by the particulars of our empirical specification.

We also consider two permutations of the baseline specification that keep the basic Jordà approach, but change the variables included in equation (1). One is to add country-specific trends, and so allow for different trends across countries in their growth rates. The results from this specification are very similar to those of the baseline specification.

A more significant change in the variables that are included is to not allow financial distress to affect output contemporaneously. One potential source of bias in estimating the effects of financial distress is within-period simultaneity, with higher distress reducing output and lower output raising distress. If within-period simultaneity were the only source of bias,

then our baseline specification would tend to overstate the impact of financial distress on output, and excluding contemporaneous financial distress would tend to understate it. We therefore consider a variant of equation (1) that replaces the $\sum_{k=1}^{4} \theta_k^i y_{j,t-k}$ term with $\sum_{k=0}^{4} \theta_k^i y_{j,t-k}$. With this change, the coefficient on $F_{j,t}$ shows the relationship between output in period $t+i$ and the component of financial distress in period t that is uncorrelated not just with financial distress and output before period t, but also with output in period t.

This change sharply reduces the estimated effects of financial crises. As before, we consider an impulse equal to 7 on our scale. By construction, the contemporaneous effect of a financial crisis on output is now zero. For industrial production, the estimated effect is very slightly negative after one half-year, and then consistently positive (though insignificant at every horizon except 2½ years, when it is marginally significant). For GDP, the response remains consistently negative, but the null hypothesis of no relationship is now never even close to being rejected, and the maximum impact of a financial crisis is only about one-fifth as large as in the baseline specification. That is, if one does not interpret any of the contemporaneous correlation between movements in output and financial distress as reflecting causation from distress to output, our results suggest that a financial crisis has little impact on output. Or, to state this finding purely in statistical terms, once one knows the path of output up through the current half-year, observing an unexpected increase in financial distress in the current half-year should have little effect on one's forecast of output going forward.

Alternative Samples. We also experiment with the sample. Another way to deal with the fact that countries may have different trends in their normal growth is to consider a shorter sample period. Since our new measure shows almost no cases of financial-market problems in the first two decades covered by our chronology, the natural shorter sample is 1987:1–2007:1. This change has very little impact on the near-term effects of a financial crisis on either industrial production or GDP. However, it has more impact on the estimated longer-run effects. For example, the effect of a crisis on GDP after 10 half-years is −3.6 (t = −1.9) in the full sample,

but just −0.5 (t = −0.4) in the shorter sample.

Because Japan may be an important outlier, we have already considered the effects of excluding Japan. As discussed in more detail in Section IV, the other country that may be having an important impact on the results is Turkey, where output fell precipitously at the time our indicator spikes up in 2001, and then rebounded strongly. We therefore examine the impact of dropping Turkey. The main effect is to lower the estimated contemporaneous impact of financial distress by about one-third. The estimated effects after one to two years are little changed, and the estimated effects at longer horizons are slightly more negative than in the baseline sample.

Standard Errors. In our baseline results, we report conventional OLS standard errors. However, there are two reasons that the residuals might not be i.i.d. First, they may be heteroskedastic. For example, output is typically more volatile in the less developed members of our sample, such as Greece and Turkey, and in the smaller countries, such as Luxembourg, Iceland, and New Zealand. Second, for horizons beyond one half-year, there is an overlapping structure to the residuals. For the case of i = 10, for example, the dependent variable for observation t for a country is its growth from t to t + 10, and for observation t + 1 it is growth from t + 1 to t + 11. This pattern could lead to serial correlation of the residuals.

We therefore consider three alternatives to conventional standard errors. The first is heteroskedasticity-consistent standard errors. The other two allow for serial correlation as well as heteroskedasticity. Specifically, for horizons i = 2 to 10, we allow for serial correlation over up to i − 1 periods; for example, for i = 2, we allow for first-order serial correlation of the residuals. We consider both Newey-West standard errors (which damp the off-diagonal elements of the covariance matrix of the residuals) and Hansen-Hodrick standard errors (which have no damping).

Considering the alternative standard errors has little impact on the overall pattern of the results. The three alternative sets of standard errors are generally very similar to one another.

For horizons of a year and longer, the alternatives are typically 10 to 20 percent larger than the conventional standard errors when we use the full sample, and 10 to 20 percent smaller when we use the no-Japan sample. The only large differences between the alternative and conventional standard errors are at very short horizons (where heteroskedasticity-consistent standard errors are the only relevant alternative to the conventional ones, since there is no reason to expect serial correlation). For these horizons, the alternative standard errors are 25 to 100 percent larger than the conventional ones. As a result, the near-term relationships between financial distress and output become less statistically significant. For example, the *t*-statistic for the contemporaneous relationship falls from −6.1 to −3.5 for GDP for the full sample, and from −2.5 to −1.8 for industrial production for the no-Japan sample.[14]

Possible Nonlinearities. A final robustness issue involves the scaling of our distress variable. In constructing our measure, we attempted to choose the gradations so that each step (such as credit disruption−regular to credit disruption−plus, or credit disruption−plus to minor crisis−minus) is of roughly equal significance. However, since the descriptions in the *OECD Economic Outlook* are qualitative rather than quantitative, we may not have been completely successful in this effort. Moreover, even if each step is of equal importance in its implications for the cost of credit intermediation, the response of economic activity to increases in the cost of intermediation may not be linear.

As a simple way of shedding light on the possibility that the effects of financial distress are nonlinear in our measure, *F*, we estimate a variant of our baseline specification that allows the effects of distress to be quadratic in *F*. That is, we estimate the system of equations:

$$(3) \qquad y_{j,t+i} = \alpha_j^i + \gamma_t^i + \beta^i f(F_{j,t}) + \sum_{k=1}^{4} \varphi_k^i f(F_{j,t-k}) + \sum_{k=1}^{4} \theta_k^i y_{j,t-k} + e_{j,t}^i,$$

[14] The effects of the heteroskedasticity correction at short horizons come entirely from the observations from Turkey. As mentioned above, Turkey in 2001 stands out as having a large spike in distress and an exceptionally large fall in output. When Turkey is omitted from the sample, the heteroskedasticity-consistent standard errors are actually slightly smaller than the conventional ones.

with $f(F) = F + bF^2$. $b > 0$ corresponds to the case where the gaps between successive steps of our distress measure increase as one moves up the scale, or where the output effects of equal increases in distress rise as distress rises. $b < 0$ corresponds to the opposite case. Our baseline specification corresponds to the case $b = 0$. We estimate (3) using nonlinear least squares. For simplicity, we focus on the results using GDP as the output measure.

The results suggest little nonlinearity. The estimate of b is -0.025 with a standard error of 0.017. Thus, the null hypothesis that the linear specification is correct cannot be rejected. Moreover, the point estimate implies that the variation across categories is only slightly different from what we assume in our baseline scaling. For example, in our baseline scaling, a regular moderate crisis is 4 times as consequential as a regular credit disruption (8 versus 2); in the quadratic specification, it is 3.4 times as consequential ($8 - 0.025(8^2)$ versus $2 - 0.025(2^2)$).

E. Comparison with Results Using Alternative Chronologies

Given that our new series on financial distress differs in important ways from existing crisis chronologies, it is useful to compare our findings for the output response with ones estimated using the other series. As in Section II, we consider both the Reinhart and Rogoff (2009) chronology and the dates from the IMF database, as described by Laeven and Valencia (2014). We again only consider the alternative chronologies for the 24 advanced countries in our sample. To incorporate the alternative chronologies into the empirical framework described above, we simply convert each to a dummy variable equal to 1 in all half-years from the start to the end of a crisis in a given country. This allows us to create a panel dataset similar to that created using our new continuous measure of financial distress.

We estimate equation (1) using each of the alternative crisis series in place of our financial distress variable (the F terms).[15] We again consider horizons up to five years after the innovation in the crisis variable. The resulting impulse response functions show the effect of a

[15] For the Reinhart and Rogoff series, the sample period is 1967:1–2006:2; for the Laeven and Valencia series, it is 1970:1–2006:2.

realization of a 1 in the alternative chronology. Since the impulse response functions we show for our new series are for a realization of a 7 on our scale from 0 to 15 (a moderate crisis–minus), the experiments considered are roughly comparable.

The impulse response functions for real GDP for both of the alternative crisis measures are shown in Figure 6. Panel (a) shows that GDP falls much more slowly when the Reinhart and Rogoff crisis series is used instead of the new series. There is almost no response of output to a crisis in the contemporaneous half-year, and the impact builds gradually over the next year and a half. Panel (b) shows that the response of GDP to an innovation in the Laeven and Valencia crisis variable is substantially faster than that estimated using the Reinhart and Rogoff series, but still much more gradual than when our new measure is used. With the Laeven and Valencia series, the contemporaneous impact is about one-third of the maximum impact and only marginally significant. With the new series, the contemporaneous impact is three-quarters of the maximum, and highly significant.

The difference in the speed of response is unsurprising given the comparison of the various series in major episodes shown in Figure 2. Reinhart and Rogoff date crises much earlier than does the new series. Similarly, Laeven and Valencia also date crises somewhat before the new series, and typically well before the new series shows an acute rise. The consideration of additional real-time evidence discussed in Section II suggests that the new measure may sometimes be slightly late in registering distress. Thus, the actual speed of response of GDP to distress may lie between the results for the new measure and that for the alternative chronologies.

Another feature of the impulse response function estimated using the Reinhart and Rogoff crisis series is that the maximum impact is remarkably small. Following a crisis, real GDP falls only 2.3 percent. This is noticeably less than the 4.2 percent fall we find following a moderate crisis in our new series. The maximum fall using the Laeven and Valencia series is 3.3 percent.

One reason that the GDP effects are smaller using the alternative chronologies than using

the new series may be that the alternative chronologies include crises in Spain in the late 1970s and early 1980s and in Turkey in the early 1980s, while the new measure based on the *OECD Economic Outlook* does not. In both those episodes, real GDP rose steadily, so their inclusion tends to reduce the estimated effect of financial problems. The additional real-time evidence suggests there was significant distress in Turkey that was missed by the OECD, and that there may have been some mild distress in Spain as well. Thus, it possible that the effect of distress in advanced countries shown by our new measure, which is already modest, is an overestimate.

The small response of output in advanced countries that we find using the Reinhart and Rogoff crisis series stands in stark contrast to the usual summary of Reinhart and Rogoff as showing that the recessions following financial crises are particularly severe. Indeed, Figure 14.4 of *This Time is Different* (2009, p. 230) shows that the average peak-to-trough decline in real GDP per capita in a sample of crises is 9.3 percent.

Some of the smaller effect we find using Reinhart and Rogoff's crisis dates in our empirical framework is surely due to the fact that we are only considering crises in advanced economies in the post-1967 period. Most of the large declines behind that 9.3 percent average come from postwar crises in emerging economies (such as Argentina, Thailand, and Indonesia, where the declines in output were well over 10 percent) and from one crisis before World War II (the U.S. Great Depression, where the peak-to-trough decline in per capita GDP was close to 30 percent). Some of the difference also reflects the fact that we are considering the full universe of crises that Reinhart and Rogoff identify in the 24 OECD countries after 1967. The 9.3 percent is the average decline in a selected set of severe episodes. Finally, some of the difference reflects our regression-based empirical approach. Reinhart and Rogoff's measure of the peak-to-trough decline in GDP around crises often includes output falls that predate the crisis. In contrast, the regression procedure only looks at the behavior of output following the crises.[16]

[16] In converting both the Reinhart and Rogoff and the Laeven and Valencia crisis dates into a dummy variable, we set the crisis dummy equal to 1 in all half-years for which the chronology indicates that a country was experiencing a crisis. An alternative is to set it equal to 1 only in the half-year corresponding

A third feature of the impulse response functions for output using the alternative chronologies involves the persistence of the decline in GDP. Recall that for the new series using the full sample, the decline in GDP following a crisis is highly persistent. The same is true following a crisis in the Reinhart and Rogoff series. In contrast, using the Laeven and Valencia series, there is a substantial amount of bounce back: two-thirds of the output decline is reversed within five years after the start of the crisis.

In the results using the new series, the persistence of the effect of a crisis on GDP is due entirely to the special case of Japan. Japan is a large outlier in our scaled crisis series, and so the prolonged slowdown in GDP growth in Japan after 1990 gets tremendous weight in the estimation. When Japan is excluded, the GDP effects disappear within two years. In contrast, the results using the Reinhart and Rogoff series are not very sensitive to the inclusion of Japan. In their non-scaled series with many crises, Japan gets only a small weight, so excluding it matters relatively little. For the Laeven and Valencia series, excluding Japan lessens the maximum GDP decline in the impulse response function by about 20 percent, and causes the effect to go away entirely after five years. This happens even though the series is not scaled because there are only eight crises in advanced economies in the period 1970 to 2006 in their chronology—so the experience of Japan is still getting substantial weight in the full-sample estimates.

The bottom line of the results using alternative crisis chronologies is that for no crisis measure is the estimated impact of severe financial distress very large. Using sensible empirical techniques on the same sample of advanced economies since 1970 shows quite minor effects on output. Financial crises just do not seem to be followed by large declines in output in advanced

---

to the start of the crisis. When we use this alternative in the estimation, the impact of a crisis is slightly more negative and slightly less precisely estimated. For example, using the Laeven and Valencia chronology, the impact on GDP in the full sample after five half-years is –3.3 percent (s.e. = 1.1) for the full dummy and –3.8 percent (s.e. = 1.6) for the start-only dummy. Both of these differences make sense. The fall in output at the start of a crisis tends to be sharper than the rise when the crisis ends. As a result, including late phases of a crisis in the dummy leads to smaller estimated negative impacts. Looking only at the start of crises discards some information about how the crisis variable evolves. Thus, it would be surprising if the standard errors did not rise when the start-only dummy is used.

economies in the modern period, whichever chronology is used.[17]

Nevertheless, there are interesting differences across the results for the various crisis measures. Using our new measure of financial distress, the impact of crises appears to be much faster than when the leading alternative crisis measures are used. Also, when Japan is excluded from the sample, any impacts of crises appear to be much less persistent when the new measure of financial distress is used. Thus, using the new measure challenges the notion that the effects of crises are particularly long-lasting.

IV. VARIATION IN THE AFTERMATH OF FINANCIAL CRISES ACROSS EPISODES

So far, we have used our new measure of financial distress in modern advanced economies to examine the usual relationship between financial crises and economic activity. We turn now to the variation across episodes in the association between financial distress and economic activity. As discussed in Section III, the overall relationship we identify is strongly influenced by the experience of Japan. In this section, we look more systematically for the presence of outliers in the estimation of the usual relationship. We also examine the range of experiences in key episodes, and the degree to which variation in the severity and persistence of financial distress accounts for the different outcomes we observe.

A. Identifying Outliers

The simplest way to get a sense of the presence of outliers, and the variation in the output

[17] That the effects of crises in advanced countries using the Reinhart and Rogoff and Laeven and Valencia chronologies are modest is somewhat surprising given that Cerra and Saxena (2008) find large effects using a simple autoregressive specification similar to our equation (2) and the Caprio and Klingebiel (2003) crisis dates that were the precursor to both of these later alternatives. Some of the difference is surely due to the specific crisis dates used; the Laeven and Valencia series, in particular, modifies the Caprio and Klingebiel dates substantially. But some of the difference is due to the empirical specification. Cerra and Saxena do not include time fixed effects in their regressions, but all of our specifications do (and the *F*-tests indicate strongly that they are important). Omitting time fixed effects leads to substantially larger estimates of the fall in output following crises. Also, in our baseline specification for the alternative chronologies, we use a dummy variable equal to 1 in all half-years when there was a crisis; Cerra and Saxena use a dummy equal to 1 only at the start of a crisis. Though this difference in specification has little impact when the Jordà approach is used, it matters more in the one-equation autoregressive specification used by Cerra and Saxena—mainly through its effect on the experiment simulated.

response to crises more generally, is to examine some partial association scatter plots. Panel (a) of Figure 7 shows the contemporaneous partial association between output and the distress variable, using real GDP as the output measure.[18] The sample period is 1969:1 to 2007:1 as before, and we use the data for all 24 countries in our sample.

The figure shows that there are far more observations in the upper left and lower right of the plot than in the lower left and upper right—consistent with the fact that we find a highly significant negative contemporaneous relationship between financial distress and output. We have labeled the most extreme points. Of these, only Turkey in 2001:1 really stands out as a powerful observation. It is clear that the contemporaneous output decline associated with Turkey's crisis in 2001 was exceptionally large. Japan in 1998:1 and Sweden in 1992:2 appear to be the next two most extreme observations of a contemporaneous negative relationship between financial distress and output.

Panel (a) also shows a handful of powerful observations going in the direction of weakening the contemporaneous relationship between financial distress and output. Norway in 1991:2 stands out as a case where extreme financial distress (a very large positive distress residual) is associated with a very small GDP residual. Similarly, the United States in 1990:1 had a large distress residual but essentially no negative GDP residual. And Sweden in 1993:2 had a very large negative distress residual (as a result of an extreme fall in distress), but a GDP residual close to zero.

Panel (b) of Figure 7 reports the partial association of output 2½ years after the date of the financial distress variable.[19] That is, it shows the scatter of points that determine the value of the $i = 5$ element of the impulse response function. This panel shows a noticeably less tight

[18] Specifically, we regress both real GDP and the financial distress variable at time t on all of the control variables in equation (1): country dummies, time dummies, four lags of GDP, and four lags of the distress variable. We then plot the residuals of the GDP regression (on the y-axis) against the residuals of the distress variable regression (on the x-axis).

[19] Panel (b) is constructed analogously to panel (a). We regress GDP five half-years after time t and the financial distress variable at time t on the same control variables as before, and then plot the residuals against one another.

relationship. It also shows an interesting pattern. Turkey in 2001:1 is again a moderate outlier going in the direction of a negative impact of financial distress. But, by far the much bigger outlier is Japan. Five of the most extreme observations of positive crisis residuals and negative GDP residuals are for Japan in the 1990s. This is consistent with the finding discussed earlier that excluding Japan has a large impact on the estimated impulse response function of GDP to the new financial distress variable, particularly at longer horizons.

As with the contemporaneous partial association, there are also some noticeable examples of residuals weakening the negative relationship between output and financial distress at longer horizons. Norway in 1991:2 had a very large positive distress residual and a substantial positive GDP residual. Sweden in 1993:1 and the United States in 1990:1 show smaller, but still substantial positive distress residuals with positive GDP residuals.

The fact that there are noticeable influential observations in the relationship between financial distress and output, going both in the direction of strengthening the relationship and in weakening it, suggests that there is indeed important variation in the aftermath of financial distress. Each time may not be different, but many episodes certainly seem to be unusual.

B. Analysis of Key Episodes

To investigate the variation in episodes in more detail, we focus on the small number of cases where there was significant financial distress. More specifically, we look at all of the cases where our new measure reached at least 7 on our scale. There are seven of these: Finland in 1993:1; Japan starting in 1997:2 and again in 2002:1; Norway starting in 1991:2; Sweden in 1993:1; Turkey starting in 2001:1, and the United States in 1990:2. For simplicity, we treat the two Japanese crises as a single episode beginning in 1997:2.

We compare what actually happened in each country following its crisis with what one would have predicted based only on the previous behavior of real GDP. To derive the baseline forecast, we run the same type of regressions as before, but without any of the financial distress

variables. That is, using our panel dataset of real GDP in the 24 advanced countries, we estimate:

$$(4) \qquad y_{j,t+i} = \alpha_j^i + \gamma_t^i + \sum_{k=1}^{4} \theta_k^i y_{j,t-k} + e_{j,t}^i,$$

where the definition of all of the variables is the same as before. We estimate (4) for i = 0 to 10 half-years.[20] The sequence of equations shows how for our sample of countries and years taken altogether, output forecasts future output.

To form the forecast for each episode, we take the relevant fitted values for the particular country from the sequence of regressions, using actual GDP data up through the half-year before the acute financial distress.[21] That is, the forecast for the contemporaneous half-year is the fitted value from the regression for horizon 0; the forecast for the t+1 half-year comes from the regression for horizon 1; and so on. To the degree that financial distress builds gradually and such distress reduces output, these effects will be included in the forecast of where an economy was headed before the distress became acute. The red line in each panel of Figure 8 shows the simple univariate forecast of output. The blue line in each panel shows the actual path of output in the five years following the crisis.

In three of the cases, actual GDP is never far below the forecast: Finland, Norway, and Sweden. Indeed, in Norway and Sweden, actual output is substantially above the forecast within a year of the acute crisis. In all three of these episodes, it appears that conditional on what had happened to GDP up through the half-year before the crisis, there was little additional impact from the acute financial distress itself. Importantly, as can be seen in Figure 1, in each of these episodes the spike in distress was quite sharp. Thus, it is unlikely that financial distress working through lagged output was a major source of output behavior either.

---

[20] The starting date for the estimation is 1969:1. As before, we do not use any GDP data after 2007:1. This means that the end date of the sample moves back by one half-year for each step forward in the horizon.

[21] Because the forecasting regression estimation includes all observations for every country, actual outcomes in the key episodes are influencing the forecasts. However, because there are many countries and many years, that influence is limited for any one episode.

In the United States in 1990:2 (panel f), actual GDP is modestly below the univariate forecast. From the contemporaneous half-year through a year after the acute financial distress, real GDP is 1 to 2 percent below forecasted. This could be consistent with acute financial distress having some impact, but not a large one.

Finally, in two cases, Turkey in 2001:1 and Japan in 1997:2, the difference between actual and forecasted GDP is very large. This could suggest that the effects of a financial crisis were particularly large in these two episodes. Interestingly, the time pattern of the forecast errors is very different in the two cases. For Turkey, the largest forecast errors are in the contemporaneous period and six months after the acute distress. At six months, the difference between actual and forecasted output is –12 percent. But the gap then fades rapidly; by three years after the crisis, GDP is above the univariate forecast. For Japan, the difference is initially small, but then rises dramatically. By three years after the crisis, actual output is 12 percent below the forecast based on output up through the half-year before the acute financial distress.

The key point of these forecasting exercises is that the behavior of output following financial crises is highly variable. It ranges from the same or better than predicted from a univariate autoregressive forecasting equation to dramatically worse.

C. Explaining the Variation

Having found substantial variation in the behavior of output following acute financial distress in six key episodes, the natural next step is to try to explain it. The explanation that can be tested most readily is that the variation in the behavior of output following a crisis depends on what happens to the financial distress itself. Even conditional on hitting a 7 on our scale of financial distress, some crises are more severe or more persistent than others. Perhaps those differences explain the variation in the output response.

To test this hypothesis, we expand our simple forecasting equation to include the financial distress variable up through the horizon of the output variable. That is, we estimate:

$$(5) \qquad y_{j,t+i} = \alpha_j^i + \gamma_t^i + \sum_{k=-4}^{i} \varphi_k^i F_{j,t+k} + \sum_{k=1}^{4} \theta_k^i y_{j,t-k} + e_{j,t}^i.$$

As before, we estimate the regression for i = 0 to 10 half-years.

To form the forecasts for each episode, we take the relevant fitted values from the sequence of regressions. Because the financial distress variable has the same horizon as the output variable, we are using the actual evolution of the distress variable in the prediction. As before, we only use output up through the half-year before the crisis in the forecast.

The green line in each panel of Figure 8 shows the forecast based on past output and the actual path of the financial distress variable. The results suggest that the severity or persistence (or both) of the financial distress explains much of the discrepancy between the simple univariate forecast of output and what actually happened in the key episodes.

The most striking case is Japan (panel b). The green and blue lines are very close for the first year following the distress variable reaching 7. That is, the forecast including the actual realizations of the distress variable matches actual GDP closely. After a year, the two diverge. But then the forecast including the actual evolution of the distress variable falls again and remains anemic, concurrent with the second bout of acute distress in Japan in the early 2000s. After 5 years, the forecast including the actual evolution of the distress variable deviates from actual GDP behavior by only about one-sixth as much as the forecast based on output alone. This suggests that much of Japan's dismal economic performance following its initial bout of acute financial distress is related to the fact that the distress itself was particularly extreme and long-lasting. At the same time, the shortfall of GDP from the forecast that includes the actual path of financial distress remains substantial in absolute terms.

For the United States (panel f), the forecast including the evolution of the financial distress variable is again substantially closer to the actual behavior of GDP than is the forecast based only on lagged output. Given that the forecast errors were not very large to begin with, the absolute forecast improvement is not nearly as impressive as that for Japan. But it does suggest

that the large financial disruption at the start of the 1990s may help explain the weak performance of the U.S. economy in this period.

For Turkey (panel e), the forecast including the actual evolution of the financial distress variable is certainly closer to the actual behavior of GDP than is the forecast based only on lagged output. The forecast error in the first half-year after the onset of acute distress falls from −12 percent to −8 percent. However, the fact that there is clearly still a very large discrepancy suggests that factors other than acute financial distress played a role in this episode. For example, the fact that Turkey's financial crisis was accompanied by a severe currency crisis is likely to have exacerbated the output decline. Turkey is also one of the least advanced economies in our sample. It is possible that financial crises have larger impacts or that real GDP is simply more volatile in less developed economies.

For both Finland and Sweden, the univariate forecast and actual GDP are very similar. Nevertheless, including the evolution of the financial distress variable brings the forecast even closer to the actual series, both in the near term and at longer horizons. Indeed, for Finland, the forecast including the evolution of distress is virtually identical to actual output.

Finally, Norway is striking in the degree to which the forecast including the financial distress variable is worse than the simple univariate forecast in predicting the behavior of output following the 1991 crisis. Based on the fact that Norway had a relatively severe panic, with a second acute phase a year after the first, the forecast including the evolution of the distress variable is quite weak for nearly two years. Yet actual output barely paused.

The bottom line of the analysis of this section is that there is substantial variation in the behavior of output following financial distress. The severity and persistence of this distress in particular episodes appears to explain much, but certainly not all, of this variation.

V. CONCLUSION

The conventional wisdom is that financial crises are followed by large downturns and weak

recoveries. Yet at least for advanced countries, the prior evidence supporting this conclusion is limited. It is based on chronologies for crises that are closely related to one another; that were constructed after-the-fact by researchers who were aware of subsequent outcomes; that simplify a continuum of possible degrees of financial distress to a 0-1 variable or just a handful of categories; that combine many crises in developing countries with just a few in advanced countries; and that are often examined using very simple empirical techniques, such as comparisons of recessions with and without crises.

Findings. This paper therefore revisits the aftermath of financial crises in advanced countries. We construct a new financial distress series for 24 advanced countries from 1967 to 2007 using the real-time accounts of countries' conditions prepared twice a year by the OECD. Our new series classifies financial distress into numerous gradations, which we scale from 0 to 15.

The new series displays both important similarities and important differences with previous crisis chronologies. There is agreement about the presence and approximate timing of many of the most severe crises. But we find that real-time awareness of financial distress sometimes did not come until well after the dates identified by earlier chronologies. And a few episodes that previous series identify as important crises do not show up at all in our new measure. Analysis of two other real-time sources largely confirms the accuracy of the new series.

Studying the aftermath of crises using our new series and conventional regression techniques leads to a view that is very different from the conventional wisdom. Crises in advanced countries are associated with falls in output, but the falls are only moderate. When measured using industrial production, output then quickly rebounds and returns to its pre-crisis path. When measured using GDP, output does not bounce back, but this pattern is driven entirely by the experience of Japan.

Interestingly, these findings are due only partly to our new measure of financial distress.

Using conventional regression techniques with standard chronologies of financial crises in advanced countries also does not provide strong support for the view that the aftermaths of crises are persistently grim.

We also find considerable variation in the aftermath of significant crises. In some cases, there was hardly any departure of output from its pre-crisis path. In others, output fell below the pre-crisis path but was soon noticeably above it. And in the case of Japan in the late 1990s, output fell further and further below the pre-crisis path. One factor that appears to be important to the variation is the severity and persistence of the crisis itself. Cases where the financial distress was more limited and resolved quickly were associated with more favorable outcomes.

***Understanding the Aftermath of the 2008 Crisis*.** Our analysis of the impact of financial crises in advanced economies in the postwar period was motivated in part by a desire to understand the experience following the 2008 financial crisis. Why was the downturn in 2008–2009 so extreme and the subsequent recovery so anemic in many advanced countries? Our finding that the fall in output following financial crises in such countries is typically modest suggests that the easy explanation for the poor economic performance—this is what always happens after financial crises—is unlikely to be true. Rather, factors specific to the 2008 episode are likely to have been key.

The factor most directly linked to our results involves the severity and persistence of the recent financial distress. The 2008 crisis was the worst in nearly a century in the United States and Europe. And, although it has been over half a decade since the acute phase of the crisis, banks in many European countries are still viewed as troubled and as constraining the availability of credit. Since we find that greater severity and persistence of financial distress is associated with worse performance following a crisis, the extreme nature of the recent distress is one likely reason for the continuing economic weakness.

A related factor is that the recent financial crisis was essentially worldwide, while most

other postwar episodes of severe financial distress were limited to individual countries or narrow regions. When a crisis is confined to a single country or to just a few, healthy financial institutions abroad may offset some of the reduction in credit supply from domestic lenders. A worldwide crisis means that there are few healthy lenders to step into the void. More generally, a crisis in one country may have adverse spillovers to others. One would therefore expect a given rise in domestic financial distress to be associated with worse outcomes when distress also increases in other countries.

While the greater severity and worldwide nature of the financial distress may be part of the explanation for the terrible aftermath of the 2008 crisis, there are certainly other possibilities. For example, the period before the crisis witnessed large housing booms and accumulations of household debt in many major countries. Many authors, such as Mian and Sufi (2014), argue that these developments, rather than the financial crisis itself, were the main drivers of the downturn and, especially, of the weak recovery. Similarly, worldwide fiscal retrenchment may have played a role. The shifts to tighter fiscal policy that began in continental Europe in late 2009 and early 2010, the United Kingdom in 2010, and the United States in 2011 cannot explain the severity of the downturn, but they may have been important to the slow recovery.

Finally, the zero lower bound on nominal interest rates may have played an important role in amplifying the effects of the recent financial crisis. Central banks in advanced countries often respond to financial distress by reducing their policy rates. But those rates have been at or near zero in most major advanced countries since 2009. The constraint on this usual cushioning mechanism may have magnified the effects of the financial crisis. The hypothesis that the effects of a financial crisis are greater when the zero lower bound binds is consistent with our finding that, even accounting for the actual behavior of financial distress, Japan's output performance following its crisis was particularly poor. Japan's policy rate has been close to zero almost without interruption since shortly after the peak of its financial crisis in the late 1990s, and it is the only country in our sample that encountered the zero lower bound in the period we

consider.

Limitations and Questions for Further Study. In many ways, our findings raise more questions than they answer. With respect to the recent episode, our results cast significant doubt on the simple explanation that what advanced countries have experienced in recent years is just the normal result of a financial crisis. But much research is needed to determine if one of the possibilities mentioned above, a combination of those factors, or something else actually explains the poor economic performance of the past five years.

With respect to the effects of crises more generally, one important limitation of previous work that our paper does not address involves the issue of causation. Lower output is likely to weaken financial institutions and so contribute to financial distress, and there may be factors that cause both higher distress and lower output. As a result of this omitted variable bias, the correlation between distress and output that we find is likely to overstate the adverse effect of crises. One indication that the overstatement may be considerable is that essentially all of the relationship between distress and output that we find arises from the contemporaneous relationship. Isolating the causal effects of crises is a challenging and important topic for future work.

Another limitation of our analysis is that it only looks at the experience of advanced countries. This may in fact be a virtue if the response of advanced countries to financial distress is different from that in developing countries. At the same time, it means that our results do not provide any direct information about the impact of financial distress in emerging economies. Importantly, such issues as potential bias coming from the ex post identification of crises and complications arising from the use of simple summary statistics may also be relevant to investigations of the aftermath of financial crises in developing countries. Indeed, because output in developing countries is generally more volatile and the data are typically of lower quality, these problems could be even more serious in that setting. Thus, addressing these concerns for developing countries is another important area for future research.

In addition, we consider only one source of variation in the aftermath of crises—differences in the severity and persistence of the financial distress itself. As suggested in the case of the 2008 crisis, many other factors may play a role. Thus another valuable area for future research is identifying and determining what else might explain the variation in the effects of distress across episodes.

Finally, there is still a great deal that we need to figure out about the role of policy. In the conventional view that the aftermath of financial crises is uniformly awful, the policy focus is naturally just on preventing crises. But our finding that there is substantial variation in outcomes suggests different questions. Are there policies that can be undertaken in normal times not just to reduce the chances of crises, but to make their consequences less grave if they occur? What policy actions when a crisis occurs—in addition to steps to lessen the crisis and resolve it quickly—could minimize its effects? The answers to these questions are important for ensuring not just that each time is different, but that none is terrible.

TABLE 1
Scale Used to Rank Financial Distress

Category	Numerical Value
No distress	0
Credit disruption–minus	1
Credit disruption–regular	2
Credit disruption–plus	3
Minor crisis–minus	4
Minor crisis–regular	5
Minor crisis–plus	6
Moderate crisis–minus	7
Moderate crisis–regular	8
Moderate crisis–plus	9
Major crisis–minus	10
Major crisis–regular	11
Major crisis–plus	12
Extreme crisis–minus	13
Extreme crisis–regular	14
Extreme crisis–plus	15

TABLE 2
Financial Distress in OECD Countries, 1967:1–2007:1

Finland	
1992:1	Credit disruption–regular
1992:2	Minor crisis–plus
1993:1	Moderate crisis–regular
1993:2	Minor crisis–regular
1994:1	Credit disruption–plus
France	
1991:2	Credit disruption–minus
1995:1	Credit disruption–regular
1995:2	Minor crisis–minus
1996:1	Minor crisis–regular
1996:2	Minor crisis–regular
1997:1	Credit disruption–plus
Germany	
1974:2	Credit disruption–regular
2003:1	Credit disruption–minus
Iceland	
2006:2	Minor crisis–regular
2007:1	Credit disruption–regular
Italy	
1997:1	Credit disruption–minus
Japan	
1990:2	Credit disruption–plus
1991:1	Credit disruption–minus
1991:2	Minor crisis–regular
1992:1	Credit disruption–plus
1992:2	Minor crisis–minus
1993:1	Minor crisis–minus
1993:2	Minor crisis–regular
1994:1	Credit disruption–plus
1994:2	Credit disruption–plus
1995:1	Minor crisis–minus
1995:2	Minor crisis–regular
1996:1	Minor crisis–plus
1996:2	Minor crisis–minus
1997:1	Minor crisis–regular
1997:2	Moderate crisis–minus
1998:1	Major crisis–regular
1998:2	Extreme crisis–minus
Japan (continued)	
1999:1	Moderate crisis–plus
1999:2	Minor crisis–plus
2000:1	Minor crisis–minus
2000:2	Credit disruption–plus
2001:1	Minor crisis–plus
2001:2	Minor crisis–plus
2002:1	Moderate crisis–regular
2002:2	Moderate crisis–minus
2003:1	Minor crisis–plus
2003:2	Minor crisis–regular
2004:1	Minor crisis–minus
2004:2	Credit disruption–plus
2005:1	Credit disruption–regular
Norway	
1991:2	Moderate crisis–plus
1992:1	Minor crisis–regular
1992:2	Moderate crisis–regular
1993:1	Minor crisis–plus
1993:2	Credit disruption–plus
1994:1	Credit disruption–regular
Sweden	
1992:2	Minor crisis–regular
1993:1	Moderate crisis–regular
Turkey	
2001:1	Moderate crisis–regular
2001:2	Moderate crisis–regular
2002:1	Minor crisis–plus
2002:2	Minor crisis–minus
2003:1	Minor crisis–minus
2003:2	Credit disruption–regular
United States	
1986:1	Credit disruption–minus
1990:1	Minor crisis–regular
1990:2	Moderate crisis–regular
1991:1	Minor crisis–minus
1991:2	Credit disruption–plus
1992:1	Credit disruption–regular
1998:2	Credit disruption–plus
2007:1	Credit disruption–minus

Notes: Based on the *OECD Economic Outlook*. The *Economic Outlook* does not describe any disruptions or crises in Australia, Austria, Belgium, Canada, Denmark, Greece, Ireland, Luxembourg, Netherlands, New Zealand, Portugal, Spain, Switzerland, and the United Kingdom. The sample period is 1967:1–2007:1 for all countries except Australia (1971:2–2007:1), Finland (1969:1–2007:1), and New Zealand (1973:1–2007:1). Countries that joined the OECD after 1993 are not considered.

TABLE 3
Granger Causality Tests of the Exogeneity of the New Financial Distress Measure

	Full Sample		No-Japan Sample	
	IP	GDP	IP	GDP
Lag				
1	-0.007	-0.006	-0.000	0.000
	(o.oo5)	(0.009)	(0.005)	(0.008)
2	0.009	0.005	0.006	0.002
	(0.007)	(0.012)	(0.006)	(0.011)
3	-0.016	-0.023	-0.017	-0.019
	(0.007)	(0.012)	(0.006)	(0.011)
4	0.015	0.028	0.012	0.018
	(0.005)	(0.009)	(0.005)	(0.008)
F-statistic	2.855	3.927	2.101	1.501
p-value	0.023	0.004	0.078	0.199

Notes: The coefficients reported are from regressions of the financial distress variable on four own lags, four lags of the output variable, country dummies, and time dummies. We multiply the log of the output variable by 100 so that the coefficient estimates show the impact of a 1 percent change in output. Standard errors are in parentheses. The *F*-statistic is for the hypothesis that all of the coefficients on the output variable are zero.

EXHIBIT 1
Sample Descriptions of Episodes from Online Appendix A

A. Credit Disruption–Regular

Germany, 1974:2. The OECD reported that during the summer, "considerable losses of exchange reserves and the imminent danger of a confidence crisis imposed particular strains on the banking system," and that "[s]pecial credit facilities were extended to small and medium-sized companies and reserve requirements were reduced in September and October" (p. 26; see also p. 51). And in a discussion of Germany, the United Kingdom, and the United States, it stated, "Recent strains on the banking system in all three countries have underlined the unfavourable climate in bank lending markets. ... There is evidence in all three countries that smaller companies have been particularly severely rationed or priced out of bank lending markets" (p. 50). Notably, there was no mention of financial-market difficulties in the section that was specifically devoted to Germany. Given that omission, it is clear that the OECD did not view financial-market problems as being a major factor in the behavior of the German economy. On the other hand, it identified strains on the banking system, and Germany had perceived a need for special facilities to support lending to certain types of businesses.

This disruption seems similar to that in the United States in 1992:1 (which we classify as a regular credit disruption), and less serious than that in the United States in 1991:2 (which we classify as a credit disruption–plus). We code this episode as a credit disruption–regular.

B. Minor Crisis–Regular

France, 1996:1. The OECD reported (p. 78):

> In 1995, the banking sector continued to suffer from low credit demand, high refinancing and operating costs and large provisions for bad debts. As a result, profitability has been very low by international comparison. The State has provided financial support to some banks and insurance companies, and several financial companies have created special corporate structures in order to assure that prudential ratios are higher than required. Lower short-term interest rates will reduce refinancing costs and help the financial sector to restore profitability. However, the current level of provisions still does not cover all doubtful credits as the real estate market has softened again and the restructuring of the banking sector is advancing only slowly.

The combination of the serious problems in the banking sector, the statement that banks faced high refinancing costs, and the fact that the banking problems were not given a central role in the OECD's discussion of the outlook causes us to identify this episode as a minor crisis–regular. This classification is consistent the fact that we classify France in 1995:2 as a minor crisis–minus, and that the description of the health of the banking sector in this issue is slightly more negative.

C. Moderate Crisis–Regular

United States, 1990:2. In the opening sentence of its entry, the OECD cited "fragility of financial markets" as one of three forces acting "to brake domestic demand" (p. 47). The conclusion cited as a downside risk, "Financial markets could encounter significant difficulties in absorbing even the moderate tightening assumed in this projection, leading to less lending and greater weakness in aggregate demand" (pp. 51–52). Elsewhere, the OECD made repeated references to financial-market difficulties, often in stronger terms. For example, in the material preceding the country entries, it said, "Concerns about debt-servicing difficulties of heavily-indebted companies have caused banks to tighten lending criteria, contributing to this apparent credit 'crunch' " (p. 14). It cited the fact that "the restructuring of the

savings and loan industry and stricter regulation of other financial institutions have apparently led to tighter credit conditions" as one of two main factors contributing to low output growth (p. 48). And in an extended discussion, it said (p. 50):

> the squeeze on the assets of financial institutions, combined with tighter capital and prudential regulations, has raised concerns that credit may be unduly restricted even in the absence of active Federal Reserve policy and that the economy may be unusually sensitive to a tightening of the Federal Reserve's stance. Several factors have undermined the financial system: the difficulties faced by savings and loan institutions; the collapse of the junk bond market; sharp declines in real estate prices, especially in the North East; and declines in stock prices and bond ratings of banks.

Finally, there were several references to the costs of "the Savings and Loans rescue programme" (p. 16; see also pp. 3 and 48).

This episode clearly fits our criteria for a moderate crisis. There had been problems requiring significant government intervention, and the OECD referred to an apparent credit crunch and to the financial system as having been undermined. Perhaps more importantly, financial-market problems were cited not just as a risk to the outlook or as a minor influence, but as one of the major influences on the path of the economy. We therefore classify this episode as a moderate crisis–regular.

D. Extreme Crisis–Minus

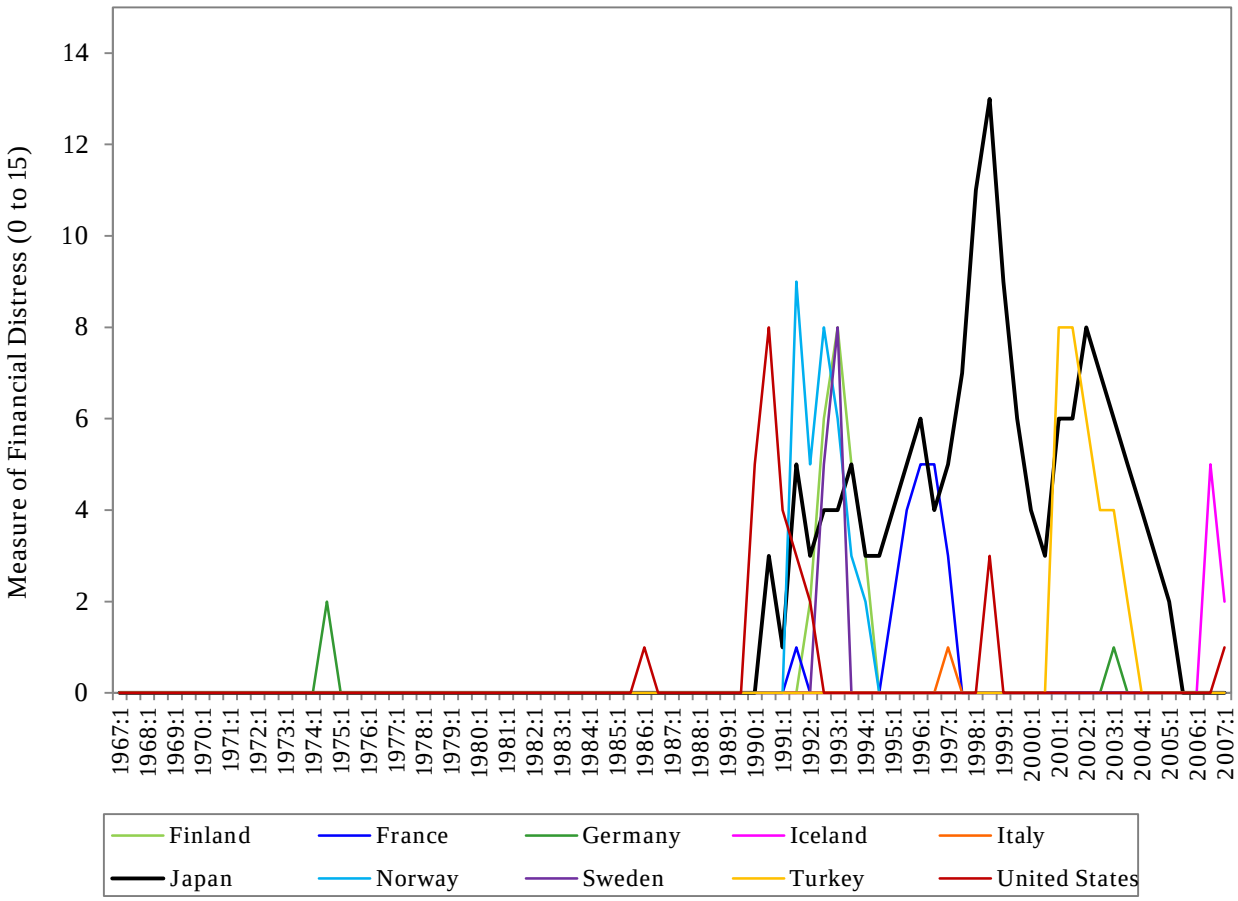
Japan, 1998:2. As just discussed, we classify Japan in 1998:1 as a major crisis–regular. Here, the OECD described a situation that was notably worse. Among its stronger phrases were "financial paralysis" (p. 20); the "breakdown in the credit creation mechanism, and the resulting widening of creditor risk premia" (p. 44); "banks remain in dire straits as risk premia widen" (p. 45); "the increasingly serious situation in the banking sector" (p. 45); and "credit crunch" (which it used repeatedly). In addition, it discussed major government interventions in the financial system: "a broad agreement was achieved in the Diet to revitalise the financial system. The new legislation includes important measures to deal with financial sector problems. To support this, the Government has made an unprecedentedly large sum of public funds available to recapitalize the banking system, amounting overall to around ¥ 60 trillion, or about 12 per cent of GDP" (p. ix).

The OECD made it clear that those developments were having an important impact on the economy. For example, it said, "a profound lack of confidence, in large part due to the severe and prolonged crisis in the banking system, has depressed private spending" (p. 20); reported that "the balance sheet problems of the banking sector remain unresolved, and the resulting uncertainty has led to diminished confidence among consumers and investors, leading to sharp declines in private spending" (p. 42); and referred to "risks of a deflationary spiral arising in part from the unresolved problems in the banking sector" (p. 44).

However, although there had clearly been a nontrivial deterioration from 1998:1, the OECD did not describe the situation as qualitatively changed. For example, it said, "banking sector problems were not improving" (p. ix), and referred to "continued concerns about the health of the financial system" (p. 12). And in the summary of its entry, it stated, "The credit crunch is continuing" (p. 42). Also, as noted above, it commented that "banks remain in dire straits" (p. 45).

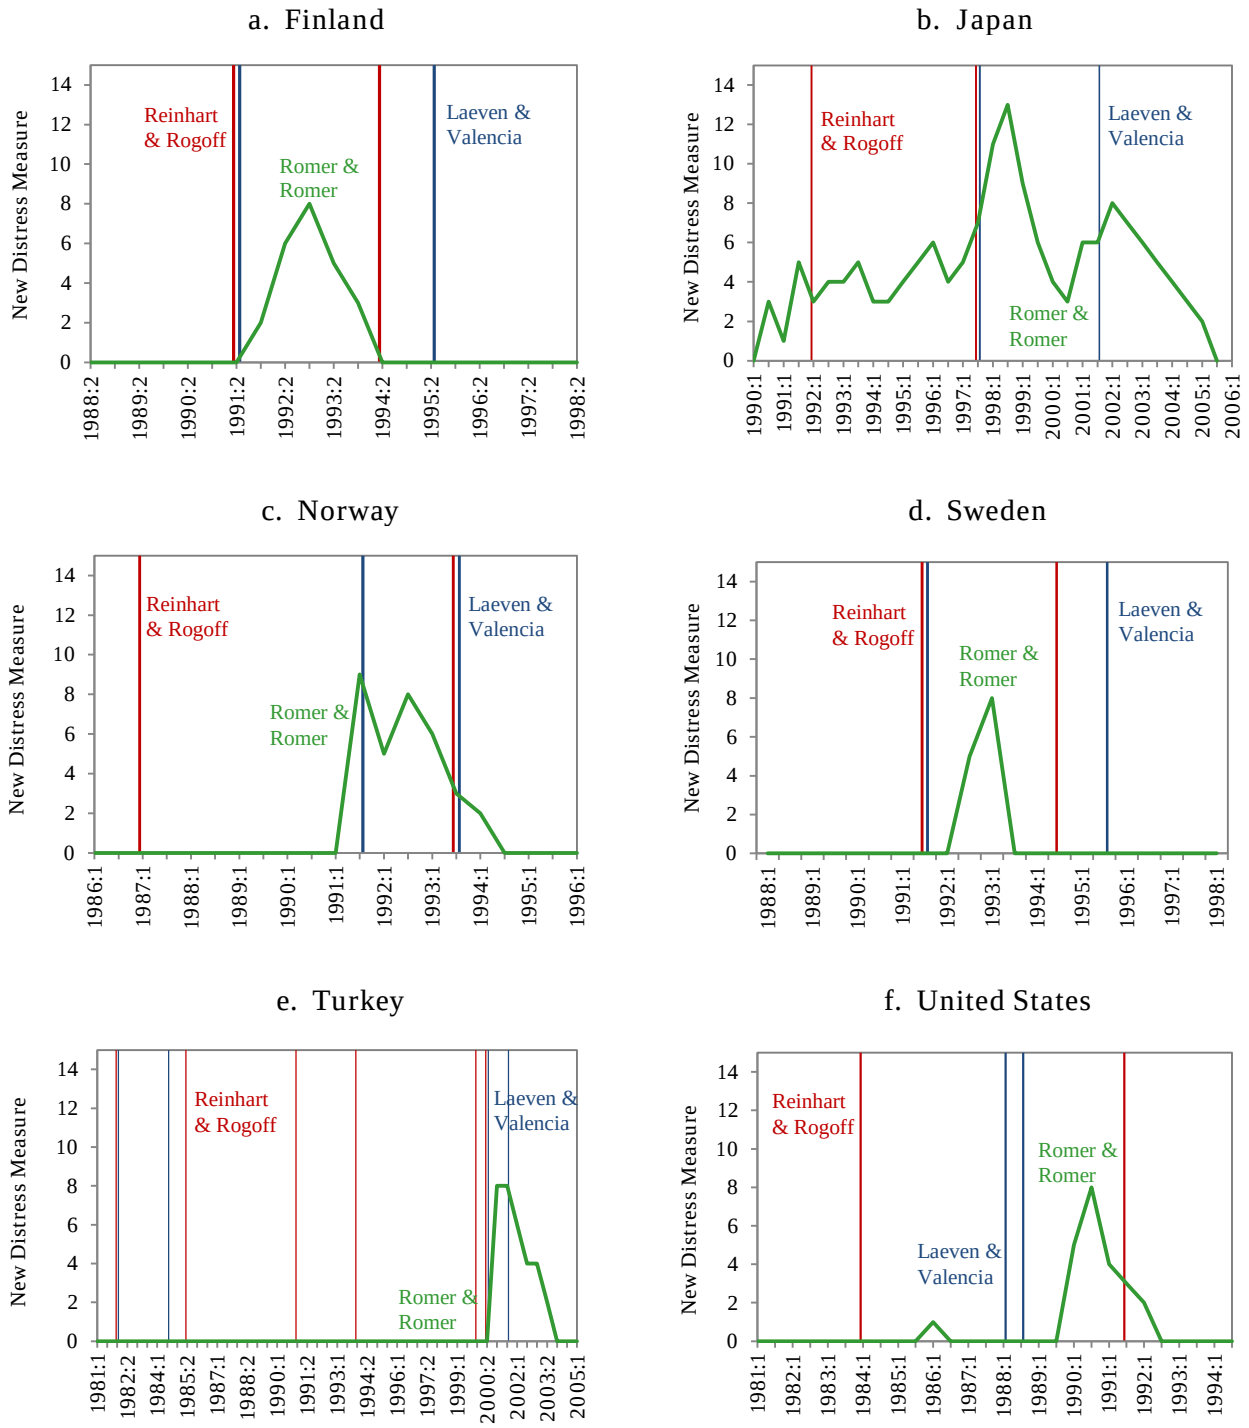
Thus, the financial-sector problems had become significantly but not dramatically worse. We therefore classify this episode as two steps more serious than in 1998:1, which corresponds to an extreme crisis–minus.

FIGURE 1
New Measure of Financial Distress for Advanced Countries



Notes: See text and appendix for details about the derivation of the new measure. The data are available semiannually from 1967:1 to 2007:1. The figure only shows the data for the ten OECD countries that had some nonzero values of our measure.

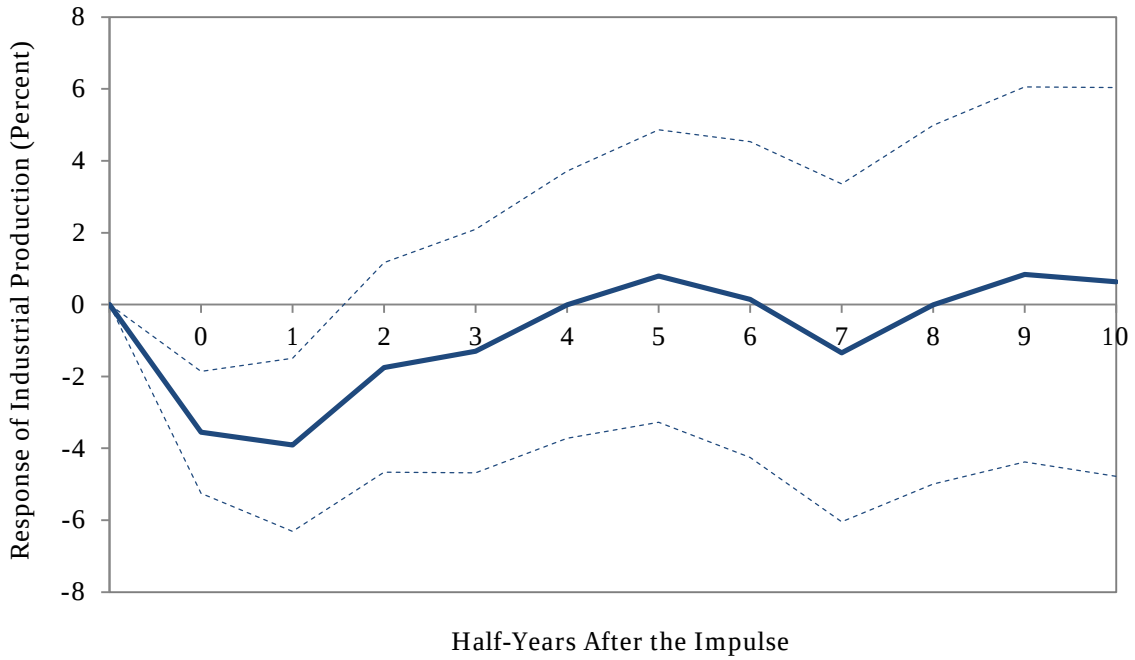
FIGURE 2
Comparison of the New Measure and Other Crisis Chronologies for Key Episodes



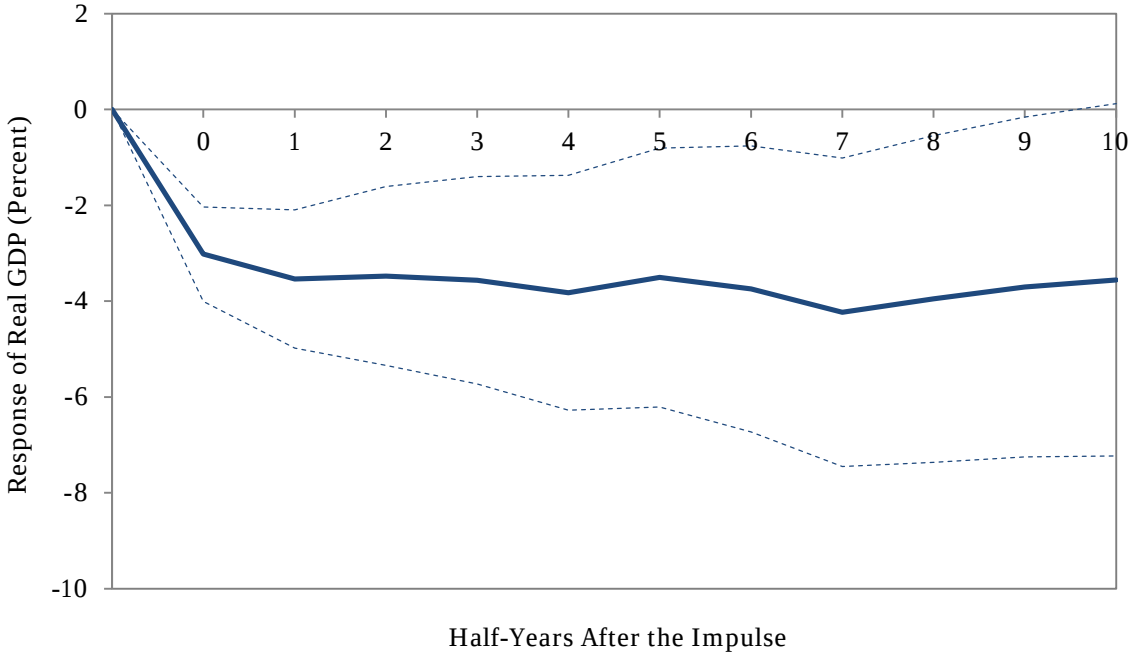
Notes: The vertical lines represent the start and end date of financial crises in the Reinhart and Rogoff (2009) and Laeven and Valencia (2014) chronologies, converted to semiannual observations as described in the text.

FIGURE 3
Impulse Response Functions, Output to Financial Distress, Full Sample

a. Industrial Production



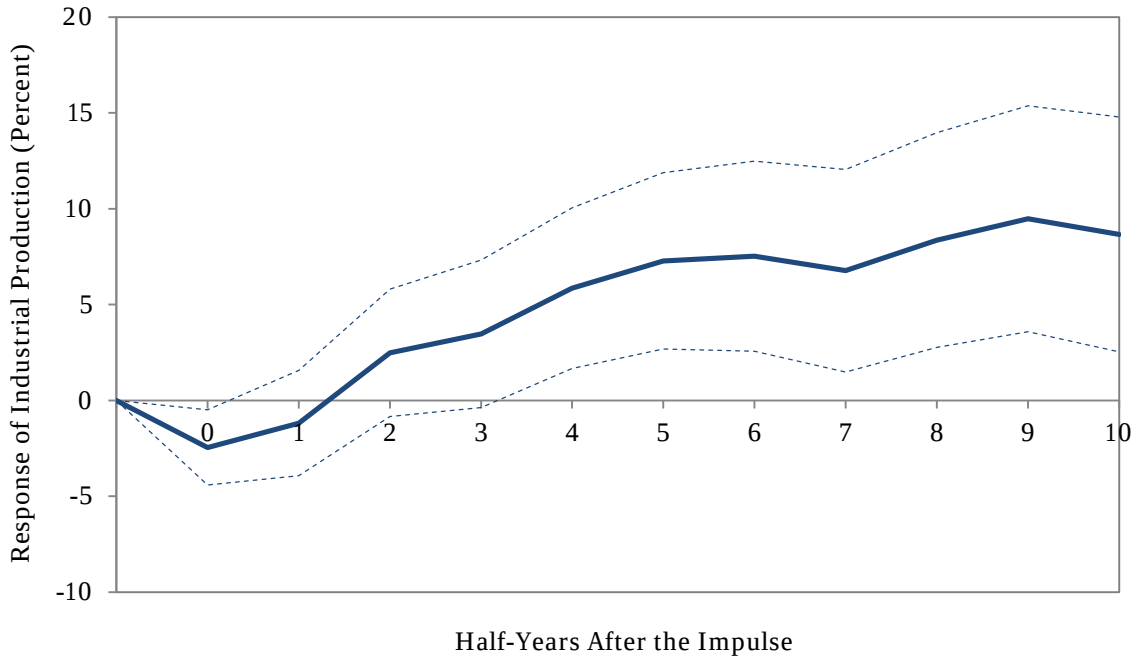
b. Real GDP



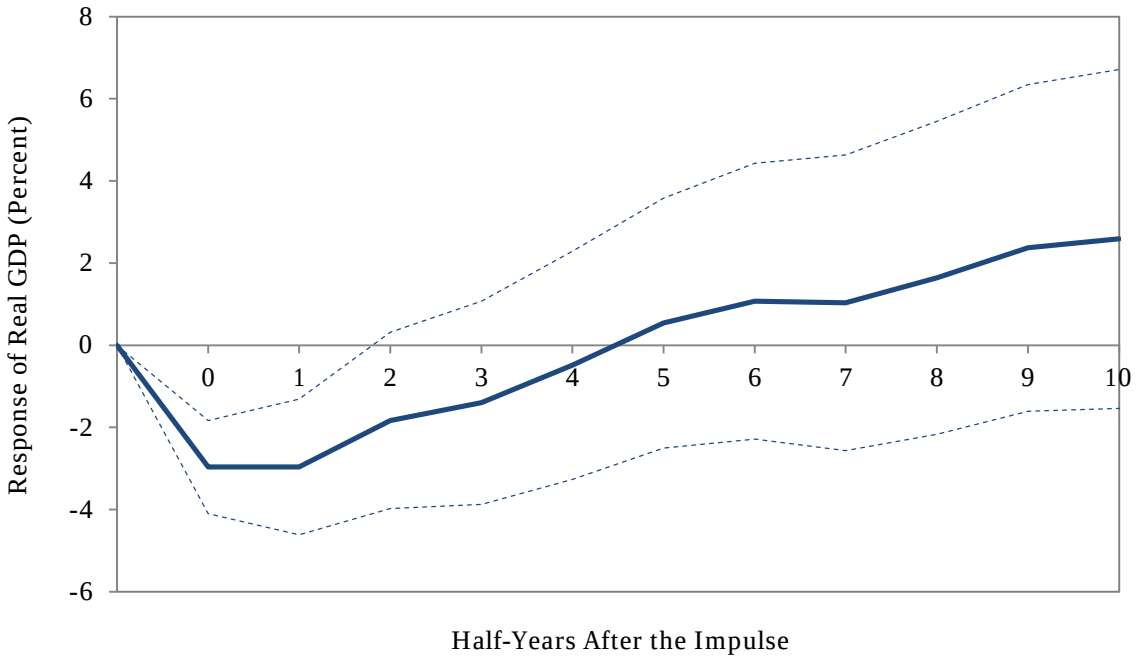
Notes: The figures show the response of output to an impulse of 7 in our new measure of financial distress. The estimates show the results for the full sample of 24 OECD countries. The dashed lines show the two-standard-error confidence bands.

FIGURE 4
Impulse Response Functions, Output to Financial Distress, Excluding Japan

a. Industrial Production



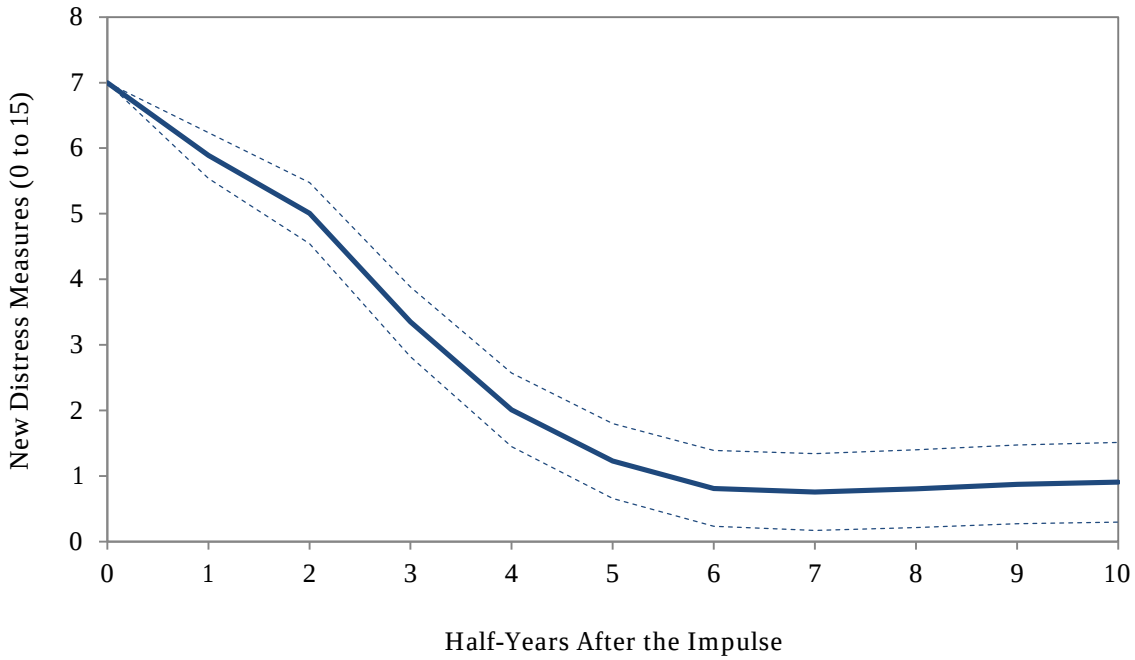
b. Real GDP



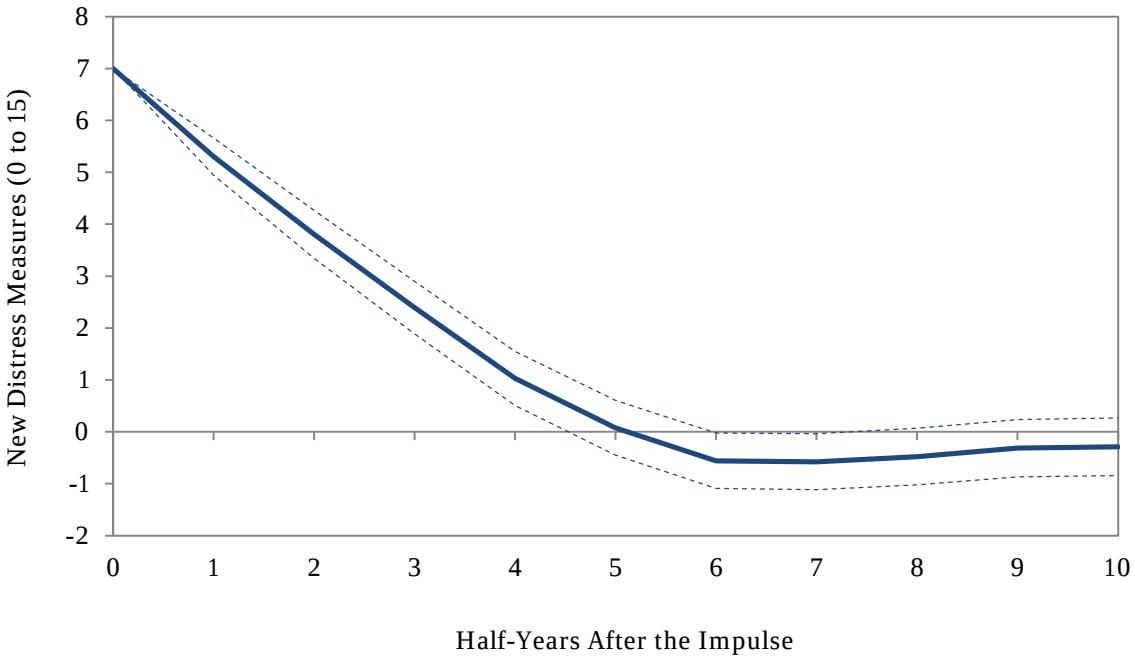
Notes: The figures show the response of output to an impulse of 7 in our new measure of financial distress. The estimates show the results for the sample excluding Japan. The dashed lines show the two-standard-error confidence bands.

FIGURE 5
Impulse Response Function, Distress to Distress

a. Full Sample



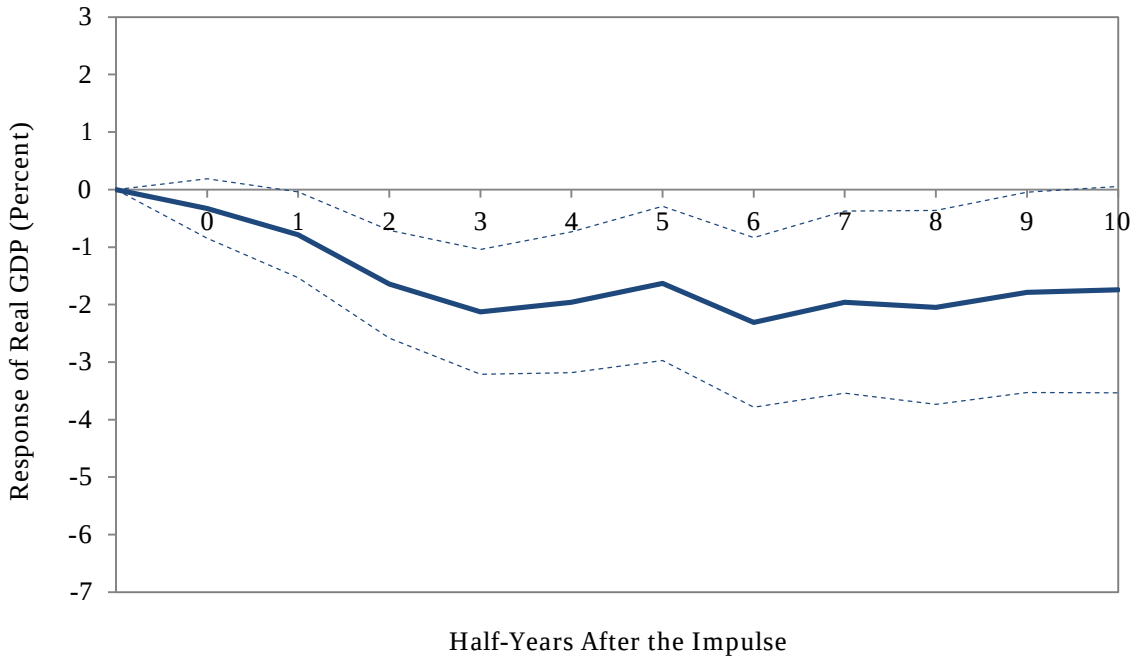
b. Excluding Japan



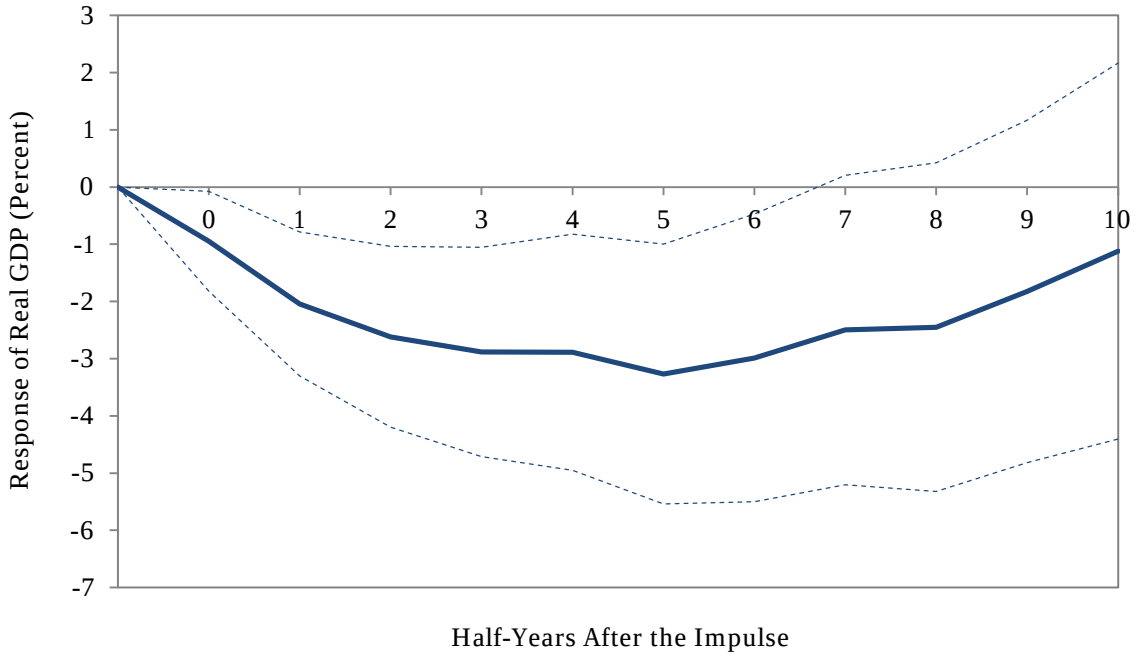
Notes: The figures show the response of the financial distress variable to itself. The initial impulse is taken to be a realization of 7 on our scale of 0 to 15. The dashed lines show the two-standard-error bands.

FIGURE 6
Impulse Response Functions, GDP to Crisis, Other Chronologies

a. Reinhart and Rogoff



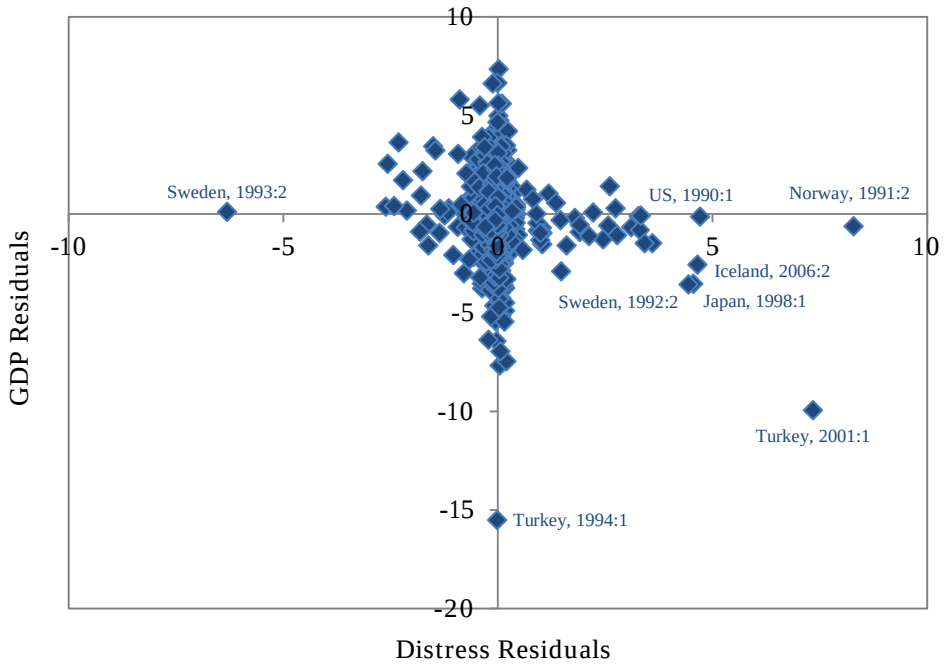
b. Laeven and Valencia



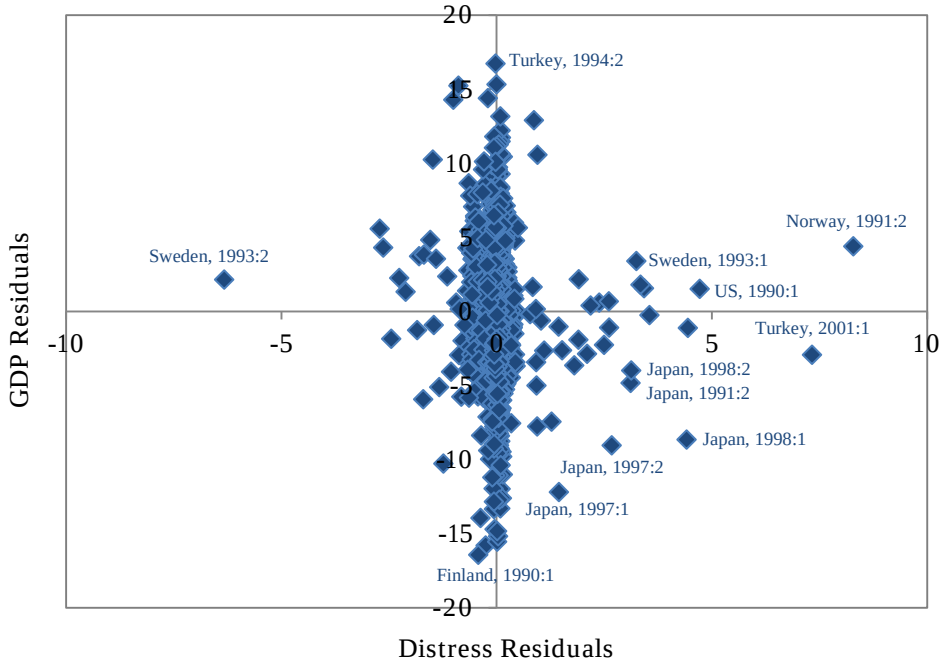
Notes: The figures show the response of GDP to an impulse of 1 in the Reinhart and Rogoff (2009) and Laeven and Valencia (2014) crisis series. The dashed lines show the two-standard-error confidence bands.

FIGURE 7
Partial Association of Financial Distress and Real GDP

a. Contemporaneous

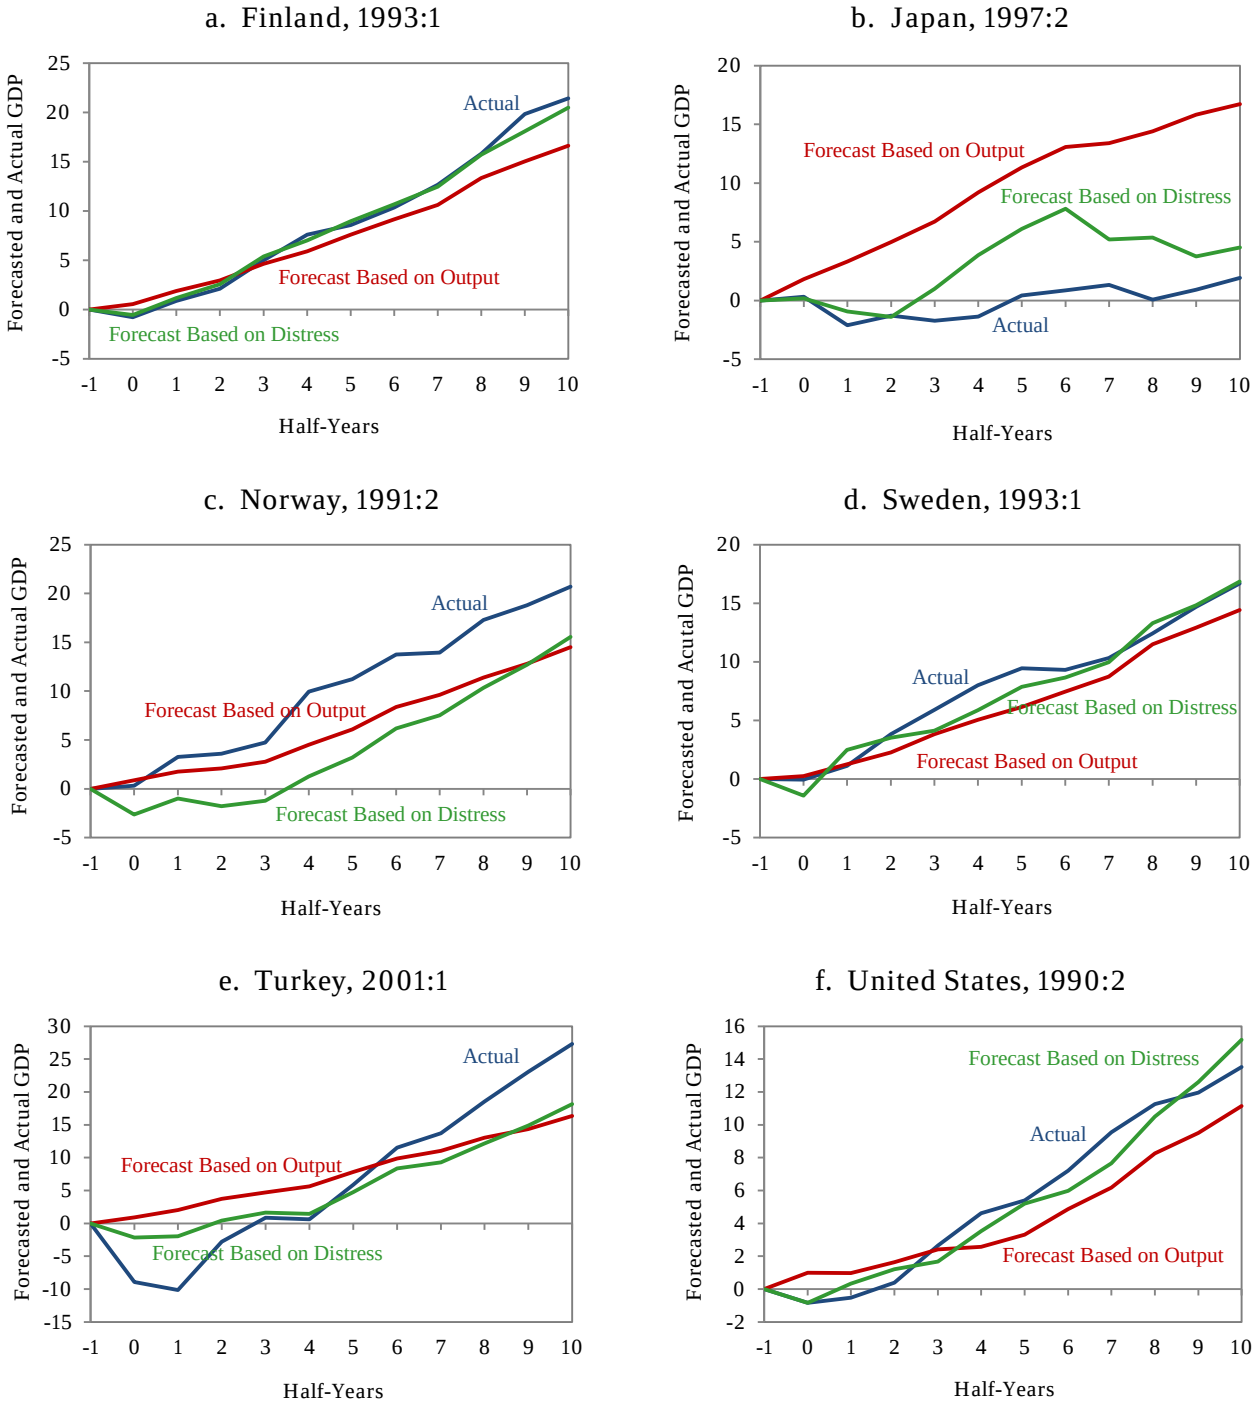


b. Five Half-Years After



Notes: The figures graph the residuals of a regression of the new financial distress series on all of the other right-hand-side variables in the baseline specification against the residuals of a regression real GDP on the same variables. We use the full sample of countries.

FIGURE 8
Actual and Forecasted GDP Following Crises



Notes: All values are expressed as 100 times the difference between the log of the series and the log of actual GDP in the half-year before the distress variable hit 7. Thus it is an index equal to zero in the half-year before the crisis. The forecast based on output uses actual data up through the half-year before the distress variable hit 7. The forecast based on distress uses output through the half-year before the distress variable hit 7 and the actual financial distress series through the date being forecast.

REFERENCES

Bernanke, Ben S. 1983. "Nonmonetary Effects of the Financial Crisis in the Propagation of the Great Depression." *American Economic Review* 73 (June): 257–276.

Bordo, Michael, Barry Eichengreen, Daniela Klingebiel, and Maria Soledad Martinez-Peria. 2001. "Is the Crisis Problem Growing More Severe?" *Economic Policy* 16 (April): 51–82.

Caprio, Gerard, Jr., and Daniela Klingebiel. 1996. "Bank Insolvencies: Cross-Country Experience." World Bank Policy Research Working Paper No. 1620 (July).

Caprio, Gerard, Jr., and Daniela Klingebiel. 1999. "Episodes of Systemic and Borderline Financial Crises." Mimeo, World Bank (October).

Caprio, Gerard, Jr., and Daniela Klingebiel. 2003. "Episodes of Systemic and Borderline Financial Crises." Mimeo, World Bank (January).

Caprio, Gerard, Daniela Klingebiel, Luc Laeven, and Guillermo Noguera. 2005. "Appendix: Banking Crisis Database." In Patrick Honohan and Luc Laeven, eds., *Systemic Financial Crises: Containment and Resolution*. Cambridge: Cambridge University Press, pp. 307–340.

Cecchetti, Stephen G., Marion Kohler, and Christian Upper. 2009. "Financial Crises and Economic Activity." In *Financial Stability and Macroeconomic Policy*. Kansas City: Federal Reserve Bank of Kansas City, pp. 89–135.

Cerra, Valerie, and Sweta Chaman Saxena. 2008. "Growth Dynamics: The Myth of Economic Recovery." *American Economic Review* 98 (March): 439–457.

Claessens, Stijn, M. Ayhan Kose, and Marco E. Terrones. 2009. "What Happens During Recessions, Crunches, and Busts?" *Economic Policy* 24 (April): 653–700.

Claessens, Stijn, M. Ayhan Kose, and Marco E. Terrones. 2014. "The Global Financial Crisis: How Similar? How Different? How Costly?" In Stijn Claessens, M. Ayhan Kose, Luc Laeven, and Fabián Valencia, eds., *Financial Crises: Causes, Consequences, and Policy Responses*. Washington, D.C.: International Monetary Fund, pp. 209–237.

Gourinchas, Pierre-Olivier, and Maurice Obstfeld. 2012. "Stories of the Twentieth Century for the Twenty-First." *American Economic Journal: Macroeconomics* 4 (January): 226–265.

Hoggarth, Glenn, Ricardo Reis, and Victoria Saporta. 2002. "Costs of Banking System Instability: Some Empirical Evidence." *Journal of Banking and Finance* 26 (May): 825–855.

International Monetary Fund. 2009a. "From Recession to Recovery: How Soon and How Strong?" Chapter 3 of *World Economic Outlook: Crisis and Recovery* (April): 97–132.

International Monetary Fund. 2009b. "What's the Damage? Medium-Term Output Dynamics after Financial Crises." Chapter 4 of *World Economic Outlook: Sustaining the Recovery* (October): 121–151.

Jalil, Andrew. 2013. "A New History of Banking Panics in the United States, 1825-1929: Construction and Implications." Working Paper, Occidental College (March). *American Economic Journal: Macroeconomics*, forthcoming.

Jordà, Òscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review* 95 (March): 161–182.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45 (December, Supplement): 3–28.

Kaminsky, Graciela L., and Carmen M. Reinhart. 1999. "The Twin Crises: The Causes of Banking and Balance-of-Payments Problems." *American Economic Review* 89 (June): 473–500.

Laeven, Luc, and Fabián Valencia. 2013. "Systemic Banking Crises Database." *IMF Economic Review* 61 (June): 225–270.

Laeven, Luc, and Fabián Valencia. 2014. "Systemic Banking Crises." In Stijn Claessens, M. Ayhan Kose, Luc Laeven, and Fabián Valencia, eds., *Financial Crises: Causes, Consequences, and Policy Responses*. Washington, D.C.: International Monetary Fund, pp. 61–137.

Mian, Atif, and Amir Sufi. 2014. *House of Debt: How They (and You) Caused the Great Recession, and How We Can Prevent It from Happening Again*. Chicago: University of Chicago Press.

Organisation for Economic Co-Operation and Development. Various years. *OECD Economic Outlook*. Paris: OECD Publishing.

Reinhart, Carmen M., and Kenneth S. Rogoff. 2009. *This Time Is Different: Eight Centuries of Financial Folly*. Princeton, NJ: Princeton University Press.

Romer, Christina D., and David H. Romer. 1989. "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz." *NBER Macroeconomics Annual* 4: 121–170.

Schularick, Moritz, and Alan M. Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008." *American Economic Review* 102 (April): 1029–1061.