

ADVANCE LAYOFF NOTICES AND LABOR MARKET FORECASTING*

Pawel M. Krolikowski[†] Kurt G. Lunsford[‡]

January 31, 2020

Abstract

We collect rich establishment-level data about advance layoff notices filed under the Worker Adjustment and Retraining Notification (WARN) Act since January 1990. We present in-sample evidence that the number of workers affected by WARN notices leads state-level initial unemployment insurance claims, changes in the unemployment rate, and changes in private employment. The effects are strongest at the one- and two-month horizons. After aggregating state-level information to a national-level “WARN factor” using a dynamic factor model, we find that the factor substantially improves out-of-sample forecasts of changes in manufacturing employment in real time.

JEL codes: E27, J65, K31

Keywords: WARN act, mass layoffs, plant closings, unemployment, employment, initial UI claims

*The idea for this project grew out of early conversations with Andrew Foote. We thank Todd Clark, Bruce Fallick, Domenico Giannone, Juan Rubio-Ramirez, Peter Rupert, and Kenneth West for helpful advice and conversations. We also thank seminar participants at the Federal Reserve Bank of Cleveland and conference participants at the Midwest Macroeconomics Meetings. Anna Weixel and Meifeng Yang provided excellent research assistance. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors, the Federal Reserve Bank of Cleveland, or the Federal Reserve System.

[†]pawel.krolikowski@clev.frb.org, +1(216)774-2543, Federal Reserve Bank of Cleveland, 1455 East 6th St, Cleveland, OH 44114

[‡]kurt.lunsford@clev.frb.org, Federal Reserve Bank of Cleveland, 1455 East 6th St, Cleveland, OH 44114

1 Introduction

We collect rich establishment-level data about advance layoff notices filed with state employment offices under the Worker Adjustment and Retraining Notification Act (WARN or the act). The act aims to provide affected workers with a smoother transition to new employment by requiring larger employers to provide at least 60 days' notice in advance of a potential mass layoff. We compile WARN data from several sources, including historical records from state employment offices and state websites. Our data are at the daily frequency, begin in January 1990, and we currently update them twice a month.¹ The states included toward the end of the sample period include a broad geography and cover roughly 90 percent of national initial unemployment insurance (UI) claims. Each collected WARN notice typically includes the name and address of the employment site, the notification date, the anticipated layoff date, and the estimated number of affected workers, among other information. To our knowledge, we are the first to create a national database of WARN notices that may be studied over time.

WARN notices contain information about impending labor market dynamics because establishments typically give their workers 60 days' notice, and this notice duration varies little over the business cycle. In particular, we find that the median number of days between the notice date and the anticipated layoff date is close to 60 days and this median did not materially change before, during, and since the Great Recession. The 25th percentile of the notice duration distribution is about 50 days and the 75th percentile is about 65 days.

After aggregating the number of workers affected by WARN notices to an unbalanced state-month panel, we find that WARN notices can help predict state-level initial UI claims, changes in private employment, and changes in the unemployment rate. In particular, when the number of workers affected by WARN notices rises in a given state, we expect initial claims and the unemployment rate to rise over the next couple of months and we expect employment to fall. For example, when the number of workers affected by WARN notices in a given state rises by one standard deviation, we expect state-level private employment to fall by about one standard deviation over the next two months. The timing of these effects is consistent with the 60 days' notice required by law.

We use a dynamic factor model (Bańbura and Modugno, 2014) to aggregate our unbalanced state-panel data to a national-level "WARN factor." The WARN factor is counter-cyclical, rising sharply just before or in recessions and falling during the initial parts of expansions. The factor also suggests that the labor market deteriorated more quickly during the Great Recession than during the 2001 recessions, which is consistent with other prominent

¹We plan to release our micro-data with a future revision of this paper. An aggregated version of the data is available in [Krolikowski, Lunsford, and Yang \(2019\)](#).

labor market measures like the unemployment rate or initial UI claims. In accompanying work (Krolikowski, Lunsford, and Yang, 2019) we provide in-sample evidence that movements in our WARN factor lead national initial UI claims, changes in the unemployment rate, and changes in private employment, and these effects are strongest at the one-month horizon. For example, a one-standard-deviation increase in the WARN factor this month is typically followed by a decrease in the average change in national private employment over the next three months of about 100,000.

In this paper, we conduct a pseudo real-time forecasting exercise and show that the WARN factor improves out-of-sample forecasts of changes in manufacturing employment, even controlling for other leading labor market indicators, such as the Institute for Supply Management's (ISM's) manufacturing employment and new orders indices and initial UI claims. Vector autoregressions (VARs) suggest that in the current month, and over the next three months, our WARN factor improves the accuracy of forecasts between 10 and 14 percent relative to a standard AR(2) model. In comparison, including ISM manufacturing employment only improves forecast accuracy between 3 and 6 percent. We focus on manufacturing employment for two reasons. First, manufacturing employment and hours are strongly pro-cyclical. Second, manufacturing establishments tend to be larger than establishments overall so that manufacturing establishments are more likely to be covered by the WARN act.

Our work informs and complements recent proposed legislation (the *Fair Warning Act of 2019*) that seeks to update the WARN Act to increase coverage (U.S. Congress, 2019), and we think future projects can leverage the richness of our data to better understand the effect of the act on worker outcomes and regional distress. We envision at least three ways to use these data in the future. First, because WARN data identify the employer, researchers may use it to estimate the causal impact of WARN notices on worker outcomes. Second, our data could be used to study the effects of state amendments to the federal WARN Act. For example, New York changed the notice period from 60 days to 90 days in February 2009. More recently, New Jersey's legislature has considered a similar change (Cryan and Pou, 2019). Our database provides evidence about how these legislative changes impact WARN notice provisions and their timing. Finally, our database could be used to measure economic distress at the regional and state level. This information could be valuable to monetary and fiscal policymakers and is similar in spirit to the requirement of the *Fair Warning Act of 2019* that the Department of Labor create and release a database of all WARN notices. Our WARN notice information could also be used to measure the effects of concentrated local demand shocks, similar to the use of mass layoffs in Foote, Grosz, and Stevens (2019), and related to work by Autor, Dorn, and Hanson (2013) and Notowidigdo (2019). Our present analysis suggests that our WARN data are high quality and relevant in aggregate, and would

benefit projects like these.

Our evidence suggests that, despite the limited coverage of the WARN Act and several exceptions, WARN notices help forecast state- and national-level labor markets. For example, as we document below, the WARN Act only applies to large establishments and to certain events, and even in these circumstances, employers can appeal to several exceptions, including “unforeseeable business circumstances.” Moreover, the WARN Act intends that workers transition to new employment before their anticipated layoff. As a result, we might not expect WARN notices to contain information about initial UI claims or the unemployment rate. Despite these caveats, we find that WARN notices contain valuable information and help improve out-of-sample forecasts.

Our paper is most closely related to the literature that uses labor market flows to help forecast the unemployment rate, as in [Barnichon and Nekarda \(2012\)](#), [Tasci \(2012\)](#), and [Meyer and Tasci \(2015\)](#). WARN notices measure anticipated layoffs, and are therefore likely useful for anticipating inflows into unemployment. [Elsby, Michaels, and Solon \(2009\)](#) and [Barnichon \(2012\)](#) present empirical evidence that job separations account for a substantial part of movements in unemployment, which suggests that forecasting inflows into unemployment may help with forecasting the unemployment and other labor market indicators. [Fujita and Ramey \(2012\)](#) and [Coles and Moghaddasi Kelishomi \(2018\)](#) find that models with endogenous job separations deliver more realistic fluctuations in unemployment and worker flows over the cycle.

Our work also complements an older literature about the WARN Act that focuses on the incidence of notice provision and the outcomes for notified workers, although none of this work uses establishment-level WARN notices, and none evaluates how these notices can help with high-frequency forecasting. [Addison and Blackburn \(1994a\)](#) and [Addison and Blackburn \(1994c\)](#) show that the incidence of lengthy written notices, at most, slightly increased around the passage of WARN, but that formal notices of less than one month became more likely after WARN. The authors suggest that the most likely explanation for this only slight increase in formal notices is that most displacement events are simply not covered by the act, a topic we return to in [Section 2.4.1](#). Advance notices reduce post-displacement non-employment ([Swaim and Podgursky, 1990](#)) by reducing the likelihood of a non-employment spell and, conditional on a non-employment spell, not affecting unemployment duration ([Addison and Portugal, 1987](#); [Addison and Blackburn, 1997](#)), but the size of these effects varies and tends to be small overall ([Addison and Portugal, 1992](#); [Ruhm, 1992](#)). Advance notices seem to improve post-displacement earnings outcomes, especially for workers receiving the longest spells of notice ([Nord and Ting, 1991, 1992](#); [Ruhm, 1994](#)), and lower turnover in subsequent jobs ([Addison and Fox, 1993](#)).

[Section 2](#) briefly describes the WARN Act and our process for collecting the data and

aggregating it to a monthly panel of states, and presents summary statistics, including the coverage of the WARN notices. This section also describes how much advance notice employers actually give workers and how this notice duration varies over the cycle, and presents some preliminary evidence for WARN notices as a leading indicator. Section 3 provides in-sample evidence at the state level about how the number of workers affected by WARN notices leads other state-level labor market indicators. Section 4 aggregates the unbalanced state panel to a national WARN factor using a dynamic factor model, and presents and describes the factor. Section 5 shows how the WARN factor can improve out-of-sample forecasts of manufacturing employment. Section 6 concludes.

2 The WARN data

2.1 The WARN Act

The WARN Act seeks to provide workers with sufficient time to begin new job searches or to obtain necessary training by requiring employers with 100 or more employees to provide their workers with at least 60 days' written notice prior to layoff. The WARN notice must be provided to the affected workers, to the state dislocated worker unit (SDWU),² and to the chief elected official of the unit of local government in which the employment site is located (e.g., the mayor). The penalty for non-compliance can be severe: The employer is liable to each aggrieved employee for an amount including back pay and benefits for the period of violation, up to 60 days. In addition, the employer is subject to a civil penalty of at most \$500 per day, but this fee may be waived if the employer settles liabilities with employees promptly. Relevant cases are brought to federal courts by workers, their representatives, or units of local government. The U.S. Department of Labor is not responsible for enforcing the act. The remainder of this section describes details of the act. [Addison and Blackburn \(1994b,c\)](#) also provide succinct summaries of the legislation, and [Ehrenberg and Jakobson \(1988\)](#) and [Addison and Portugal \(1991\)](#) provide more historical context.

The act was passed on August 4, 1988, and became effective on February 4, 1989. The act was legislated partly in response to a report by the U.S. Government Accountability Office ([GAO, 1987](#)), which found that few laid-off workers had enough notice to obtain job-related assistance. The act covers employers who have 100 or more employees, not counting employees who have worked less than 6 months in the last 12 months and not counting employees who work an average of less than 20 hours a week (part-time workers), or 100 or more workers who work at least a combined 4,000 hours a week, exclusive of overtime. The act covers private and quasi-public employers, including for-profit and non-profit employers,

²SDWUs aim to minimize disruptions associated with job loss, often providing on-site services, such as career counseling, assistance with job search, and job training ([DOLETA, 2019a](#)).

but regular federal, state, and local government employers are not covered. The act also does not cover job loss from temporary facilities, if employment occurred with the understanding that it was temporary, or facilities that are affected by strike or lockout.

WARN notification is triggered by large, permanent reductions (layoffs exceeding 6 months) in the labor force at individual employment sites, defined in two ways. First, a “plant closing” is triggered when an employment site is shut down and 50 or more full-time workers lose their jobs over any 30-day period. Second, a “mass layoff” is triggered if employment is reduced by 500 or more full-time workers at a given site over any 30-day period, or by 50 to 499 full-time workers if they make up at least one-third of the employer’s full-time workforce.³ The act is also triggered if the plant closing or mass-layoff definitions are met by the number of employment losses for two or more groups of workers during any 90-day period. The act applies when hours of work for 50 or more individual employees are reduced by more than 50 percent for each month in any 6-month period. While part-time employees are not considered in determining whether plant closing or mass layoff thresholds are reached, such workers are due notice. Workers who are offered a transfer to another site within a reasonable commuting distance before the plant closing or mass layoff occurs are not counted as having suffered an employment loss under WARN.

Notices must be received at least 60 days before a closing or layoff, with three exceptions, which are purposefully construed narrowly. First, “faltering company” applies when a company is in the process of obtaining capital or business that would allow it to avoid or postpone a shutdown and believe that the advance notice would prevent it from obtaining this capital or business. Second, “unforeseeable business circumstances” applies when business circumstances change suddenly and unexpectedly, such as the unexpected cancellation of a major order. Third, “natural disaster” applies when the workforce reduction is the direct result of, among others, hurricane, flood or earthquake. Even under these exceptions, notice must be given as soon as is “practicable” and the employer must provide a statement of the reason for reducing the notification period.

WARN notices provided to the SDWU should include, among other information, the name and address of the affected employment site, the expected date of the first separation, and the anticipated number of affected employees.

The act does not impose the condition that firms carry out their expected layoffs, but there are likely non-pecuniary costs to issuing unrealized WARN notices. It is possible that firms issue WARN notices and then do not undergo a mass layoff. Nevertheless, issuing a WARN notice might reveal something about the employer’s financial health to employees,

³The one-third rule exempts many large layoffs from coverage by the act according to the [GAO \(1993\)](#) and, as a result, employers are more likely to issue notices of plant closures than mass layoffs ([GAO, 2003](#)). Of course, for reductions in workforce of 500 or more employees, which can only occur at larger firms, the one-third rule does not apply.

which might encourage some workers, especially the most productive ones with outside options, to seek alternative employment. This adverse selection, together with only a slight increase in lengthy written notices around the enactment of WARN, suggests that employers likely do not issue WARN notices without some cause.

Some states and localities have passed their own legislation that expands the federal WARN Act, but we largely abstract from these variations in legislation for the purposes of this paper. For example, the state of New York, since February 2009, requires that employers give their workers 90 days' notice instead of the federally mandated 60 days' notice, and it requires WARN notices from smaller firms and in anticipation of smaller reductions in workforce ([NY WARN, 2019](#)). New Jersey's legislature has recently considered a similar change ([Cryan and Pou, 2019](#)). Some states had voluntary advance notice programs prior to the act ([Ehrenberg and Jakubson, 1990](#)) and before WARN took effect approximately half of all displaced workers received some sort of advance notice, most of which took place "informally" ([Addison and Blackburn, 1994b](#)).

2.2 Revisions and creating a monthly panel

We collect establishment-level WARN notices from local SDWUs, via websites and contacting state officials. As mentioned in the previous section, each notice typically includes the name and address of the employer, the date the notice was filed, the anticipated layoff date, and the estimated number of affected workers, but there are some exceptions. For example, we can obtain the month of notice but not the exact day of notice for several states, including KS, NJ, PA, and UT. For some notices we do not have the estimated number of affected workers, likely due to uncertain business conditions. For some notices we have an associated city or county, but we do not have a street address. There is also some variation in the amount of additional information provided with notices, although we do not exploit this information in the present paper. For example, in some states notices differentiate between a layoff and a plant closure and sometimes they provide the cause of the layoff, such as "economic downturn," "financial," or "September 11th tragedy." Sometimes notices will include the number of workers affected by occupation and establishment industry code.

Employers may submit revisions of original WARN notices to SDWUs and we remove revisions to original WARN notices from our data. We remove revised notices from the data for consistency because in some states (for example, Ohio) we only have original notices. From our sample of states CT, IA, ID, IN, NY, PA, UT, WI, and WV differentiate revised notices from original notices and we remove revisions to original notices. Because of its long history and size, we use original and revised notices from the state of WI to develop an algorithm that identifies notice revisions for states that do not identify revisions from original notices. We find that two criteria work well for identifying revisions to WARN

notices. First, we identify as a revision any notice that is filed by the same company on the second of two subsequent days. This procedure removes any notices that alter the number of affected workers or the anticipated date of the layoff. Second, we identify as a revision any notice that has exactly the same amount of affected workers as a previous notice from the same establishment. This criterion removes any notices that change the anticipated layoff date but leaves the number of affected workers the same. Our algorithm does not remove more than 3 percent of any state’s notices.

To obtain a monthly state panel we sum the estimated number of affected workers across individual establishments in each state-month for every state in our sample and seasonally adjust this series separately for each state. The timing of these notices is based on the notice date. The panel data are unbalanced, because the number of states in our sample changes over time, as discussed in Section 2.4.1. Table 1 lists the set of states in our sample. We use the Census Bureau’s X-12-ARIMA process to seasonally adjust the state-level data.

We drop several states from our analysis due to restrictions on WARN information. First, we have difficulty obtaining and updating complete WARN notice information from several (mostly) smaller states – AR, DE, HI, ME, MT, NH, NV, SD, VT, and WY – so we omit these states from our sample. Second, GA’s and NE’s WARN records do not have a notice date (or month) so we could not include these states in the analysis, and a similar issue arises for SC before 2009 and after 2012. Third, the panel of WARN information for AK, ND, NM, and RI are too short to allow for seasonal adjustment so we drop these states. Finally, we could not obtain historical WARN data for MA.

We also make use of several other data sets in our analysis. For state-level data, we obtain not seasonally adjusted (NSA) data and seasonally adjust with the Census Bureau’s X-12-ARIMA process and for national-level data we obtain published seasonally adjusted (SA) data. We seasonally adjust the state-level data ourselves to maintain consistency as some of the SA series we use are not published, although using NSA data throughout does not alter our conclusions. The data sets we use are as follows. First, we use NSA data from the now-defunct Mass Layoff Statistics (MLS, 2019) program run by the Bureau of Labor Statistics (BLS). The MLS program published statistics between April 1995 and May 2013 from establishments that had at least 50 initial claims for UI filed against them during a 5-week period. Second, we use SA data from the BLS on layoffs and discharges at private employers from the Job Openings and Labor Turnover Survey (JOLTS, 2019b). Third, we obtain NSA state-level unemployment rates from the Local Area and Unemployment Statistics (LAUS) program at the BLS (LAUS, 2019). Fourth, we use estimates of NSA state-level private employment from the State and Metro Area Employment program (SAE, 2019), although using total employment data from LAUS does not change our conclusions. Finally, we use NSA weekly initial UI claims from the Employment and Training Administration

(DOLETA, 2019b). To aggregate the weekly initial UI claims data to the monthly level, we assume a uniform distribution over the week.

2.3 Real-time properties of the WARN data

We began collecting real-time data in March 2019, and we currently update our WARN database in the middle of the month and at the end of the month. We stop counting the number of workers affected by WARN notices in a given state and a given month when we observe a WARN notice for that state in a subsequent month. For some smaller states, this means that we may not have a measure of workers affected by WARN notices in a given month until several months later when a subsequent WARN notice is filed. In this situation, we treat the most recent data for these states as unavailable. In practice, we have found that data for most states for month $t-1$ are available by the middle of month t . For example, data for December 2019 were available for 25 of 33 states as of January 15, 2020, and included essentially all of the large states: CA, FL, IL, MI, NJ, NY, OH, PA, TX, VA and WA.

Our updating schedule is timely relative to other prominent labor market information. Our update in the middle of the month lags the employment report (BLS, 2020) by 5 to 10 days; but our data cover the entire previous month, whereas the employment report covers the reference week that includes the 12th of the month (Current Population Survey) and the pay period that includes the 12th of the month (Current Employment Statistics). Our data are slightly less timely than weekly initial UI claims data. We could make our series more frequent than monthly, although states release their WARN data on different (unannounced) schedules.

2.4 Summary statistics

2.4.1 Monthly panel summary statistics

In this section we describe the WARN state-level data, provide some summary statistics of these data, and provide a sense of its coverage. We also provide summary statistics of the more conventional data measuring state-level unemployment rates, private employment, and initial UI claims, as well as national layoffs and discharges.

As of December 31, 2019, the final WARN data have almost 55,000 WARN notices affecting over six million workers, which implies that the average WARN notice affects about 110 workers.⁴ Over the years 2014 to 2019, when the sample of states has not changed, there

⁴We find that about 30 percent of employers filing WARN notices in our sample anticipate fewer than 50 layoffs, which is below the federally mandated threshold for issuing a WARN notice. This is partly due to variation in legislation by locality (some states require WARN notices when fewer than 50 workers are laid off), as mentioned in Section 2.1, and partly because employers may be confused about when a WARN notice is required, as documented by the GAO (1993, 2003). Ehrenberg and Jakubson (1990) find that, in

were roughly 280,000 workers affected by 2,600 WARN notices each year.

Table 1 presents our baseline sample of states and some summary statistics about the number of workers affected by WARN notices in these states. The data set consists of 33 states, with an average of over 16 years of monthly observations for each state, yielding 6,589 state-month observations. Our first WARN notifications come from the state of Michigan in January 1990, and by January 2000, we have 14 states; by January 2006, we have 21 states; and by January 2019, we have 32 states in our sample. Since data retention policies vary from state to state, we are unable to collect historical data for some states and no state seems to have WARN notice data back to February 1989, when the act took effect. WARN data for the seven most populous states – CA, TX, FL, NY, PA, IL, and OH – start in 08/2005, 01/1999, 01/1998, 10/2001, 01/2011, 01/1999, and 07/1996, respectively. Larger states tend to have more workers affected by WARN notices. Many states have at least one monthly observation that is zero, which means that no workers were affected by WARN notices in that state during the month.

Table 2 presents summary statistics for the number of workers affected by WARN notices in our state-month panel, along with summary statistics for several other state-level labor market indicators. Table 2 (row 1) shows that the average number of workers affected by WARN notices in a month is about 900, with a standard deviation of about 1,150. The median number of workers affected by WARN notices in a month is about 550.

For state-month observations in our sample, using seasonally adjusted data, the median number of workers affected by WARN notices is about 10 percent of the median number of initial UI claimants from plants in the MLS (Table 2, row 2). This suggests that around one-tenth of all workers affected by mass layoffs and plant closures are subject to, and comply with, the act, which is consistent with a [GAO \(2003\)](#) report that finds that about 25 percent of all mass layoffs and plant closures are subject to the act and that employers provide WARN notices for around one-third of these events.

These WARN data cover a small but non-trivial fraction of overall labor market activity, as measured by various other related indicators. First, according to the Business Dynamics Statistics from the U.S. Census Bureau ([Census, 2016](#)), between 1990 and 2014, 60 to 65 percent of employment has been located in firms with at least 100 workers, which is the firm size restriction in the federal WARN Act. Thus, the act covers almost two-thirds of all employment relationships. Second, if all WARN notices end in actual layoff, then WARN notices cover about 1.5 percent of all private-sector layoffs and discharges in the United States as measured by the Job Openings and Labor Turnover Survey ([JOLTS, 2019b](#)). Over the years 2014 to 2018, annual layoffs and discharges among private employers were around 20 million per year. Third, WARN notices cover about 2 percent of all initial UI claims, Pennsylvania, employers filed notices with the SDWU even when WARN did not apply.

which were roughly 13 million over the years 2014 to 2018 (Claims, 2019). Finally, the initial UI claims in our sample of states in 2018 cover almost 90 percent of national initial UI claims.

Table 2 also provides summary statistics for state-level unemployment rates, private employment, and initial UI claims over the unbalanced sample for which we have available WARN data. The average unemployment rate was 5.7 percent over this period, with monthly initial UI claims in an average state being about 40,000 and employment at about 3.25 million. The standard deviation of unemployment rate changes is about 0.2 pp.

2.4.2 How much advance notice do employers give?

Our WARN data will be useful for assessing current and future labor market conditions only if employers issue timely WARN notices. So in this section we investigate how much advance notice employers give their workers and how this notice period varies over the business cycle. To assess the amount of advance notice, we use the number of days between the date a notice was filed and the anticipated layoff date for each establishment-level WARN notification. Figure 1 shows the fraction of WARN notices by the number of days between the date the notice was filed and the anticipated layoff date.

The figure has three features that suggest that WARN notices provide substantial advance layoff notice, largely consistent with the structure of the WARN Act. First, there is a large spike at 60 days. This is not surprising since the law mandates that employers give their workers 60 days advance notice. Second, almost two-thirds of all notices are filed between 40 to 80 days (one to three months) prior to the anticipated layoff date. The 25th (75th) percentile of the notice duration distribution is about 50 (65) days. Third, there is another large increase in the fraction of issued WARN notices around 90 days of notice. This is because, as mentioned in Section 2.1, New York is one of a few (but large) states that deviate from the federal mandate of 60 days' advance notice and, as of February 2009, requires employers in its state to give their workers 90 days' advance notice. There is also a spike at 0 days, which seems to represent employers that file the WARN notice the same day they start laying off some of their workers.

The notice period does not vary much over the business cycle, as shown in Figure 2. The figure shows the median number of days between the layoff announcement and the anticipated layoff date for all WARN notices in a particular month. We begin in January 2006 and include only those states that have available WARN notices during that month to maintain a balanced sample of states, as in Section 2.5.⁵ The figure suggests that the median notice duration is about 60 days, congruent with the WARN Act. Importantly, this lead

⁵We start in January 2006 because CA enters our sample in late 2005 and we have 21 states in our sample at this date. Starting the figure after the 2001 recession does not alter our conclusions. The data become noisy for start dates before the 2001 recession.

time does not vary much before, during, and after the Great Recession, with perhaps only a slight increase in notice duration since 2014. There is a spike in notice duration in August 2015 that is driven by the closure of the Great Atlantic & Pacific Tea Company, which gave its workers 117 days' notice at many locations across the country. As a whole, the evidence in Figure 2 suggests that the temporal structure of our data does not vary substantively over the business cycle.

2.5 WARN as a leading indicator: Preliminary evidence

We present preliminary evidence that the WARN notices may lead initial UI claims from the MLS program and layoffs and discharges in the JOLTS data. The remainder of the paper uses state-level and aggregate data, together with more sophisticated analysis that confirms and expands this central message.

Figure 3 shows the (not seasonally adjusted) aggregate number of workers affected by WARN notices and the (seasonally adjusted) aggregate number of initial UI claimants as covered by the MLS program for each month since January 2006. These data are noisy, but they support the notion that the number of individuals affected by WARN notices tends to move with the MLS initial claimants. Moreover, during the Great Recession, the number of workers affected by WARN notices rose sharply in October 2008, three months before the sharp rise in MLS initial claims in January 2009. Toward the end of the Great Recession, the number of workers affected by WARN notices falls somewhat before MLS. For example, the number of individuals affected by WARN notices peaked in December 2008 and was on a clear decline before MLS initial claims peaked in May 2009.

Figure 4 shows a similar pattern between the (not seasonally adjusted) aggregate number of workers affected by WARN notices and the (seasonally adjusted) aggregate number of layoffs and discharges from JOLTS (total private). Again, the number of individuals affected by WARN notices tends to move with layoffs and discharges in JOLTS, and, during the Great Recession, the number of workers affected by WARN notices rose sharply in October 2008, two months before the sharp rise in JOLTS layoffs in December 2008. The number of layoffs and discharges in JOLTS rises through April 2009, whereas the number of individuals affected by WARN notices peaks in December 2008 and declines considerably by April 2009.

3 State-level evidence

In this section we assess the predictive content of WARN notice information for several labor market indicators at the state level. In particular, the number of workers affected by WARN notices in state s in month t may be informative about the number of imminent initial UI claims in state s . Since the WARN notices measure flows, they may also lead changes in the

unemployment rate and private employment. In this section, we use final data and perform an in-sample analysis.

To assess the predictive content of WARN information, we estimate the following equation:

$$y_{s,t} = \alpha_s + \sum_i \beta_i \text{WARN}_{s,t-i} + \sum_i \eta_i X_{s,t-i} + \epsilon_{s,t} \quad (1)$$

in which s denotes state and t denotes month, $y_{s,t}$ is initial UI claims, the change in the unemployment rate, or the change in private employment, α_s are state fixed effects, $\text{WARN}_{s,t-i}$ is the number of workers in state s affected by WARN notices in month t and its lags, and $X_{s,t-i}$ controls for lags of the left-hand-side variable, $y_{s,t-i}$, and the other variables of interest. Since WARN requires 60 days notice, and in section 2.4.2 we find that most notices occur around one to three months before the anticipated layoff event, we include three lags, so that $i = 1, 2, 3$. This means that the number of observations used for estimation is 99 fewer than in our sample (33 states with the first three observations dropped). We have considered specifications that allow for interactions between lagged variables, such as initial UI claims and state fixed effects, but these deliver similar quantitative results. We do not consider JOLTS layoff and discharge data at the state level because these data are “experimental,” are limited to 3-month moving averages, and rely on statistical models (JOLTS, 2019a).

The number of workers affected by WARN notices leads state-level labor market indicators by one to two months. Table 3 presents the results. Column (1) suggests that if the number of workers affected by WARN notices rises by 1,000 this month (almost a one-standard-deviation increase) in state s , we expect initial claims in state s to rise two months from now by about 425, after controlling for lags of initial UI claims and changes in state-level unemployment and private employment. Although this effect is statistically significant, this magnitude is small as the monthly average of state-level initial UI claims in our sample is about 40,000. If the number of workers affected by WARN notices rises by 1,000 this month in state s , we expect the unemployment rate change to increase by about 0.025 pp and employment growth to slow by about 2,100 over the next two months. The standard deviation of changes in the unemployment rate is 0.2 pp, so the magnitude of the effect of WARN information is meaningful. Similarly, the mean change in state-level private employment over our whole sample is about 2,700, so the predictive content of WARN notices appears to be economically meaningful.

4 National-level aggregation with a dynamic factor model

In this section, we aggregate our state-level panel data to a national-level labor market indicator. The challenge for aggregating our data is that the state-level panel is unbalanced in two distinct ways. First, as discussed above and as shown in Table 1, the state-level data

do not all begin at the same time. For example, MI enters our sample in 1/1990 but CA enters in 8/2005. Second, states do not publish their most recent data at the same time. Some state-level data are available within a few days of the end of the previous month, while data from some smaller states may not be available for several months, as explained in Section 2.3. Hence, the real-time data flow is unsynchronized, leading to a “jagged” or “ragged” edge problem (Bańbura, Giannone, Modugno, and Reichlin, 2013).

Our unbalanced panel prevents us from being able to sum up our WARN data across states as one could, for example, with UI claims. We could restrict our sample to those states that are available at a point in time, such as January 2006, and simply sum those particular states going forward. Indeed, this is what we do for Figures 3 and 4. While this is a useful expository approach, it limits both the historical time-series and the cross-sectional number of states that can be used for more formal statistical analysis. Also, the jagged/ragged edge problem remains and limits the real-time usefulness of the data.

To get around these problems, we model our unbalanced panel with a dynamic factor model (DFM), which we estimate with an expectation maximization (EM) algorithm, as in Dempster, Laird, and Rubin (1977), Shumway and Stoffer (1982), and Bańbura and Modugno (2014). This modeling and estimation approach is commonly used to handle unbalanced panels with jagged/ragged edge problems. The DFM takes the following structure:

$$z_t = d + \Lambda_t f_t + e_t, \quad (2)$$

in which $z_t = [z_{1,t}, \dots, z_{N,t}]'$ with $z_{s,t} = \ln(WARN_{s,t})$ as the vector of observed data, f_t is an unobserved scalar factor, $d = [d_1, \dots, d_N]'$ is a vector of means, $\Lambda_t = [\lambda_{1,t}, \dots, \lambda_{N,t}]'$ is an N -dimensional process of factor loadings, and $e_t = [e_{1,t}, \dots, e_{N,t}]'$ is an N -dimensional process of disturbances. In this model, a given state’s log-level of WARN notices is given by $z_{s,t} = d_s + \lambda_{s,t} f_t + e_{s,t}$.

We interpret f_t as being the national-level “WARN factor” that affects all states through the factor loadings, Λ_t . We assume that f_t follows an AR(1) process:

$$f_t = A f_{t-1} + u_t, \quad (3)$$

in which A is the autoregression slope coefficient and u_t is the autoregression innovation. This AR(1) structure allows the WARN factor to be persistent, similar to other national-level labor market variables. However, we will estimate A and, hence, do not impose persistence.

We make the following assumptions for e_t and u_t :

$$e_t \stackrel{iid}{\sim} N(0, R), \quad u_t \stackrel{iid}{\sim} N(0, Q), \quad (4)$$

in which R is a diagonal matrix with all diagonal elements being strictly positive and $Q > 0$. Equations (2), (3), and (4) compose our DFM.

We use our data sample $\{z_t\}_{t=1}^T$ to estimate the parameters of the model, $\theta = \{\{\Lambda_t\}_{t=1}^T, A, R, Q\}$, and a sequence for the WARN factor, $\{f_t\}_{t=1}^T$. We estimate θ and $\{f_t\}_{t=1}^T$ with maximum likelihood by using the EM algorithm in [Bańbura and Modugno \(2014\)](#). This algorithm, which builds on [Dempster, Laird, and Rubin \(1977\)](#) and [Shumway and Stoffer \(1982\)](#), is designed for data with an arbitrary pattern of missing observations. The algorithm is iterative. Beginning with a guess of θ , the algorithm computes expectations of the DFM's log-likelihood and of $\{f_t\}_{t=1}^T$ conditional on the available data by using the Kalman filter and smoother. Then, the algorithm estimates an update of θ by maximizing the expected log-likelihood. The estimated log-likelihood of the DFM is non-decreasing with each update of θ ([Dempster, Laird, and Rubin, 1977](#); [Shumway and Stoffer, 1982](#)). We end the algorithm when the increase in the log-likelihood is sufficiently small and keep the expected values of $\{f_t\}_{t=1}^T$ from the final iteration as our estimate of the WARN factor.

We make three remarks about our estimation. First, without further restrictions, the scale of the factor loadings and the factor itself are not separately identified. That is, we can rewrite equation (2) to be $z_t = d + \Lambda_t c^{-1} c f_t + e_t$ for any constant scalar $c \neq 0$ and define $\tilde{f}_t = c f_t$, $\tilde{\Lambda}_t = \Lambda_t c^{-1}$, and $\tilde{Q} = c^2 Q$. The expected log-likelihood of the DFM is the same when using $\{\Lambda_t\}_{t=1}^T$ and Q as it is when using $\{\tilde{\Lambda}_t\}_{t=1}^T$ and \tilde{Q} conditional on $\{z_t\}_{t=1}^T$. Intuitively, because the WARN factor is not observed, its units are unknown and not identified from the data.⁶ To make our estimates of the WARN factor interpretable, we normalize its variance to be 1 and we discuss it in terms of its standard deviation.⁷

Second, we assume that $\Lambda_t = \kappa \Psi_t$ in which κ is a scalar, $\Psi_t = [\ln(E_{1,t-1}), \dots, \ln(E_{N,t-1})]' / \sum_{s=1}^{50} \ln(E_{s,t-1})$, and $E_{s,t}$ is employment in state s at time t . We impose this structure on the loadings so that the DFM puts more weight on larger states. We do this following the spirit of [Solon, Haider, and Wooldridge \(2013\)](#) with the intent to make the WARN factor representative of the United States, which is our target population.

Third, building on the previous two remarks, we allow κ to be re-estimated in each iteration of the EM algorithm so that a unit variance on f_t is imposed. However, Ψ_t does not change within the EM algorithm. This structure implies that the ratios of elements in Λ_t , $\lambda_{s,t} / \lambda_{s^*,t}$ for states s and s^* are constant for all iterations.

We estimate the parameters of the DFM and our WARN factor with a sample of July 1996 to December 2019. We start our sample in July 1996 in order to balance two issues.

⁶[Dempster, Laird, and Rubin \(1977\)](#) note that the log-likelihood of a factor model has a ridge of local maxima. [Bańbura and Modugno \(2010\)](#) also have a brief discussion of identification in DFMs.

⁷Formally, we impose the condition that the unobserved factor f_t has a variance equal to 1. The estimate of $\{f_t\}_{t=1}^T$, which we display below and get from the Kalman smoother, has a variance less than 1. For ease of discussion, we discuss our WARN factor in terms of standard deviation units.

First, we want the longest sample possible to produce our WARN factor. This will give us the best sense of the WARN data's aggregate properties and aid in estimating and evaluating the forecasting models in the next section. Second, we want a moderate number of states in our sample in order to get a good estimate of the WARN factor. In July 1996, we observe data for five states: MI, NC, OH, VA, and WI. As we move forward in time, we add state-level data into the sample as it becomes available, and our cross section of states increases (see Table 1). In short, we begin in July 1996 in order to balance having a long time-series against having many states in the initial cross-section.

We make two further remarks about our DFM. First, as shown in the definition of z_t , we take logs of the WARN data before estimating the DFM. We do this because the WARN data are right skewed and bounded below by zero. Taking logs makes the data more approximately normal as assumed in the DFM. Second, as shown in Table 1 some states have observations equal to zero. We do not take logs of these observations; rather, we treat them as missing observations. Similar to the jagged/ragged edge problem at the end of our sample, the EM algorithm handles these missing observations within our sample. Because only smaller states have observations equal to zero, treating these observations as missing puts more weight on larger states for estimating the WARN factor in the corresponding months.

Figure 5 shows the WARN factor and its 68 and 95 percent confidence intervals. We use the parametric bootstrap in [Pfeffermann and Tiller \(2005\)](#) to compute the mean-squared errors for the confidence intervals. We note three important features of our WARN factor. First, the WARN factor rises just before or in recessions, falls during the initial parts of expansions, and is relatively flat during later parts of expansions. While the factor is typically near or below zero outside of recessions, it is significantly above zero for the 2001 recession, the subsequent slow labor market recovery, and the 2008-09 recession. Second, the peak of the WARN factor is higher in the 2008-09 recession than in the 2001 recession. This is consistent with the larger decreases in employment and the larger increase in the unemployment rate in 2008-09 compared to 2001. Third, although the WARN factor has some underlying noise, it has few false positives. When it is one standard deviation or more above zero, it indicates increasing slack in the labor market.

The WARN factor has no economically interpretable units so we supplement it with another aggregate measure. While we can normalize the WARN factor and interpret it in terms of its standard deviation, it may be useful to have an aggregate measure of WARN notice data with economically interpretable units. Toward this end, we estimate a second aggregate measure of the number of workers affected by WARN notices. Given the maximum likelihood estimates of $\hat{\theta}$ and $\{\hat{f}_t\}_{t=1}^T$, we can compute $\hat{z}_{s,t} = \hat{d}_s + \hat{\lambda}_{s,t}\hat{f}_t$ for all s and all t ,

giving us a balanced panel. We then define

$$\widehat{WARN}_t = \sum_{s=1}^N \exp(\hat{z}_{s,t}), \quad (5)$$

which is a scalar process of estimated national WARN notices implied by our DFM, and this estimate can be interpreted in terms of the number of workers affected by WARN notices.

Figure 6 shows \widehat{WARN}_t from July 1996 to December 2019. By construction, it is highly correlated with the WARN factor (correlation above 0.98) and inherits the WARN factor's business cycle properties. One important difference is that the WARN factor is assumed to be normally distributed and so is roughly symmetric around its mean. In contrast, \widehat{WARN}_t displays more skew with larger peaks in the recessions. This asymmetry better matches U.S. labor market data (for example, see Ferraro, 2018), which may help explain why \widehat{WARN}_t forecasts employment changes better than the WARN factor in the next section.

5 Pseudo real-time forecasting

In this section, we demonstrate the WARN data's usefulness in assessing the state of the national labor market by showing that they help improve pseudo real-time forecasts of changes in manufacturing employment. In Krolikowski, Lunsford, and Yang (2019), we use a VAR to show that the estimated factor, $\{\hat{f}_t\}_{t=1}^T$, has predictive power for national UI claims, changes in the national unemployment rate, and changes in national private employment for the sample of July 1996 to October 2019. However, those results are based on an in-sample regression and not on pseudo out-of-sample forecasts.

We denote changes in national manufacturing employment with $\Delta E_{man,t}$. To produce baseline forecasts, we use least squares to estimate a univariate AR(2) model with direct multistep forecasts,

$$\Delta E_{man,t+h} = \beta_{h,0} + \beta_{h,1}\Delta E_{man,t-1} + \beta_{h,2}\Delta E_{man,t-2} + \xi_{t+h}. \quad (6)$$

Throughout this section, h denotes the forecast horizon.

Next, we produce forecasts from competing models. We first use bivariate VAR models with direct multistep forecasts,

$$\begin{bmatrix} \Delta E_{man,t+h} \\ x_{t+h} \end{bmatrix} = \Gamma_{h,0} + \Gamma_{h,1} \begin{bmatrix} \Delta E_{man,t-1} \\ x_{t-1} \end{bmatrix} + \Gamma_{h,2} \begin{bmatrix} \Delta E_{man,t-2} \\ x_{t-2} \end{bmatrix} + \zeta_{t+h}. \quad (7)$$

In these models, x_t is a scalar process. We consider five different choices of x_t : \hat{f}_t , \widehat{WARN}_t , the monthly average of initial UI claims, the ISM's manufacturing employment index, and

the ISM’s manufacturing new orders index. Other than \hat{f}_t and \widehat{WARN}_t , which we introduce in this paper, these other control variables are timely and commonly followed economic indicators. [Barnichon and Nekarda \(2012\)](#) use the monthly average of UI claims in their forecasting model of the unemployment rate, and the Conference Board includes average initial UI claims in its Leading Economic Index (LEI) ([Levanon, Manini, Ozyildirim, Schaitkin, and Tanchua, 2015](#); [CB, 2019](#)). [Lahiri and Monokroussos \(2013\)](#) highlight the timeliness of the ISM indexes for nowcasting gross domestic product. We use the new orders index because the Conference Board includes it in the LEI. We also use the employment index because our target variable is the change in employment.

Next, we use larger-dimensional VARs with direct multistep forecasts,

$$\begin{bmatrix} \Delta E_{man,t+h} \\ X_{t+h} \end{bmatrix} = \Xi_{h,0} + \Xi_{h,1} \begin{bmatrix} \Delta E_{man,t-1} \\ X_{t-1} \end{bmatrix} + \Xi_{h,2} \begin{bmatrix} \Delta E_{man,t-2} \\ X_{t-2} \end{bmatrix} + \eta_{t+h}. \quad (8)$$

In these VARs, we always include the monthly average of initial UI claims, the ISM’s manufacturing employment index, and the ISM’s manufacturing new orders index in X_t . Hence, the VAR is at least 4-dimensional. We also consider two cases where the VAR is 5-dimensional by adding either \hat{f}_t or \widehat{WARN}_t to X_t . We estimate all VARs in this section by least squares.

For all forecasting models, we use 11 years of data to estimate the model parameters and produce the forecasts. We first take July 1996 to June 2007 as our data sample and use these data to forecast the monthly change in manufacturing employment from July 2007 to July 2008. We then take August 1996 to July 2007 as our data sample and use these data to forecast the monthly change in manufacturing employment from August 2007 to August 2008. We move forward one month at a time until we use our final sample of December 2008 to November 2018 to forecast the change in manufacturing employment from December 2018 to December 2019. This gives us 138 forecasts at each forecast horizon for each forecasting model. We take this approach for three reasons. First, we use a rolling sample with a fixed number of observations to facilitate forecast evaluation with nested models ([Giacomini and White, 2006](#)). Second, our first sample begins in July 1996 so that we make full use of our WARN factor, which also begins in July 1996. Third, we choose an 11-year rolling sample so that the first set of forecasts is produced with information through June 2007, which is prior to the start of the 2008-09 recession. Hence, our forecasting exercise indicates how WARN notice data might have changed forecasts from just before the 2008-09 recession to the present.

We make three additional remarks about our forecasts. First, as mentioned in [Section 2.3](#), we have WARN data for a sufficient number of states for month $t - 1$ by the middle of the subsequent month to compute \hat{f}_t and \widehat{WARN}_{t-1} . All other variables for the forecasting models in [\(7\)](#) and [\(8\)](#) are available for month $t - 1$ by the middle of month t , and we assume

that a forecaster has these variables through month $t - 1$ when forecasting $\Delta E_{man,t+h}$ for $h = 0, 1, \dots, 12$. In particular, weekly UI claims are published every Thursday. We use monthly averages of UI claims to compute our forecasts, implying that data for month $t - 1$ are always available within the first week of month t . Also, the ISM indexes for month $t - 1$ also become available in the first week of month t . Finally, changes in manufacturing employment for month $t - 1$ generally become available on the first Friday of month t .

Second, in an effort to replicate a hypothetical forecaster’s information set, we only estimate the DFM with data through month $t - 1$, when forecasting $\Delta E_{man,t+h}$ for $h = 0, 1, \dots, 12$. That is, we re-estimate the DFM every time we move one month forward in our forecasting exercise. One caveat here is that we do not know the exact historical pattern of which states had data for month $t - 1$ by the middle of month t . Because of this, we use all states through month $t - 1$ when estimating the DFM. This gives us slightly more information than what a forecaster would have had in real time. However, as discussed in Section 2.3, most state-level data are available by the middle of the month and only data for small states are typically unknown. Hence, we do not view this information difference as having a big impact on our results. Further, by re-estimating the DFM for each subsequent set of forecasts, we ensure that data for months $t, t + 1, \dots$ are not included when computing the maximum likelihood estimates or the forecasts.

Third, we use the real-time vintages of changes in manufacturing employment. These data are from the Federal Reserve Bank of St. Louis’s ALFRED database ([ALFRED, 2020](#)) with the series code MANEMP. We make one adjustment to the manufacturing employment data. The BLS identified large effects on manufacturing employment from two strikes at General Motors plants in June and July of 1998 ([BLS, 1998](#)). These strikes were resolved in August 1998, introducing a large one-time fluctuation in the manufacturing employment data. Because of this we add 9 thousand and 134 thousand back to manufacturing employment in June and July 1998.

We show our results in Table 4. Row (1) of Table 4 shows the root mean squared prediction errors (RMSPEs) of the baseline forecasting model in (6). We compute the RMSPEs as follows. First, let $\Delta \hat{E}_{man,t+h}$ denote the forecasted change in manufacturing employment at horizon h . Then, $\hat{d}_{h,t} = \Delta \hat{E}_{man,t+h} - \Delta E_{man,t+h}$ is the estimated forecast error, in which $\Delta E_{man,t+h}$ is the value from the January 2020 employment report. Then, we compute the RMSPE with $\sqrt{K^{-1} \sum_{t=1}^K \hat{d}_{h,t}^2}$, in which $K = 138$ is the number of forecast errors. The units in row (1) can roughly be interpreted as the number of employees in thousands. For forecast horizon $h = 0$, this implies that the baseline model’s forecast errors average about 26 thousand employees.⁸ Row (1) shows that these errors grow with the forecast horizon.

⁸Formally, the RMSPE will be slightly larger than the mean absolute prediction error due to Jensen’s inequality.

Rows (2) through (6) of Table 4 show the RMSPEs from the VARs in (7) divided by the RMSPEs from the baseline model in row (1). Values less (greater) than 1 indicate that the VAR in (7) is more (less) accurate than the baseline model. Each of the rows (2) through (6) indicates forecasts produced with the difference choices of x_t . All of the variables that we consider generally improve forecast accuracy at short horizons. However, \hat{f}_t and especially \widehat{WARN}_t yield the largest improvements. \widehat{WARN}_t is the only variable that yields statistically significant improvements, and these improvements range from 11 to 14 percent at horizons zero through three.⁹ These horizons with the biggest forecast improvements are consistent with WARN notices generally having anticipated layoff dates around 60 days in the future as shown in Figure 1.

Row (7) of Table 4 shows the RMSPEs from the model in (8) divided by the RMSPEs from the baseline model in row (1) when only UI claims and ISM indexes are included in the VAR. This model yields reasonably large but not statistically significant forecast improvements at short horizons. Row (8) adds \hat{f}_t and row (9) adds \widehat{WARN}_t to the VAR in (8). These lines show that adding these variables yields larger reductions in RMSPE compared to the baseline model. The VAR with \widehat{WARN}_t has RMSPE reductions between 13 and 19 percent at horizons zero through three compared to the baseline AR(2) model. However, none of these reductions are statistically significant.

To build on the results in rows (7) through (9), we have panel B of Table 4. Row (10) shows the RMSPEs of the VAR model in (8) when only UI claims and ISM indexes are included in the VAR. Rows (11) and (12) show the relative RMSPEs of the VAR model in (8) when either \hat{f}_t or \widehat{WARN}_t is included to the RMSPEs in row (10). Row (12) shows that \widehat{WARN}_t reduces the RMSPEs of the VAR model by about 10 percent at horizons zero through two and these reductions are statistically significant. These results show that WARN notice data can improve the forecast accuracy of a VAR even after controlling for UI claims and ISM indexes. Further, as in row (6), the horizons of the statistically significant RMSPE reductions in row (12) are generally consistent with WARN notices having anticipated layoff dates around 60 days in the future.

Overall, we find that WARN notice data provide useful leading information for manufacturing employment over and above other commonly used leading indicators. We note that these are not obvious results. As we discuss above, the WARN Act has several exceptions that allow firms to issue notices with less than 60 days' notice. There is also evidence that not all firms comply with the WARN Act. For example, the [GAO \(1993\)](#) provides evidence

⁹For statistical significance, let $d_{h,t}^{base}$ and $d_{h,t}^{alt}$ denote the forecast errors from the baseline model and some alternative model, respectively. We test the null hypothesis $\mathbb{E}[(d_{h,t}^{base})^2 - (d_{h,t}^{alt})^2 | \mathcal{I}_t] = 0$ against the alternative $\mathbb{E}[(d_{h,t}^{base})^2 - (d_{h,t}^{alt})^2 | \mathcal{I}_t] \neq 0$, in which \mathcal{I}_t denotes the forecaster's information set when making the forecasts in month t . Following the recommendation in [Clark and McCracken \(2013\)](#), we test the null hypothesis with the modified [Diebold and Mariano \(1995\)](#) test proposed by [Harvey, Leybourne, and Newbold \(1997\)](#).

of noncompliance and delayed compliance in the years shortly after the WARN Act went into effect. Further, some states do not report the exact date that the WARN notice was issued, making the exact lead time unknown, and some notices do not include the number of workers affected or later revise the number of workers affected, potentially introducing noise into our measurement. Nonetheless, our forecasting results show that WARN notice data are of sufficiently high quality and relevance that they provide useful leading labor market information. Hence, WARN data can be helpful for economists who use near-term labor market forecasts, such as business economists and policymakers. In addition, our results also validate the quality and relevance of WARN data for other research topics, such as the impact of WARN notices on worker outcomes, the effect of state amendments to the WARN Act, and the measurement of economic distress at the regional and state level.

6 Conclusion

We compile a unique database with establishment-level WARN notifications starting in January 1990 by collecting information from state officials and websites. We aggregate these data to the state-level and find that the number of workers affected by WARN notices leads other labor market indicators, such as initial UI claims and changes in the unemployment rate and private employment. The lead relationship is strongest in the two months following WARN announcements and is statistically and economically significant. Using a dynamic factor model, we show that an aggregated version of these state-level notices (the “WARN factor”) behaves in intuitive ways, rising sharply before or during recessions and indicating a sharper increase in labor slack during the Great Recession than the 2001 recession. In-sample results suggest that this aggregated measure leads national initial UI claims, changes in the unemployment rate, and changes in private employment. An out-of-sample forecasting exercise that controls for prominent leading indicators, such as the ISM manufacturing employment and new orders indices, suggests that the WARN factor can improve the forecast accuracy of changes in manufacturing employment by as much as 15 percent at various forecast horizons.

Our results suggest that our WARN database contains valuable information for the state of the aggregate economy. We think the richness and timeliness of our data make them useful for other researchers and economic analysts and policymakers.

References

- Addison, John T. and McKinley L. Blackburn (1994a). “Has WARN Warned? The Impact of Advance-Notice Legislation on the Receipt of Advance Notice.” *Journal of Labor Research*, 15(1), pp. 83–90. doi:10.1007/bf02685677. URL <https://link.springer.com/journal/volumesAndIssues/12122>.
- Addison, John T. and McKinley L. Blackburn (1994b). “The Worker Adjustment and Retraining Notification Act.” *Journal of Economic Perspectives*, 8(1), pp. 181–190. doi:10.1257/jep.8.1.181.
- Addison, John T. and McKinley L. Blackburn (1994c). “The Worker Adjustment and Retraining Notification Act: Effects on Notice Provision.” *Industrial and Labor Relations Review*, 47(4), pp. 650–662. doi:10.1177/001979399404700409.
- Addison, John T. and McKinley L. Blackburn (1997). “A Puzzling Aspect of the Effect of Advance Notice on Unemployment.” *Industrial and Labor Relations Review*, 50(2), pp. 268–288. doi:10.1177/001979399705000205.
- Addison, John T. and Douglas A. Fox (1993). “Job Changing after Displacement: A Contribution to the Advance Notice Debate.” *Southern Economic Journal*, 60(1), pp. 184–200. doi:10.2307/1059942. URL <https://www.jstor.org/stable/pdf/1059942.pdf>.
- Addison, John T. and Pedro Portugal (1987). “The Effect of Advance Notification of Plant Closings on Unemployment.” *Industrial and Labor Relations Review*, 41(1), pp. 3–16. doi:10.1177/001979398704100101.
- Addison, John T. and Pedro Portugal (1991). “Advance Notice.” In John T. Addison, editor, *Job Displacement*, chapter 8, pp. 203–243. Detroit: Wayne State University Press.
- Addison, John T. and Pedro Portugal (1992). “Advance Notice and Unemployment: New Evidence from the 1988 Displaced Worker Survey.” *Industrial and Labor Relations Review*, 45(4), pp. 645–664. doi:10.1177/001979399204500402.
- ALFRED (2020). “Archival Economic Data.” <https://alfred.stlouisfed.org/>. Accessed: January 21, 2020.
- Autor, David H., David Dorn, and Gordon H. Hanson (2013). “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review*, 103(6), pp. 2121–68. doi:10.1257/aer.103.6.2121.

- Bañbura, Marta, Domenico Giannone, Michele Modugno, and Lucrezia Reichlin (2013). “Now-Casting and the Real-Time Data Flow.” In Graham Elliott and Allan Timmermann, editors, *Handbook of Economic Forecasting*, pp. 195–237. Elsevier. doi:[10.1016/B978-0-444-53683-9.00004-9](https://doi.org/10.1016/B978-0-444-53683-9.00004-9).
- Bañbura, Marta and Michele Modugno (2010). “Maximum Likelihood Estimation of Factor Models on Data Sets with Arbitrary Pattern of Missing Data.” European Central Bank Working Paper Series no. 1189.
- Bañbura, Marta and Michele Modugno (2014). “Maximum Likelihood Estimation of Factor Models on Datasets with Arbitrary Pattern of Missing Data.” *Journal of Applied Econometrics*, 29(1), pp. 133–160. doi:[10.1002/jae.2306](https://doi.org/10.1002/jae.2306).
- Barnichon, Regis (2012). “Vacancy Posting, Job Separation and Unemployment Fluctuations.” *Journal of Economic Dynamics and Control*, pp. 315–330. doi:[10.1016/j.jedc.2011.09.006](https://doi.org/10.1016/j.jedc.2011.09.006).
- Barnichon, Regis and Christopher J. Nekarda (2012). “The Ins and Outs of Forecasting Unemployment: Using Labor Force Flows to Forecast the Labor Market.” *Brookings Papers on Economic Activity*, pp. 83–117. doi:[10.1353/eca.2012.0018](https://doi.org/10.1353/eca.2012.0018). Available at <https://www.jstor.org/stable/41825365>.
- BLS (1998). “General Motors Strikes Affect Employment Counts.” <https://www.bls.gov/opub/ted/1998/sep/wk5/art01.htm#bls-print>. Accessed: January 21, 2020.
- BLS (2020). “Employment Situation Summary.” <https://www.bls.gov/news.release/empsit.nr0.htm>. Accessed: January 21, 2020.
- CB (2019). “Global Business Cycle Indicators.” <https://www.conference-board.org/data/bcicountry.cfm?cid=1>. Accessed: January 21, 2020.
- Census (2016). “Business Dynamics Statistics: Firm Characteristics Data Tables.” https://www.census.gov/ces/dataproducts/bds/data_firm.html. Accessed: September 10, 2019.
- Claims, UI (2019). “Initial Claims, NSA.” <https://fred.stlouisfed.org/series/ICNSA>. Accessed: September 10, 2019.
- Clark, Todd and Michael McCracken (2013). “Advances in Forecast Evaluation.” In Graham Elliott and Allan Timmermann, editors, *Handbook of Economic Forecasting*, chapter 20, pp. 1107–1201. Elsevier B.V. doi:[10.1016/B978-0-444-62731-5.00020-8](https://doi.org/10.1016/B978-0-444-62731-5.00020-8).

- Coles, Melvyn G. and Ali Moghaddasi Kelishomi (2018). “Do Job Destruction Shocks Matter in the Theory of Unemployment?” *American Economic Journal: Macroeconomics*, 10(3), pp. 118–136. doi:[10.1257/mac.20150040](https://doi.org/10.1257/mac.20150040).
- Cryan, Joseph P. and Nellie Pou (2019). “Increases Prenotification Time and Requires Severance Pay in Certain Plant Closings, Transfers, and Mass Layoffs.” URL https://www.njleg.state.nj.us/2018/Bills/S3500/3170_S2.HTM, December 5.
- Dempster, Arthur P., Nan Laird, and Donald B. Rubin (1977). “Maximum Likelihood from Incomplete Data via the EM Algorithm.” *Journal of the Royal Statistical Society Series B (Methodological)*, 39(1), pp. 1–38. doi:[10.1111/j.2517-6161.1977.tb01600.x](https://doi.org/10.1111/j.2517-6161.1977.tb01600.x).
- Diebold, Francis X. and Roberto S. Mariano (1995). “Comparing Predictive Accuracy.” *Journal of Business & Economic Statistics*, 13(3), pp. 253–263. doi:[10.2307/1392185](https://doi.org/10.2307/1392185).
- DOLETA (2019a). “Rapid Response Services.” <https://www.doleta.gov/layoff/>. Accessed: September 9, 2019.
- DOLETA (2019b). “Unemployment Insurance Weekly Claims Data.” <https://oui.doleta.gov/unemploy/claims.asp>. Accessed: September 12, 2019. Retrieved from Haver.
- Ehrenberg, Ronald G. and George H. Jakubson (1988). *Advance Notice Provision in Plant Closing Legislation*. Kalamazoo: Upjohn Institute. doi:[10.17848/9780880995344](https://doi.org/10.17848/9780880995344). URL https://research.upjohn.org/cgi/viewcontent.cgi?referer=&httpsredir=1&article=1110&context=up_press.
- Ehrenberg, Ronald G. and George H. Jakubson (1990). “Why WARN? Plant Closing Legislation.” *Regulation*, Summer, pp. 39–46. URL <https://www.jstor.org/stable/pdf/2109896.pdf>.
- Elsby, Michael W. L., Ryan Michaels, and Gary Solon (2009). “The Ins and Outs of Cyclical Unemployment.” *American Economic Journal: Macroeconomics*, 1(1), pp. 84–110. doi:[10.1257/mac.1.1.84](https://doi.org/10.1257/mac.1.1.84).
- Ferraro, Domenico (2018). “The Asymmetric Cyclical Behavior of the U.S. Labor Market.” *Review of Economics Dynamics*, 30, pp. 145–162. doi:[10.1016/j.red.2018.05.005](https://doi.org/10.1016/j.red.2018.05.005).
- Foote, Andrew, Michel Grosz, and Ann Stevens (2019). “Locate Your Nearest Exit: Mass Layoffs and Local Labor Market Response.” *Industrial and Labor Relations Review*, 72(1), pp. 101–126. doi:[10.1177/0019793917753095](https://doi.org/10.1177/0019793917753095).

- Fujita, Shigeru and Garey Ramey (2012). “Exogenous versus Endogenous Separation.” *American Economic Journal: Macroeconomics*, 4(4), pp. 68–93. doi:[10.1257/mac.4.4.68](https://doi.org/10.1257/mac.4.4.68).
- GAO (1987). “Plant Closings: Limited Advance Notice and Assistance Provided Dislocated Workers.” URL <https://www.gao.gov/products/HRD-87-105>, GAO/HRD-87-105, Washington, D.C., July 17.
- GAO (1993). “Worker Adjustment And Retraining Notification Act Not Meeting Its Goals.” URL <https://www.gao.gov/products/HRD-93-18>, GAO/HRD-93-18, Washington, D.C., February 23.
- GAO (2003). “The Worker Adjustment And Retraining Notification Act: Revising the Act and Educational Materials Could Clarify Employer Responsibilities and Employee Rights.” URL <https://www.gao.gov/products/GAO-03-1003>, GAO-03-1003, Washington, D.C., September 19.
- Giacomini, Raffaella and Halbert White (2006). “Tests of Conditional Predictive Ability.” *Econometrica*, 74(6), pp. 1545–1578. doi:[10.1111/j.1468-0262.2006.00718.x](https://doi.org/10.1111/j.1468-0262.2006.00718.x).
- Harvey, David, Stephen Leybourne, and Paul Newbold (1997). “Testing the Equality of Prediction Mean Squared Errors.” *International Journal of Forecasting*, 13(2), pp. 281–291. doi:[10.1016/S0169-2070\(96\)00719-4](https://doi.org/10.1016/S0169-2070(96)00719-4).
- JOLTS (2019a). “JOLTS Experimental State Estimates.” https://www.bls.gov/jlt/jlt_statedata.htm. Accessed: September 12, 2019.
- JOLTS (2019b). “Layoffs and Discharges: Total Private, NSA.” <https://fred.stlouisfed.org/series/JTU1000LDL>. Accessed: September 10, 2019. Retrieved from Haver.
- Krolikowski, Pawel M., Kurt G. Lunsford, and Meifeng Yang (2019). “Using Advance Layoff Notices as a Labor Market Indicator.” doi:[10.26509/frbc-ec-201921](https://doi.org/10.26509/frbc-ec-201921). Federal Reserve Bank of Cleveland *Economic Commentary*, no. 2019-21.
- Lahiri, Kajal and George Monokroussos (2013). “Nowcasting US GDP: The Role of ISM Business Surveys.” *International Journal of Forecasting*, 29(4), p. 644–658. doi:[10.1016/j.ijforecast.2012.02.01](https://doi.org/10.1016/j.ijforecast.2012.02.01).
- LAUS (2019). “Local Area Unemployment Statistics.” <https://www.bls.gov/lau/>. Accessed: September 12, 2019. Retrieved from Haver.

- Levanon, Gad, Jean-Claude Manini, Ataman Ozyildirim, Brian Schaitkin, and Jennelyn Tanchua (2015). “Using Financial Indicators to Predict Turning Points in the Business Cycle: The Case of the Leading Economic Index for the United States.” *International Journal of Forecasting*, 31(2), pp. 426–445. doi:[10.1016/j.ijforecast.2014.11.004](https://doi.org/10.1016/j.ijforecast.2014.11.004).
- Meyer, Brent and Murat Tasci (2015). “Lessons for Forecasting Unemployment in the U.S.: Use Flow Rates, Mind the Trend.” doi:[10.26509/frbc-wp-201502](https://doi.org/10.26509/frbc-wp-201502). Federal Reserve Bank of Cleveland, Working Paper no 15-02.
- MLS (2019). “Mass Layoff Statistics.” <https://www.bls.gov/mls/>. Accessed: September 10, 2019.
- Nord, Stephen and Yuan Ting (1991). “The Impact of Advance Notice of Plant Closings on Earnings and the Probability of Unemployment.” *Industrial and Labor Relations Review*, 44(4), pp. 681–691. doi:[10.1177/001979399104400405](https://doi.org/10.1177/001979399104400405).
- Nord, Stephen and Yuan Ting (1992). “The Impact of Advance Notice: A Rejoinder.” *Industrial and Labor Relations Review*, 45(4), pp. 674–682. doi:[10.1177/001979399204500404](https://doi.org/10.1177/001979399204500404).
- Notowidigdo, Matthew (2019). “The Incidence of Local Labor Demand Shocks.” doi:[10.1086/706048](https://doi.org/10.1086/706048). Forthcoming in the *Journal of Labor Economics*.
- NY WARN (2019). “New York State Worker Adjustment and Retraining Notification.” <https://www.labor.ny.gov/workforcenypartners/warn/warnportal.shtm>. Accessed: September 10, 2019.
- Pfeffermann, Danny and Richard Tiller (2005). “Bootstrap Approximation to Prediction MSE for State–Space Models with Estimated Parameters.” *Journal of Time Series Analysis*, 26(6), pp. 893–916. doi:[10.1111/j.1467-9892.2005.00448.x](https://doi.org/10.1111/j.1467-9892.2005.00448.x).
- Ruhm, Christopher J. (1992). “Advance Notice and Postdisplacement Joblessness.” *Journal of Labor Economics*, 10(1), pp. 1–32. doi:[10.1086/298276](https://doi.org/10.1086/298276). URL <https://www.jstor.org/stable/pdf/2535127.pdf>.
- Ruhm, Christopher J. (1994). “Advance Notice, Job Search, and Postdisplacement Earnings.” *Journal of Labor Economics*, 12(1), pp. 1–28. doi:[10.1086/298341](https://doi.org/10.1086/298341). URL <https://www.jstor.org/stable/pdf/2535118.pdf>.
- SAE (2019). “State and Metro Area Employment, Hours, and Earnings.” <https://www.bls.gov/sae/>. Accessed: September 12, 2019. Retrieved from Haver.

Shumway, Robert H. and David S. Stoffer (1982). “An Approach to Time Series Smoothing and Forecasting Using the EM Algorithm.” *Journal of Time Series Analysis*, 3(4), pp. 253–264. doi:[10.1111/j.1467-9892.1982.tb00349.x](https://doi.org/10.1111/j.1467-9892.1982.tb00349.x).

Solon, Gary, Steven J. Haider, and Jeffrey M. Wooldridge (2013). “What Are We Weighting For?” *Journal of Human Resources*, 50(2), pp. 301–316. doi:[10.3368/jhr.50.2.301](https://doi.org/10.3368/jhr.50.2.301).

Swaim, Paul and Michael Podgursky (1990). “Advance Notice and Job Search: The Value of an Early Start.” *Journal of Labor Economics*, 25(2), pp. 147–178. doi:[10.2307/145752](https://doi.org/10.2307/145752).

Tasci, Murat (2012). “Ins and Outs of Unemployment in the Long-Run: Unemployment Flows and the Natural Rate.” doi:[10.26509/frbc-wp-201224](https://doi.org/10.26509/frbc-wp-201224). Federal Reserve Bank of Cleveland Working Paper No. 12-24.

U.S. Congress (2019). “Fair Warning Act of 2019.” URL <https://www.congress.gov/bill/116th-congress/senate-bill/2938>, November 21.




Figure 1: Fraction of WARN notices by the number of days of advance notice

Note: Establishment-level WARN notice data from January 1990 to August 2019. Histogram excludes Pennsylvania, New Jersey, Utah, Kansas, Oklahoma (partial), Minnesota (2000m6-2015m12), and Arizona (2019m7 and on). Total number of observations is 41,901. Outliers are dropped at the 1 percent level on both sides of the distribution. See Section 2.4.2 for details.




Figure 2: Median notice duration in days

Note: The median number of days between the notice date and the anticipated layoff date for a balanced sample of states since January 2006. See Section 2.4.2 for details.




Figure 3: Number of workers affected by WARN leads MLS initial claimants

Note: The number of not seasonally adjusted individuals affected by WARN notices leads the number of seasonally adjusted MLS initial claimants. WARN data are for a balanced panel of states as of January 2006. MLS data are for the same set of states. Initial claimant data from MLS stop in May 2013. MLS data are seasonally adjusted using the Census Bureau's X-12-ARIMA process. Shading represents NBER recession dates. See Section 2.5 for details.




Figure 4: Number of workers affected by WARN leads JOLTS layoffs and discharges

Note: The number of not seasonally adjusted individuals affected by WARN notices leads the number of seasonally adjusted layoffs and discharges in JOLTS (total private). Both series are in thousands. WARN data are for a balanced panel of states as of January 2006. JOLTS data are for the nation. JOLTS seasonally adjusted data are from the BLS. Shading represents NBER recession dates. See Section 2.5 for details.




Figure 5: The WARN factor

Note: Our estimated WARN factor from July 1996 to December 2019 with 68 and 95 percent confidence intervals. The factor is estimated using the method in [Bańbura and Modugno \(2014\)](#). Standard errors are computed using the parametric bootstrap approach in [Pfeffermann and Tiller \(2005\)](#). Vertical grey bars indicate recessions. See Section 4 for details.




Figure 6: The time-series of \widehat{WARN}_t

Note: \widehat{WARN}_t is the aggregate number of workers affected by WARN notices in month t implied by our dynamic factor model in Figure 5 (equation 5). The sample is July 1996 to December 2019. Vertical grey bars indicate recessions. See Section 4 for details.

State	Start	Avg. # of workers affected	Median # of workers affected	Std. dev	# of months with 0 # of workers affected
AL	1998m7	618	473	583	10
AZ	2010m4	374	249	398	6
CA	2005m11	5397	5105	2063	0
CO	2006m1	409	324	344	7
CT	1998m1	381	292	348	4
FL	1998m1	1463	1297	785	0
IA	2011m1	259	181	244	9
ID	2009m1	116	31	299	27
IL	1999m1	1490	1173	1266	0
IN	1997m1	948	693	860	5
KS	1998m8	392	233	492	15
KY	1998m10	428	361	402	25
LA	2007m1	305	231	295	15
MD	2010m1	373	290	362	6
MI	1990m1	954	696	907	10
MN	2000m6	261	140	288	11
MO	2000m7	584	497	444	8
MS	2010m7	364	273	335	7
NC	1996m5	1217	1008	838	5
NJ	2004m1	1190	900	944	3
NY	2001m10	1982	1750	997	0
OH	1996m7	1436	1153	1003	0
OK	1999m11	236	131	304	28
OR	2008m10	267	193	288	4
PA	2011m1	1026	910	711	0
SC	2009m1	678	552	438	0
TN	2012m1	643	568	443	3
TX	1999m1	1885	1573	1272	0
UT	2009m1	150	99	169	12
VA	1994m7	862	664	669	6
WA	2004m1	524	401	469	7
WI	1996m1	843	710	531	4
WV	2011m3	287	140	476	7

Table 1: Summary statistics of state-month panel of the number of workers affected in WARN notices

Note: There are 33 states in our final sample. We drop AR, DE, HI, ME, MT, NH, NV, SD, VT, and WY due to data accessibility. We drop GA and NE because they did not have notice dates. We drop MA because we could not obtain historical data. We drop AK, ND, NM, and RI because the time series could not be seasonally adjusted. SC and MN are the only states that have data for some period and then drop out of our sample. SC drops out because subsequent to December 2012, no notice month is available. MN drops out because of missing data after April 2015. All other states have data through August 2019. See Section 2.4.1 for details.

	Mean	Median	Std. dev	Min	Max	Observations
(1) $\text{WARN}_{s,t}$ (lvl)	919	553	1,168	0	13,229	6,589
(2) MLS initial claims $_{s,t}$ (lvl)	6,268	4,597	6,612	0	69,157	2,761
(3) JOLTS (LDs) $_{s,t}$ (k)	1,763	1,718	175	1,482	2,523	167
(4) UI claims $_{s,t}$ (lvl)	41,387	28,935	41,310	3,869	351,466	6,589
(5) $U_{s,t}$ (%)	5.7	5.2	2.1	2.1	14.8	6,589
(6) $\Delta U_{s,t}$ (pp)	-0.01	-0.02	0.20	-0.92	2.2	6,589
(7) $E_{s,t}$ (k)	3,264	2,428	2,578	483	14,956	6,589
(8) $\Delta E_{s,t}$ (k)	2.8	2.1	11.9	-157.3	181.6	6,588

Table 2: Summary statistics of state-month panel of labor market indicators

Note: Summary statistics for seasonally adjusted data since January 1990, when we first have state WARN notice information, to December 2019 for the states in our sample. “ $\text{WARN}_{s,t}$ ” stands for the number of individuals affected by WARN notices in state s in month t . “MLS” stands for Mass Layoff Statistics, “JOLTS” stands for Job Openings and Labor Turnover Survey, and “LDs” stands for layoffs and discharges. The last row has one fewer observation than the three rows above because state-level employment data begin in January 1990, whereas the other series have longer histories. As a result, the first difference for this series is missing in January 1990. The MLS program started in April 1995 and ended in May 2003 and therefore there are fewer observations in row (2). JOLTS data are for the nation (total private), and all other series are for our sample of states when WARN information is available. JOLTS seasonally adjusted data are from the BLS, and all other series are seasonally adjusted using the Census Bureau’s X-12-ARIMA process. “lvl” stands for levels, “pp” stands for percentage points, and “k” stands for thousands. See Section 2.4.1 for details.

	(1)	(2)	(3)
	ui claims _{s,t}	$\Delta U_{s,t}$	$\Delta E_{s,t}$
$WARN_{s,t-1}$	210 (167)	0.020*** (0.0041)	-1,168*** (201)
$WARN_{s,t-2}$	423*** (77)	0.0056* (0.0032)	-957*** (238)
$WARN_{s,t-3}$	-20 (144)	0.0010** (0.0042)	-137 (223)
3 lags of UI claims _{s,t}	✓	✓	✓
3 lags of $\Delta U_{s,t}$	✓	✓	✓
3 lags of $\Delta E_{s,t}$	✓	✓	✓
State FEs	✓	✓	✓
p-value $WARN_{s,t-i} = 0$	0.00001	0.0004	0.000006
Observations	6,490	6,490	6,490
R-squared	0.98	0.13	0.42

Table 3: WARN as a leading state-level indicator

Note: WARN notices lead state-level initial UI claims, changes in the unemployment rate, and changes in private employment. The number of observations in columns (1) and (2) is 99 less than in Table 2 (rows 4 and 6) because equation (1) includes 3 lags and we have 33 states in our sample. Standard errors clustered at the state level. ***p<0.01, **p<0.05, *p<0.1. See Section 3 for details.

	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$	$h = 5$	$h = 6$	$h = 7$	$h = 8$	$h = 9$	$h = 10$	$h = 11$	$h = 12$
Panel A: Forecast Results of and Comparisons to the Baseline AR(2) Forecasting Model													
(1) Baseline RMSPE	26.1	31.2	35.1	39.5	43.0	45.8	48.5	49.9	51.2	52.6	54.1	55.4	56.2
Relative RMSPEs of Bivariate VARs:													
(2) UI Claims	0.99	0.97	0.97	0.95	0.97	0.98	0.98	0.97	0.99	0.98	0.99	1.00	0.99
(3) ISM Emp	0.94	0.94	0.95	0.97	0.98	0.98	0.98	0.99	1.00	1.00	1.00	1.00	1.00
(4) ISM New Order	0.97	0.97	0.99	0.99	1.01	1.02	1.02	1.03	1.01	1.01	1.01	1.01	1.01
(5) \hat{f}_t	0.92	0.91	0.91	0.92	0.91	0.90	0.90	0.91	0.91	0.91	0.93	0.95	0.95
(6) \widehat{WARN}_t	0.86**	0.86*	0.86*	0.89*	0.90	0.90	0.92	0.94	0.94	0.95	0.97	0.98	0.99
Relative RMSPEs of Large-Dimensional VARs:													
(7) VAR4	0.91	0.89	0.91	0.93	0.97	0.99	0.98	0.99	1.00	1.00	0.99	1.01	1.01
(8) VAR5, \hat{f}_t	0.88	0.84	0.86	0.88	0.92	0.92	0.91	0.93	0.94	0.93	0.93	0.97	0.97
(9) VAR5, \widehat{WARN}_t	0.84	0.81	0.83	0.87	0.92	0.92	0.93	0.96	0.97	0.97	0.98	1.01	1.01
Panel B: Forecast Results of and Comparisons Among the Larger-Dimensional VAR Forecasting Models													
(10) VAR4 RMSPE	24.5	28.1	31.8	36.7	40.9	44.3	47.5	49.1	51.4	53.1	54.6	56.1	57.1
Relative RMSPEs of Large-Dimensional VARs:													
(11) VAR5, \hat{f}_t	0.95	0.93	0.92	0.93	0.92	0.90	0.91	0.92	0.92	0.92	0.93	0.95	0.95
(12) VAR5, \widehat{WARN}_t	0.91*	0.90**	0.89**	0.93	0.93	0.92	0.95	0.97	0.97*	0.98	0.99	1.00	1.01

Table 4: Forecast results for changes in manufacturing employment

Note: Row (1) shows the root mean squared prediction errors (RMSPEs) of the baseline AR(2) forecasting model for each forecasting horizon, h . The units can be interpreted as the number of employees in thousands. The sample period for the forecast errors is July 2006 to December 2019. Rows (2) to (9) show the ratios of the RMSPEs from the corresponding model to the baseline model. Values less than 1 indicate lower RMSPEs than the baseline model. Rows (2) to (6) show relative RMSPEs from the bivariate VAR(2) forecasting model. This model includes the change in manufacturing employment and the variable indicated in rows (2) through (6). Row (7) shows the relative RMSPEs for the 4-dimensional VAR(2) model, which includes change in manufacturing employment, change in initial UI claims, and the employment and new orders ISM indexes. Rows (8) and (9) show the relative RMSPEs for the 5-dimensional VAR model, which adds either \hat{f}_t or \widehat{WARN}_t to the 4-dimensional VAR model. Row (10) shows the root mean squared prediction errors (RMSPEs) of the 4-dimensional VAR model. Rows (11) and (12) show the relative RMSPEs for the 5-dimensional VAR model, which adds either \hat{f}_t or \widehat{WARN}_t to the 4-dimensional VAR model. Stars, *, **, ***, indicate statistical significance at the 10, 5 and 1 percent levels. See Section 5 for details.