

Economic Time

By Oscar Jordá

Introduction

Economics is a tricky science. We have to investigate the underpinnings of economic behavior without the luxury of controlled experimentation, building upon axioms based on the fickle nature of the human psyche, and relying on data collected at the convenience and stricture of calendars rather than in "economic time."

Despite these challenges, we have certainly come a long way. We have learned to rely on "natural experiments" - something like studying the resilience of a roofing structure by investigating the aftermath of a hurricane - to compensate for the unavailability of controlled experimentation. The rational expectations revolution and recent developments on bounded rationality and experimental economics continue to perfect our ability to construct sensible models. However, incorporating the timing with which economic events take place into economic models and reconciling the time scale of economic behavior with that of available data, remains a difficult barrier to our comprehension of economics.

To understand what I mean, go back in time to April 13, 1970. After the explosion of an oxygen tank brought the Apollo XIII Mission to the brink of catastrophe, astronauts James Lovell, John Swigert and Fred Haise had to execute manual, controlled, engine burns to keep their spacecraft in an optimal trajectory to return to earth. The timing and the length of these burns were critical if they were to reenter the atmosphere safely. Store managers face a similar problem everyday in trying to manage their inventory levels. Given expected sales, stock-out costs and the financial costs of excess inventory, one can compute the optimal level of inventories (the optimal trajectory to return to earth) and adjust any deviations from this level (the manual burns). However, unlike the famous astronauts of our story, store managers often do not get to know their inventory position until the end of the day. Had the heroes of Apollo XIII not been able to continuously

monitor their exact coordinates, we would have likely celebrated their odyssey in the form of a memorial at Arlington cemetery rather than with a Hollywood blockbuster.

This essay reviews recent developments in the analysis of dynamic econometric models when the economic data generating process evolves in a different time-scale than that of available sampled data. These two parallel time-scales will rarely coincide and in fact can differ quite dramatically from one another in practice. We will denominate as *operational time* the time scale native to the economic process under investigation while *observational time* will refer to the time scale of the sampled data. An example will clarify the meaning of this nomenclature.

The literature on time aggregation typically assumes that the true economic decision interval is finer than the data sampling interval in addition to assuming that economic decision-making is done at fixed intervals. An example of this type of aggregation occurs when the economic process evolves at a monthly frequency but only quarterly data is available to the investigator. Numerous papers exist on this topic and I shall not dedicate any effort to review the main results, which are well established.

The view that the economic process evolves at fixed intervals of time is rather restrictive however. A more interesting and realistic scenario is to consider situations in which operational time evolves at stochastic intervals of time while data collection is done at fixed, calendar-time intervals. This type of aggregation is common in finance, for example, where available data on financial assets is reported weekly, or daily at best, while transactions in these markets take place anywhere from seconds to hours and days apart. This is also an example of the type of problem faced by our store manager and space enthusiast. More importantly, the timing of economic events is likely to be endogenous to the economic process under observation. The implication of this observation is that the information contained in the manner the data is aggregated is very likely to be helpful in the choice of econometric model.

The remainder of this essay will elaborate on ways to model and estimate problems that involve this type of disconnection between the timing of economic behavior and the available data and the implications that this disconnection has on structural economic analysis.

Transformations of the Time Scale

Suppose that the economic process under consideration follows a generic stochastic process that evolves in *operational time* τ , namely, $x = \{x_t\}_{t=1}^{\infty}$. By contrast, the available data will be the realizations of a different process, $z = \{z_t\}_{t=1}^{\infty}$, whose elements are functions of those of x . z is said to evolve in *observational time* t . The transformation of the operational time scale τ , into the observational time scale t , is given by

$$\tau = \mathbf{j}(t) = \mathbf{j}(k(t)) = \sum_{j=1}^t k_j \quad \text{for } k = \{k_t\}_{t=1}^{\infty}$$

k is termed the frequency of aggregation which can be thought of as a sequence of numbers or more generally, a stochastic process itself. The easiest way to understand the notation is by noting that the number of operational time observations aggregated per sampling interval is given by $\mathbf{j}(t) - \mathbf{j}(t-1) = k_t$. Therefore, if the internal time scale to the process is monthly but the data is only recorded quarterly, then $k = 3$, as in the old fashion thinking on time aggregation. If one considers asset prices in the stock market for example, then k is more naturally thought of as a stochastic process.

Given the frequency of aggregation, k , and knowledge of whether the data was averaged over the interval of aggregation or only the last observation in the interval recorded, it is possible to generate the mechanism that transports us from the operational time process x to the observational time process z . At this point, understanding the properties of the

process z given the process x becomes mathematically complex although the main points can be stated succinctly.¹

The most important point is this: Even when one assumes that the operational time process x is a simple linear dynamic process (such as an ARIMA process) drawn from a distribution with a constant conditional variance, the observational time process z will have, in general, time-varying parameters and a conditional variance that depend on the frequency of aggregation k . Furthermore, the resulting process, z , will have additional persistence that was not present originally.

Structural Bias

Why should one care about these claims and the discrepancies between time scales? There are at least three good reasons: (1) Any empirical test of an economic model will require that the frequency of aggregation be appropriately modeled and incorporated into the econometric specification to avoid structural bias (after all we would not want to see our astronauts burning up upon reentry into the atmosphere); (2) forecasting performance can be improved by conveniently exploiting information on the aggregation frequency; and (3) understanding that the stylized statistical properties of many economic variables may be artifacts of a transformation of the time scale may lead to better economic models.

Time matters. To understand reason (1), suppose that we try to model William Baumol's and James Tobin's holdings of cash. According to the Baumol-Tobin model of cash management, holding money is convenient but costly. People hold cash to avoid having to go to the bank each time they want to make a purchase but the cost of this convenience is the forgone interest in their accounts. To test this theory, we thus record the total cash holdings of our famed economists every month. That, along with the interest rate paid on their checking accounts, allows us to estimate their willingness to hold cash.

¹ The reader is referred to Jordá and Marcellino (2000) for a further discussion and proofs.

This estimate will likely be flawed. Here is why. William lives in New York, more specifically in the Bronx, which is the only place where he can afford to live in given his meager professor salary. Fully aware of the risks of withdrawing money in public from a cash machine, he goes there only once a month (while less law-abiding segments of society are scouting the local gun show for new gear) and withdraws \$300.00. By contrast, James lives in New Haven, but as a Nobel Prize winner, he is able to afford a cottage in a peaceful community in Connecticut. Therefore, he hardly ever carries more than \$30.00 since he can always go to the cash machine by the university library, thus withdrawing cash an average of 10 times over the month. From the econometrician's point of view, both economists behave similarly since the only data available indicates that their total cash withdrawals amounted to \$300.00 over the month although it is clear that the behavior of both economists is radically different.

In the notation used in this paper, the number of times each economist went to the cash machine is what we termed k , the frequency of aggregation. Armed with this information, our fellow econometrician would have had no problem describing the behavior implied by the Baumol-Tobin model, given the available data. We will reserve the next section of the paper to discuss an example that extrapolates the methods sketched here in an investigation of the foreign exchange market. This roughly corresponds to reason (2) raised above.

Modeling Strategies

In the previous section I argued that knowledge of the frequency of aggregation might be very useful in empirical analysis. When we are trying to estimate a model of economic behavior, this frequency of aggregation gives us an idea of the intensity of economic decision-making that pervaded throughout each interval of the sampled data. This section illustrates methods that can be used to improve the analysis of economic phenomena and

the forecasting performance of econometric models when the aggregation frequency is observed. Situations in which the aggregation frequency is unobserved are perhaps more frequent but require of rather complex methods. The reader is referred to Jordá and Marcellino (2000) for a discussion. Here, I present a simple example designed to investigate the behavior of the bid-ask spread in the U.S. Dollar/Deutsche Mark foreign exchange rate (USD/DM, FX) market.

The FX market is a 24 hours, global market although the activity pattern throughout the day is dominated by three major trading centers: East Asia, with Tokyo as the major trading center; Europe, with London as the major trading center; and America, with New York as the major trading center. Figure 1 displays the activity level in a regular business day as the number of quotes received per half hour interval. The seasonal pattern presented is calculated nonparametrically with a set of 48 time-of-day dummies. Figure 2 illustrates the weekly seasonal pattern in activity levels by depicting a sample week of raw data.

The original data corresponds to the HFDF-93 dataset available from Olsen & Associates.² The sub-sample I consider here contains 3,500 observations of half-hour intervals (approximately 300,000 ticks) constructed by counting the number of quotes in half-hour intervals throughout the day. For each individual half-hour observation I then record the corresponding bid-ask spread. The average intensity is approximately 120 quotes/half-hour during regular business days, although during busy periods this intensity can reach upwards of 250 quotes/half-hour. The activity level drops significantly over the weekend although not completely.

² These data span one year beginning October 1, 1992 and ending September 30, 1993, approximately 1.5 million observations. The data have a 2 second granularity and are pre-filtered for possible coding errors and outliers at the source (approximately 0.36% of the data is therefore lost).

Figure 1
Seasonal Intraday Pattern
Diurnal Dummies

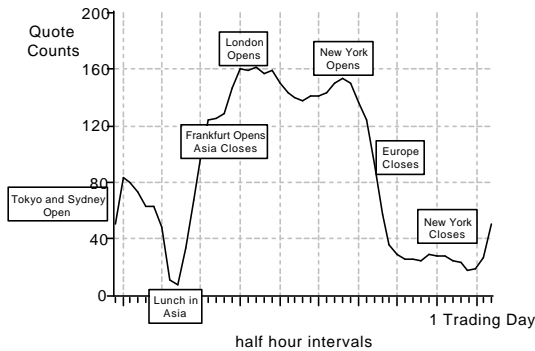



Figure 2
A Typical Week of Trading
Exchange Rate Quotes by 30 Minute Interval-Counts



The explanations for the width of the bid-ask spread vary widely (see O'Hara, 1995), ranging from "market failure" and "market power" explanations to more transactions, cost-related "dealer risk-aversion" and "gravitational pull" theories. The simple approach I take here is to investigate the dynamics of the bid-ask spread as a function of information flows measured by the level of activity in the market.

Let k_t denote the number of quotes per half hour interval (a count variable which we have denominated here as the aggregation frequency) and let z_t denote the bid-ask spread that corresponds to the half-hour interval t . Thus, this is an example in which the variable z_t is the result of aggregating the original tick-by-tick data but for which we observe the frequency of aggregation (k_t). The problem thus consists of jointly modeling these two random variables. However, since I am only interested in knowing the effect of information flow on the bid-ask spread, I will concentrate on the conditional distribution of z_t given k_t .

A researcher interested in modeling z_t but ignorant that the data has been aggregated into half-hour intervals, may proceed by using an off-the-shelf ARIMA model. To illustrate the pitfalls of his approach and based on our knowledge that the dynamics of the bid-ask spread probably depend on the information flow, I specify the following model,

$$z_t = \text{seasonals}(1 + F_0(k_t)) + \Phi(L, k_t)z_{t-1} + \mathbf{e}_t \quad (1)$$

$$\Phi(L, k_t) = \mathbf{f}_1(1 + F_1(k_t)) + \mathbf{f}_2(1 + F_2(k_t))L + \mathbf{f}_3(1 + F_3(k_t))L^2$$

where the *seasonals* variables are a collection of time-of-day effect, day-of-week effect, and holiday effect dummies, L is the lag operator and $F_i(k_t)$ for $i = 0, 1, 2, 3$ is a non-parametric estimate based on a sixth order polynomial. Note that as a particular case, if all the $F_i(k_t)$ are zero, one recovers a conventional linear AR(3) model.

Figure 3 depicts the estimated autoregressive parameters as a function of the transaction intensity, k_t . In the limit, notice that as the transaction intensity approaches 0 (i.e. nobody trading during that particular half-hour interval), then \mathbf{f}_1 goes to 1 while \mathbf{f}_2 and \mathbf{f}_3 go to 0 since z_t attains the same value as z_{t-1} . However, as the transaction intensity picks up, the estimated parameters exhibit a fair amount of non-monotonic variation, ranging from high persistence to negative correlation and back into high levels of persistence. Figure 4 reports the fluctuations in the average, seasonally adjusted, residual spread as a function of the transaction intensity. After accounting for the intra-day trading patterns, the spread exhibits two well defined peaks: One at low levels of activity, which is consistent with the view that when the market is inactive, the best position is to maintain a sizeable bid-ask spread to hedge against unforeseen and sudden large transactions. The second peak takes place when the intensity reaches 140 quotes per half-hour (recall that the average trading intensity in a regular business day is around 120 quotes per half hour). However, as the trading intensity increases, the filtration and price learning processes in the market evolve much faster and therefore, it becomes easier to price the FX rate. Consequently the bid-ask drops.

Figure 3

Nonparametric Estimates of the Autoregressive Parameters of the Spread Model as a Function of the Intensity





Figure 4

Nonparametric Estimate of the Mean Spread as a Function of the Intensity of Quote Arrivals



A Word on Forecasting

The results in the preceding section demonstrate the advantages of incorporating knowledge on the transaction intensity into a model of the bid-ask spread or in the wording of this paper, the model of the aggregated data z_t depends on the frequency of aggregation k . However, given knowledge of the dynamic behavior of the spread, an agent in this market may be interested in forecasting the width of the spread in the coming half-hour for the purposes of minimizing the costs of his next transaction. Therefore, one needs to construct a forecast of the trading intensity k that is expected in the next half-hour interval. Figures 3 and 4 illustrate that depending on the expected transaction intensity, the best forecast of the bid-ask spread will be different.

The transaction intensity k is a count variable so it is natural to think of it as a Poisson distributed random variable.³ Hence, the conditional distribution of k_t given all past information (that is, lags of k_t as well as lags of z_t) can be written as,

$$P(k_t = j | \tilde{z}_{t-1}, \tilde{k}_{t-1}, \mathbf{q}) = \frac{e^{-I_t} I_t^j}{j!}$$

³ Recall that the Poisson distribution is often used in problems such as those describing the arrival of customers to a line per interval of time. Here we are investigating the arrival of quotes per half-hour interval of time.

(2)

where \tilde{z}_{t-1} and \tilde{k}_{t-1} denote all past values of z and k as of time $t-1$. λ_t is the conditional expected transaction intensity at time t . Since the Poisson distribution is a one parameter density, the variance coincides with the conditional mean, λ_t . A natural way to parameterize λ_t is to use a formulation similar to that of an ARCH⁴ model. In particular, to account for the fact that λ_t must remain non-negative, I use the following E-GARCH⁵ type specification,

$$\log(I_t) = \text{seasonals} + \Theta(L) \log(I_{t-1}) + \Psi(L)k_{t-1} + \Pi(L)z_{t-1} \quad (3)$$

The variables in *seasonals* are the same dummies described in equation (1) while $\Theta(L)$, $\Psi(L)$, and $\Pi(L)$ are conventional, finite order, autoregressive, lag polynomials. The model described by equations (2) and (3) is called the autoregressive intensity model (ACI) introduced in Jordá and Marcellino (2000) and can be easily estimated by maximum likelihood.

To get a sense of the properties of this formulation, I estimated the ACI model on the transaction intensity data that I used to estimate the bid-ask model in equation (1). Table 1 below summarizes the salient results and compares them to a benchmark Poisson regression based on the *seasonals* variables alone. The likelihoods are directly comparable since the Poisson model is a nested particular case of the ACI model.

⁴ ARCH stands for autoregressive conditional heteroskedasticity. The ARCH model was first introduced by Engle (1982).

Table 1 - The ACI model versus a benchmark Poisson Model

	Poisson	ACI
<i>Log-Likelihood</i>	-28562.37	-18145.52
<i>Parameters</i>	55	80
<i>Akaike Information C.</i>	19.729	12.565
<i>Schwarz Information C.</i>	19.842	12.730
<i>Ljung-Box Q_5</i>	1608	99
<i>Ljung-Box Q_{10}</i>	1903	128
<i>Ljung-Box Q_{50}</i>	2398	372
<i>LR test ACI vs. Poisson</i>	0.000 (p-value)	

These are rather striking results by any measure. The ACI model does a very good job at explaining the enormous amount of persistence indicated by the Ljung-Box statistic. The ACI has 15 additional parameters relative to the Poisson model, which could explain the dramatic improvement in the likelihood. However, as the Akaike and Schwarz information criteria⁶ show, the ACI model is overwhelmingly preferred to the Poisson. Overall, the ACI model appears to be a good way to forecast movements in transaction intensity.

Final Thoughts

Recently a student asked me why was I not teaching overlapping generations models of monetary economies in my monetary theory course. My response was that overlapping generations models are very useful in thinking about long-run issues, such as social security, taxation, government debt and the like, but that they are less well suited to explain money. Money usually arises in these economies as a medium for a generational transfer. That is simply the wrong time scale (then again, my need to have cash for my

⁵ E-GARCH stands for exponential generalized autoregressive conditional heteroskedasticity. The E-GARCH model was first introduced by Nelson (1991).

afternoon coffee may skew my view toward a transactions-cost based view of the world). My argument in this paper is that different economic decisions have different time scales, and that the timing of economic events is not gratuitous but rather endogenous to the economic problem. By contrast, available data is limited to the calendar traditions of data recording.

The implications of the divergence between "economic time" and "observational" time are twofold. On one hand, we should devote more effort to constructing models in which each economic decision has the right time scale. It can not be that my consumption decisions, my financial decisions and my retirement decisions coincide in timing, not only with each other, but also with the hiring decisions, investment decisions and production decisions of the firm that is employing me, along with the public sector's decisions on taxes and the deficit. On the other hand, we need to become more adept at understanding the subtleties that data availability impose on empirical analysis. This essay concentrated on the second implication to argue that: (1) Empirical analysis of economic hypotheses requires that we carefully model the transformation of the operational time scale into the observational time scale; (2) from a purely econometric perspective, forecasting can benefit from explicit modeling of the aggregation scheme; and (3) the data will inherit properties that originated from the economic process as well as properties that relate to the manner it was recorded. An economic model should only try to match the first set of properties.

References

- Engle, Robert F. (1982) "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 50/4, 987-1006.
- Jordá, Oscar and Marcellino, Massimiliano (2000) "Stochastic Processes Subject to Time-Scale Transformations: An Application to High Frequency, FX Data," U.C. Davis, Working Paper 00-02.
- Nelson, Daniel B. (1991) "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59/2, 347-370.

⁶ The Akaike and the Schwarz information criteria are measures that balance the quality of the fit against overparameterized models.