Journal of Economic Dynamics & Control 83 (2017) 107–148

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Surprise, surprise – Measuring firm-level investment innovations

Rüdiger Bachmann[a,1,*], Steffen Elstner[b], Atanas Hristov[c]

[a] University of Notre Dame, CEPR, CESifo and ifo, United States
[b] RWI Essen, Germany
[c] European Commission, Belgium

ARTICLE INFO

Article history:
Received 17 March 2017
Revised 25 July 2017
Accepted 26 July 2017
Available online 4 August 2017

JEL classification:
E20
E22
E30
E32

Keywords:
Survey data
Expectation errors
Firm data
Investment
Idiosyncratic shocks
Higher moments

ABSTRACT

Using the IFO Investment Survey for the German manufacturing sector, we construct a forty-year panel of firm-level investment innovations (surprises). We document the cross-sectional and time-series properties of this panel, and our main findings are: the cross-sectional dispersion of investment surprises is countercyclical, as is their average within-firm time series volatility, and both are highly correlated. This justifies, in part, strategies in the literature to use cross-sectional moments for the calibration of heteroscedasticity. At the same time, the cross-sectional dispersion of investment is procyclical, suggesting a nonsmooth capital adjustment friction at the micro level. There is substantial dispersion in within-firm volatility, a feature consistent with a recent literature on information frictions at the firm-level. Finally, the aforementioned second moments of investment innovations are Granger-caused by recessions, but not vice versa, rendering simple exogenous uncertainty shocks less plausible as drivers of business cycles.

1. Introduction

This is a facts paper. It studies investment innovations (or surprises, terms we will use interchangeably in this paper) for West German manufacturing firms. We define an investment innovation as an investment expectation error, i.e., the difference between an investment realization for a given year and the prior investment expectation for that same year, for both of which we have direct, quantitative observations through confidential survey data.

A growing body of recent literature has studied the time series properties of the second (and higher) moments of a number of firm-level outcome variables such as sales and employment growth rates, investment rates and price changes. Arguably the most pervasive fact that this literature has uncovered is the countercyclicality of the cross-sectional second moment, dispersion, of most of these firm-level variables. Part of the literature has used this fact to infer, through the lens of structural firm or more reduced-form econometric models, that firm-level shocks are countercyclically risky.

* Corresponding author.
E-mail addresses: rbachman@nd.edu (R. Bachmann), Steffen.Elstner@rwi-essen.de (S. Elstner), Atanas-Dimitrov.Hristov@ec.europa.eu (A. Hristov).

[1] We thank the JEDC editor, three anonymous referees and seminar/conference participants at the Bonn 2014 Conference "The European Crisis - Causes and Consequences", the 2014 SED in Toronto, the NBER Summer Institute 2014 (Macroeconomics and Productivity), the 2014 EEA meeting in Toulouse, the Federal Reserve Board, the University of California at Santa Cruz, the University of Düsseldorf, as well as Michael Bryan and Petr Sedlacek for their comments. We are grateful to Annette Weichselberger from the IFO Institute for helping us with the data and introducing us to the institutional backgrounds.

http://dx.doi.org/10.1016/j.jedc.2017.07.009
0165-1889/

The novelty of our approach is to use quantitative survey data on firm investment expectations and data on how much in a given year a firm actually invested, combing them to compute investment surprises in a model- and econometrics-free way.[2] We then construct an (unbalanced) panel of these investment surprises and analyze their long-run, cross-sectional and cross-sectional cum time-series properties.

With this approach, which we view as complementary to the empirical studies based on firm-level outcome variables, we contribute to a number of debates about the nature of firm risk, that is, the conditional distribution of future firm-level shocks. For example, do data on the cross-sectional dispersion of firm variables identify within-firm volatility processes? Our answer is yes, their cyclical dynamics, but not their long-run average behavior. Second, do researchers need to look beyond the second moment to understand the time-varying nature of firm risk? Our answer is: likely no. And finally, do second moments of investment surprises Granger-cause recessions or vice versa. Our answer is: the latter. We also provide an example in this paper where results based on the dispersion of firm-level outcome variables – the dispersion of firm-level investment is procyclical – are very different from those based on the dispersion of surprises – the dispersion of firm-level investment surprises is countercyclical. This is consistent with nonsmooth capital adjustment frictions at the firm level. Finally, the fact that there is substantial heterogeneity in within firm volatility is supportive of recent models of firms with information frictions.

Our analysis comes with an important caveat: we do not measure surprises in objects that are usually considered exogenous to the firm, such as idiosyncratic revenue productivity. To the best of our knowledge, such firm-level expectational data do not exist and would be difficult to obtain. Thus, we view our approach as the next best possible, even though we are not measuring deep structural shocks.

Since we measure surprises with respect to an endogenous variable, firm-level investment – indeed, in technical lingo, investment is usually a control variable of the firm –, the question arises: why does the firm not invest according to plan? The answer may be found in the time horizon of the investment expectations in the survey, the longest of which spans over more than a year. Over such a time horizon, firms may obtain new information and/or experience a shock and thus change their investment activity accordingly. These shocks will materialize as the investment surprises we observe. Moreover, as we will show, firms deviate from their investment plans in a plausible way: the longer the expectation horizon, the larger on average and the more procyclical the investment surprise; during periods of lower overall economic activity investment surprises tend to be negative. All this structure in investment surprises makes it implausible that they are merely the result of measurement error in investment realizations (or expectations) or of random survey answers - why should the importance of these issues decline with the expectation horizon? This makes us comfortable using investment surprises as informative about underlying shocks, that is, underlying firm risk.[3]

Our primary data source is the West German part of the semiannual IFO Investment Survey (IFO-IS). It covers all manufacturing industries plus mining, and provides quantitative information on expected and realized investment for various horizons.[4] The data goes back to 1970, so we can compute a panel of investment surprises for a large number of firms for roughly forty years, a period that includes five recessions. Another advantage of our data is that they are highly confidential, which makes it unlikely that investment expectations are the result of strategic behavior of the sort that can be expected in, e.g., public earnings announcements, which firms might use to signal to the market.

Using the IFO-IS data, we establish the following facts: first, the pooled distribution of investment surprises, which summarizes their long-run properties, displays excess kurtosis for all expectation horizons, but no significant skewness. This is consistent with the view that firms are exposed to symmetric, but somewhat fat-tailed risk, to the extent that, as we have argued, data on investment surprises identify underlying firm risk.[5]

Second, the cross-sectional average of investment innovations is procyclical, which is to be expected, but it provides a good plausibility check for our survey data. Moreover, the procyclicality of average investment innovations increases with the expectation horizon, which means the longer the expectation horizon, the less plausible is the “investment plans should

[2] Jurado et al. (2015) use a saturated factor approach to estimate expectations both for aggregate and firm-idiosyncratic data to control for the predictable part of changes in economic outcome variables.

[3] In addition, our data only gives us firms’ expected investment, not the probability distribution of firms’ investment expectations or even just their subjective uncertainty around these expectations. Such data, however, do not exist over a relatively long time horizon. The one, albeit imperfect exception to this is the Italian Survey on Investment in Manufacturing, used by Bontempi et al. (2010) and Guiso and Parigi (1999), which asks firms not only about their expected sales growth, but also about their max-min sales growth range. However, the time series of this survey is too short to derive business cycle properties of firm-level sales surprises. Of course, the situation on the household side, owing to the pioneering work of Charles Manski (see, for instance, Bellemare and Manski, 2011) is different in this regard.

[4] This distinguishes the IFO-IS from other business surveys that often feature qualitative, up-down-unchanged expectation and realization data. Notable exceptions are the Canadian Capital Expenditures Survey for the manufacturing sector (see Dave, 2011), the aforementioned Italian survey, a very recent survey in New Zealand (see Kumar et al., 2015), and the annual German IAB Establishment Panel (see Müller, 2011), all of which, however, cover a much shorter period than our data. Another, now defunct example in the U.S. is the BEA survey of business expenditures on plant and equipment (see de Leeuw and McKelvey, 1981; 1984) with quantitative annual expectation data about aggregate prices. The IFO institute itself has another survey with qualitative expectations about many more firm-level variables, which has been used in the literature (see, e.g., Bachmann et al., 2013b). In the U.S., there are the Philadelphia FED Business Outlook Survey and the Small Business Economic Trends Survey by the National Foundation of Independent Businesses, each of which contains a plethora of qualitative firm-level expectations. However, to derive quantitative expectation errors from these surveys, quantification assumptions are needed, see Bachmann and Elstner (2015) and Müller and Köberl (2007).

[5] Bachmann and Bayer (2013) and Midrigan (2011) have used data on, respectively, firm-level Solow residual growth rates and price changes to infer these properties of firm risk.