Weather Adjustment of Economic Output[*]

Sven Schreiber[†]

Abstract

While recurring and regular variations of weather conditions are implicitly addressed by standard seasonal adjustment procedures of economic time series, extraordinary weather outcomes are not. We propose a way of measuring aggregate abnormal weather conditions based on available local measurements and a straightforward regression-based framework to analyze their impact on German monthly total industrial and construction-sector production data, and find noticeable effects. In the historical –and seasonally adjusted– construction sector growth data the extra explanatory power of the weather regressors over a benchmark univariate autoregressive model even exceeds 50% of the variation. The estimated effects of weather deviations can be subtracted from the already seasonally adjusted data to obtain (seasonally as well as) weather adjusted series, which might capture economic developments better. Given the timely availability of the weather data compared to the publication lag of economic measurements, we also point out how measuring the weather impact may help short-term forecasting or nowcasting of industrial production in real time.
JEL codes: E32 (Business fluctuations), E27 (Forecasting production)

Key Words: weather, business cycle, nowcasting

1. Introduction

Whenever new observations of macroeconomic aggregates such as production or (un-) employment are published by statistical agencies, it is often heard that some part of the changes in the respective variables is due to some extraordinary weather effect, such as a mild winter or an unusually snowy spring. However, a precise magnitude of this effect is typically not provided. Therefore, the aim of this paper is to fill this gap by analyzing the impact of unusual weather conditions on economic output measures. In this study we restrict ourselves to the case of Germany, but our aim is to develop a tractable framework which could be universally applied in the future.

Impacts of weather phenomena on economic variables are usually associated with seasonal patterns and therefore treated as regular. Statistical agencies address this pattern by providing seasonally adjusted series. Nevertheless, one might expect that deviations of weather conditions from their seasonal average may affect economic activities and partly conceal the underlying structural dynamics. For example, Bloesch and Gourio (2015, p.2) pointed out that whether the economic slowdown in the winter 2013/2014 in the U.S. was due to harsher winter weather or instead due to an underlying economic trend would have

---

[†]Macroeconomic Policy Institute Duesseldorf (IMK, in the Hans Boeckler Foundation), and Free University Berlin; Hans-Boeckler-Str. 39, 40476 Duesseldorf, E-mail mail@sven.schreiber.name.

had implications for monetary policy. A slowdown of the U.S. economy due to weather effects rather than a negative economic trend might have implied less of a need for adjusting monetary policy.

Depending on the primary objective, controlling for abnormal weather effects and extracting the real economic trend can be accomplished in two different ways: Wright (2013) suggested to include –and Boldin and Wright (2015) then included– weather variables in the seasonal adjustment process for U.S. employment and GDP data, resulting in a weather as well as seasonally adjusted time series. They argue that abnormal weather effects may influence the seasonal adjustment procedure. Ouwehand and van Ruth (2014) provided a quite differentiated analysis for Dutch GDP data on the national and sectoral level. Estimating an ARIMA model they concluded that no significant weather effects could be identified for the majority of the sectors. A similar approach was used by the Bundesbank for German GDP data (Deutsche Bundesbank, 2014).

In a second type of approach the seasonally adjusted series is taken as given, relying on asymptotic orthogonality between the seasonal component and the unusual weather effects. Bloesch and Gourio (2015) for example found an overall weak but significant weather effect on the nonfarm employment growth rate using a fixed-effects regression model. Our approach follows this two-step approach, because it allows us much more flexibility in modelling month-specific and nonlinear weather terms. While we do not doubt the theoretical possibility that omitted weather effects might bias the first-step seasonal adjustment in finite samples, we believe that other weather modelling issues are more relevant in practice.

A previous important contribution is Hummel, Vosseler, Weber, and Weigand (2015) who analyzed the effect of several weather variables like temperature, snowfall, or snow height on German national-level employment, based on 310 representative weather stations. They identified several weather and catch-up effects in the following months. For instance, a one degree temperature increase in January raises employment by 14,000 persons on average between 2006 and 2014. Also for Germany Döhrn and an de Meulen (2015) showed that including weather variables in a business-cycle oriented forecasting procedure improves the model, but not in a significant way in their setup.

There are also attempts to identify longer-run weather (or climate) effects on economic outcomes, see Dell, Jones, and Olken (2014), but in this paper we focus on the shorter-run dynamics of occurrences of abnormal weather. The longer-run impact of climatic trends on economic activity raises difficult questions about endogenous adaptation and restructuring of production, as well as the adequacy of national accounts measurements that (almost) do not take into account environmental damages.

We confirm and follow Hummel, Vosseler, Weber, and Weigand (2015) in their choice of relevant weather measurements, namely air temperature, snowfall, and snow height. We focus on economic output instead of labor inputs, however. (In this study we focus on monthly industrial production, but in the future a similar analysis for quarterly GDP would also be interesting.) We also provide a separate analysis for the construction sector since any weather effects will be felt there most. A further difference is that we stick to a straightforward model framework that is linear in the parameters, instead including non-linearities through polynomial terms and by modelling heterogeneous month-specific effects. Finally, we discuss the use of the weather observations for forecasting purposes in a (pseudo) real-time setting, when the current production data as well as their immediate lags would not have been published yet (often called "nowcasting").

2. Data and empirical approach

The dependent variables that we analyze are the monthly growth rate of German real total industrial production (IP) and the production in the construction sector, shown in figure 1. Total output represents an important cyclical indicator, while production in the construction sector is the part of economic activity which is most likely to depend on weather conditions. An overview about the different production indices and their hierarchial structure is given in Statistisches Bundesamt (2015), data are taken from the Bundesbank website, and both indices are calendar and seasonally adjusted.

It can be seen that the production growth series with this monthly frequency are quite noisy, but the great recession at the outbreak of the financial crisis is clearly visible especially in total production which includes export sectors. In the estimated equations we remove these effects by a small number of impulse dummies added to the regressions. The salient feature of the construction sector growth distribution is its heavy tails, with a considerable number of observations that exceed $\pm 10\%$ monthly growth, leading to an empirical excess kurtosis of 6.7.

Weather data for Germany have recently begun to be provided on the internet and are freely available.[1] The construction of the weather dataset was initially inspired by the approach of Hummel, Vosseler, Weber, and Weigand (2015), that is we aggregated the weather data of the available 251 weather stations to the state levels of the sixteen German federal states (including the three city-states Berlin, Hamburg, and Bremen), then weighing them by the state-level number of employees to obtain aggregated data at the national level. See below for further details on the aggregation method. The locations of the weather stations are displayed in figure 2, and the time series sample used in this paper is January 1991 through November 2015. We consider three measurable weather aspects, namely air temperature, snow height and snow fall per week in centimeters, all time-averaged from daily to monthly series. Other weather variables would also be possible in principle; the Deutsche Bundesbank (2014) for example used the sum of ice-days in a specific time interval (quarter or month), but that information does not differ much from the combined content of snow fall and (cumulated) height.

Given that the weather data are published already about a week after the end of the month – in contrast to the production data that suffer from a publication delay of at least one month – this allows one to predict or "nowcast" the weather effect on a real-time basis. A potential disadvantage, however, is that the most recent data are mostly not yet checked for measurement errors.

What we have in mind as a first approximation is a simple additive framework that distinguishes between different components that together yield the observed realization of the economic variable of interest:

$$y_t = struc_t + weatherdev_t + \varepsilon_t, \tag{1}$$

where y_t will be a seasonally adjusted growth rate of the underlying economic variable, and $struc_t$ is interpreted as a component which is structural in the sense that it indicates the underlying tendency attributable to purely economic forces and intrinsic dynamics. In contrast, $weatherdev_t$ is an irregular component which measures influences that stem from

---

[1]Original data series provided by Deutscher Wetterdienst and freely available at `http://www.dwd.de/`.
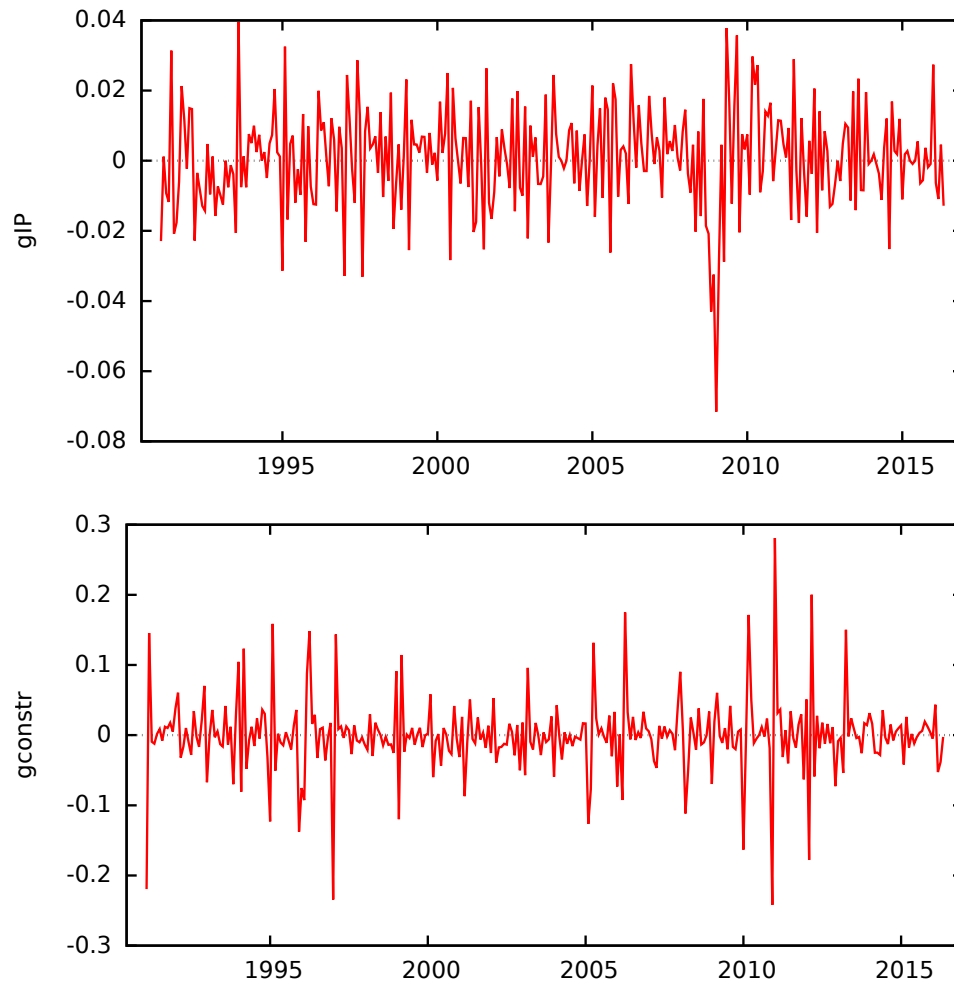
Figure 1: Total industrial (upper panel, gip) and construction sector (lower panel, gconstr) production growth. Seasonally adjusted at the official source and transformed to month-on-month growth rates (log differences).
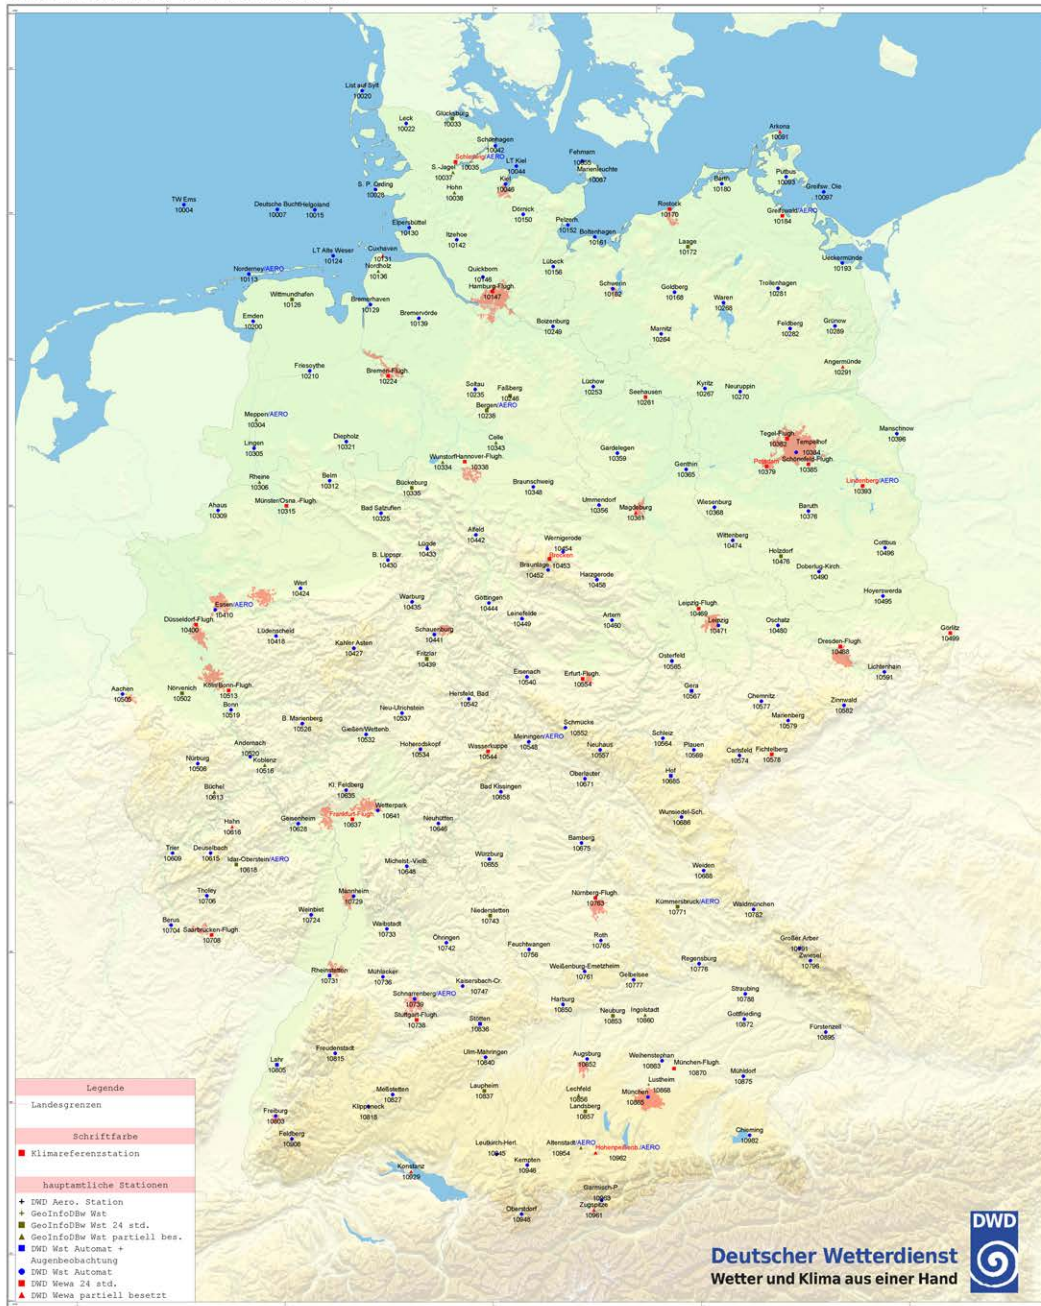
Figure 2: Weather measurement stations in Germany (map provided by the German Weather Agency DWD, www.dwd.de).

weather realizations beyond the systematic and regular seasonal cycles. We allow these components to be dynamic, such that they will include lags as well. Finally, ε_t is a purely random error component which should be (close to) white noise. As a consequence, $y_t - \widehat{weatherdev}_t$ will be a weather (and seasonally) adjusted series.

We proceed by defining the extent of "abnormal weather" as the absolute deviations of the observation $X_{j,t}$ from a month-specific ($m = 1...12$, corresponding to January...December) time average for an individual weather station j:

$$x_{t,j} = X_{j,t} - \bar{X}_{j,m(t)}, \tag{2}$$

where $X \in \{temperature, snowfall, snowheight\}$.[2] The next step is to aggregate the time deviations across stations to the corresponding federal state level by a simple average, $x_{t,s} = \bar{x}_{t,j \in s}$, followed by a weighted average (by state employment numbers $e_{t,s}$, with $e_t = \sum_{s=1}^{16} e_{t,s}$) to the national level:

$$x_t = e_t^{-1} \sum_{s=1}^{16} e_{t,s} x_{t,s}, \tag{3}$$

This construction of aggregate weather time series might be called "deviate-then-aggregate". The advantage is that regional station-specific seasonal patterns are captured.[3] Figure 3 shows the resulting time series. The lower panel of that figure displays the snow-related variables, and while the two series are highly correlated –for example because in the summers snow is mostly absent and thus the deviation series are both zero– there are some marked differences in the spikes which could be especially important when considering non-linear effects.

From now on, the deviation of a weather variable and the name of a weather variable are used synonymously. For example, the deviation of temperature (from its month-specific average) and temperature are used synonymously, and the absolute level of a weather variable never enters any estimated model.

Our econometric framework is a straightforward dynamic regression. The benchmark specification is an AR(6) model of the growth rate of the respective production index, augmented with two impulse dummies (2008M11, 2009M1) that capture obvious outliers in the great recession episode:

$$y_t = c + \sum_{i=1}^{6} a_i y_{t-i} + \delta_1 d_{2008M11,t} + \delta_2 d_{2009M1,t} + u_t \tag{4}$$

In the full specification the following weather terms are included, where D_t^m is a dummy variable for month m with mean zero (centered):

- Month-specific direct weather regressors:

$$\sum_{x \in \{temp, sfall, sheight\}} \sum_{m=1}^{12} (b_{x,1}^m x_t D_t^m) \tag{5}$$

---

[2]We have also experimented with relative deviations (where possible), with inferior results. Taking into account also a potential time trend in the regular weather series is left for future research.

[3]In Haustein and Schreiber (2016) the inverse ordering was used, "aggregate-then-deviate". The revision was partly inspired by comments from Matthias Hertweck.
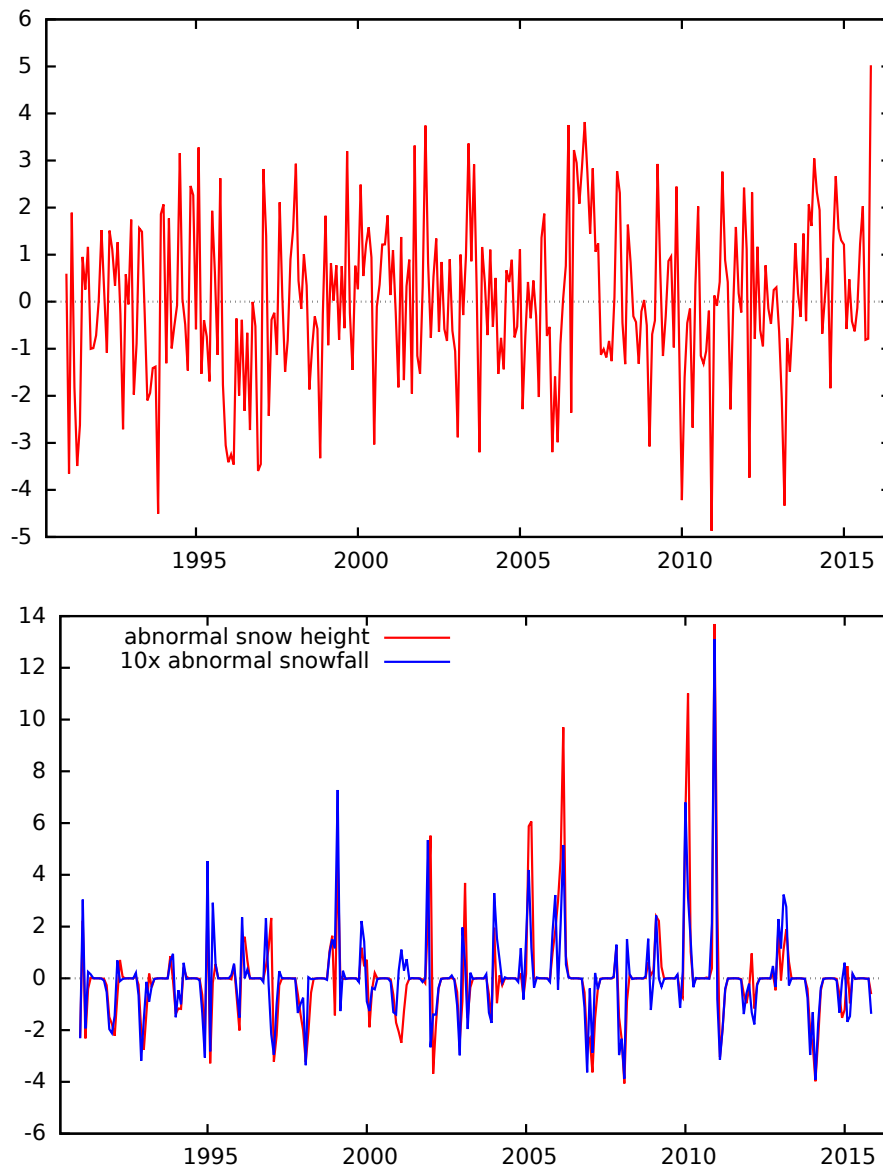
Figure 3: Observed weather deviations (aggregates of station-specific deviations). Upper panel air temperatures in centigrades, lower panel snow fall (times 10) and snow height in centimeters.

- Squared month-specific weather terms, where $x_t^2 \, sgn(x_t)$ represents a signed quadratic function which unlike a pure parabola is negative for $x_t < 0$:[4]

$$\sum_{x \in \{temp, sfall, sheight\}} \sum_{m=1}^{12} \left(b_{x,2}^m \, x_t^2 \, sgn(x_t) \, D_t^m \right) \tag{6}$$

- Additional auxiliary dummy regressors that merely serve to balance the month-specific non-zero means of the squared terms:

$$\sum_{m=1}^{12} d_m D_t^m \tag{7}$$

- Lagged weather terms (not month-specific to limit the total number of regressors):

$$\sum_{x \in \{temp, sfall, sheight\}} \sum_{k=1}^{6} b_{x,k} \, x_{t-k} \tag{8}$$

These details are accurate for the estimation of the historical adjustments; in the case of forecasting or nowcasting the direct (and month-specific) terms enter as lags, and the first lag(s) of the output growth are excluded, see below for further explanations.

The lags of each weather variable were also included to control for possible catching-up effects in the following months. However, we impose homogeneity across months for the lagged effects because the number of parameters would otherwise explode relative to the available observations. By catching-up effects we mean a shift of production in point of time; for example orders and contracts which could not be carried out in February and March due to a harsh winter might be completed one or two months later.

A further concern in time series analysis might be the existence of some structural break. During their analyses, Hummel, Vosseler, Weber, and Weigand (2015) found some evidence for a structural break in 2006, which prompted them to use a smooth transition regression model. The advantage of that model is that weather effects can be flexibly modeled over time. However, the authors assign the structural break in 2006 mainly to the introduction of seasonal short time work benefits (*Saison-Kurzarbeitergeld*), given that they focus on labor market variables. This is not relevant for our focus on industrial production. Our modelling of month-specific effects also limits the possibility of splitting the sample, as we effectively require up to 12 times more observations. In any case, apart from the mentioned inclusion of impulse dummies for the great recession there was no indication of structural changes in our specifications.

The overall sample size in this monthly data is $T \approx 290$, and we employ a simple general-to-specific search to obtain a sparse model. More than 80% of the roughly hundred regressors are typically removed as insignificant by the procedure. A more sophisticated method could also be employed, but we expect that variations there would only have marginal effects.

---

[4]We have also experimented with a threshold model as an alternative nonlinear specification, but with disappointing results.

3. Estimating weather influences

We can now report the estimated $\widehat{weatherdev}_t$ component by adding together all terms from the estimated regressions containing a weather-related variable, using the estimated parameters in place of the unknown truth. The result is shown in figure 4. It is clear that the observed output growth can most clearly be associated with weather developments in the case of the construction sector (lower panel), where many of the extraordinarily large realizations are explained quite well; the resulting $\bar{R}^2$ of that regression is 72%, compared to a mere 15% in the benchmark specification (4). As expected, for total industrial production (upper panel) the explanatory power of abnormal weather is more modest with an $\bar{R}^2$ of 36%, rising from 19% in the benchmark AR. Most movements in total production are not attributable to weather but to other types of shocks.

Notice that some nonlinearity indeed remains in the sparse model, for example in the case of the total IP equation we retain the sign-squared temperature deviations in May, yielding a composite abnormal temperature impact of

$$\left[0.020\,temp_t - 0.013\,temp_t^2\,sgn(temp_t)\right] \times D_t^{May},$$

or the sign-squared snow height in January:

$$\left[-0.0087\,sheight_t + 0.0018\,sheight_t^2\,sgn(sheight_t)\right] \times D_t^{Jan}.$$

Furthermore, the weather coefficients are estimated with a certain amount of sampling uncertainty, and in Figure 5 we take the associated standard error of $\widehat{weatherdev}_t$ into account, using the estimated covariance matrix of the weather-related coefficients.[5] This yields interval estimates of the weather (as well as seasonally) adjusted series at a nominal (asymptotically justified) 95% level of confidence. The intervals in that figure are depicted as shaded grey areas, while for easier visibility we display the un-adjusted original observations as red circles instead of lines.

One way to read this figure is that whenever the red circle is *not* touched by the grey area the weather adjusted datapoint is significantly different from the original observation. This is quite often the case, of course especially for the construction sector (lower panel). If it happens with larger observations (in absolute value), the weather adjustment is "inwards", closer to zero structural growth, which we observe in 87 of the roughly 290 observations. But notice that also a significant "outward" adjustment is quite frequent, where the confidence interval of the weather adjusted value is farther away from zero than the original observation. In the construction sector this happens for 93 observations. Furthermore in about 10% of all cases the weather adjustment in the construction sector flips the sign of the observation significantly. (These cases may be overlapping with the significant inward or outward adjustments.)

All in all, it appears that the impact of abnormal weather conditions affects the majority of observations in the construction sector in a statistically significant way. Even though the situation is less extreme for total industrial output, in order to assess the economic situation it seems that weather adjustments should play a more prominent role than is currently the practice among macroeconomists.

---

[5]Let ξ_t be the consolidated vector of all weather regressors in the full regression described in (5) through (8), and β the corresponding stacked coefficient vector. Then we have $\widehat{weatherdev}_t = \hat{\beta}'\xi_t$, and its variance is given as $\xi_t'\widehat{Cov}(\hat{\beta})\xi_t$.
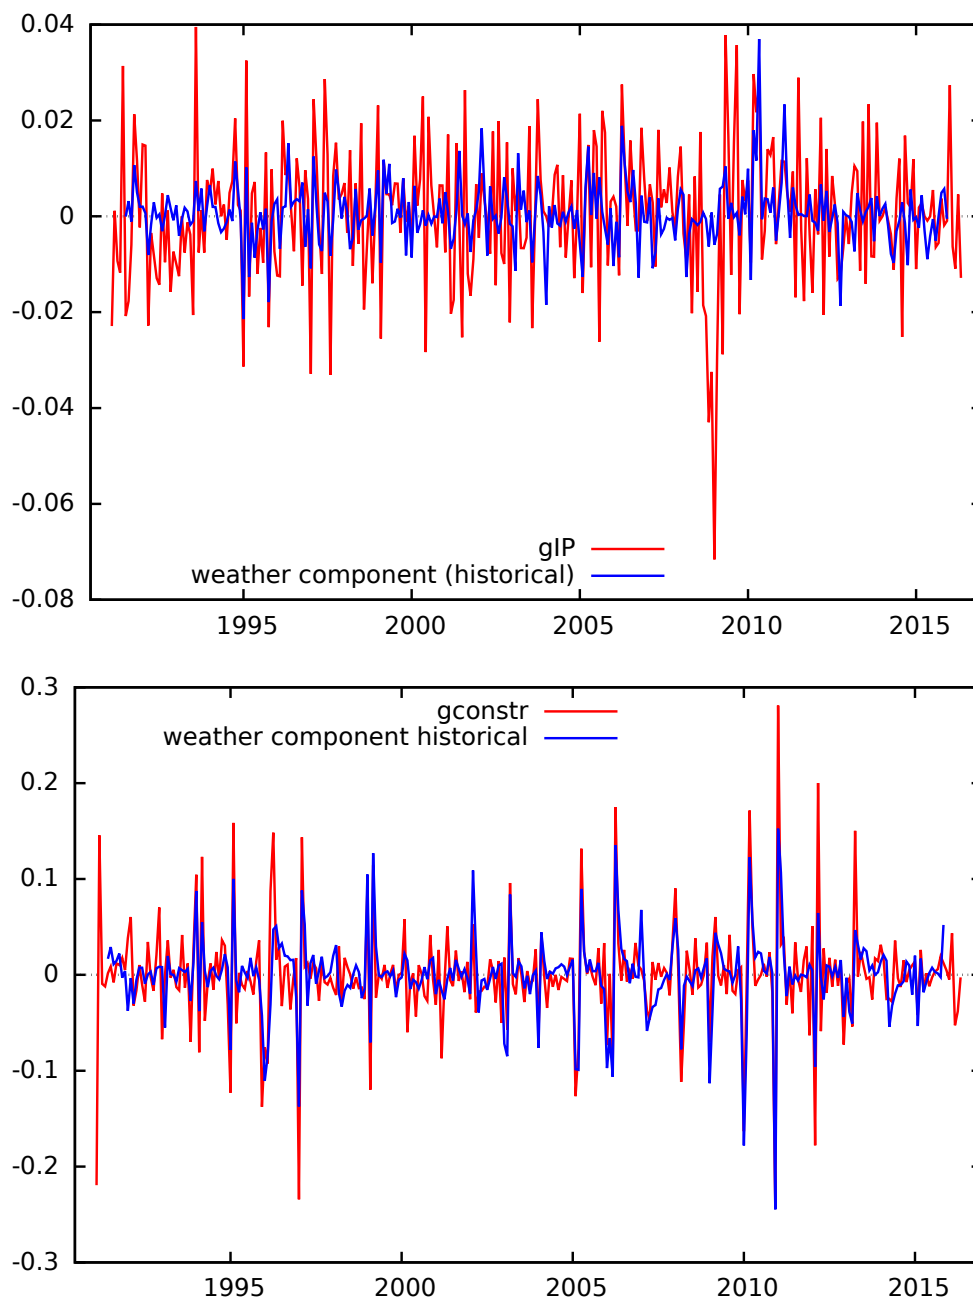
Figure 4: Observed German industrial production (same as in figure 1) and estimated historical weather components. Upper panel total production growth (gIP), lower panel construction sector growth (gconstr).
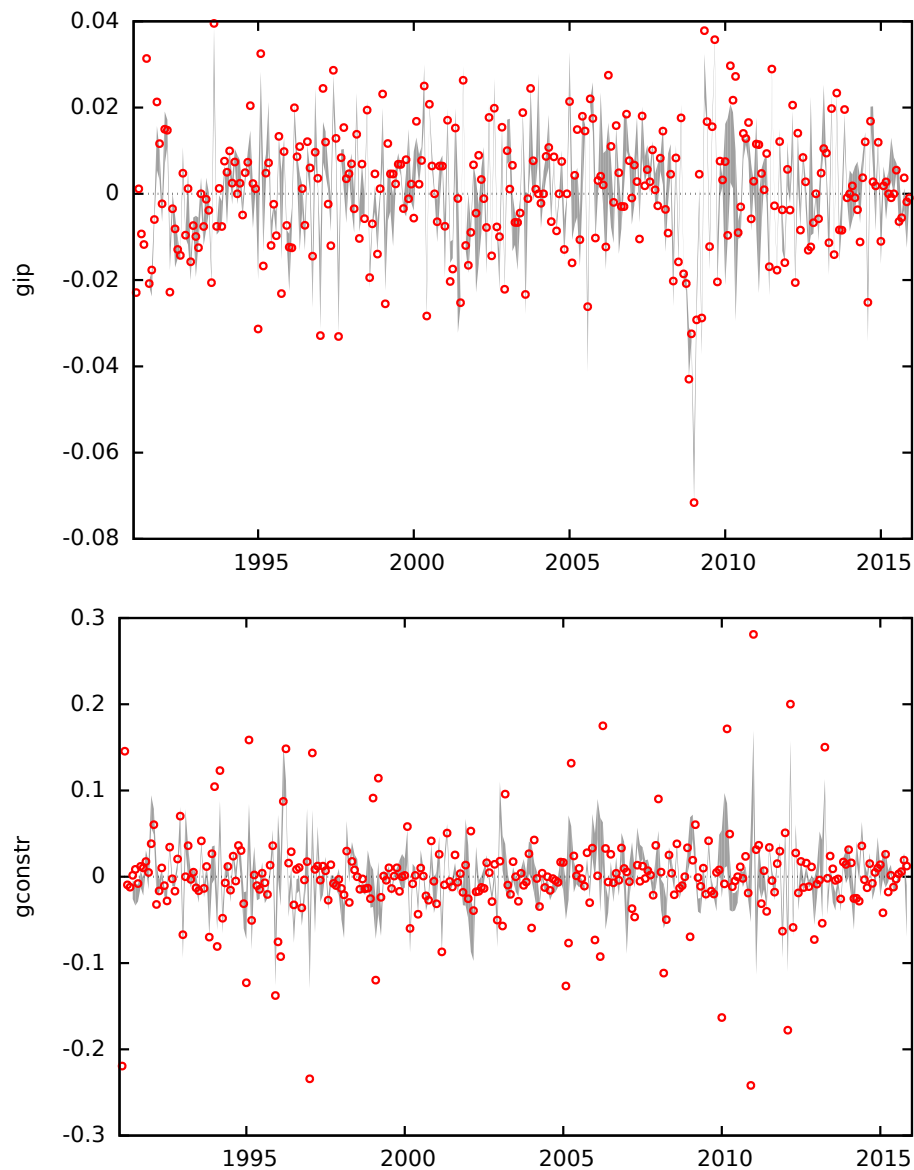
Figure 5: Estimation uncertainty of weather-adjusted industrial production (pointwise 95% confidence intervals for the adjusted series $y_t - \widehat{weatherdev}_t$; $\bigcirc$: non-adjusted observations, i.e. the same series as in figure 1). Upper panel total production growth (gIP), lower panel construction sector growth (gconstr).

4. Improving real-time analysis of output developments

In the previous section we performed a historical adjustment of German production time series by estimating the dynamic influences of irregular or abnormal weather conditions. In this section we want to exploit the fact that observations of the weather measures are available much more quickly than the first publications of production values by statistical agencies – in Germany the publication delay for the first and tentative official figures on industrial production is around 38 days, more than one month. Given that we found some significant contemporaneous impact of the weather (deviations), it is natural to take these effects into account when the aim is to produce a short-term forecast of economic activity.

However, even the weather data are not available contemporaneously in real time, but take about a week after the end of the month to appear on the agency website. Therefore we distinguish between the following two scenarios.

First we work with the information set of the middle of any given month, where the aim is to produce a forecast for the current month. This timing could be called "nowcast", because the current month is affected, or it could also be called "semi-forecast", because the second half of the current month still lies in the future. In the middle of the current month t we have available the weather data describing the previous month $t - 1$ (and earlier), and the industrial production data relating to $t - 2$ (and earlier). This means that the weather data in this scenario may only enter with a lag, and the first lagged endogenous term must be removed from the regressors in order to replicate the real-time information set.

Secondly we also consider a scenario of the beginning of the next month $t + 1$, meaning that the weather data for period t are already available, but the output data for $t - 1$ have not yet been released, because that typically takes until the middle of the month. The target quantity is still output in period t, so this is called an (early) "backcast". Again no first lagged endogenous term must be included in the equation, but now the period-t weather terms are allowed as regressors again.

However, our present aim is merely to check whether this is a promising route for future research, and hence we employ some shortcuts. First we do not work with a full real-time dataset but instead continue to use our dataset on industrial and construction-sector production which effectively contains only a single vintage (from 2015). We therefore do not take into account the data revisions occurring after the respective first publications. Secondly, in order to gauge the value added of the weather effects for forecasting, in principle one would have to use a full-fledged forecasting model as the starting point. Instead we will here restrict the non-weather predictive variables to be the lagged endogenous variables available on a pseudo real-time basis. Both the real-time aspect and the use of other predictive variables can be found in Proaño and Theobald (2014) or in Schreiber and Soldatenkova (2016), but only for total industrial production.

As indicated above, the present pseudo real-time exercise for nowcasting and backcasting thus effectively boils down to removing from the predictive regressions a lag of the dependent variable, and for the nowcasting exercise also a lag of the weather terms. The empirical strategy is unchanged otherwise with respect to the historical analysis in section 3.

In table 1 we report the simple $\bar{R}^2$ values (fit adjusted for number of retained regressors) that are attained in the various scenarios, each compared with the respective AR benchmark. (The numbers in the historical adjustment columns were already mentioned in the text above and are repeated for convenience and comparison.) Moving from the historical full-information scenario to the backcasting case without the first endogenous lag, the extra

Table 1: Explanatory and nowcasting/backcasting power

($\bar{R}^2$ in %)	nowcasting / middle-of-month info set	backcasting / early-next-month info set	historical adjustment
total industry output growth			
benchmark AR	15	15	19
with (abnormal) weather	27	32	36
construction sector output growth			
benchmark AR	0	0	15
with (abnormal) weather	49	67	72

Notes: The historical adjustment column corresponds to results in section 3. The benchmark equations are univariate AR(6) models, where the first lag is removed in the nowcasting and backcasting cases.

predictive power of the abnormal weather terms –measured as the simple difference between the $\bar{R}^2$ values– cannot decrease by definition of the regressor set changes. However, even removing the contemporaneous weather terms in the nowcasting scenario reduces the extra predictive power by less than one third in the total industrial production equation (of the percentage point difference: $36 - 19 = 17$ vs. $27 - 15 = 12$), and less than one sixth for the construction sector ($72 - 15 = 57$ vs. $49 - 0 = 49$).

For the construction sector it seems very unlikely that enlarging the information set with economic indicators like new orders received or financial variables would change that picture dramatically. On the other hand, for total industrial production we mainly expect some robust (albeit perhaps modest) gains in the backcasting case when the contemporanous weather information can be used, not so much for the nowcasting scenario. This is clearly a natural next step for future research.

5. Conclusions

We conclude that abnormal weather conditions in Germany affect the construction sector and aggregate production. Generally and not surprisingly, the impact as well as the estimation precision are larger for the construction sector than for total industrial production. Controlling for measurable weather effects using freely available datasets thus helps to determine the underlying economic dynamics and should lead to a more accurate assessment of the business cycle, ultimately also implying more appropriate stabilization policy advice.

By relying on the (approximate) orthogonality between regular seasonal effects and irregular random weather outcomes we were able to keep the econometric methods simple, using straightforward regression models that are linear in parameters while being non-linear in some of the variables. Within this framework we found it important to allow the effects of the weather variables such as air temperature or snow height (in deviations from seasonal averages) to be month-specific. The specification also had to account for serially correlated production and dynamic reactions to past weather incidents.

We provided initial evidence that the weather effects could also be used to improve the "nowcasting" or "backcasting" of monthly output growth realizations that are still unknown

because of the publication delay of such macroeconomic data. This latter part of the analysis used the simplifying shortcut of a pseudo real-time setup, whereas a refined fully real-time estimation would work with a different data vintage for each datapoint. Also, the marginal value added of the weather variables as predictors would have to be assessed relative to a broader information set. Nevertheless, this first set of results was encouraging especially in the case of the construction sector. It remains to be seen to what extent the measured influences hold up at a lower measurement freuqency, most importantly in quarterly GDP data. This is also left for future research.

Finally, we expect that such effects of abnormal weather apply to most other economies as well, not only to Germany, and our framework is intended to be easily adaptable given that it only requires three widely available weather measurements. In the currently standard approach of conducting structural macroeconomic analysis seasonally adjusted data is used, which means that weather variations are implicitly seen as an uninteresting nuisance for economic trends. If taken seriously, this position would imply that structural macroeconomic analysis needs series that are also adjusted for other, non-seasonal and exogenous, weather variations as presented in this paper.

References

BLOESCH, J., AND F. GOURIO (2015): "The Effect of Winter Weather on U.S. Economic Activity," *Journal of Economic Perspectives*, 39(1), 1–20.

BOLDIN, M., AND J. H. WRIGHT (2015): "Weather-Adjusting Economic Data," *Brookings Papers on Economic Activity*, (Fall), 227–260.

DELL, M., B. F. JONES, AND B. A. OLKEN (2014): "What Do We Learn from the Weather? The New Climate-Economy Literature," *Journal of Economic Literature*, 52(3), 740–798.

DEUTSCHE BUNDESBANK (2014): "Wettereffekte auf das Bruttoinlandsprodukt im Winterhalbjahr 2013/2014," *Monatsbericht*, 66(5), 58–59.

DÖHRN, R., AND P. AN DE MEULEN (2015): "Weather, the forgotten factor in business cycle analyses," Ruhr Economic Papers No. 539 05, Rheinisch-Westfälisches Institut für Wirtschaftsforschung (RWI).

HAUSTEIN, E., AND S. SCHREIBER (2016): "Adjusting Production Indices for Varying Weather Effects," Working Paper 171, IMK.

HUMMEL, M., A. VOSSELER, E. WEBER, AND R. WEIGAND (2015): "Wie das Wetter den Arbeitsmarkt beeinflusst," *IAB-Kurzbericht*.

OUWEHAND, P., AND F. VAN RUTH (2014): "How Unusual Weather Influences GDP," Working paper, CBS.

PROAÑO, C., AND T. THEOBALD (2014): "Predicting Recessions with a Composite Real-Time Dynamic Probit Model," *International Journal of Forecasting*, 30(4), 898–917.

SCHREIBER, S., AND N. SOLDATENKOVA (2016): "Anticipating business-cycle turning points in real time using density forecasts from a VAR," *Journal of Macroeconomics*, 47, 166–187.

STATISTISCHES BUNDESAMT (2015): "Indizes der Produktion und der Arbeitsproduktivität im Produzierenden Gewerbe, Fachserie 4 Reihe 2.1," .

WRIGHT, J. H. (2013): "Unseasonal Seasonals?," *Brookings Papers on Economic Activity*, pp. 65–110.