MISSING AGGREGATE DYNAMICS:
ON THE SLOW CONVERGENCE OF LUMPY ADJUSTMENT MODELS

David Berger
Northwestern

Ricardo Caballero
MIT and NBER

Eduardo Engel
U. of Chile and NBER

January 25, 2013[1]

Abstract

When the microeconomic adjustment underlying an aggregate variable is lumpy, conventional VAR procedures often imply less persistence than there really is. This is relevant for non-, semi- and structural models in macroeconomics. The extent to which persistence is underestimated decreases with the level of aggregation, yet convergence is very slow and the bias is likely to be present for sectoral data in general and, in some cases, for aggregate series as well. Paradoxically, while idiosyncratic productivity and demand shocks smooth away microeconomic non-convexities and are often used as a justification for approximating aggregate dynamics with linear models, their presence exacerbates the bias. We propose procedures to correct for the bias and provide various applications. In one of them, the difference in the speed with which inflation responds to sectoral and aggregate shocks disappears once we account for the "missing aggregate dynamics" bias.

JEL Codes: C22, C43, D2, E2, E5.

Keywords: Aggregate dynamics, persistence, lumpy adjustment, idiosyncratic shocks, aggregation, Calvo model, *Ss* model, inflation, sectoral shocks, aggregate shocks, investment, labor demand, sticky prices, biased impulse response functions.

[1]We are grateful to Filippo Altissimo, William Brainard, Xavier Gabaix, Pablo García, Robert Hall, Fabiano Schiaverdi, Harald Uhlig and seminar participants at Humboldt Universität, MIT, Universidad de Chile (CEA), University of Maryland, University of Pennsylvania, Yale University, NBER EFG Meeting, and the 2nd ECB/IMOP Workshop on Dynamic Macroeconomics, Hydra, for their comments. Financial support from NSF is gratefully acknowledged. This paper is an extensively revised version of "Adjustment is Much Slower than You Think," NBER WP #9898.

1 Introduction

The dynamic response of aggregate variables to shocks is one of the central concerns of applied macroeconomics. The main procedure used to measure these dynamics consists in estimating a vector autoregression (VAR). In non- or semi-structural approaches, the characterization of dynamics stops there. In other, more structural approaches, researchers try to uncover underlying parameters from the estimated VAR and use the implied response to shocks as the benchmark against which the success of the calibration exercise, and the need for further theorizing, is assessed.

The main point of this paper is that when the microeconomic adjustment underlying an aggregate variable is lumpy, conventional VAR procedures often lead the researcher to conclude that there is less persistence than there really is. The extent to which persistence is underestimated decreases with the level of aggregation, linear models capture no persistence when applied to an individual series and they capture all the persistence when applied to a series that aggregates infinitely many agents (Rotemberg, 1987). Interestingly, convergence is very slow: the bias is likely to be present for sectoral data quite generally and, in some cases, for aggregate series as well. For example, for U.S. investment, even after aggregating across all establishments in the non-farm sector, the estimated half-life of a shock is only a small fraction of the half-life that would be obtained with infinitely many establishments.

We propose three procedures for correcting this bias. The first approach uses a proxy for the shocks while the second one is based on instrumental variables. The third procedure estimates an ARMA specification that captures the true dynamics in the AR component. We provide two detailed applications. In the first application, we show that Bils and Klenow's (2004) finding of much slower speeds of adjustment when looking at individual price adjustment frequencies than when estimating the speed of adjustment with linear time-series models can be explained entirely by our missing aggregate dynamics bias. The difference between adjustment speeds disappears once we correct for the bias.

Our second application revisits Boivin, Giannoni and Mihov's (2009) finding that sectoral inflation responds much faster to sectoral shocks than to aggregate shocks. Once we correct for the bias highlighted in this paper the responses of inflation to both types of shocks look very similar.

The basic intuition underlying our main results begins by noting that linear models are unable to capture the nonlinear relation between shocks and aggregates when adjustment at the micro level is lumpy. To see this, consider the series for one agent. If the agent does not adjust in a given period, its response in that period to past shocks will be zero. On the other hand, if every time the agent adjusts it responds one-for-one to the shocks that took place since the previous adjustment, the average response to a shock that took place k periods ago will be equal to the probability of not adjusting in periods 0 through $k-1$ and adjusting in k. The average nonlinear impulse response at lag k will then be equal to the probability of first adjusting k periods after the shock.

By contrast, the impulse response inferred from a linear model only considers correlations be-

tween the agent's actions in different points in time. If the agent did not adjust in one of the periods under consideration, there is no correlation. And even if the agent adjusted at both points in time, since the agent's actions necessarily reflect shocks in non-overlapping time periods, the correlation will also be zero if shocks are uncorrelated. This implies that the impulse response obtained via linear methods will be zero at all positive lags, suggesting immediate adjustment to shocks.

The bias falls as aggregation rises because the correlations at leads and lags of the adjustments across individual units are non-zero. That is, the common components in the adjustments of different agents at different points in time provides the correlation that allows us to recover the nonlinear impulse response. The more important this common component is —as measured either by the by the variance of aggregate shocks relative to the variance of idiosyncratic shocks or the frequency with which adjustments take place— the faster the estimate converges to its true value as the number of agents grows. In practice, the variance of aggregate shocks is significantly smaller than that of idiosyncratic shocks, and convergence takes place at a very slow pace. That is, while idiosyncratic productivity and demand shocks smooth away microeconomic non-convexities and are often used as a justification for approximating aggregate dynamics with linear models, their presence exacerbates the bias. Since in practice idiosyncratic uncertainty is many times larger than aggregate uncertainty, we conclude that the problem of missing aggregate dynamics is prevalent in empirical and quantitative macroeconomic research.

The remainder of the paper is organized as follows. Section 2 presents the Rotemberg (1987) equivalence result that justifies using linear time-series methods to estimate the dynamics for aggregates with lumpy microeconomic adjustment, as long as the number of units in the aggregate is infinite. Section 3 presents the missing-aggregate-dynamics bias that arises when the number of units considered is finite. This section also provides results establishing the slow convergence to the Rotemberg limit. Section 4 describes various approaches to correct for the bias while Section 5 considers various extensions of the baseline model and shows that the bias is present in all of them. Section 6 considers two detailed applications and Section 7 concludes. Various appendices follow.

2 ARMA Dynamics and the Calvo-Rotemberg Limit

Regardless of whether the final goal is to have a reduced form characterization of aggregate dynamics, or whether this is an intermediate step in identifying structural parameters, or whether it is just a metric to assess the performance of a calibrated model, at some key stage researchers estimate an equation of the form:

$$A(L)\Delta y_t = \epsilon_t, \tag{1}$$

where Δy represents the change in the log of some aggregate variable of concern, such as a price index, the level of employment, or the stock of capital; ϵ is an i.i.d. innovation; and $A(L) = 1 + a_1 L + a_2 L^2 + ...$, where L is the lag operator and the a_is are fixed parameters.