

Schreiber, Sven

Conference Paper

Weather-induced Short-term Fluctuations of Economic Output

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2018: Digitale Wirtschaft - Session: Environmental Economics IV, No. D06-V3

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Schreiber, Sven (2018) : Weather-induced Short-term Fluctuations of Economic Output, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2018: Digitale Wirtschaft - Session: Environmental Economics IV, No. D06-V3, ZBW - Leibniz-Informationszentrum Wirtschaft, Kiel, Hamburg

This Version is available at:

<https://hdl.handle.net/10419/181622>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Weather-induced Short-term Fluctuations of Economic Output*

Sven Schreiber[†]

this version: February 2018

Abstract

We contribute to the recent literature on the economic effects of those weather conditions that deviate from their regular seasonal pattern. To this end we use local temperature and snow measurements across Germany to analyze their impact on German monthly total industrial and construction-sector production. We find noticeable effects of the various (linear and nonlinear, contemporaneous and dynamic) weather regressors, which in the –seasonally adjusted– construction sector growth data imply an extra explanatory power of more than 50% of the variation, compared to benchmark predictive regressions. As expected, the impact is quite a bit less in total industrial production. From our estimates we obtain (seasonally as well as) weather adjusted production series, and our regression-based approach also yields confidence intervals for these adjustments. The estimated adjustments are quantitatively relevant also for broad output (quarterly GDP). In a mixed-frequency framework we find some value of the estimated monthly weather impact for quarterly GDP nowcasts in (quasi) real time.

Keywords: weather, business cycle, nowcasting, MIDAS

JEL codes: E32 (Business fluctuations), E27 (Forecasting production)

*I thank Erik Haustein for building the database with the underlying German weather series at the station-based micro level. An interim version of this paper appeared in the conference proceedings of the Joint Statistical Meetings Chicago 2016. I wish to thank the session participants for helpful comments.

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

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. The aim of this paper is to provide a quantitative assessment of the impact of unusual weather conditions on economic output. 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 wherever systematic weather measurement is recorded.

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 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 non-farm employment growth rate using a (fixed-effects) regression model. Our approach follows this two-step approach, because it allows us much more flexibility in modeling month-specific and nonlinear weather terms. Furthermore, because the second step is performed with a straightforward OLS regression which includes nonlinear terms but is linear in the parameters, in contrast to other methods suggested in the literature we can assess the parameter uncertainty and thus report confidence intervals for the estimated weather adjustment. 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 these and other weather modeling issues are more relevant in practice.

A short but inspiring 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.

In contrast to Hummel, Vosseler, Weber, and Weigand (2015) we focus on economic output instead of labor inputs, however. We also provide a separate analysis for the construction sector since any weather effects will be felt there most. We include non-linearity through polynomial terms and by modeling heterogeneous month-specific effects, see the following section 2. Empirical results for monthly production indices as well as some implications for quarterly GDP figures are reported in section 3. 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”. To this end section 4 also covers the nowcasting of quarterly GDP with the help of a mixed-frequency (MIDAS) approach, before offering concluding comments in section 5.

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. Later in section 3.2 we analyze the implications of these adjustments in the industrial sector for aggregate output (GDP). Total industrial 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 hierarchical 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 data set 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.

¹Original 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).

Messnetzkarte DWD - Hauptamtliches Stationsnetz und GeoInfoDBw
Karte vom: 02.05.2016, 10:24 mit 211 Stationen

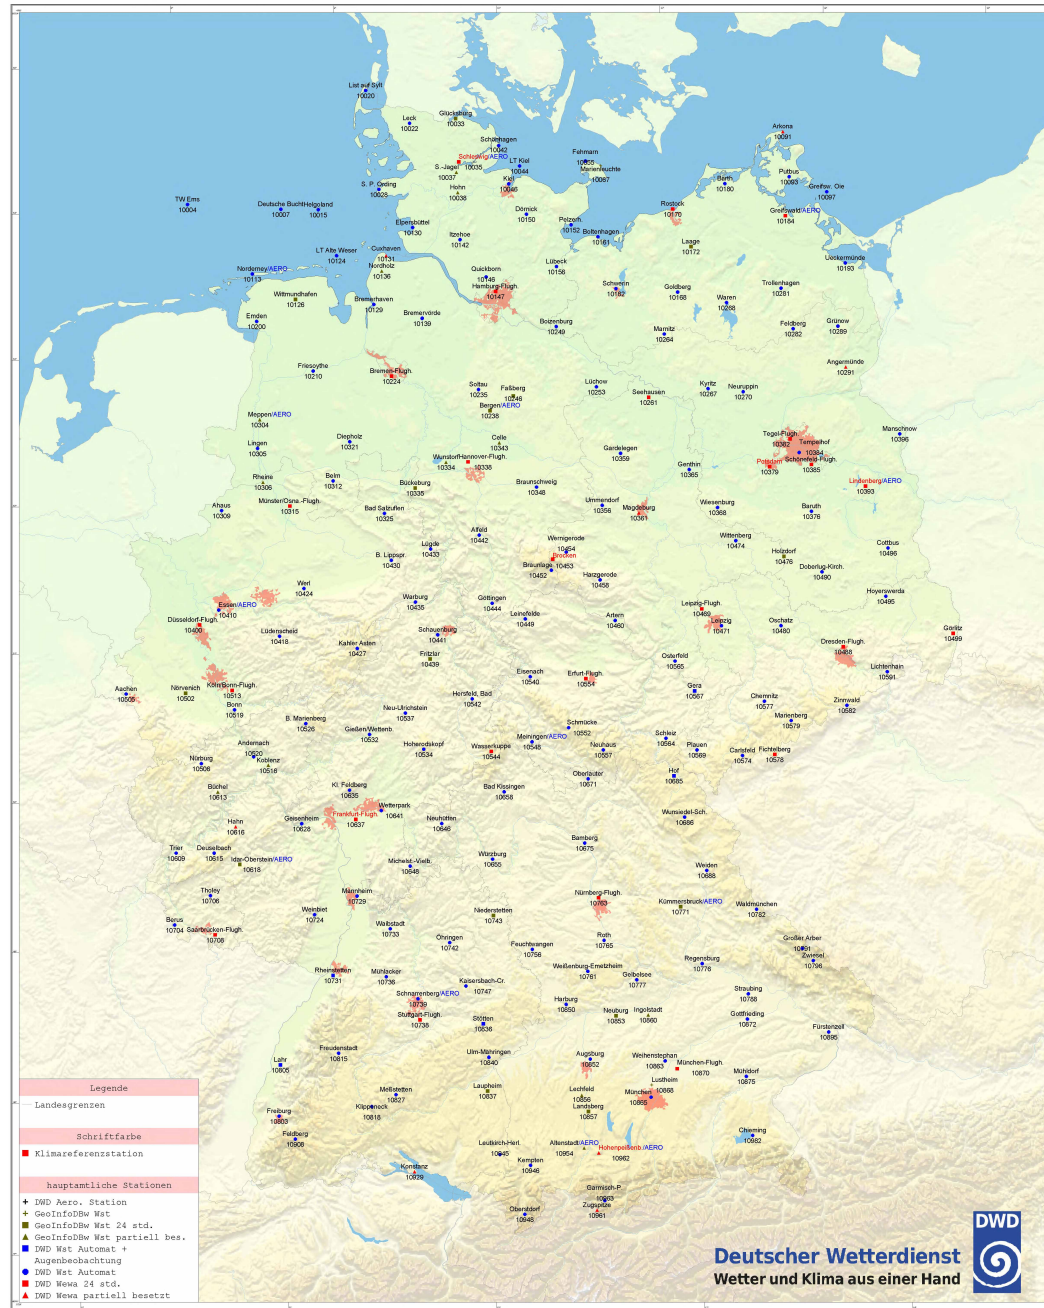


Figure 2: Weather measurement stations in Germany (map provided by the German Weather Agency DWD, www.dwd.de).

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, see section 4 below for further discussion. A potential disadvantage is that the most recent data are mostly not yet checked for measurement errors.

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

$$y_t = struc_t + weatherdev_t + \varepsilon_t, \quad (1)$$

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 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 ($m = 1 \dots 12$, corresponding to January...December) time average for an individual weather station j :

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

where $X \in \{temperature, snowfall, snowheight\}$.² The next step is to aggregate

²We 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.

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:

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

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 a model where the growth rate of the respective production index depends on own lags plus lagged indices of incoming orders o_t and lags of the term spread of interest rates s_t , augmented with two impulse dummies (2008M11, 2009M1) that capture obvious outliers in the great recession episode:

$$y_t = c + \sum_{i=1}^6 a_{y,i} y_{t-i} + \sum_{j=1}^3 (a_{o,j} o_{t-j} + a_{s,j} s_{t-j}) + \delta_1 d_{2008M11,t} + \delta_2 d_{2009M1,t} + u_t \quad (4)$$

In the full specification the following weather terms are included, where D_t^m is

³Before –in Haustein and Schreiber (2016)– we had used the reversed approach “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.

a dummy variable for month m with mean zero (centered):

- Month-specific direct weather regressors:

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

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

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

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

$$\sum_{m=1}^{12} d_m D_t^m \quad (7)$$

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

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

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 oth-

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

erwise 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 modeling 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.⁵

3 Estimating weather influences

3.1 Monthly industrial production

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

⁵We have also cross-checked some results with a more sophisticated best-subset method, but without any noticeable changes.

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 adjusted \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 48%, rising from 33% in the benchmark specification without weather terms. Most movements in total production are not attributable to weather but to other types of shocks; nevertheless the explained variation is noticeable.

Notice that some non-linearity 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

$$[0.017temp_t - 0.012temp_t^2 \text{sgn}(temp_t)] \times D_t^{May},$$

or the sign-squared snow height in January:

$$[-0.0098sheight_t + 0.0019sheight_t^2 \text{sgn}(sheight_t)] \times D_t^{Jan}.$$

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 co-variance matrix of the weather-related coefficients.⁶ This yields interval estimates of the weather (as well as seasonally) adjusted series at a nominal 95% level of confidence. The intervals in that figure are depicted as shaded gray areas, while for easier visibility we display the

⁶Let 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 $\xi_t' \widehat{Cov}(\hat{\beta}) \xi_t$.

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 gray area the weather adjusted data-point 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 about 90 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 macro-economists.

3.2 Implied adjustment of quarterly GDP

Our estimates are based on industrial production data. First the IP series is available at a higher (monthly) frequency, secondly it is well known that its short-run volatility is considerably higher than that of GDP, and finally a higher dependency on weather conditions is expected at least for some industrial sectors. However, for macroeconomic analysis such as in monetary policy we are usually interested in aggregate output, and therefore we have also applied the estimated weather adjustments to the GDP series. The monthly IP growth rates are cumulated to the

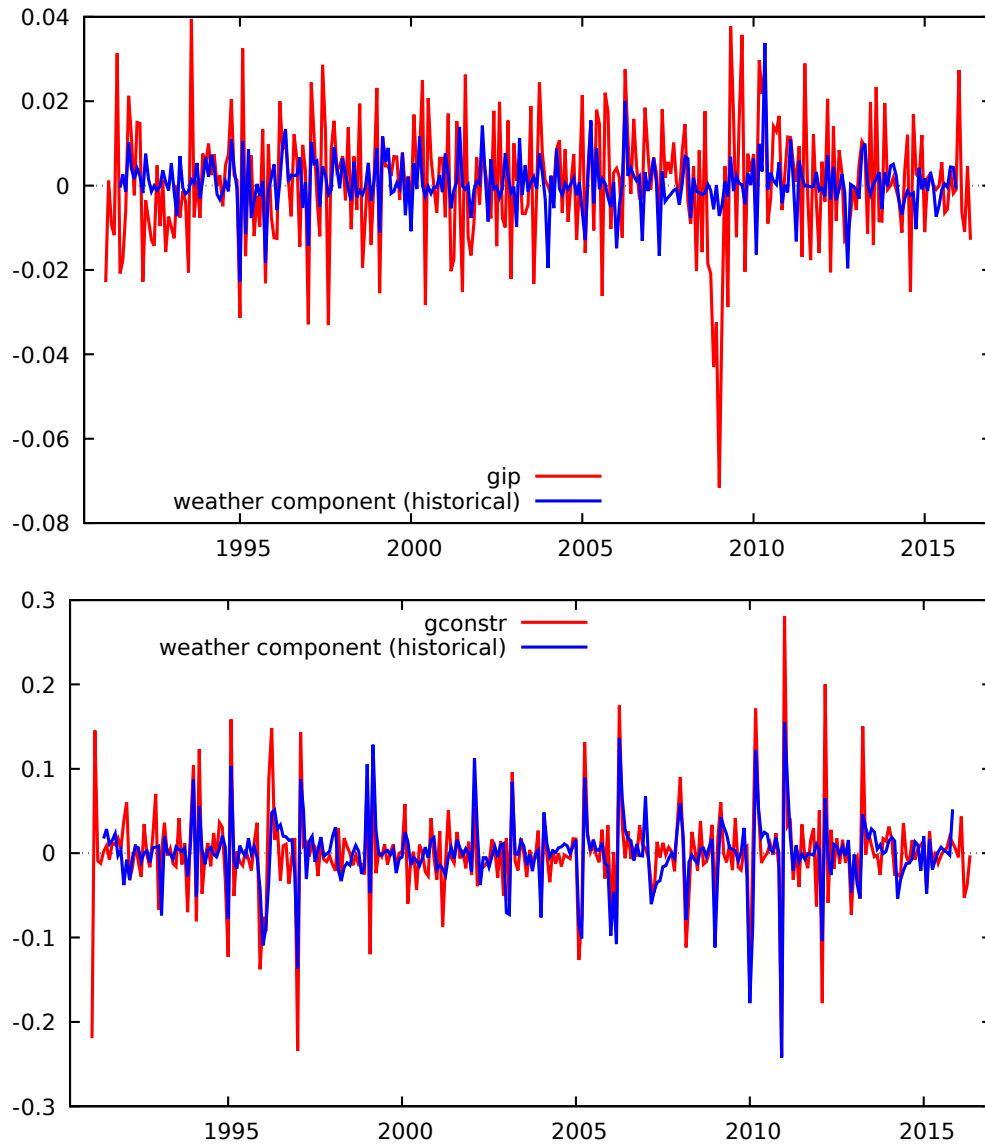


Figure 4: Observed German industrial production (same as in figure 1) and estimated historical weather components. Upper panel total production growth (gIP), lower panel construction sector growth ($gconstr$).

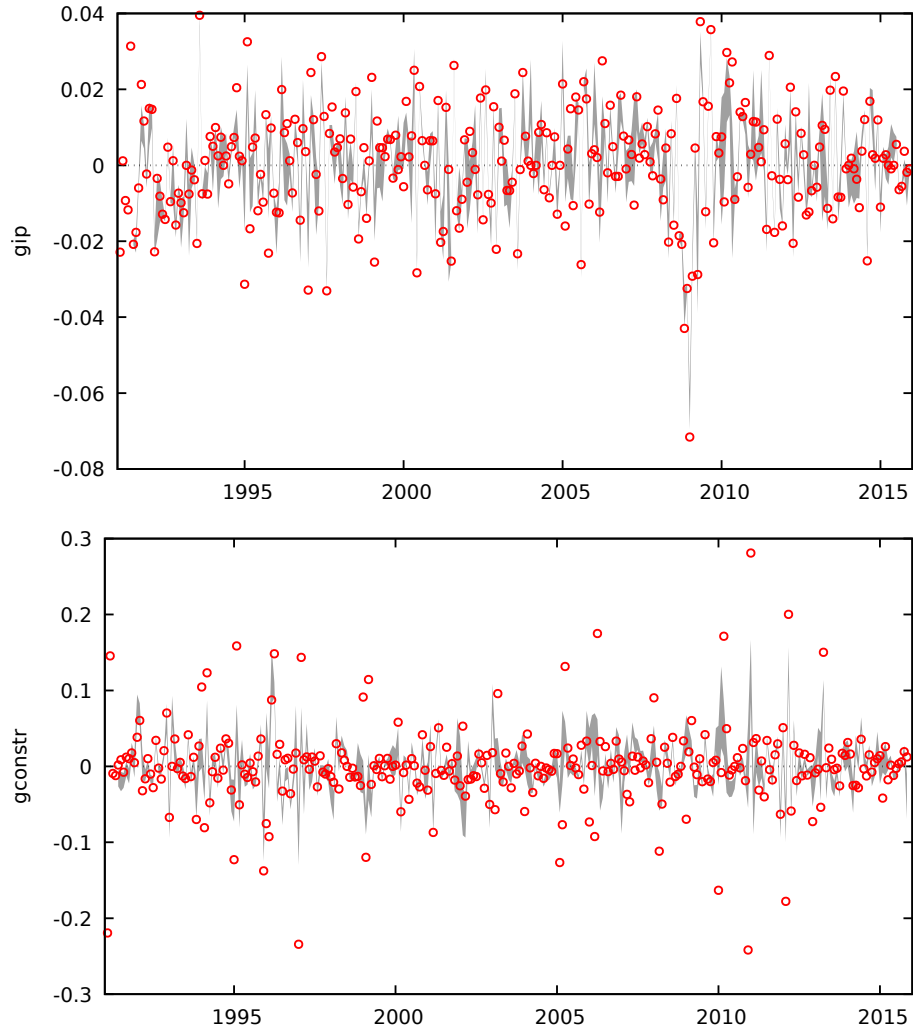


Figure 5: Estimation uncertainty of weather-adjusted industrial production (point-wise 95% confidence intervals for the adjusted series $y_t - \widehat{weatherdev}_t$; \circ : non-adjusted observations, i.e. the same series as in figure 1). Upper panel total production growth (gIP), lower panel construction sector growth ($gconstr$).

quarterly frequency, then the growth differential between un-adjusted and adjusted IP growth is re-scaled with the (annually varying) share of industry output in aggregate output. This procedure implicitly assumes that any abnormal weather effects in the non-industrial service and public sectors are negligible, yielding the results in Figure 6.⁷ We find that even though the share of industry in aggregate output is much less than 50% and declining in Germany, the fluctuations are important enough to matter at the economy-wide level at least in some periods. The largest implied adjustment occurs in 2010Q2 with a whopping figure of almost -4% (at annualized rate, percentage points). However, we should bear in mind that the average effect over all four quarters of 2010 is much smaller (about -1.1%), and that reported growth in 2010 was extraordinarily high in Germany (about 4.5%), catching up from the huge dip in the Great Recession. Thus the implied negative adjustment for 2010 would still leave German annual GDP growth at a remarkable positive rate above 3%.

4 Real-time growth assessment – nowcasting

4.1 Monthly industrial production itself

In the previous section we performed historical adjustments of German production time series by estimating the dynamic influences of irregular or abnormal weather conditions. In this section we want to investigate whether it helps 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

⁷The agricultural sector is obviously expected to also be affected by the weather, but its share in aggregate output is very small for Germany. For other economies it might be necessary to take that effect into account as well.




Figure 6: GDP quarterly growth adjustment for abnormal weather (at annualized rate differential in %), based on estimated weather deviation effects in total industrial production. Calculated as $\widehat{gGDP}_t - gGDP_t$, i.e. positive numbers mean that weather-adjusted growth would have been higher than reported growth.

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.

Because our present aim is merely to check whether this constitutes a promising route for future research, we employ the shortcut that we do not work with a full real-time data set but instead continue to use our data set on industrial and construction-sector production which effectively contains only a single vintage (from the end of the sample). We therefore do not take into account the data revisions occurring after the respective first publications.

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 of the explanatory variable of orders received, and for the nowcasting exercise to exclude 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 non-weather benchmark. (The numbers in the historical adjustment columns were already mentioned in the text above and are repeated for convenience and comparison.) Without the contemporaneous weather terms in the nowcasting scenario the extra predictive power of the remaining weather terms amounts to eleven percentage points in the total industrial production equation ($32 - 21 = 11$); on the one hand this is quite a limited value for practical nowcasting purposes, but on the other hand it reflects a rise of the overall fit by about one half.

Furthermore, the quantitative picture is certainly different for the construction

Table 1: Explanatory and nowcasting/backcasting power

$(\bar{R}^2 \text{ in } \%)$	nowcasting / middle-of-month info set	backcasting / early-next- month info set	historical adjustment
<i>total industry output growth</i>			
benchmark predictive regression	21	21	33
with (abnormal) weather	32	35	48
<i>construction sector output growth</i>			
benchmark predictive regression	0	0	15
with (abnormal) weather	49	67	72

Notes: The historical adjustment column corresponds to results in section 3. The benchmark equations are ARDL type models (see equation 4); the first lag of output growth is removed in the nowcasting and backcasting cases for the benchmark as well as the full model. Furthermore the contemporaneous weather terms are removed in the nowcasting case.

sector, for which Table 1 reveals that the extra explanatory power of weather fluctuations is still almost fifty percentage points even under the nowcasting scenario (having dropped from $72 - 15 = 57$ to $49 - 0 = 49$).

4.2 Nowcasting GDP in a mixed-frequency framework

Our estimates are based on monthly industrial production for econometric reasons, but quarterly GDP figures are often more important in public economic discourse. Therefore we also assessed whether the inclusion of abnormal weather realizations yields additional predictive power. Given that one important advantage of the weather data is its timeliness, we focus on a nowcasting exercise that takes place right in the middle of the quarter. The middle of the quarter is also the middle of the second month of the quarter, or February/M2, May/M5, August/M8, November/M11. At this point of time the value of the previous quarter is just published, and the value of industrial production as well as incoming orders of two months

back is also just becoming available (thus: M12, M3, M6, M9). The realizations of the immediately preceding month (i.e., the first month of the quarter) are not yet known, but they can be backcast using our dynamic regressions with weather terms from 3.

We employ a mixed-frequency MIDAS approach as proposed and popularized by Ghysels, Santa-Clara, and Valkanov (2004). The idea is to run a regression with a quarterly left-hand side variable (GDP), including a candidate monthly variable with all three monthly observations per quarter as three separate econometric terms. Although it has recently also been suggested to simply run this somewhat “inflated” regression unrestrictedly (Forni, Marcellino, and Schumacher, 2015), the standard in the literature is to impose a certain polynomial connection between the various terms pertaining to a high-frequency variable. The overall impact of the variable can then be defined by a low-dimensional parameter that prevents overfitting, especially with a view to forecasting.

More formally, a generic MIDAS equation can be written as follows:

$$y_t = \alpha + \sum_{k=1}^K \beta_k y_{t-k} + \sum_{i=0}^M X_{t-i} \mathbf{b}_i + \sum_{j=1}^J \gamma_j W_j(x_{j,\tau-p_j}, \dots, x_{j,\tau-p_j}; \theta_j) + \varepsilon_t \quad (9)$$

In this equation, t denotes the low-frequency (quarterly) time periods whereas τ stands for the high-frequency (monthly) time. The first standard block comprises autoregressive terms up to lag K , where publication lags imply that lower-order coefficients β_1, \dots may have to be restricted to zero because the corresponding terms y_{t-1}, \dots are not yet observable at the time of the nowcast. Next is a block of distributed lags of N exogenous variables measured at the low (quarterly) frequency, such that X_t is an N -column matrix with associated $N \times 1$ coefficient vector \mathbf{b}_i for each lag i . Note that some of the indicators may be available quickly enough such

that contemporaneous values at lag $i = 0$ may also appear.

Finally we have the different high-frequency indicators $x_{1,\tau}, \dots, x_{J,\tau}$ that enter the predictive equation in a special way. Each indicator group j is loaded with a coefficient γ_j , but the individual coefficients are governed by a weighting function W_j that depends on a low-dimensional $L_j \times 1$ hyperparameter θ_j . Notice that in the high-frequency time axis we use the convention that $\tau \pm 0$ denotes the first month of the quarter.⁸ Since it is possible that a monthly indicator for the second or third month of the t -quarter is already observable at the time of the forecast, this means that the initial high-frequency term may be a lead relative to $\tau \pm 0$, thus $p_j < 0$ is possible. Given a maximal high-frequency lag P_j , the requirement of having a low-dimensional hyperparameter means $L_j < P_j^*$, where $P_j^* = P_j - p_j + 1$ counts the number of high-frequency terms of the j -th variable. Typical dimensionality values are $L = 2$ or $L = 3$.

Regarding the weighting functions W_j , the simplest case is the unrestricted U-MIDAS specification, where $\gamma_j W_j(x_{j,\tau-p_j}, \dots, x_{j,\tau-P_j}; \theta_j) = \sum_{i=1}^{P_j^*} w_{j,i} x_{j,\tau-i}$ with all $w_{j,i}$ freely varying. That is, γ_j is implicitly normalized to unity, and there is no role for θ_j . In contrast, the various restricted MIDAS variants rely on polynomial functions as follows. Let each weight be

$$w_{j,i} = f(i, \theta_j) / \sum_{k=1}^{P_j^*} f(k, \theta_j) \quad (10)$$

such that $\sum_{i=1}^{P_j^*} w_{j,i} = 1$.

The “normalized exponential Almon” polynomial weighting function is then given by plugging $f(i, \theta_j) = \exp(\sum_{l=1}^L \theta_{j,l} i^l)$ into (10). For example, the widespread case of $L = 2$ leads to $w_{j,i} = \exp(\theta_{j,1} i + \theta_{j,2} i^2) / \sum_{k=1}^{P_j^*} \exp(\theta_{j,1} k + \theta_{j,2} k^2)$.

⁸This convention is common to Ghysels’ toolbox for Matlab as well as the MIDAS implementation in the gretl program that we use (Cottrell and Lucchetti, 2017). The relevant R package `midasr` appears to follow a different timing rule. Also note that MIDAS specifications commonly contain only one set of high-frequency terms instead of J different ones, so our setup is slightly more general.

The “plain Almon” weighting function is represented by $w_{j,i} = \sum_{l=1}^L \theta_{j,l} i^{l-1}$.

The next possibility is the normalized beta case, where the hyperparameter is two-dimensional and yields:

$$f(i, \theta_j) = ((i-1)/(P_j^* - 1))^{\theta_{j,1}-1} (1 - (i-1)/(P_j^* - 1))^{\theta_{j,2}-1} \quad (11)$$

Loosely speaking, if $\theta_{j,2} \gg \theta_{j,1}$ then the weights are decreasing in i and relatively smoothly approach the final value zero. Therefore this parametrization is also labeled “normalized beta, last lag zero”. The commonly available implementations do not seem to impose this restriction, however. With a third hyperparameter element $\theta_{j,3}$ the final value can be allowed to depart from zero. Under the plain label “normalized beta” the weights $w_{j,i}^{bnz}$ are given by starting from the previous weights $w_{j,i}$ obtained from plugging (11) into (10), and then transforming them via: $w_{j,i}^{bnz} = (w_{j,i} + \theta_{j,3}) / (1 + P_j^* \theta_{j,3})$, which ensures that the sum of weights is still unity.

As before, we include as predictors in the benchmark specification the growth of new orders received, a term spread of interest rates, and lags of industrial production along with (few) lags of GDP. Table 2 reports the results of these MIDAS regressions.

The enhanced specifications that include weather terms –comprising the backcast for the IP value of the preceding month taking advantage of the weather information– are capable of reducing the RMSE as a measure of forecast accuracy in this MIDAS application. Only the unrestricted U-MIDAS variant appears to suffer from overfitting and associated parameter uncertainty. For the first three MIDAS variants reductions of RMSEs between 10 and 15% are achieved. With respect to an even simpler quarterly AR(1) benchmark the compound gain is 30%, but this latter number contains the effects of standard leading indicators such as new orders,

Table 2: MIDAS nowcasts of German GDP

(all values RMSE $\times 10000$)	MIDAS parametrization				
	Normalized	Almon	Normalized	Norm.	memo
	exp. Almon		Beta	Beta, zero last lag	item: U-MIDAS
Benchmark	48	44	44	45	43
Specification with weather terms	41	38	39	43	54

Notes: AR(1) root mean square error (RMSE) $\times 10000$ for comparison: 54. Forecast evaluation sample: 2010Q1-2015Q4. The assumed information set is all published values up to the middle of the quarter (middle of the second month of the respective quarter, e.g. mid-May for Q2). “U-MIDAS”: unrestricted polynomial coefficients. For the specifications see the text.

albeit on a monthly basis.

5 Conclusions

We conclude that weather conditions that deviate from their regular seasonal pattern affect the construction sector and aggregate production in Germany. 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 data sets 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. This allowed some further

insights such as a quantification of the estimation uncertainty. We could confirm earlier findings in the literature that the effects of the weather variables such as air temperature or snow height (in deviations from seasonal averages) are month-specific. The specification also had to account for serially correlated production and dynamic reactions to past weather incidents.

Weather effects could also be used to improve the backcasting of monthly output growth realizations that are still unknown because of the publication delay of such macroeconomic data. It was also shown that the estimated weather adjustments do not always cancel within a quarter and often are large enough to be noticeable even when compared to economy-wide GDP measures.

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.
- COTTRELL, A., AND R. LUCCHETTI (2017): *Gretl User’s Guide*.
- 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).
- FORONI, C., M. MARCELLINO, AND C. SCHUMACHER (2015): “Unrestricted mixed data sampling (MIDAS): MIDAS regressions with unrestricted lag polynomials,” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 178(1), 57–82.
- GHYSELS, E., P. SANTA-CLARA, AND R. VALKANOV (2004): “The MIDAS touch: Mixed data sampling regression models,” Discussion paper, Série Scientifique CIRANO.
- 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.
- 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.