Gauging Hourly Economic Activity in Your Neighborhood*

Fabrizio Ghezzi† UC San Diego

Allan Timmermann‡ UC San Diego

Max Yang§ UC San Diego

March 26, 2026

Abstract

Imagine gauging hourly economic activity at the local (zip-code) level. We argue that this is now feasible and propose a new granular economic modeling approach (GEM) for measuring, modeling, and predicting economic activity at high levels of granularity. Our approach uses bridging equations and aggregation constraints to link highly granular spatio-temporal data on variables such as foot traffic, traffic flows, electricity usage and credit card spending with aggregate data on payrolls and income. We extract a common dynamic factor that summarizes local, hourly economic activity and is consistent with observable measures observed at far coarser levels of granularity such as city-wide payroll data and GDP.

Keywords: Hourly local economic activity; dynamic factor modeling; Cascade Kalman filter; high-frequency economics; alternative data

1 Introduction

"Everything is related to everything else, but near things are more related than distant things"

W. R. Tobler, The first law of geography

*We are grateful to Davide Pettenuzzo for discussions of the modeling approach in our paper.

†Rady School of Management, UC San Diego. **Email:** fghezzi@ucsd.edu

‡Rady School of Management and Department of Economics, UC San Diego. **Email:** atimmermann@ucsd.edu

§Rady School of Management, UC San Diego. **Email:** ygnmax@ucsd.edu

Economic activity is local in nature: activity at service providers such as restaurants, dry cleaners, gyms, or retail stores varies with the state of the local economy, sometimes more so than with the broader economy. Office attendance and industrial manufacturing—while typically catering to a broader market—leave clear local imprints on the economy through workers' commuting patterns and local spending. In turn, local economic activity (from consumption to retail shopping and office attendance) fluctuates significantly during the day and follows systematic hour-of-day and day-of-week patterns related to school schedules, office operating hours, and plant production cycles. Through externalities such as traffic congestion and delays in the provision of services at peak hours, spatio-temporal patterns in economic activity introduce frictions that affect business decisions on investment in capacity and household consumption decisions and can have important consequences for economic welfare.

Variation in economic activity at very high levels of granularity is typically not examined in economic analysis which focuses on aggregate outcomes such as GDP growth, unemployment, or inflation at national or regional levels. In large part, this omission has been dictated by limits on granular data, especially the absence of direct, observable measures similar to quarterly national or annual state-level GDP. Local measures of economic activity such as payroll at the city level are available only on a monthly basis, often with lengthy delays.

In this paper, we show that modeling and predicting hourly local economic activity is now feasible. New data types measured at unprecedented high spatial and temporal resolution combined with vast improvements in computing power make it possible to track and predict local economic activity at levels of granularity previously thought impossible, namely hourly zip-code data. We develop a new granular economic modeling (GEM) approach that is anchored on the highest level of granularity at which we observe data and links this with observed data at coarser levels of the spatio-temporal grid through aggregation constraints.

Our GEM approach merges new data sources, such as foot traffic, traffic flows, electricity usage, and credit card sales data, with variables that are correlated with economic activity, such as county-level payroll or state-wide GDP. We start at the finest level of granularity (hourly, zip-code) and aggregate both temporally (to the daily, monthly, quarterly, and annual frequencies) and spatially (to county and state levels). To ensure coherence with more traditional disaggregate measures of economic activity, we merge these data with (low-frequency) government data sources such as monthly city-wide payrolls, annual data on county- and state-wide GDP, and unemployment figures.

Individual variables measured at the most granular levels differ in how strongly they are correlated with economic activity. We capture this aspect through a dynamic factor model that summarizes local economic activity through a common factor extracted from our granular proxies for economic activity. The central idea is that observable proxies for local economic activity can be decomposed into exposures to the common economic activity factor along with an idiosyncratic local activity component. Our dynamic factor model captures both a slow-moving mean (including a trend component) as well as a calendar component that adjusts to different hours-of-day, days-of-week, holidays and seasons. Variation over time and across local areas is then captured through heterogeneity

in individual zip-code areas' factor loadings, which also determine spatial correlation patterns in economic activity, and their idiosyncratic components.

The ability to track and predict economic activity at high levels of resolution presents three key advantages. First, from a business planning perspective, high levels of local economic activity present both opportunities and costs for governments and firms, including questions such as how to operate dynamic pricing schemes and when to schedule potentially disruptive activities (e.g., infrastructure work). For activities generating negative externalities such as congestion, better models for temporal variation in local activity can be used to design optimal timing strategies. Second, augmenting conventional data series measured at more aggregate levels (e.g., annual state-wide GDP) with more granular proxies for economic time-series facilitates more accurate forecasts—and more timely nowcasts—even for policy makers, businesses or researchers who are predominantly interested in activity measured at a more aggregate level. Third, proxies for economic activity at high levels of spatio-temporal resolution can be used to test economic theories on preferred consumption patterns and measure the local economic impact of events such as concerts, conventions, or sport games. As part of these objectives, it is important to ensure that the "bottom up" perspective offered by the more granular data is anchored by the aggregate data.

Because every year consists of nearly 9,000 hourly periods, with potentially hundreds or thousands of zip codes and data on several variables, conventional dynamic factor algorithms quickly run into computational challenges. We deal with these issues by combining spatial pooling, a new Cascade Kalman filter, which leverages the aggregation constraints, and a computationally fast filtering approach. Specifically, our algorithm uses the prediction error at the lowest frequency (annual) to update the factor estimate through the Kalman gain and then cascades down to the next-lowest frequency (quarterly or monthly) until stopping at the hourly level. The resulting algorithm is designed to handle both the high-dimensional data sets that we use and to capture the rich dynamics of the underlying data. The Kalman gain varies across the individual series and data frequencies (e.g., daily credit card spending), attributing different weights to each series in the filtering of the factor.

Throughout our analysis, we impose aggregation constraints to ensure that our measures of economic activity are consistent across different levels of the spatial and temporal hierarchy, i.e., zip-code activity aggregates to city-wide activity which in turn sums to state-wide activity. Similarly, hourly activity within a given zip code sums to daily activity within the same zip code which in turn aggregates to lower frequencies such as monthly, quarterly, or annual. Bridging equations help us anchor our factors to observed economic activity, which is usually measured at a much lower granularity (e.g., annual GDP at the state level).

We implement the GEM to provide a detailed analysis of local economic activity in San Diego. Our main findings can be summarized as follows. First, local economic activity follows a distinct diurnal pattern with clearly marked hour-of-day and day-of-week effects reflecting commuting times, work hours, holidays, and week-ends. Economic activity also varies greatly across seasons and we identify a pronounced underlying trend component. Along the spatial dimension, we document novel patterns that would not be discoverable

by approaches based on data observed at lower levels of granularity. We develop a new spatio-temporal activity radar (STAR) similar to radar maps used in meteorology which provides hourly snapshots of economic activity in individual zip codes. We show how hourly activity is much higher at night in zip codes with an airport, hospital facilities, or restaurants and bars—in the latter case particularly during weekends. Conversely, economic activity in residential zip codes fluctuates more along school and office hours.

We also demonstrate how our dynamic factor modeling approach can be used to predict local (zip-code level as well as city-wide) economic activity for the next hour, day, or week. Measured against a proxy for local economic activity constructed using payroll data, we find that our GEM approach generates accurate forecasts compared with simpler benchmarks that ignore the dynamics in the data or do not fully exploit the heterogeneity across zip-codes.

Finally, we develop a framework for spatio-temporal scenario analysis that allows us to compute the effect of changes to the trajectories of the variables from which we extract our local economic activity factor. We use this framework to study the impact of a shutdown in foot traffic, traffic flow, or electricity for a 24-hour period in three residential or commercial zip-codes. The cumulative effect of shutting down one of these variables for a 24-hour period in three zip codes is economically large, ranging from 1% to 5% of San Diego's daily GDP and varies along three key dimensions: which variable is shut down (e.g., foot traffic versus traffic flows), which hour-of-day the shutdown occurs (e.g., daytime versus night) and the type of zip code affected (residential versus commercial). For example, shutting down traffic flow in the residential zip codes has twice as large a negative effect on San Diego's daily economic activity (-2%) as shutting down foot traffic (-1%) and the effect of either type of shutdown in commercial zip codes is significantly larger (-4%). Similarly, shutdowns occurring in the late night or early morning hours have a much smaller economic impact than those that occur in the middle of the day.

Related literature. Our paper contributes to empirical macroeconomics by proposing a framework to measure, model and predict economic activity at a highly granular spatio-temporal level. Within this field, we make two major contributions. First, by measuring economic activity at an unprecedented level of spatio-temporal granularity, we document new empirical regularities in local economic activity. Second, we contribute to the methodological literature on economic activity measurement by developing a novel modeling and estimation approach that combines series observed at different spatio-temporal frequencies and bridges the gap between alternative data and standard economic aggregates through economic restrictions.

Decision makers, such as central banks, need timely information to adjust policies to economic shocks, so the paucity of high-frequency measures of economic activity has spurred the development of different approaches to extract the best estimate of the current state of the economy.[1] This interest in timely predictions of economic activity is reflected

[1]For examples of central banks attempts to track economic activity in real time, see the Atlanta Fed's GDPNow: `https://www.atlantafed.org/cqer/research/gdpnow`; the ADS Index maintained by the Philadelphia FED: `https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/ads` built on Aruoba et al. (2009); and the New York Fed's nowcasts: `https://www.newyorkfed.org/research/policy/nowcast/#nowcast/2025:Q4`.

in the vast nowcasting literature, see, e.g., Giannone et al. (2008), Altissimo et al. (2010), Bańbura et al. (2013), and Bok et al. (2018).

Most of the literature on measuring economic activity in real time has focused on aggregate spatio-temporal levels, typically at the monthly country level. For example, Aruoba et al. (2009) and Aruoba and Diebold (2010) propose modeling approaches for tracking U.S. economic activity, while Camacho and Perez-Quiros (2010) adopt a similar approach for the euro area. More recent contributions exploit data at finer spatial and temporal levels. For example, Lewis et al. (2021) and Lewis et al. (2022) develop a weekly economic activity indicator for the aggregate U.S. economy; Owyang et al. (2009) and Koop et al. (2025b) estimate economic activity at the monthly-state level; Arias et al. (2016) and Koop et al. (2025a) consider granular spatial areas such as metropolitan areas at the lower monthly and quarterly frequencies, respectively; Baumeister et al. (2024) develop a weekly-state level economic conditions indicator.

We push the boundaries of these studies by moving to a much finer granularity in space and time, namely the zip-code-hourly level. In doing so, we contribute to the literature that seeks to exploit information from data series measured at different spatio-temporal levels. Approaches such as the MIDAS regression of Ghysels et al. (2007) have been used to deal with temporally-mixed data frequencies to measure economic activity (Carriero et al. (2015) and Marcellino et al. (2016)). A key difference with our work is that we are extracting economic activity from series observed at different spatio-temporal levels. For example, we seek to jointly exploit hourly-zip code electricity usage and yearly-county Gross Domestic Product. This dual heterogeneity in spatial and temporal measurement is outside the scope of standard MIDAS approaches.

In relation to the large body of work employing dynamic factor models in empirical macroeconomics, our contribution stems from specifying a high-frequency dynamic factor and using spatio-temporal aggregation restrictions to link variables at different levels of aggregation.[2] By using granular series to measure local economic activity, we contribute to the emerging literature that exploits the information content of alternative data (see, e.g., D'Amuri and Marcucci (2017), Chen et al. (2019), Chetty et al. (2024), Klopack and Luco (2025), and De Polis et al. (2025)).

Finally, our work is related to the literature on spatial general equilibrium models, e.g., Ahlfeldt et al. (2015), Redding and Rossi-Hansberg (2017), Desmet et al. (2018), and Kleinman et al. (2023). Our focus is on granular patterns at a much higher temporal frequency than assumed in these papers which requires developing and applying different analytical tools.

The remainder of the paper proceeds as follows. Section 2 introduces our granular data sets and documents new stylized facts. Section 3 presents our modeling approach, and Section 4 the estimation methods. Section 5 reports key empirical findings. Section 6 proposes methodologies for forecasting and scenario analysis. Section 7 corroborates the economic validity of the results. Finally, Section 8 concludes. A set of appendices report details on the datasets, algorithms and provide additional results.

[2]See, e.g., Stock and Watson (2002), Kose et al. (2003), Moench et al. (2013), and Stock and Watson (2016).

2 Data

We begin this section by introducing our data sources and describing the construction of the granular dataset that we use in our empirical analysis. Since our focus is on the highest level of spatio-temporal granularity (hourly-zip codes), and economic activity at this frequency has not been extensively analyzed, we then illustrate some stylized features of our data and provide formal justification for using such variables to measure local economic activity.

2.1 Spatio-Temporal Granular Dataset

We assemble our dataset by sourcing a multitude of series sampled at different spatio-temporal levels of granularity from several data providers.[3] Many of the variables we include in our dataset are part of the so-called 'alternative' data family. These data have become available only recently, but their success in uncovering economic dynamics is already evident.[4] Our analysis focuses on San Diego County which comprises 62 separate zip codes.

All variables included in our dataset are motivated by either prevailing narratives or established economic reasoning.[5] Starting with data recorded at the hourly zip-code level, we include four types of data. First, we use data on foot traffic. This variable tracks the aggregate counts of visits to different points of interest (POI) gathered using geolocational data obtained from mobile phones which monitor when a mobile device visits a particular zip code within a given hour. Second, we obtain data on traffic flow, defined as the number of vehicles passing by a given location within a given time interval. This data is obtained from individual detectors situated on highways across all major metropolitan areas of California and is again aggregated to the hourly interval for each zip-code. Third, we obtain data on electricity usage based on an anonymized sample of utility accounts, although we know if the account belongs to an agricultural, commercial, industrial, or residential customer. Finally, we use hourly zip-code level weather data from the National Centers for Environmental Information (NCEI) which records weather at numerous monitoring stations. We focus on dry bulb temperature, which is the closest approximation to humans' thermal perception and thus most likely to be correlated with economic activity. At the daily frequency, we have access to spending data, i.e., credit and debit card transactions aggregated to the level of zip codes. Turning to the lower frequency, we have weekly unemployment insurance claims and monthly payroll data at the county level from the California Employment Development Department (CAEDD), as well as annual zip-code payroll data obtained from the County Business Patterns (CBP)

[3] We refer to Appendix A for a detailed explanation of the data, data sources, and transformations.

[4] See, e.g., Chetty et al. (2024), Klopack and Luco (2025), Batch et al. (2025), and Buda et al. (2025). Bloomberg's second measure illustrates how alternative data can be used to construct consumption measures.

[5] For example, we include weather data drawing on a growing literature in economics and finance that demonstrates how temperature and weather influence investor behavior and economic activity (Bassi et al. (2013), Colacito et al. (2019)). In Section 2.3, we provide additional evidence that these variables are related to local economic activity by demonstrating their relationship with local payrolls.

dataset. Finally, we observe county-level GDP for San Diego at the annual frequency.

Figure 1 depicts the spatio-temporal resolution of the series included in our dataset. The horizontal axis represents the temporal dimension (hourly, daily, weekly, monthly, quarterly, and annual), while the vertical axis corresponds to the spatial dimension (zip code, MSA/county, and state-level).

Moving upward in the table reduces the level of spatial granularity while moving from left to right reduces the level of temporal granularity. The most granular series in our dataset—traffic flow, foot traffic, electricity consumption, and weather—are available at high temporal (hourly) and spatial (zip code) resolutions.

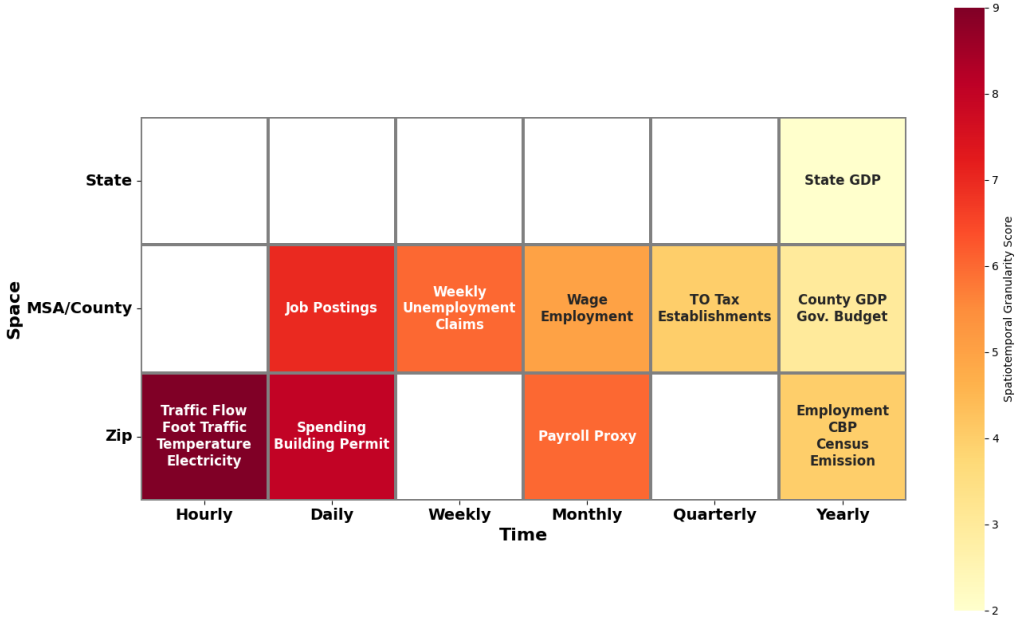


Figure 1: **Spatio-Temporal Granularity of the Data.** This figure shows the spatio-temporal granularity of the data. We assign to each variable a spatio-temporal granularity score, adding one point for each level of increased granularity in time and space, starting from the most aggregate level (i.e., yearly and state-level).

We apply a bottom-up approach linking data observed at the most granular levels to their counterparts observed at lower levels of granularity through aggregation constraints. Hence, the most granular observations are first assigned to their respective hour and zip code.[6]

[6]For recording purposes, the values of variables available only at coarser temporal resolutions are assigned to the corresponding zip code at the first hour when the data become available. For example, spending data for day d is recorded at 00:00 on day $d+1$, ensuring that all spending on day d is correctly matched to its full time interval. This method is applied consistently across other temporally aggregated variables. Similarly, we attribute the value of variables available only at higher spatial aggregation levels to all corresponding zip-code datasets, treating these aggregated variables as common information across zip codes.

2.2 Features of Hourly Variables

In this subsection, we document the empirical features of the hourly series that we employ to measure local economic activity. Our hourly variables are measured in very different units. For example, traffic flows and foot traffic are counted in trips made by cars or people, whereas energy is measured in kilowatt-hours (kWh). To make the series easier to compare, we therefore apply logarithmic transforms.[7]

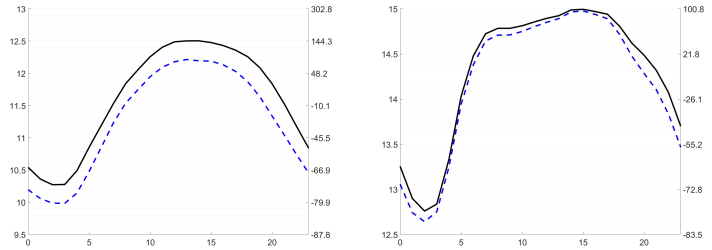
Figure 2 displays hour-of-day values of the (log) foot traffic, traffic flow, electricity, and (seasonally demeaned) temperature series, averaged across all summer (June-August) or winter (December-February) days in our sample at the aggregate San Diego level.[8] Foot traffic reaches a minimum in the early morning hours around 2-3 a.m. from which it increases sharply towards noon, stabilizing until 3:30 p.m. before starting a marked descent that lasts for the remainder of the day. This pattern is the same for the summer and winter months although the level is systematically higher during the summer. Traffic flow also reaches its lowest point in the early morning hours, around 2-3 a.m., rises sharply between 4 a.m. and 8 a.m., at which point it continues climbing at a slower pace until 4 p.m. before descending for the remainder of the day. Once again, the patterns are very similar for summer and winter months with the former being slightly higher on average.

Electricity usage follows a distinct hour-of-day pattern that differs markedly during the summer and winter months. During summer months, electricity usage reaches a low point at 3 a.m. after which it gradually increases until around 4 p.m., before stabilizing until 8 p.m. after which it declines. During winter months, the low point is reached an hour earlier (2 a.m.), from which it rises sharply to 6 a.m., stays constant until 3 p.m., rises sharply again until 5 p.m., and then declines from 6 p.m. and throughout the night. Again, electricity usage is generally higher during summer months than in the winter. The hourly temperature curve is more symmetric and sharply peaked in the winter months than the economic activity variables.

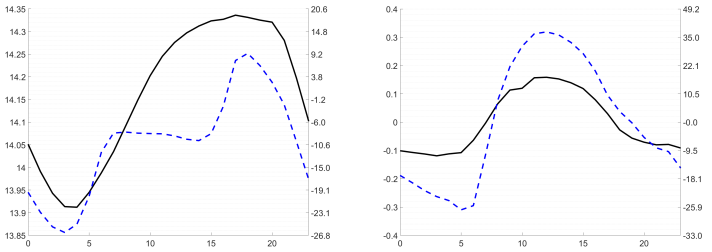
In percentage terms, the ranges of intra-day values are very wide. Measured relative to their average daily level, we see the widest ranges for foot traffic and traffic flow ($(-80\%, 150\%)$ and $(-80\%, 100\%)$, respectively) and a narrower range for electricity usage ($(-25\%, 20\%)$).

[7] For ease of interpretation we report temperature in deviation from its seasonal mean. We refer to data Appendix A for details of data transformations.

[8] We use dual vertical axes with the left axis showing logs and the right axis showing percentage deviations from the mean.

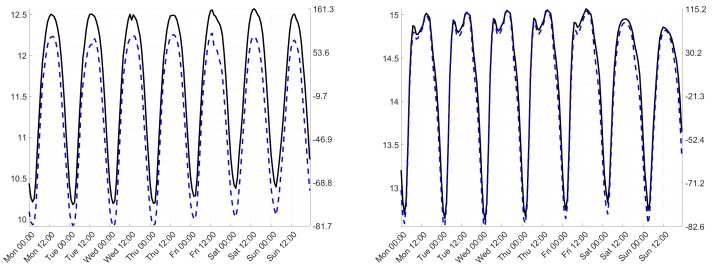


(a) **Foot Traffic.** (b) **Traffic Flow.**

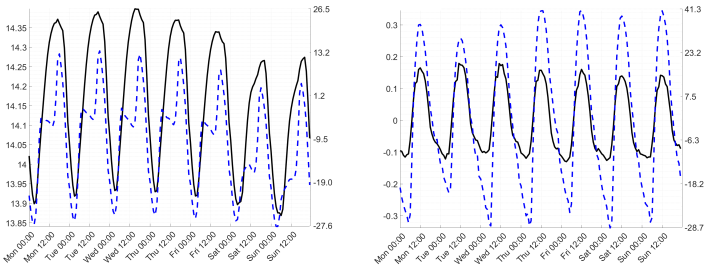


(c) **Electricity.** (d) **Temperature.**

Figure 2: **Hour-of-Day Variation.** This figure plots the hourly time series of foot traffic, traffic flow, electricity consumption, and temperature at the aggregate San Diego level for the average diurnal cycle across summer (June-August, solid black line) and winter months (December-February, broken blue line). The left scale reports log-level values of the series, while the right scale reports percentages deviation from the mean of the series.



(a) **Foot Traffic.** (b) **Traffic Flow.**



(c) **Electricity.** (d) **Temperature.**

Figure 3: **Day-of-Week Variation.** This figure plots the hourly time series of foot traffic, traffic flow, and electricity consumption at the aggregate San Diego level for the averages across summer (June-August, solid black line) and winter months (December-February, broken blue line). The left scale reports log-level values of the series, while the right scale reports percentages deviation from the mean of the series.

To study weekly (day-of-week) patterns in our series, Figure 3 plots the seasonal weekly averages for the (log) foot traffic, traffic flow, electricity consumption, and temperature series. The figure illustrates key empirical features of how our hourly data vary across days of the week, again split by summer (June-August) and winter (December-February) months.

While the series differ in levels, heteroskedasticity, and temporal dynamics, there are also some notable commonalities which, we argue below, any modeling approach must account for to extract economic activity at a granular level. These features prompt our common factor approach to modeling economic activity.

For example, a common feature is that all series have their peaks during the central hours of the day (6 a.m. - 6 p.m.) and their troughs during the overnight period (7 p.m. - 5 a.m.). Specifically, during most days of the week, foot traffic (top panel) peaks around noon, suggesting that people leave their offices to have lunch. On Wednesdays, there are two distinct peaks of similar magnitude, while on Fridays the peak occurs earlier and the descent is sharper, consistent with people leaving the office a little earlier. As expected, foot traffic is higher in summer months than in the winter.

Hourly intra-day patterns in traffic flows vary notably across different days of the week. During weekdays, traffic flow shows two peaks: one in the early morning, around 9 a.m., and one in the afternoon, around 3 p.m. This 'double-peak' behavior is absent during weekends and holidays.

The weekday patterns align with common commuting patterns as people drive to school or work in the morning and return home (or pick up children) in the late afternoon, along with recreational activities in the afternoon and early evening. Activity patterns tend to be less regular on weekends. Traffic flow is notably lower on Sundays than on Saturdays, consistent with reduced economic activities on Sundays.

Electricity consumption (third panel) typically peaks early in the evening around 5 p.m. During business days, this occurs one or two hours later than the second peak in traffic flows, consistent with an increase in electricity usage as people return to their homes. Electricity consumption is single-peaked during the summer but double-peaked during the winter with a smaller peak in the morning followed by a higher peak in the late afternoon. Electricity consumption is also notably lower during weekends than during weekdays and during the winter than in the summer. Compared to foot traffic and traffic flow, electricity consumption displays smaller intra-day variation, consistent with electricity demand being relatively rigid.

Differently from the diurnal movement in these series, the hourly temperature follows a more symmetric pattern with a single mid-day peak that, of course, does not vary across week days and weekends, while the level and range vary across seasons.

Taken together, our results highlight both common and variable-specific components in the hour-of-day and day-of-week patterns in variables that can be used to track variation in economic activity. We further show below that there are important differences linked to holidays and seasons of the year.

2.3 Can Granular Data Proxy for Local Economic Activity?

To what extent does our granular data provide a good proxy for local economic activity? The answer to this question is not obvious. On the one hand, our analysis combines conventional measures directly reflective of economic activity such as electricity consumption and credit card spending, with local mobility data such as foot traffic and traffic flows.[9]

On the other hand, economic activity is inherently unobserved at every level of spatio-temporal granularity from the local, hourly to the national, annual level. A well-established literature on national accounts has been successful in constructing proxies of low-frequency, aggregate economic activity such as annual GDP. The same does not hold at the most granular end of the spectrum where there is no observable outcome which we can measure our proxies against.

As a partial answer to the proxy question, we next relate our granular data series to a local measure of economic conditions, namely payroll data. Section 7 contains a more detailed analysis. We start from the relationship between the time series of the first difference of the log of our estimate of local (zip-code level) monthly payroll ($PAY_{zip,m}$) and the first differences of the log of the local series aggregated to the monthly frequency:[10]

$$\Delta PAY_{zip,m} = b_0 + b_{TF}\Delta TF_{zip,m} + b_{FT}\Delta FT_{zip,m} + b_{ElU}\Delta ElU_{zip,m} + b_{SPEND}\Delta SPEND_{zip,m} + \phi\Delta PAY_{zip,m-1} + \nu_{zip,m} + e_{zip,m}.$$

We estimate this panel regression using a mean-grouped estimator with monthly data from June 2020 to December 2023 for the 62 San Diego zip codes in our sample.[11]

We begin by reporting a set of univariate regressions that consider the variables one at a time while controlling for the lagged monthly payroll.[12] The results of these individual regressions are reported in columns (1)-(5) of Table 1. Each of the variables generates a significantly positive coefficient. Months with higher foot traffic, traffic flows, electricity usage and spending in a given zip code are therefore associated with significantly higher payrolls.[13]

The final column of Table 1 reports the results of the multivariate payroll regression in (1) and the associated t-statistics. All coefficient estimates are again positive, although their statistical significance varies, with the coefficient on foot traffic now becoming insignificant

[9]The use of local mobility data is analogous to earlier approaches focused on global economic activity. See, e.g., Kilian (2009) for a prominent example that uses shipping traffic and freight rates to construct a measure of global activity.

[10]We construct monthly zip-level payroll series by multiplying the monthly San Diego-wide payroll data by the payroll fraction of each zip code, updated annually, since zip-code payroll data is only available annually. An autoregressive term is included to control for persistence in the payroll data. Weather data is omitted from the multivariate payroll regression because of its high correlation with electricity usage which leads to multicollinearity.

[11]We start our analysis from June 2020 to avoid the immediate effects of the lock-down caused by the COVID-19 pandemic.

[12]See Appendix B for details on these specifications and additional results.

[13]We find a clear positive association between variables such as foot traffic and spending and this correlation tends to increase when the variables are aggregated to monthly values. This creates an issue with multicollinearity and makes interpretation of the individual coefficient estimates difficult.

due to the issue with multicollinearity mentioned above.

R^2 values vary from 51% for the model that uses temperature data to 73% for the spending regression and 81% for the multivariate model. In total, our granular economic data aggregated to the monthly horizon explains more than 80% of the monthly variation in local payroll growth. Only a small part of this explanatory power comes from the lagged payroll growth, whose coefficient is only 0.18 in the multivariate regression.

We conclude from these results that there is a significantly positive link between the local monthly values of granular variables such as foot traffic, traffic flow, electricity usage and spending on the one hand and payrolls on the other. This helps build a strong case that our granular activity variables can be used to extract a common component related to economic activity.

	Payroll Regression					
	Individual Payroll Regressions					Multivariate
	Foot Traffic	Traffic Flow	Electricity	Spending	Temperature	
b_{FT}	0.357***					0.181***
t-stat	6.892					3.677
b_{TF}		0.296***				0.160***
t-stat		3.301				2.412
b_{Elu}			0.522***			0.338***
t-stat			4.341			2.689
b_{SPEND}				0.818***		0.364**
t-stat				4.196		2.142
b_{Temp}					0.114*	—
t-stat					1.291	—
ϕ	0.338***	0.450***	0.296***	0.431***	0.493***	0.235***
t-stat	8.795	21.511	10.335	18.394	567.441	8.570
$R^2(\%)$	64.02	51.02	67.85	52.99	46.36	74.72
Obs.	2666	2666	2666	2666	2666	2666

Table 1: **Payroll Regressions.** This table reports results of payroll regressions over the sample from June 2020 to December 2023 for the 62 zip codes within San Diego county included in our granular dataset. The t-statistics are computed using clustered-robust standard errors.

3 Granular Economic Modeling

In this section, we introduce our methodology for modeling hourly local economic activity designed to match the features of the data described in the previous section. We start by introducing notations and the main model specification. Next, we detail the specifics of the model for the local economic activity components. Finally, we discuss how we deal with the mixed spatio-temporal scale of our dataset. In particular, we discuss the role of spatio-temporal aggregation restrictions to ensure the coherency of our local economic activity factor across space and time.

3.1 A Dynamic Common Factor Model

We begin by introducing our notations. Let t and h refer to the daily time interval and hours within the day, respectively, so t_h denotes hour h of day t.[14] Further, we use s to refer to a particular spatial location while i refers to individual variables. Finally, we use a superscript τ to denote variables that, even if recorded at time t_h, refer to a given time interval τ. Using these notations, y^{τ}_{s,i,t_h} refers to variable i at location s observed at hour h of day t at the $\tau-$frequency.

The set of individual variables available for each location can vary with time and space and is denoted by N^{τ}_{s,t_h} and stacked into the vector y^{τ}_{s,t_h}. For example, for a given location s such as a zip code, y^{h}_{s,t_h} may contain hourly data on foot traffic (FT^{h}_{s,t_h}), traffic flows (TF^{h}_{s,t_h}), electricity consumption (ElU^{h}_{s,t_h}) and weather (temperature) ($Temp^{h}_{s,t_h}$).

Our empirical analysis in Section 2 suggests the presence of a common factor for the high-frequency variables.[15] Motivated by this observation, consider the following factor model for the granular spatial units in some set $\mathcal{S}_G$:[16]

$$y^{h}_{s,t_h} = \varrho_s + \beta_s f^{h}_{t_h} + \epsilon^{h}_{s,t_h},\ s \in \mathcal{S}_G. \tag{1}$$

Here $f^{h}_{t_h}$ is a latent (unobserved) component that is common across spatial units and β_s is the loading on this common factor for location s, ϱ_s is a spatial fixed effect and ϵ^{h}_{s,t_h} is a location-specific mean-zero component that can be used to capture additional idiosyncratic dynamics including autoregressive components and exogenous regressors.[17]

In addition to the granular spatio-temporal observations modeled in (1), we also consider data on spatially aggregated variables $s^{\mathcal{A}} \in \mathcal{S}_A$. The model for the spatially-aggregated

[14] We drop the hour subscript from t whenever it is not needed.

[15] Early studies such as Stock and Watson (1989), Stock and Watson (2002), and Mariano and Murasawa (2003) extract diffusion indices or common factors from panels of economic data.

[16] Although granularity at the spatial dimension is context specific, a useful narrative to have in mind here is the postal zip-code level.

[17] When more precision is needed, we index the most granular locations by a $\mathcal{G}$ superscript, $s^{\mathcal{G}}$, where $s^{\mathcal{G}} = 1, \ldots, G$, and G is the number of granular locations.

variables is specified analogously to (1):[18]

$$y^h_{s,t_h} = \varrho_s + \beta_s f^h_{t_h} + \epsilon^h_{s,t_h},\ s \in \mathcal{S}_A. \tag{2}$$

The individual components in (2) can be interpreted equivalently to those in (1), except that s and $\mathcal{S}_A$ denote the set of spatially aggregated variables. Let $L = \{G, A\}$ denote the set of spatial aggregation levels.[19] For each $l \in L$, let $\mathcal{S}_l$ denote the set of spatial units at level l, with cardinality $S_l = |\mathcal{S}_l|$. The total number of spatial units is then $S = \sum_{l \in L} S_l$, and the set of all spatial units is given by $\mathcal{S} = \bigcup_{l \in L} \mathcal{S}_l$.

As a first step to deal with the spatio-temporally-mixed data structure of our dataset, we stack the granular and aggregate spatial variables, ordered hierarchically from the aggregate to the granular level, in the observation vector:

$$y^h_{t_h} = \varrho + \beta f^h_{t_h} + \epsilon^h_{t_h},$$

where

$$y^h_{t_h} \equiv \begin{bmatrix} y^h_{\mathcal{S}_A,t_h} \\ y^h_{\mathcal{S}_G,t_h} \end{bmatrix}, \quad \varrho = \begin{bmatrix} \varrho_{\mathcal{S}_A} \\ \varrho_{\mathcal{S}_G} \end{bmatrix}, \quad \beta = \begin{bmatrix} \beta_{\mathcal{S}_A} \\ \beta_{\mathcal{S}_G} \end{bmatrix}, \quad \epsilon^h_{t_h} = \begin{bmatrix} \epsilon^h_{\mathcal{S}_A,t_h} \\ \epsilon^h_{\mathcal{S}_G,t_h} \end{bmatrix}.$$

Most of the spatially aggregated variables are available at frequencies lower than hourly. For example, GDP for San Diego County is available only at the annual frequency.[20] The observation vector at frequency τ is

$$y^\tau_{t_h} \equiv \begin{bmatrix} y^{\tau_{\mathcal{S}_A}}_{\mathcal{S}_A,t_h} \\ y^{\tau_{\mathcal{S}_\mathcal{G}}}_{\mathcal{S}_\mathcal{G},t_h} \end{bmatrix} = \left[y^{\tau_{\mathcal{A}},1}_{s^{\mathcal{A}}_1,t_h}, y^{\tau_{\mathcal{A}},2}_{s^{\mathcal{A}}_2,t_h}, \cdots, y^{\tau_{\mathcal{A}},\mathcal{S}_A}_{s^{\mathcal{A}}_A,t_h}, y^{\tau_{\mathcal{G}},1}_{s^{\mathcal{G}}_1,t_h}, y^{\tau_{\mathcal{G}},2}_{s^{\mathcal{G}}_2,t_h}, \cdots, y^{\tau_{\mathcal{G}},\mathcal{S}_G}_{s^{\mathcal{G}}_G,t_h} \right]',$$

where τ_s denotes the frequency of variable s.[21] We write $y^\tau_{t_h}$ when all variables in the observation equation are recorded at frequency τ and omit the superscript (y_{t_h}) when the variables are related to different frequencies. Hence, our model for the observation vector is given by

$$y_{t_h} = A + BF_{t_h} + \epsilon_{t_h}, \tag{3}$$

where

$$A = \begin{bmatrix} \varrho_1 \\ \vdots \\ \varrho_s \\ \vdots \\ \varrho_S \end{bmatrix}, \quad B = \begin{bmatrix} \beta_1 \\ \vdots \\ \beta_s \\ \vdots \\ \beta_S \end{bmatrix}, \quad \beta_s = \begin{bmatrix} \beta_{s,1} \\ \beta_{s,2} \\ \vdots \\ \beta_{s,K} \end{bmatrix}', \quad F_{t_h} = \begin{bmatrix} f^{yr}_{t_h} \\ \vdots \\ f^h_{t_h} \end{bmatrix}, \quad \epsilon_{t_h} = \begin{bmatrix} \epsilon_{1,t_h} \\ \vdots \\ \epsilon_{s,t_h} \\ \vdots \\ \epsilon_{S,t_h} \end{bmatrix},$$

[18] If the most granular level is the zip code, then spatially aggregated variables can be thought of as measures at the county, MSA, or state level.

[19] For simplicity, we assume only two spatial levels, granular and aggregate, but our setup generalizes to multiple layers of aggregation. For instance, by slicing $\mathcal{S}_A$ into $\mathcal{S}_{A_1}$ and $\mathcal{S}_{A_2}$, we can obtain two distinct levels of spatial aggregates.

[20] To account for this heterogeneity, we consider the lower frequency variables within the hourly observation equation by assigning them to the hour of the day at which they get observed. For example, a variable related to the annual data of year $yyyy$, released at 8 a.m. on February 1st, $yyyy + 1$, would be incorporated into the observation equation as of 8 a.m. on that same date.

[21] Note that τ_A and τ_G are allowed to be different.

are the vector of spatial intercepts (A), the matrix of spatio-temporal factor loadings (B), the vector of unobserved components related to the different frequencies (F_{t_h}), and the vector of disturbances (ϵ_{s,t_h}), respectively. K denotes the number of components of the vector F_{t_h}, with S the total number of locations. We do not make any strong assumptions on ϵ_{t_h} other than it has mean zero. In particular, ϵ_{t_h} can be serially correlated, and we allow it to be heteroskedastic across variables.

Before specifying the model for the unobserved component, it is worth highlighting that our hourly model can accommodate both global and local factors. As an example, suppose we have aggregate and granular factors f^h_{A,t_h} and f^h_{G,t_h}, respectively. Re-writing $f^h_{t_h}$ and the coefficients β_s as

$$f^h_{t_h} \equiv [f^h_{A,t_h}, f^h_{G,t_h}]',$$
$$\beta_s \equiv [\gamma_s, b_s],$$

we obtain the specification

$$\begin{aligned} y^h_{s,t_h} &= \varrho_s + \beta_s f^h_{t_h} + \epsilon^h_{s,t_h} \\ &= \varrho_s + \gamma_s f^h_{A,t_h} + b_s f^h_{G,t_h} + \epsilon^h_{s,t_h},\ s \in \mathcal{S}, \end{aligned}$$

which can be used, e.g., to model hierarchical spatial units.[22]

3.2 Hourly Local Economic Activity

We next specify the dynamics of the latent factor $f^h_{t_h}$, the vector containing our hourly local economic activity components. We treat this as an unobserved component specified as

$$f^h_{t_h+1} = \begin{bmatrix} \mu^h_{t_h+1} \\ \psi^h_{t_h+1} \end{bmatrix},$$

where $\mu^h_{t_h+1}$ is a smoothly evolving mean component and $\psi^h_{t_h+1}$ is a calendar component. We refer to *local economic activity* (LEA) as the sum of these components:

$$LEA^h_{t_h} := \iota_2' f^h_{t_h} = \mu^h_{t_h+1} + \psi^h_{t_h+1}, \tag{4}$$

where $\iota_2 = (1,1)'$.

We specify the dynamics for the mean component as a local linear trend (Harvey (2006)):

$$\begin{aligned} \mu^h_{t_h+1} &= \mu^h_{t_h} + \ell^h_{t_h} + \epsilon^\mu_{t_h+1}, \\ \ell^h_{t_h+1} &= \ell^h_{t_h} + \epsilon^\ell_{t_h+1}, \end{aligned} \tag{5}$$

or, equivalently, in matrix form, as

$$\begin{bmatrix} \mu^h_{t_h+1} \\ \ell^h_{t_h+1} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \mu^h_{t_h} \\ \ell^h_{t_h} \end{bmatrix} + \begin{bmatrix} \epsilon^\mu_{t_h+1} \\ \epsilon^\ell_{t_h+1} \end{bmatrix}.$$

[22] A natural extension of our model would allow the economic activity to have two hierarchical dimensions: space and time: $f_{t_h} = [f^{yr}_{S_A,t_h}, f^{yr}_{S_G,t_h}, \ldots, f^h_{S_A,t_h}, f^h_{S_G,t_h}]'$.

We complete the model by specifying the distributions for the initial values and innovations:

$$\begin{bmatrix}\mu_1^h \\ \ell_1^h\end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix}\overline{\mu}_1 \\ \overline{\ell}_1\end{bmatrix}, \begin{bmatrix}\sigma^2_{\mu_1} & 0 \\ 0 & \sigma^2_{\ell_1}\end{bmatrix}\right) \text{ and } \begin{bmatrix}\epsilon^{\mu}_{t_h+1} \\ \epsilon^{\ell}_{t_h+1}\end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix}0 \\ 0\end{bmatrix}, \begin{bmatrix}\sigma^2_{\epsilon^{\mu}} & 0 \\ 0 & \sigma^2_{\epsilon^{\ell}}\end{bmatrix}\right),$$

where $\overline{\mu}_1$ and $\overline{\ell}_1$ are the means of the initial state.

In the local linear trend model, $\mu^h_{t_h}$ is the slowly-varying mean component, while $\ell^h_{t_h}$ is the trend component. The dynamics of the model is affected by the specification of the variance terms of the unobserved components, $\sigma^2_{\epsilon^{\mu}}$ and $\sigma^2_{\epsilon^{\ell}}$. For example, when $\sigma^2_{\epsilon_\ell}$ is set to zero, the model reduces to a random walk with drift.

The calendar component, $\psi^h_{t_h}$, takes the form

$$\psi^h_{t_h} = HoD_{t_h} + DoW_t + Hol^-_t + Hol^+_t + \varphi_{t_h}, \tag{6}$$

with components given by

$$\begin{aligned}
HoD_{t_h} &= \sum_{\mathrm{day}\in\{WD,WE\}} \sum_{i=1}^{23} \delta^{HoD}_{i,\mathrm{day}} \mathbb{1}\{hour(h) = i, day(t) = \mathrm{day}\}, \\
DoW_t &= \sum_{i=1}^{6} \delta^{DoW}_i \mathbb{1}\{day(t) = day(i)\}, \\
Hol^-_t &= \sum_{i=1}^{HOL^-} \delta^{hol}_i \mathbb{1}\{day(t) = holiday(i)\}, \\
Hol^+_t &= \sum_{i=1}^{Hol^+} \delta^{MH}_i \mathbb{1}\{day(t) = maj.hol.(i)\}, \\
\varphi_{t_h} &= \sum_{k=1}^{K} \left[\delta^{\varphi}_{1,k} \cos\left(\frac{2\pi k t_h}{8670}\right) + \delta^{\varphi}_{2,k} \sin\left(\frac{2\pi k t_h}{8670}\right)\right],
\end{aligned}$$

where HoD_{t_h} is an *hour-of-day* effect that differs across weekdays (WD) and weekends (WE), DoW_t captures a *day-of-week* effect, φ_{t_h} is a smooth component modeled as a Fourier series that can capture seasonal (time-of-year) variations, Hol^-_t captures a *minor holiday* effect, and Hol^+_t captures a *major holiday* effect.[23]

In our model, observations load separately onto the mean and calendar components:

$$\begin{aligned}
y^h_{s,t_h} &= \varrho_s + \beta_s' f^h_{t_h} + \epsilon^h_{s,t_h} \\
&= \varrho_s + \beta_{s,\mu}\mu^h_{t_h} + \beta_{s,\psi}\psi^h_{t_h} + \epsilon^h_{s,t_h}.
\end{aligned}$$

[23]Note that DoW_t, Hol^-_t, and Hol^+_t are indexed by day and thus are the same across hours within a given day. Furthermore, we note that the calendar effects could be allowed to be time-varying by adding a disturbance term in their specifications. In our empirical implementation, we specify as minor holidays all the U.S. calendar holidays with the exception of the Memorial Day, the Independence Day, Thanksgiving, and Christmas, which we specify as major holidays.

This feature is important for potential heterogeneity between the two components.[24]

3.3 Spatio-Temporal Aggregation Restrictions

We next explain how we impose spatio-temporal aggregation restrictions to ensure that the unobserved factor recovered from our heterogeneous dataset is *coherent*.[25] A *coherent factor* can be interpreted as the same object across space and time, although it is extracted from series measured at different frequencies, spatial scales and for different types of data such as weather, traffic flows, and credit card spending.

Consider the following factor structure:

$$f_{t_h} \equiv \begin{bmatrix} f^{yr}_{t_h} \\ f^{h}_{t_h} \end{bmatrix},$$

where $f^{yr}_{t_h}$ and $f^{h}_{t_h}$ denote the yearly and hourly economic activity factors at time t_h, respectively. If the factor structure is coherent, then $f^{yr}_{t_h}$ must be the annual aggregate of $f^{h}_{t_h}$, that is, it must satisfy the following temporal aggregation restriction (TAR):

$$f^{yr}_{t_h} \overset{TAR}{=} \log\left(\sum_{j=0}^{364\times 24} \exp\left(f^{h}_{t_h-j}\right)\right). \tag{7}$$

Hence, the annual economic activity factor must be the (backward) sum of its hourly values during the year. The backward sum means that we only measure economic activity up to time t_h, ensuring that our approach can be implemented in real time based on data available at time t_h.[26]

In our model, we incorporate the temporal aggregation restrictions by specifying economic activity at the highest frequency allowed by the observation vector (hourly) and imposing the aggregation restrictions on the factors extracted from the variables that do not refer to hourly intervals. This approach allows us to handle series observed at any lower frequency, making the model flexible across temporal scales. For a generic time period, $\tau > h$, under the log-transformation that we apply, the corresponding temporal aggregation restriction is given by

$$f^{\tau}_{t_h} = \log\left(\sum_{t_h \in \tau} \exp\left(f^{h}_{t_h}\right)\right).$$

As we show below, under the Cascade Kalman filter that we propose, these restrictions are fully enforced in the computation of the Kalman gain via the prediction error, allowing us to exploit information from variables sampled at different frequencies.

On top of these spatio-temporal restrictions, we impose two further economic identification restrictions using the annual, county-level GDP and monthly, county-level payroll

[24] In fact, as we will later show (Figure 7 and Figure D.2), the empirical analysis finds that the two components play distinct roles.

[25] Recall that our dataset is characterized by three levels of heterogeneity: time, space, and domain.

[26] In Section 6.1 we discuss how to forecast future economic activity at time $t_h + H$, for $H > 0$.

series. First, we anchor the loading of the factor on local log GDP to be unity:

$$y^{yr}_{S_A,t_h} = f^{yr}_{t_h} + \epsilon^{yr}_{S_A,t_h}, \tag{8}$$

where $y^{yr}_{S_A,t_h} = \log(1 + GDP^{yr}_{SD,t_h})$. Next, from the relationship between GDP and payroll (both in logs) at the annual, county-level:

$$\log(1 + GDP^{yr}_{SD,t}) = a + b \log(1 + PAY^{yr}_{SD,t}) + e_t. \tag{9}$$

we derive and impose a restriction on the payroll specification in our model:

$$y^{m}_{S_A,t_h} = \varrho^m + \gamma^m f^m_{t_h} + \epsilon^m_{S_A,t_h},$$

where $y^{m}_{S_A,t_h} = \log(1 + PAY^{m}_{SD,t_h})$. Using annual GDP and payroll data at the county level, we estimate the coefficient γ^m to be $1.39\iota_2$.[27] To summarize, our restrictions that link our economic factor to the observable state-wide GDP and payroll figures are $\gamma^{GDP} = 1$, $\gamma^{PAY} = 1.39$, and $\mu_0 = \gamma^{PAY} y^m_{S_A,1}$.[28]

In summary, our model for the hourly local economic activity under temporal aggregation restrictions takes the form:

$$\begin{aligned}
y_{t_h} &= A + B f_{t_h} + \epsilon_{t_h}, \\
f_{t_h} &= [f^{yr}_{t_h}, f^{m}_{t_h}, \ldots, f^{h}_{t_h}]', \\
f^{\tau}_{t_h} &= \log\left(\sum_{t_h \in \tau} \exp\left(f^h_{t_h}\right)\right), \quad h < \tau, \\
f^h_{t_h+1} &= [\mu^h_{t_h+1}, \psi^h_{t_h+1}]', \\
\mu^h_{t_h+1} &= \mu^h_{t_h} + \ell^h_{t_h} + \epsilon^{\mu}_{t_h+1}, \\
\ell^h_{t_h+1} &= \ell^h_{t_h} + \epsilon^{\ell}_{t_h+1}, \\
\psi^h_{t_h} &= HoD_{t_h} + DoW_t + h_t + H_t + \varphi_{t_h}.
\end{aligned} \tag{10}$$

The first line in (10) is the observation equation followed by the vector of factors at the different frequencies, the temporal aggregation restrictions governing the factor dynamics, the hourly economic activity factor, and finally the three equations specifying the dynamics of its underlying components. We refer to these restrictions as *coherence equations*.[29]

4 Estimation

In this section, we introduce the estimation methodology for our local economic activity model. Specifically, we propose a two-step methodology that consists of: i) spatial pooling and a sequential hierarchical Kalman filter for recovering the latent factor; ii) a second-stage least squares step to recover spatial heterogeneity.

[27] We also use this relationship to initialize the latent factor in the Kalman filter.

[28] See Appendix C.1-C.2 for additional details on the derivations of these restrictions.

[29] Note that although f_{t_h} contains multiple components (such as $f_{t_h} = [f^{yr}_{t_h}, f^h_{t_h}]'$), its dynamics is entirely determined by $f^h_{t_h}$.

4.1 Step I: Spatial Pooling and the Cascade Kalman Filter

4.1.1 Spatial Pooling

The first step in the estimation of our model consists of pooling the variables to the common (aggregate) spatial level. To understand why spatial pooling as a first step is necessary, consider the example of Section 3.1 with two spatial levels. Working at the granular spatio-temporal level requires us to work with a high-dimensional vector whose T and S dimensions rapidly grow large. For example, each time series measured at the hourly frequency contains approximately 9,000 ($365\times$ 24) data points per year. This dimension has to be multiplied by the number of spatial locations (S) and by the number of series available for each location (N). In total, this entails working with approximately $9000 \times S \times N$ data points (per year). For 100 spatial location and 5 series for each location, this is equivalent to more than 4 million data points which leads to computational challenges.

The first step of our estimation approach therefore involves reducing the spatial dimension by pooling observations across s. Specifically, for each point in time, t_h, and assuming that $\bigcup_{s\in\mathcal{S}_G} = S_A$, we aggregate the observations spatially to a common aggregate level:[30]

$$y^h_{S_A,t_h} = \sum_{s\in S_\mathcal{G}} y^h_{s,t_h}.$$

Following this spatial pooling step, the observation equation is homogeneous in the spatial dimension, i.e., all its elements are observed at the same spatial level. Next, we arrange the variables in the observation vector in decreasing order starting from the lowest frequency (annual) and ending with the highest frequency (hourly):

$$\begin{bmatrix} y^{yr}_{S_A,t_h} \\ \vdots \\ y^h_{S_A,t_h} \end{bmatrix} = \begin{bmatrix} \varrho^{yr} \\ \vdots \\ \varrho^h \end{bmatrix} + \begin{bmatrix} \beta^{yr} & 0_\tau & \cdots & 0_h \\ \vdots & \vdots & \vdots & \vdots \\ 0_{yr} & 0_\tau & \cdots & \beta^h \end{bmatrix} \begin{bmatrix} f^{yr}_{t_h} \\ \vdots \\ f^h_{t_h} \end{bmatrix} + \begin{bmatrix} \epsilon^{yr}_{S_A,t_h} \\ \vdots \\ \epsilon^h_{S_A,t_h} \end{bmatrix}, \tag{11}$$

where 0_τ denotes the zero vector of the same dimension as the factor at frequency τ.

For each frequency, τ, the vector $y^\tau_{S_A,t_h}$ in (11) stacks all observations for the spatial aggregate S_A at the same frequency τ:

$$y^\tau_{S_A,t_h} = \varrho^\tau_{S_A} + \beta^\tau_{S_A} f^\tau_{t_h} + \epsilon^\tau_{S_A,t_h}, \tag{12}$$

where $f^\tau_{t_h} = \log(\sum_{t_h\in\tau} \exp(f^h_{t_h}))$ from the temporal aggregation restriction. An option is to conduct the filtering on the full system in (11). However, because this is specified jointly across all frequencies, τ, this is a large-dimensional problem. To deal with this challenge, we filter each block $y^\tau_{S_A,t_h}$ sequentially, starting from the top (most aggregate) level and moving down to the granular, hourly, frequency as we next explain.

[30] An advantage from spatial pooling is that the aggregate series is more outlier robust, i.e., less affected by outliers that impact the individual zip-level data more frequently but averages out in the aggregate.

4.1.2 Cascade Kalman Filter

We begin by casting our hourly model for local economic activity in state space form. The model uses the equations in (10) that govern the economic activity factor along with the observation equation:

$$y^h_{t_h} = \varrho^h + \beta^h_\mu \mu^h_{t_h} + \beta^h_\psi \psi^h_{t_h} + \epsilon^h_{t_h}.$$

To facilitate the analysis, we define the state vector $\alpha^h_{t_h+1} \equiv [\mu^h_{t_h+1}, \ell^h_{t_h+1}, \psi^h_{t_h+1}]'$. We first consider the case where the observation vector contains *all* hourly variables. For this case, we have the following linear time-invariant state space representation:

$$y^h_{t_h} = \varrho^h + Z^h \alpha^h_{t_h} + \epsilon^h_{t_h}, \quad \epsilon^h_{t_h} \sim \mathrm{N}(0, H^h), \tag{13}$$

$$\alpha^h_{t_h+1} = T\alpha^h_{t_h} + \eta^h_{t_h+1}, \quad \eta^h_{t_h+1} \sim \mathrm{N}(0, Q), \tag{14}$$

where

$$\varrho^h = \begin{bmatrix} \varrho^h_1 \\ \vdots \\ \varrho^h_{N^h} \end{bmatrix}, \quad Z^h = \begin{bmatrix} \beta^h_\mu & 0 & \beta^h_\psi \end{bmatrix}, \quad H^h = Var(\epsilon^h_{t_h}),$$

$$T = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \eta^h_{t_h+1} = \begin{bmatrix} \epsilon^\mu_{t_h+1} \\ \epsilon^\ell_{t_h+1} \\ 0 \end{bmatrix}, \quad Q = \begin{bmatrix} \sigma^2_{e^\mu} & 0 & 0 \\ 0 & \sigma^2_{e^\ell} & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

The equations summarizing the Kalman filter recursions used to evaluate the mean of the hourly state vector $\alpha^h_{t_h+1}$, conditional on the observations $\mathcal{Y}^h_{t_h}$, are now given by

$$v^h_{t_h} = y^h_{t_h} - \varrho^h - Z^h \alpha^h_{t_h}, \quad F^h_{t_h} = Z^h P^h_{t_h} Z^{h\prime} + H^h, \quad K^h_{t_h} = P^h_{t_h} Z^{h\prime},$$

$$\alpha^h_{t_h+1} = T(\alpha^h_{t_h} + K^h_{t_h} {F^h}^{-1}_{t_h} v^h_{t_h}), \quad P^h_{t_h+1} = T(P^h_{t_h} - K^h_{t_h} {F^h}^{-1}_{t_h} {K^h}'_{t_h})T' + Q,$$

where $F^h_{t_h} = Var(y^h_{t_h}|\mathcal{Y}^h_{t_h-1}) = Var(v^h_{t_h})$, and $K^h_{t_h} = Cov(\alpha^h_{t_h}, y^h_{t_h}|\mathcal{Y}^h_{t_h-1})$, and $\mathcal{Y}^h_{t_h}$ is the set of (hourly) observations up to time t_h.

Next, building on the hierarchical structure of y_{t_h}, we follow the same logic of the hourly state space in equations (13)-(14) and construct, for each block τ, the state space:

$$y^\tau_{t_h} = \varrho^\tau + Z^\tau \alpha^\tau_{t_h} + \epsilon^\tau_{t_h}, \quad \epsilon^\tau_{t_h} \sim \mathrm{N}(0, H^\tau),$$

$$\alpha^\tau_{t_h} = \log\left(\sum_{t_h \in \tau} \exp\left(\alpha^h_{t_h}\right)\right). \tag{15}$$

The system matrices for the $\tau-$block are

$$Z^\tau = \begin{bmatrix} \beta^\tau_\mu & 0 & \beta^\tau_\psi \end{bmatrix}, \quad H^\tau = Var(\epsilon^\tau_{t_h}), \tag{16}$$

with corresponding prediction error and updating matrices given by

$$v_{t_h}^{\tau} = y_{t_h}^{\tau} - \varrho^{\tau} - Z^{\tau}\alpha_{t_h}^{\tau}, \quad F_{t_h}^{\tau} = Z^{\tau}P_{t_h}^{\tau}Z^{\tau\prime} + H^{\tau}, \quad K_{t_h}^{\tau} = P_{t_h}^{\tau}Z^{\tau\prime}. \tag{17}$$

The filtering equations for the state recursion can be written as

$$\underbrace{\alpha_{t_h}^{*,\tau-1}}_{\text{updated estimate}} = \underbrace{\alpha_{t_h}^{\tau-1}}_{\text{TAR}} + \underbrace{K_{t_h}^{\tau}(F_{t_h}^{\tau})^{-1}v_{t_h}^{\tau}}_{\text{updating term}},$$

and, for the update to the state variance,

$$P_{t_h}^{\tau-1} = P_{t_h}^{\tau} - K_{t_h}^{\tau}(F_{t_h}^{\tau})^{-1}K_{t_h}^{\tau\prime},$$

The transition from time t_h to time $t_h + 1$ is now achieved by the relations

$$\alpha_{t_h+1}^{h} = T(\alpha_{t_h}^{h} + K_{t_h}^{h}(F_{t_h}^{h})^{-1}v_{t_h}^{h}), \quad P_{t_h+1}^{h} = T(P_{t_h}^{h} - K_{t_h}^{h}(F_{t_h}^{h})^{-1}K_{t_h}^{h\prime})T' + Q.$$

Starting from the highest level of aggregation (the lowest frequency, $\max(\tau)$) available at time t_h, we construct a (yearly) proxy for $\alpha_{t_h}^{\max(\tau)}$ using the estimates of $\alpha_{t_h}^{h}$ and the temporal aggregation restriction. We use this proxy to evaluate the state space of the block $\max(\tau)$ to obtain a filtered state vector $\widehat{\alpha}_{t_h}^{\max(\tau)}$ that uses information contained in the series at the $\max(\tau)$ frequency (annual). The filtered estimate of $\widehat{\alpha}_{t_h}^{\max(\tau)}$ is next used in the filtering equations to obtain an estimate of the state vector of the subsequent hierarchical block, $\widehat{\alpha}_{t_h}^{\max(\tau)-1}$, which we use to evaluate the state space of the $\max(\tau) - 1$ block (quarterly). We follow the same strategy down to the lowest level of the temporal hierarchy, i.e., the hourly block h. In this last step, we produce the final estimate of the hourly economic activity $\widehat{\alpha}_{t_h}^{h}$ at time t_h. Using the transition equations from time t_h to time $t_h + 1$, we obtain $\widehat{\alpha}_{t_h+1}^{h}$, and we start again from $\max(\tau)$ corresponding to the block at time t_{h+1}.

This methodology allows us to incorporate *all* the information from the various hierarchical levels available at t_h. When no information other than hourly is available, our approach simply proceeds to filter $\alpha_{t_h}^{h}$ using the information in the hourly block. However, when information from lower frequency series is available, we take advantage of this through the cascading recursions and the temporal aggregation restrictions. This ability to leverage simultaneously information available at several layers of aggregation is one of the key strengths of our approach.[31] Moreover, by imposing the identification restrictions discussed in Section 3.3, which anchors the loadings of the factors on the aggregated economic series (GDP and payroll), we can identify the latent factor as the underlying driver of economic activity.[32]

[31] An alternative estimation strategy would be to first obtain an estimate of the calendar component, use it to obtain a calendar-adjusted observation vector, and finally extract the unobserved mean component from the adjusted series. However, as Wallis (1974) noted, adjusting the data series for calendar components can lead to distortions in the results.

[32] We refer to Appendix C.1 for a discussion of the identification restrictions.

4.1.3 Fast Filtering

The granularity of our dataset introduces challenges due to the large dimensions involved. After spatial pooling of the variables, we still have two dimensions: time, evolving daily, and the cross-section—the number of variables examined in the layers of the hierarchy.

The time series implies a large number of recursions to compute the likelihood of the model, while the multidimensional series involve computing inverse matrices at each iteration, each with the same dimension as the τ block. To avoid matrix inversions, we adopt the approach of Koopman and Durbin (2000) and filter the system equation by equation within each block. This greatly simplifies the scale of the problem as it allows us to deal with scalars instead of matrices.

We extend this approach to our framework in the following way. Consider the τ block, with observation vector given by $y^{\tau}_{t_h} = [y^{\tau}_{1,t_h}, \ldots, y^{\tau}_{N^{\tau},t_h}]'$, where N^{τ} denotes the number of series available at time t_h in the τ block.[33] By considering each element in $y^{\tau}_{t_h}$ individually, we can re-write the state space for each series as

$$y^{\tau}_{i,t_h} = \varrho^{\tau}_i + Z^{\tau}_i \alpha^{\tau}_{t_h} + \epsilon^{\tau}_{i,t_h}, \quad \epsilon^{\tau}_{i,t_h} \sim \mathrm{N}(0, h^{\tau}_i)$$
$$\alpha^{\tau}_{t_h} = \log\left(\sum_{t_h \in \tau} \exp\left(\alpha^{h}_{t_h}\right)\right).$$

The system matrices take the form

$$Z^{\tau}_i = \begin{bmatrix} \beta^{\tau}_{i,\mu} & 0 & \beta^{\tau}_{i,\psi} \end{bmatrix}, \quad h^{\tau}_i = Var(\epsilon^{\tau}_{i,t_h}), \tag{18}$$

while the prediction error and the updating matrices are given by

$$v^{\tau}_{i,t_h} = y^{\tau}_{i,t_h} - \varrho^{\tau}_i - Z^{\tau}_i \alpha^{\tau}_{t_h}, \quad F^{\tau}_{i,t_h} = Z^{\tau}_i P^{\tau}_{i,t_h} Z^{\tau\prime}_i + h^{\tau}_i, \quad K^{\tau}_{i,t_h} = P^{\tau}_{i,t_h} Z^{\tau\prime}_i. \tag{19}$$

It is immediately evident where the computational efficiency with respect to the system in (16)-(17) is gained. The filtering equations in (18)-(19) are univariate equations, while the filtering equations in (16)-(17) involve $N^{\tau} \times N^{\tau}$ dimensional matrices.

In this case the updating equation for the state vector $\alpha^{\tau}_{t_h}$ and its conditional variance $P^{\tau}_{t_h}$ are sequentially updated through the following recursions

$$\alpha^{\tau}_{i+1,t_h} = \alpha^{\tau}_{i,t_h} + K^{\tau}_{i,t_h}(F^{\tau}_{i,t_h})^{-1} v^{\tau}_{i,t_h}, \quad P^{\tau}_{i+1,t_h} = P^{\tau}_{i,t_h} - K^{\tau}_{i,t_h}(F^{\tau}_{i,t_h})^{-1} K^{\tau}_{i,t_h}{}'.$$

Under mild conditions, Koopman and Durbin (2000) show that this procedure makes use of the full information contained in the dataset.

4.2 Step II: Recovering the Granular Loadings

The second step of our approach recovers the spatial heterogeneity among the variables considered in the system. From the first step, we obtain a filtered estimate, $\widehat{\alpha}^{h}_{t_h}$, which

[33] Because our first step consists of spatially aggregating the series, we drop S_A from the subscript.

is the hourly state vector common to the spatial units. To recover spatial heterogeneity, recall that our granular model is defined in equation (1) as

$$y^h_{s,t_h} = \varrho_s + \beta_s f^h_{t_h} + \epsilon^h_{s,t_h},\ s \in \mathcal{S}_{\mathcal{G}}.$$

From the first step of our methodology, we obtain a factor estimate, $\widehat{f}^h_{t_h}$. Treating this as a proxy for $f^h_{t_h}$, we estimate the following regressions of the granular series:

$$y^h_{s,i,t_h} = \varrho_{s,i} + \beta_{s,i}\ \widehat{f}^h_{t_h} + \epsilon^h_{s,i,t_h},\ s \in \mathcal{S}_{\mathcal{G}}, \tag{20}$$

from which we obtain the spatial intercepts and loadings, $\varrho_{s,i}$ and $\beta_{s,i}$, respectively. The coefficients $\beta_{s,i}$ capture spatial heterogeneity and allow each of the series (e.g., traffic flows) to have different loadings on the local economic activity factor.[34] We refer to $\beta_{s,i}$ as *granular loadings*.

5 Empirical Results

We begin this section by explaining how we implement the granular economic model empirically. We then present the results in the following order: time-series, spatial, and spatio-temporal. Next, we develop a radar-like economic activity tracker. We conclude the section comparing the results from our model with the results from a surrogate proxy, which we construct using the hourly series.

5.1 Modeling Granular Economies

We consider San Diego county as the spatial aggregate, S_A, and zip codes within San Diego as the granular locations S_G. To enhance clarity, we write SD in place of S_A to refer to the spatial aggregate, and zip_j instead of s to refer to the granular, zip-code locations. Specifically, we consider 62 San Diego zip codes over the period January 1^{st} 2021 to December 31^{st} 2023, for a total of approximately 26,280 hourly time points and 1,629,422 panel observations. When we work with the full spatial panel, we have a total of approximately 86,359,366 entries.

Consider a zip code within San Diego County, zip_j. The hourly-zip_j observation vector contains the hourly series observed at time t_h for zip_j, and includes traffic flows ($TF^h_{zip_j,t_h}$), foot traffic ($FT^h_{zip_j,t_h}$), electricity usage ($ElU^h_{zip_j,t_h}$), and temperature ($Temp^h_{zip_j,t_h}$). Merging these, the hourly observation vector for zip code j is given by:

$$y^h_{zip_j,t_h} = \left[TF^h_{zip_j,t_h}, FT^h_{zip_j,t_h}, ElU^h_{zip_j,t_h}, Temp^h_{zip_j,t_h}\right]'.$$

By stacking all the zip-code observation vectors in an encompassing vector, $y^h_{zip_{SD},t_h}$, we obtain the hourly observation vector:

$$y^h_{zip_{SD},t_h} = \left[y^h_{zip_1,t_h}, y^h_{zip_2,t_h}, \cdots, y^h_{zip_{S_G},t_h}\right]',$$

[34] This modeling feature allows series such as electricity consumption in a zip code to have a different sensitivity to the LEA factor than the sensitivity of traffic flow.

which contains all the granular observations at the hourly level.

Next, we consider the following variables measured at a granular spatial level but available only at lower frequencies: daily zip-level spending ($Spend^{d}_{zip_j,t_h}$), and monthly zip-code level total employee pay ($PAY^{m}_{zip_j,t_h}$). Finally, we consider yearly-*county* level GDP (GDP^{yr}_{SD,t_h}). The GDP and payroll series serve as anchoring equations that allow us to interpret the unobserved factor as an economic activity factor, see Section 3.3.

Following the hierarchical approach of Section 3, we stack the observation equations at the different levels of aggregation—county and zip codes—obtaining the San Diego economy observation vector:[35]

$$y_{SD,t_h} = \beta_{SD} f_{SD,t_h} + \epsilon_{SD,t_h},$$

where

$$y_{SD,t_h} = \begin{bmatrix} GDP^{yr}_{SD,t_h} \\ y_{zip_{SD},t_h} \end{bmatrix}. \tag{21}$$

The observation equation in (21) represents the empirical counterpart of equation (10).

5.2 Time-Series Dynamics

Beginning with intra-day patterns, Figure 4 displays hourly patterns in local economic activity. Specifically, the left panel plots hourly values of our local economic activity factor $f^{h}_{t_h}$ averaged separately across weekdays and weekends. We show results both in logs (left axes), the unit adopted by our model, and in percentage deviation from the average (right axes).

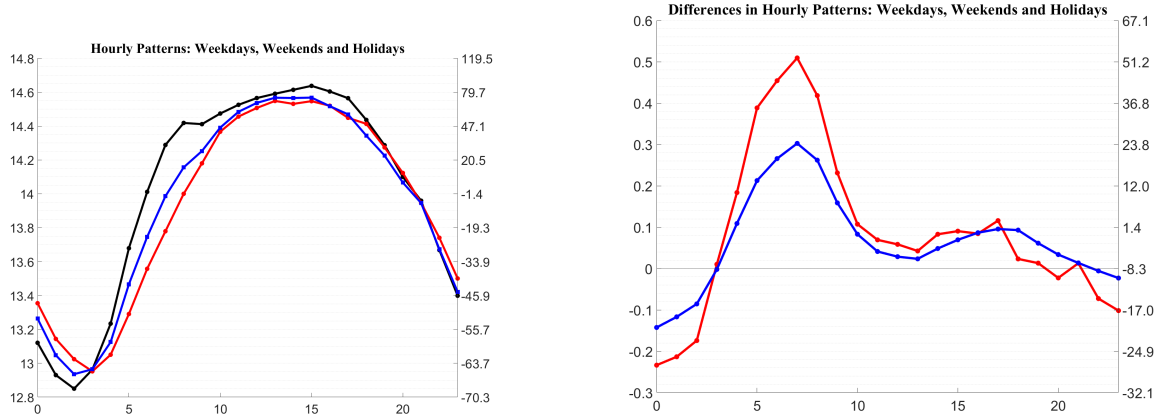
Economic activity on weekdays starts out very low at midnight, falling further until 2 a.m. which is the nadir for economic activity on weekdays. From that point, average economic activity rises sharply until 7 a.m.. Between 7 a.m. and noon, economic activity continues to rise, though at a slower rate and it remains persistently high until 5 p.m. after which it declines systematically until midnight.

During weekends, economic activity follows a broadly similar shape, although it starts from a higher level at midnight and remains higher until 3 a.m. (consistent with people staying out longer during weekends). During weekends, local economic activity also rises at a slower pace during the early morning hours as people start the day at a slower pace.

The right panel in Figure 4 plots the difference between the hour-of-day pattern in local economic activity for weekdays versus weekends. Between midnight and 3 a.m., economic activity is notably higher on weekends while conversely, weekdays register the highest economic activity between 4 a.m. and 10 a.m.. Differences in economic activity are relatively small during the late morning and afternoon, while economic activity tends to stay higher in the late evening on weekends. Figure 4 also shows that the hourly pattern of holidays behave similarly to that observed during weekends.

[35] Note that we drop the superscript h, as the observation vector contains series at different frequencies.

Turning to differences in economic activity patterns across different days of the week, Figure 5 plots hourly values of the local economic activity factor for San Diego over the seven-day period from Monday March 18^{th} 2024 00:00 to Sunday March 24^{th} 2024 at 23:59. The solid line tracks the sum $LEA^h_{t_h} = \mu^h_{t_h} + \psi^h_{t_h}$, while the broken line plots the mean component $\mu^h_{t_h}$. The vast majority of the variation in the factor $f^h_{t_h}$ is picked up by the local linear trend component $(\mu^h_{t_h})$ with the calendar component $(\psi^h_{t_h})$ playing a smaller role. This suggests that the dynamic stochastic factor component is capable of quickly adapting to diurnal patterns both within and across days. Economic activity is more muted on weekends than during week days, but there are also interesting differences between individual week days with Wednesdays being more bi-modal than other days and Fridays showing a sharper peak in the middle of the day.[36]



(a) **Hour-of-Day Patterns.**

(b) **Hour-of-Day Difference.**

Figure 4: **Hour-of-Day Dynamics.** This figure plots the hour-of-day dynamics across weekdays, weekends and holidays. **Left panel:** the black line represents weekdays, the red line represents weekends, the blue line holidays. **Right panel:** the red line represents weekday-weekend difference, the blue line represents weekday-holiday difference. In both plots, the left axis reports the value of the factor in log units, whereas the right axis reports values in percentage deviations from the mean.

[36] We map the factor to percentage units by applying the transformation $f^h_{t_h}(\Delta\%) = (e^{f^h_{t_h} - \bar{f}^h} - 1) \times 100$, and we add the notation $(\Delta\%)$ when denoting economic activity expressed in these units.

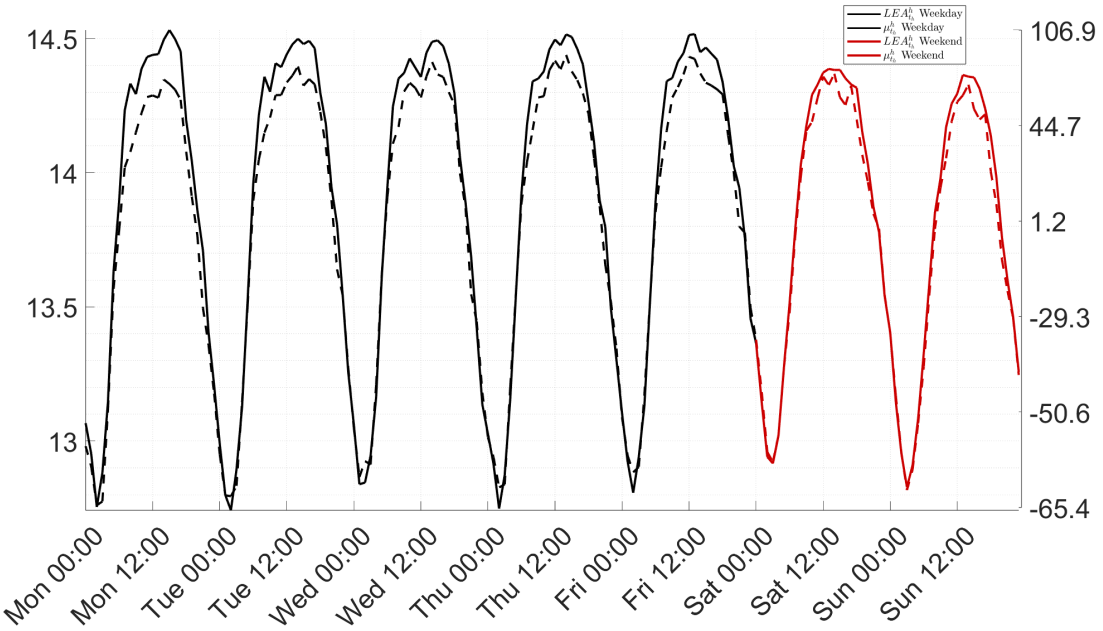


Figure 5: **Hourly Local Economic Activity Factor.** This figure plots the San Diego hourly local economic activity factor. The solid line tracks $LEA^h_{t_h} = \mu^h_{t_h} + \psi^h_{t_h}$. The broken line plots $\mu^h_{t_h}$. The plot starts on Monday November 13^{th} 2023 at 00:00 and ends on Monday November 20^{th} 2023 at 00:00. The left axis reports values of the factor in log units, whereas the right axis reports values in percentage deviations from the mean.

Expressed in percentage terms, the intra-daily variation in local economic activity is very large. Measured relative to its average level, the hourly local economic activity factor drops by two-thirds (-65%) at night during most week days, and is close to 100% higher at its peak on Monday and Friday afternoon.

Figure 6 plots daily values of our LEA factor over our full sample 1/1/2021 - 12/31/2023. The figure illustrates a strong weekly cycle along with a seasonal cycle as local economic activity is notably higher during the summer months (June - August) than during the winter months (December-February). We also see notable drops in economic activity in the immediate aftermath of holidays.

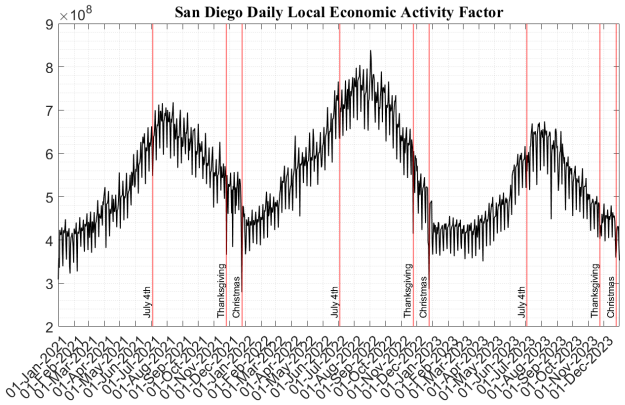


Figure 6: **Daily Local Economic Activity Factor.** This figure plots the daily economic activity (GDP) factor measured in daily nominal dollar terms for the period from January 1^{st} 2021 to December 31^{st} 2023. The vertical red lines indicate the three major U.S. holidays: Independence Day (July 4th), Thanksgiving (floating date), and Christmas Day (December 25th).

5.3 Spatial and Spatio-Temporal Heterogeneity

The central assumption underpinning our modeling approach is that each of the underlying granular series (e.g., traffic flow, foot traffic, electricity consumption, and spending) in a given location can be decomposed into its factor exposure times a common economic activity factor plus an idiosyncratic term that is largely variable-specific.

How a given series splits into these two components is informative about how closely that variable is linked with economic activity: the larger the factor loading and the smaller the variance of the idiosyncratic term, the more important the common factor will be for the dynamics of a given series.

We begin our analysis of the importance of the common economic factor by examining the factor loadings for the individual series on the unobserved mean component of the local economic activity factor, $\mu^h_{t_h}$, denoted by $\widehat{\beta}_{s,\mu,FT}$, $\widehat{\beta}_{s,\mu,TF}$, and $\widehat{\beta}_{s,\mu,ElU}$, respectively.[37] These estimates are recovered using the second-stage step (equation (20)).

Figure 7 shows the histograms of loadings of foot traffic, traffic flow, and electricity consumption on our LEA factor across the 62 zip codes. There is significant heterogeneity in the factor loadings, both across zip codes (within the same panel) as well as across the series (across different plots). Foot traffic and traffic flow display the largest magnitude in the $\beta_{s,\mu,\cdot}$ coefficients, which range between [1.19,4.75] and [1.13,2.56], respectively. The range of factor loadings on electricity usage is narrower ([-0.36, 0.48]). All loading estimates are positive except for a single loading on electricity, which reflects data from a zip code with a major government facility.

[37] In Appendix D.2 (Figure D.2), we report the loadings of the calendar component.

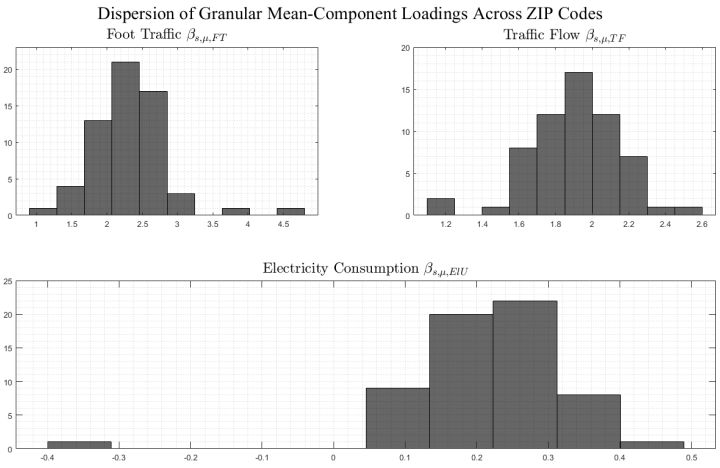
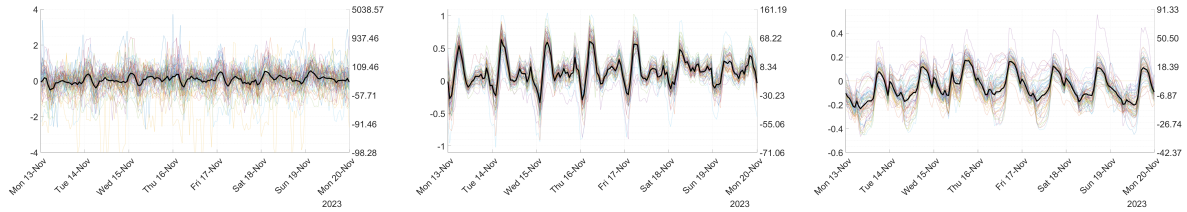


Figure 7: **Factor Loadings on the Granular Mean-Component** This figure plots histograms of granular betas on the mean component of the hourly LEA factor, $\beta_{s,\mu,i}$, for hourly foot traffic (upper left plot), hourly traffic flow (upper right panel), and hourly electricity consumption (bottom panel) for the 62 San Diego zip codes.

Given the substantial heterogeneity in factor loadings displayed in Figure 7, we might also expect to find a high degree of variation in the magnitude of the idiosyncratic components, $\epsilon_{i,s,t}$. Figure 8 confirms that this is indeed the case. The figure plots the residuals of the hourly series for foot traffic, traffic flow, and electricity usage for all zip codes after accounting for the effect of the common factor.[38]

Three features stand out from these plots: first, as demonstrated by the difference in ranges, the magnitude of the idiosyncratic term varies significantly across the three series, being largest for foot traffic and much smaller for electricity consumption, with traffic flow in the middle. Second, there is substantial variation in the idiosyncratic terms across the individual zip codes as indicated by the span between the 62 individual series. Third, there is clearly some persistent dynamics in the idiosyncratic components but the patterns are specific to foot traffic, traffic flow and electricity and so are not attributed to the common factor, consistent with most of the variation in these residuals being idiosyncratic. In fact, we find that the average AR(1) coefficients for the zip-code level residuals' series is around 0.5 for foot traffic, 0.75 for traffic flow and 0.9 for electricity consumption.

[38]Specifically, for each hourly series in the 62 San Diego zip codes we construct fitted values as $\widehat{y}^h_{s,i,t_h} = \widehat{\varrho}_{s,i} + \widehat{\beta}_{s,i}\widehat{f}^h_{t_h}$ and compute residuals as $\widehat{e}^h_{s,i,t_h} = y^h_{s,i,t_h} - \widehat{y}^h_{s,i,t_h}$.



(a) **Foot Traffic** (b) **Traffic Flow.** (c) **Electricity Consumption.**

Figure 8: **Model-Implied Residuals.** This figure plots the hourly residuals implied by the GEM for the logs of foot traffic, traffic flow, and electricity consumption at the aggregate San Diego level from midnight, Monday November 13^{th} 2023 to midnight of Monday, November 20^{th} 2023. The solid black line plots the mean across the 62 zip codes, plotted in the shaded lines. The left axis reports values of the factor in log units, whereas the right axis reports values in percentage deviations from the mean.

5.4 San Diego Spatio-Temporal Activity Radar

Our granular economic modeling approach enables us to generate a radar-like visualization of local economic conditions, which we refer to as the *spatio-temporal activity radar* (STAR).

We start from the conditional expectations of the series implied by our granular model:

$$\widehat{y}^h_{s,i,t_h} = \widehat{\varrho}_{s,i} + \widehat{\beta}_{s,i}\widehat{f}^h_{t_h}. \tag{22}$$

To obtain an index for the local economic activity importance ($LEAI$) across zip codes, we use the weighted GEM-implied conditional expectations from equation (22) normalized by their sum across all zip codes:[39]

$$LEAI_{s,t_h} = \frac{\sum_i \omega_i \widehat{y}^\tau_{s,i,t_h}}{\sum_s \sum_i \omega_i \widehat{y}^\tau_{s,i,t_h}},$$

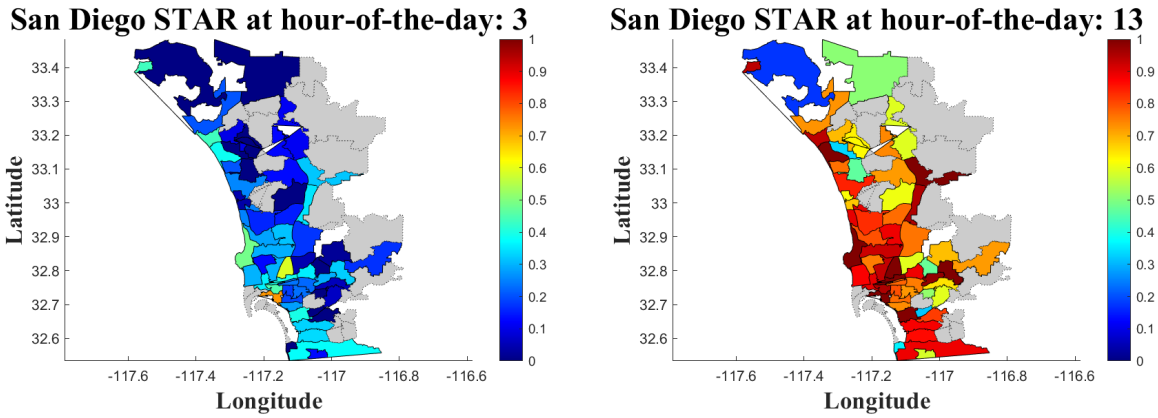
Higher values of $LEAI_{s,t_h}$ indicate relative stronger economic activity in zip code s at time t_h. We construct the STAR as the spatial plot of $LEAI_{s,t_h}$ and display it for two different hour-of-day averages in Figure 9.[40]

Comparing the activity radars for the daytime (1 p.m.) and nighttime (3 a.m.) periods in Figure 9, we note distinctly higher levels of economic activity during the day, with areas near the coast and Mexican border showing particularly high LEA values during the day and inland areas generally being less active. At night, most zip codes show very low levels of economic activity. One area lights up sharply, namely the zip code where San Diego International Airport is located, which remains active during much of the night. Other areas with high nighttime economic activity include a zip code (92037) with a large educational institution (UCSD) and large medical facilities which operate

[39] We use equal weights for ω_i here, leaving more general specifications for future research.

[40] The radar in Figure 9 is constructed using the time-series average of the local economic activity factor, $\widehat{f}^h_{t_h}$, computed over the estimation sample for a given hour.

throughout the night. The figure shows that economic activity fluctuates notably both in the time and spatial dimensions, with several areas switching intensity at different times of the day. This is consistent with the patterns shown in Figures 5-4, which documented high excursion in economic activity across the hours of the day.



(a) **STAR at 3 a.m.** (b) **STAR at 1 p.m.**

Figure 9: **Spatio-Temporal Economic Radar Hour-of-Day Variation.** This figure plots the spatio-temporal activity radar for San Diego county at different hour-of-day. The left panel plots the STAR at 3 a.m., the right panel plots the STAR at 1 p.m. Grey areas are zip codes within San Diego county excluded due to incomplete data.

The STAR can shed light on local economic activity dynamics across zip codes. To illustrate this, suppose we classify the zip codes as predominantly commercial versus predominately residential using their share of commercial and residential electricity usages:[41]

$$\frac{ElU_{zip_j}^{\text{Commercial}}}{ElU_{zip_j}^{\text{Commercial}} + ElU_{zip_j}^{\text{Residential}}}.$$

Using this measure, we select the top and bottom 15 zip codes and compute the following quantity:[42]

$$\sum_{\{\text{Fr,Sat}\}} \left(\sum_{t_h=5\text{p.m.}}^{10\text{p.m.}} \widehat{y}_{s,i,t_h} - \sum_{t_h=9\text{a.m.}}^{3\text{p.m.}} \widehat{y}_{s,i,t_h} \right) - \sum_{\{\text{Sun}-\text{Thur}\}} \left(\sum_{t_h=5\text{p.m.}}^{10\text{p.m.}} \widehat{y}_{s,i,t_h} - \sum_{t_h=9\text{a.m.}}^{3\text{p.m.}} \widehat{y}_{s,i,t_h} \right),$$

which we interpret as an entertainment effect. Intuitively, this quantity captures the additional economic activity in zip codes during the evening hours of Fridays and Saturdays (5 p.m. to 10 p.m.) relative to both daytime hours and the corresponding weekday baseline akin to constructing a diff-in-diff estimator. This measure is intended to isolate activity primarily associated with leisure and entertainment.

We expect the entertainment effect to be stronger for commercial zip codes than for residential ones. To verify whether this is the case, we construct the STAR of the entertainment effect by using the entertainment effect in the numerator of the $LEAI_{s,t_h}$, and

[41] Other methods used to identify zip codes as mostly residential versus mostly commercial yield a very similar ordering.

[42] We exclude 4 p.m. to create a discontinuity in time.

as denominator its sum over the 15 commercial zip codes and 15 residential zip codes. We report the STAR obtained in this way for the 15 commercial zip codes and the 15 residential zip codes in Figure 10.

Zip codes identified as commercial (left plot) display on average stronger entertainment effects than those identified as residential (right plot). The zip code that displays the strongest entertainment effect (92101) includes San Diego International Airport, as well as districts with many bars and restaurants (such as the Gaslamp Quarter, which is one of the major food hubs of Downtown San Diego). A strong entertainment effect among the residential zip codes is reported for zip code 91932 (Imperial Beach), which even if it is classified as residential, also contains a large number of restaurants.

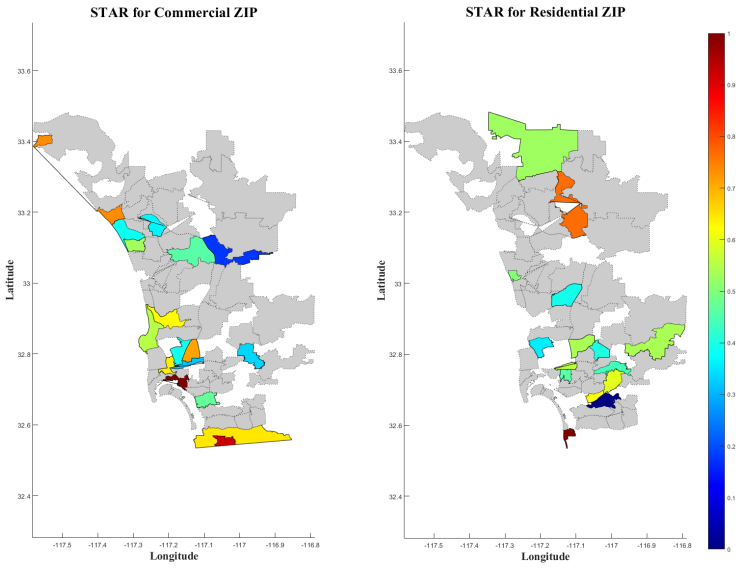


Figure 10: **STAR of the Entertainment Effect.** This figure plots the spatio-temporal activity radar for the entertainment effect. Grey areas are zip codes within San Diego county excluded due to incomplete data.

5.5 A Surrogate Proxy of Local Economic Activity

Section 2.3 established a link between our hourly variables and local economic activity using a regression on payroll. Building on that regression, we construct a surrogate proxy for local economic activity by linearly combining the high-frequency variables:

$$\widehat{\widetilde{LEA}}^{h}_{SUR,t_h} = \widehat{b}_{TF} TF^{h}_{SD,t_h} + \widehat{b}_{FT} FT^{h}_{SD,t_h} + \widehat{b}_{ElU} ElU^{h}_{SD,t_h} + \widehat{b}_{Temp} Temp^{h}_{SD,t_h}, \qquad (23)$$

where TF, FT, ElU, and $Temp$ refer to variables aggregated along the spatial dimension across the 62 local San Diego zip codes (denoted by the subscript SD), and the weights $\widehat{b}_i$ are given by the estimates of the coefficients in the univariate payroll regressions reported in Table 1, rescaled to sum to one similar to an equal-weighted combination of fitted values. Despite its simplified form, this surrogate proxy offers a simple and intuitive

link between the hourly series and payroll growth, giving us additional insights on the information contained in the hourly variables, as well as a measure of comparison.

Figure 11 shows weekly averages of hourly values of the surrogate proxy from (23) and the GEM mean component ($\mu^h_{t_h}$).[43] Although the two series look broadly similar, the figure also highlights that the granular model captures hourly patterns missed by the surrogate proxy. In the next section, we will use the surrogate proxy to generate competing forecasts and draw comparisons with the granular modeling approach.

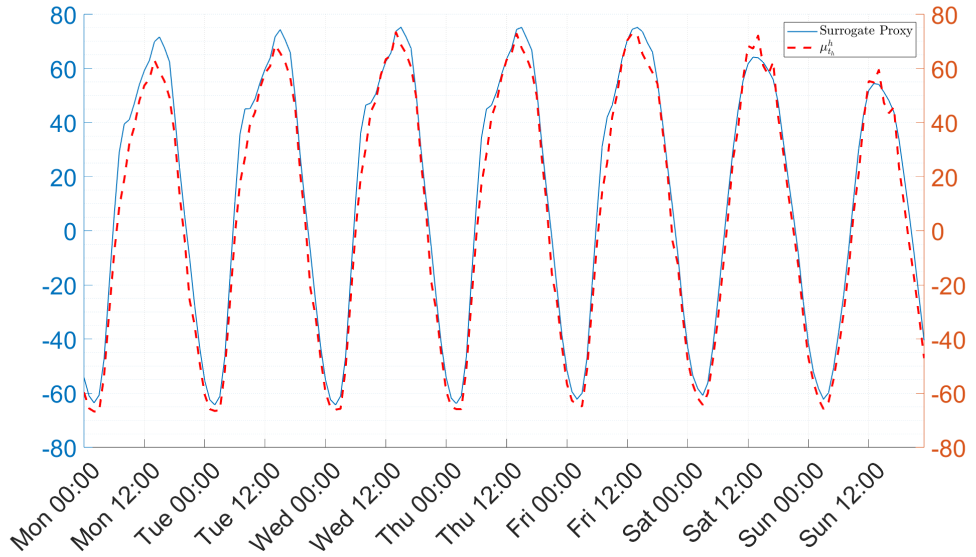


Figure 11: **Surrogate Proxy of Local Economic Activity.** This figure displays the weekly average of the high-frequency surrogate proxy of local economic activity constructed as in equation (23) (solid line, left scale) and the weekly average of the mean components from the granular economic modeling (broken line, right scale). The scales report respective percentage devotions from the mean of the series.

6 Forecasts and Scenario Analysis

We next study forecasts of local economic activity at short and long horizons and develop a framework for conducting scenario analysis based on assumptions about the future trajectories of key model variables.

6.1 Forecasting

In this subsection, we develop methodologies to construct forecasts of local economic activity at the hourly zip-code level. We first present granular forecasts based on our GEM approach. We then illustrate forecasts using the surrogate proxy for local economic ac-

[43] We report the series in percentage deviation from their mean to allow a direct comparison.

tivity. Finally, we report forecasting results for both short- (one hour) and long (multiple hours) horizons for the observable series.

6.1.1 Granular Forecasts

Our granular modeling approach naturally extends to generating forecasts of local economic activity at different horizons and for both local and more aggregate spatial levels. We begin by explaining how to generate such forecasts and by constructing a set of plausible benchmarks for comparison.

To start, consider predicting local economic activity one hour ahead. Such forecasts involve predicting the individual components of $f^h_{t_h+1}$, $\mu^h_{t_h+1}$ and $\psi^h_{t_h+1}$, given the information available at time t_h: $\widehat{f}^h_{t_h+1|t_h} = [\widehat{\mu}^h_{t_h+1|t_h}, \widehat{\psi}^h_{t_h+1|t_h}]'$.

Forecasts of $\mu^h_{t_h+1|t_h}$ are computed using the transition equation:

$$\begin{bmatrix} \widehat{\mu}^h_{t_h+1|t_h} \\ \widehat{\ell}^h_{t_h+1|t_h} \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \widehat{\mu}^h_{t_h|t_h} \\ \widehat{\ell}^h_{t_h|t_h} \end{bmatrix} = T \begin{bmatrix} \widehat{\mu}^h_{t_h|t_h} \\ \widehat{\ell}^h_{t_h|t_h} \end{bmatrix}. \tag{24}$$

Forecasts of the calendar component, $\psi^h_{t_h+1|t_h}$, are straightforward to compute since future values of the underlying diurnal dummies is known ahead of time:

$$\widehat{\psi}^h_{t_h+1|t_h} = \widehat{HoD}_{t_h+1|t_h} + \widehat{DoW}_{t_h+1|t_h} + \widehat{h}_{t_h+1|t_h} + \widehat{H}_{t_h+1|t_h} + \widehat{\varphi}_{t_h+1|t_h}. \tag{25}$$

The one-hour-ahead forecast of local economic activity is obtained as the sum of the forecasts of the two components of $\widehat{f}^h_{t_h+1}$: $\widehat{LEA}^h_{t_h+1|t_h} = \widehat{\mu}^h_{t_h+1|t_h} + \widehat{\psi}^h_{t_h+1|t_h}$.

At longer horizons, the diurnal hour-of-day cycles suggest that there are strong similarities among the same hours across different days but large differences within the day. Building on these patterns, the H-hour-ahead forecast of the unobserved mean component at time $t_h + H$ is obtained using the most recent value of $\widehat{\mu}^h_{t_h}$ observed at the same hour on the latest available day:

$$\widehat{\mu}^h_{t_h+H|t_h} = \widehat{\mu}^h_{t_h+H-24k|t_h},$$

where $k = \lceil H/24 \rceil$ is an integer selected such that $t_h + H - 24k$ denotes the hour corresponding to $t_h + H$ at the most recent available observation.[44]

Given that $t_h + H$ is a known date at time t_h, the $H-$hour-ahead forecasts of the calendar component are again predetermined at time t_h, and easily computed as

$$\widehat{\psi}^h_{t_h+H|t_h} = \widehat{HoD}_{t_h+H|t_h} + \widehat{DoW}_{t_h+H|t_h} + \widehat{h}_{t_h+H|t_h} + \widehat{H}_{t_h+H|t_h} + \widehat{\varphi}_{t_h+H|t_h}.$$

[44]The alternative of iterating on the transition equation H times

$$\begin{bmatrix} \widehat{\mu}^h_{t_h+H|t_h} \\ \widehat{\ell}^h_{t_h+H|t_h} \end{bmatrix} = T^H \begin{bmatrix} \widehat{\mu}^h_{t_h|t_h} \\ \widehat{\ell}^h_{t_h|t_h} \end{bmatrix}$$

produces inferior forecasts due to the propagation of estimation errors (Marcellino et al. (2006)).

6.1.2 Surrogate Forecasts

To the best of our knowledge, our paper is the first to construct a latent factor at the hourly-zip code granularity. Since economic activity is not directly observable at this resolution, we cannot directly compare our forecasts of $f^h_{t_h}$ to a measured outcome. As an economically grounded measure, instead we consider the approach based on the surrogate proxy constructed using the estimates $\widehat{b}_j$ from the payroll regressions in equation (23):[45]

$$\widetilde{LEA}^h_{SUR,t_h+1|t_h} = \sum_{j=1}^{N_h} \widehat{b}_j y^h_{S_A,j,t_h}.$$

We use this surrogate measure to produce competing forecasts for the underlying observable series and evaluate the relative performance of our local economic activity factor. Ideally, if our approach is successful in modeling local economic activity, it should provide better forecasts for these variables.

6.1.3 Hour-Ahead Forecasts

We now report the empirical results. We focus on forecasting the observable series that depend on local economic activity. We distinguish between two spatial levels: aggregate (San Diego) and local forecasts (zip code), labeled y^h_{SD,i,t_h+H} and y^h_{zip,i,t_h+H}, respectively.

In Section 5, we found persistent dynamics in the idiosyncratic components, ϵ^h_{s,i,t_h}, at the zip-code level.[46] Such persistence could help us improve our hourly forecasts, and we propose capturing it by modeling the idiosyncratic term as an AR(1) process which we add to the predicted factor:

$$\widehat{y}^h_{s,i,t_h+1|t_h} = \widehat{\varrho}_{s,i} + \widehat{\beta}_{s,i}\widehat{f}^h_{t_h+1|t_h} + \widehat{\phi}\,\widehat{\epsilon}^h_{s,i,t_h|t_h},\ s = SD, zip,$$

where $\widehat{\epsilon}^h_{s,i,t_h|t_h} = \widehat{y}^h_{s,i,t_h} - \widehat{\varrho}_{s,i} - \widehat{\beta}_{s,i}\widehat{f}^h_{t_h|t_h}$ are the residuals.[47]

$$\widehat{\epsilon}^h_{s,i,t_h|t_h} = \widehat{y}^h_{s,i,t_h} - \widehat{\varrho}_{s,i} - \widehat{\beta}_{s,i}\widehat{f}^h_{t_h|t_h},\ s = SD, zip.$$

Replacing forecasts of $\widehat{f}^h_{t_h+1|t_h}$ with those from $\widetilde{LEA}^h_{SUR,t_h+1|t_h}$, we obtain the competing surrogate forecasts for the individual series, where in this specification we do not include an AR(1) term for the surrogate forecasts.

[45] We normalize the coefficients to sum to one, so that they can be interpreted as weights.

[46] This component includes measurement noise affecting a specific zip code. For example, a car accident that impacts the measurement station of traffic flows in a given zip code may disrupt the measurement of traffic flow for that zip code for a given window of time.

[47] In our forecasting exercise, we compute the residuals recursively using only information up to t_h, making sure that no forward-looking information is used to construct the forecasts.

We compare the forecasts produced by these models against competing benchmarks at both the aggregate and disaggregated levels. Given the important differences between forecasting aggregate versus more granular outcomes, we employ two different benchmark models.

At the aggregate level, the benchmark used to evaluate the forecasting performance of the models based on the economic activity proxies—whether using our LEA factor or the surrogate proxy—is an hourly random walk:

$$\text{benchmark for aggregate San Diego:} \quad \widehat{y}^{h}_{SD,i,t_h+1|t_h,bmk_{SD}} = y^{h}_{SD,i,t_h}.$$

For the disaggregated zip-code level, we consider the grand mean computed across zip codes for each series:

$$\overline{y}^{h}_{\overline{zip},i,t_h} := \log\left(1 + \frac{1}{N_{zip}} \sum_{zip} \exp(y^{h}_{zip,i,t_h})\right).$$

We use this to set up the following benchmark AR(1) model at the zip-code level:

$$\text{benchmark for individual zip codes:} \quad \widehat{y}^{h}_{zip,i,t_h+1|t_h,bmk_{zip}} = \widehat{a}_{zip,i} + \widehat{b}_{zip,i}\overline{y}^{h}_{\overline{zip},i,t_h}.$$

This benchmark aligns more closely with the purpose of our LEA factor (or the surrogate proxy) than random walks for each individual zip-code series, namely to capture the common drivers underlying the series rather than their idiosyncratic dynamics.

Our forecasting exercise employs the first two years of hourly observations in our sample to estimate the parameters of the forecasting models, and we keep the full final year—January 1st to December 31st 2023—as the test sample.

Table 2 displays the one-hour-ahead performance for the GEM and SUR forecasts measured as RMSE ratios relative to those for a given benchmark, so values below one show performance better than the benchmark. We show forecasting performance across three different levels of aggregation and for the series on foot traffic, traffic flows and electricity consumption. The top row reports aggregate results for San Diego while rows (2)-(4) and rows (5)-(7) report results for the top and bottom three zip code ranked by the share of commercial electricity usage, respectively. The bottom row reports the grand mean across all the zip codes in our sample.

The GEM forecasts outperform the benchmarks in most cases. Starting from the aggregate San Diego level, the GEM outperforms both the individual benchmark, showing smaller RMSFE, as well as the SUR forecasts, which outperform the benchmark for foot traffic but not for traffic flow and electricity usage.

These results are confirmed by the grand mean result, which summarizes the forecasting performance across all the individual zip codes. Note the substantial heterogeneity in forecasting performance at the zip-code level.

We also report the results from a one-sided Diebold–Mariano test (Diebold and Mariano (1995)), which evaluates the hypotheses

$$H_0 : E(d_{s,i,t_h}) \leq 0 \quad \text{against } H_1 : E(d_{s,i,t_h}) > 0,$$

where $d_{s,i,t_h} = e^2_{s,i,t_h,bmk} - e^2_{s,i,t_h,\mathcal{M}}$ denotes the loss differential between the benchmark and the local economic activity model. Rejection of the null indicates statistical evidence for a superior predictive accuracy of the local economic activity model.[48]

In the great majority of cases, the GEM-based forecasts display significantly superior predictive ability relative to the benchmark. Forecasts based on the surrogate proxy display greater variation, both across space and across series. For example, while the null is rejected for the foot traffic series in residential zip codes, it is not rejected for traffic flows or electricity usage in the same locations.

Granular Forecast Evaluation						
	Foot Traffic		Traffic Flows		Electricity	
	SUR	GEM	SUR	GEM	SUR	GEM
SD	1.410	0.998***	1.274	0.994***	2.479	0.959***
92121	0.999**	0.372***	1.677	1.104	0.973***	0.266***
92101	1.033	0.640***	1.767	1.050	0.790***	0.377***
92123	1.016	0.675***	1.652	0.917***	1.082	0.336***
92116	0.934***	0.482***	1.712	0.967***	1.789	0.529***
91941	0.968***	0.642***	1.581	0.948***	1.579	0.482***
92139	0.944***	0.739***	1.564	0.920***	1.492	0.459***
Grand Mean	1.006	0.625***	1.616	0.965***	1.644	0.608***

Table 2: **Granular Forecasting Performance.** This table reports the forecasting performance of the one-hour-ahead forecasts for foot traffic, traffic flow and electricity usage at the aggregate (San Diego) and zip code levels. SUR denotes the forecasts based on the surrogate proxy, while GEM refers to our granular-based forecasts. Results are reported as $RMSFE$ ratios with respect to the underlying benchmarks. The test sample covers all hourly periods during 2023. The individual zip codes reported are the top 3 commercial and top 3 residential, respectively, identified by their share of commercial electricity usage. *, **, and *** denote significance at the 10%, 5%, 1% level of the one-sided Diebold-Mariano test.

6.1.4 Longer-horizon Forecasts

We next turn to the longer one-day or one-week forecasts. We adopt the same hour-of-day anchoring approach based on the latest hour-matched observation from the historical sample, $\widehat{\mu}^h_{t_h+H|t_h} = \widehat{\mu}^h_{t_h+H-24k}$, where $k = \lceil H/24 \rceil$ is selected to match the last observation by day of the week and hour of the day available. Following this approach, forecasts for $y^h_{s,i,t_h+H,}$ under the granular economic model are computed as

$$\widehat{y}^h_{s,i,t_h+H|t_h} = \widehat{\varrho}_{s,i} + \widehat{\beta}_{s,i,\mu}\widehat{\mu}^h_{t_h+H-24k|t_h} + \widehat{\beta}_{s,i,\psi}\widehat{\psi}^h_{t_h+H|t_h} + \widehat{\phi}\widehat{\epsilon}^h_{s,i,t_h+H-24k|t_h+H-24k},\ s = SD, zip.$$

We compare these forecasts with those generated by the benchmark models and with

[48]We use HAC-robust standard errors computed with a Bartlett kernel and 24 lags (1 diurnal cycle).

those obtained from the surrogate proxy constructed using a similar methodology:

$$\begin{aligned}
\text{Benchmark Aggregate Level:} \quad & \widehat{y}^{h}_{SD,i,t_h+H|t_h,bmk_{SD}} = y^{h}_{SD,i,t_h+H-24k}, \\
\text{Benchmark zip-code Level:} \quad & \widehat{y}^{h}_{zip,i,t_h+H|t_h,bmk_{zip}} = \widehat{a}_{zip,i} + \widehat{b}_{zip,i}\overline{y}^{h}_{\overline{zip},i,t_h+H-24k}, \\
\text{Surrogate Forecasts:} \quad & \widehat{y}^{h}_{s,i,t_h+H|t_h} = \widehat{a}_{s,i} + \widehat{b}_{s,i}\ \widetilde{LEA}^{h}_{SUR,t_h+H-24k|t_h},\ s = SD, zip.
\end{aligned}$$

We report the results for the day-ahead and week-ahead forecasts in Table 3. Day-ahead corresponds to a 24-hour forecast ($k = 1$), while week-ahead corresponds to a 168-hour forecast ($k = 7$).

The results show that the hour-matched GEM forecasts deliver substantial gains in one-day and one-week-ahead forecasting performance. This indicates that the LEA factor captures predictive information about the underlying series that is not contained in their own past values, consistent with the idea that it isolates signals about the underlying economic forces driving these variables.

Long-Horizon Granular Forecast Evaluation						
	Foot Traffic		Traffic Flows		Electricity	
	SUR	GEM	SUR	GEM	SUR	GEM
Panel A) one-day-ahead forecasts						
SD	1.413	0.989	1.663	0.970	1.910	1.025
92121	0.956***	0.495***	1.390	0.883***	0.948***	0.695***
92101	1.020	0.754***	1.640	0.935***	0.847***	0.586***
92123	1.050	0.939***	1.682	1.003	1.045	0.911***
92116	0.950***	0.538***	1.602	0.980***	1.588	0.633***
91941	0.970***	0.685***	1.602	0.993***	1.455	0.574***
92139	0.966***	0.907***	1.635	0.972***	1.365	0.517***
Grand Mean	1.017	0.714***	1.541	0.904***	1.503	0.724***
Panel B) one-week-ahead forecasts						
SD	1.415	1.002	3.152	0.955***	1.619	1.058
92121	0.938***	0.428***	2.117	0.624***	0.949***	0.485***
92101	1.002	0.784***	2.478	0.957***	0.829***	0.549***
92123	1.053	0.902***	3.076	0.930***	1.043	0.680***
92116	0.939***	0.499***	2.423	1.033	1.348	0.871***
91941	0.950***	0.649***	2.795	0.950***	1.288	0.833***
92139	0.959***	0.867***	2.577	0.885***	1.265	0.646***
Grand Mean	1.007	0.685***	2.299	0.785***	1.342	0.807***

Table 3: **Long-Horizon Granular Forecast Evaluation.** This table reports the forecasting performance of the day-ahead and week-ahead forecasts for foot traffic, traffic flow and electricity usage at the aggregate (San Diego) and zip-code levels. SUR denotes forecasts based on the surrogate proxy, while GEM refers to our granular-based forecasts. The evaluation metric is expressed in terms of $RMSFE$ ratios with respect to the benchmark models explained in the main text. The forecasting period extends from January 1^{st} 2023 01:00 a.m. to December 31^{st} 2023 23:00 p.m. The individual zip codes reported are the top 3 commercial and top 3 residential zip codes identified by their share of commercial electricity usage. *, **, and *** denote significance at the 10%, 5%, 1% level of the one-sided Diebold-Mariano test.

6.2 Spatio-Temporal Scenario Analysis

Suppose a municipality seeks to assess the impact on local business activity of restricting traffic flow or foot traffic within a particular zip code over a given period. How can the economic effects of such an intervention be quantified?

Impulse response functions (IRFs) are a fundamental tool for analyzing the dynamic effects of shocks and policy intervention in economic systems. They capture how variables respond to a conditioning event at a given time but are typically used for lower-dimensional models than those we consider. In our granular framework, the observation vector y_{t_h} lacks explicit VAR dynamics, eliminating the feedback effects central to standard IRF analysis.[49]

Building on our GEM approach, we introduce a framework for *spatio-temporal scenario analysis* that generalizes the intuition behind impulse response functions to spatio-temporal settings. This approach enables us to analyze how shocks propagate across individual units within a local economy.

6.2.1 General Framework

Let $\mathcal{E}_{\mathcal{T}}$ denote a scenario of interest over some time interval or at a single point in time $\mathcal{T} \equiv \{t_h \in T_h\}$. Our spatio-temporal scenario analysis studies the conditional distribution $y_{t_h}|\mathcal{E}_{\mathcal{T}}$ for $t_h \in T_h$. For example, a scenario could be a restriction on the granular variables such as shutting off traffic flow for a certain period of time:

$$\mathcal{E}_{t_h} = \mathcal{Y}_{t_h} \bigcup \{y_{s,i,t_h} = 0\},\ t_h \in T_h,$$

where $\mathcal{Y}_{t_h}$ denotes the granular system under no conditioning and $\{y_{s,i,t_h} = 0\}$ is the conditioning event (scenario), which shuts off variable i at location s for $t_h \in T_h$. We denote the observation vector under the conditioning event as $y_{t_h}(\mathcal{E}_{\mathcal{T}})$. To evaluate the effects of scenarios on local economic activity we use the normal distribution as an approximation to compute the expectation of the local economic activity factor conditional on the scenario as

$$E[f^h_{t_h+1}|\mathcal{E}_{t_h+1}] = E[f^h_{t_h+1}|\mathcal{Y}_{t_h}] + \Sigma_{f,y|\mathcal{Y}}\Sigma^{-1}_{y|\mathcal{Y}}\Big(y_{t_h+1}(\mathcal{E}_{t_h+1}) - E[y_{t_h+1}|\mathcal{Y}_{t_h}]\Big),$$

where $E[f^h_{t_h+1}|\mathcal{Y}_{t_h}] = f^h_{t_h+1|t_h}$ is the predicted factor, $y_{t_h+1}(\mathcal{E}_{t_h+1})$ is the observation vector under the restriction in $\mathcal{E}_{t_h+1}$, $E[y_{t_h+1}|\mathcal{Y}_{t_h}]$ is the conditional expectation of y_{t_h+1} given the system in place at time t_h, and $\Sigma_{f,y|\mathcal{Y}}$ and $\Sigma_{y|\mathcal{Y}}$ are the conditional covariance and

[49] See, e.g., Sims and Zha (1998) and Waggoner and Zha (1999).

variance-covariance matrices under no conditioning:[50]

$$\Sigma_{f,y|\mathcal{Y}} = Cov(f^h_{t_h+1}, y_{t_h+1}|\mathcal{Y}_{t_h}), \quad \text{and} \quad \Sigma_{y|\mathcal{Y}} = Var(y_{t_h+1}|\mathcal{Y}_{t_h}).$$

We define the expected scenario effect (ESE) as the time-series difference between the expected value of the factor under no conditioning versus under the scenario:

$$ESE(\mathcal{E}_{t_h+1}) := E[f^h_{t_h+1}|\mathcal{Y}_{t_h+1}] - E[f^h_{t_h+1}|\mathcal{E}_{t_h+1}]. \tag{26}$$

Cumulating the ESE over the scenario period, we obtain the cumulative expected scenario effect ($CESE$):

$$CESE(\mathcal{E}_{\mathcal{T}}) := \sum_{t_h \in T_h} ESE(\mathcal{E}_{t_h}).$$

$CESE$ measures the total expected impact of the scenario on local economic activity. Denoting by $w_{s,i}$ the element of $\Sigma_{f,y|\mathcal{Y}}\Sigma^{-1}_{y|\mathcal{Y}}$ corresponding to the s, i granular variable, we have

$$CESE(\mathcal{E}_{\mathcal{T}}) = \sum_{t_h \in \mathcal{T}} \sum_{s \in \mathcal{S}} \sum_{i \in N_s} w_{s,i} \times (y_{s,i,t_h+1}(\mathcal{Y}_{t_h+1}) - y_{s,i,t_h+1}(\mathcal{E}_{t_h+1})). \tag{27}$$

The weights $w_{s,i}$ provide a measure of the sensitivity of the local economic factor to a given scenario for variable i in location s.

In the special case where $\Sigma_{y|\mathcal{Y}}$ takes the form of a diagonal matrix, $D_{y|\mathcal{Y}}$, with the variances of the y_{t_h} granular units on its main diagonal, we obtain an analytical expression for the weights:

$$(\Sigma_{f,y|\mathcal{Y}} D^{-1}_{y|\mathcal{Y}})_j = \frac{\sigma_{f,y_j}}{\sigma^2_{y_j}} \equiv \beta_j \times \frac{Var(f)}{Var(y_j)}.$$

This expression for the weights provides a conservative estimate of the expected scenario effect since it neglects the covariances across the variables in the observation vector. Nevertheless, it offers a useful interpretation of the components of the scenario effect.[51] Ceteris paribus, the impact of a given scenario is expected to be larger for variables with higher factor loadings ($\beta_{s,i}$) and lower variance ($Var(y_{s,i})$).

6.2.2 Empirical Results

We next show how our framework for spatio-temporal scenario analysis can be used to assess the impact of local policy interventions on economic activity.

[50] The dimension of the variance-covariance matrix $\Sigma_{y|\mathcal{Y}}$ is 186×186 in our empirical application (62 zip codes times 3 variables). Hence, our approach requires us to invert a high-dimensional matrix at each point in time. Our implementation considers a Ridge-type regularized variance-covariance matrix:

$$\Sigma^{\dagger}_{y|\mathcal{Y}} = \Sigma_{y|\mathcal{Y}} + \lambda I,$$

where λ is the penalty parameter and I is the comforming identity matrix.

[51] Appendix D.4 reports robustness results for the scenario effect which confirm this finding.

We consider scenarios $\mathcal{E}_{\mathcal{T}}$ lasting one full calendar day (24 hours)—November 13, 2023—and affecting the top three commercial and top three residential zip codes identified in Section 5.4. We consider four different restrictions: a shutdown of foot traffic, a shutdown of traffic flow, a shutdown of electricity usage, and, lastly, a shutdown of all three series simultaneously.[52]

Figure 12 plots $CESE$ for these scenarios. To quantify how much economic output would be lost relative to under normal conditions, we report cumulative losses as percentages of the daily San Diego GDP.

Starting from the top left panel of Figure 12a, we show the $CESE$ resulting from the shutdown of foot traffic for each of the two sets of zip codes, one at a time, while holding everything else constant. We also show the $CESE$ when foot traffic is shut down in all six zip codes simultaneously. A similar analysis, now for traffic flow and electricity consumption, is presented in Panels 12b and 12c while Panel 12d shows the results of a complete shutdown in all three series across the six zip codes.

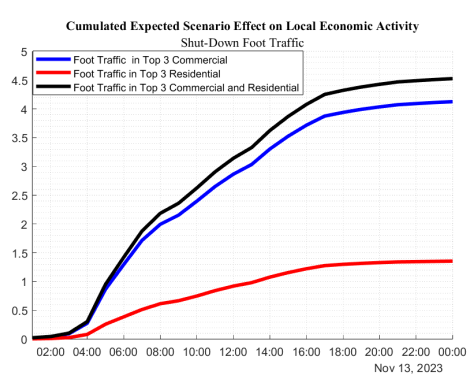
The plots show how the economic impact of a shutdown varies along three important dimensions: (i) the type of variable that is perturbed; (ii) the time of the day where the shutdown takes place; and (iii) the type of economic activity dominating the affected zip code.

Starting with point (i), we see that the cumulative reduction in San Diego-wide economic activity from a shutdown in foot traffic in the top-3 residential areas is only 1% compared to double the effect (2%) from a shutdown in electricity and three times the impact from a shutdown in traffic flow. If all three variables are shut down simultaneously, San Diego's daily GDP declines by -2.5%. The economic impact of a shutdown in the top-3 commercial zip-codes, a decline of -4.5%, is larger and more homogeneous across these variables.

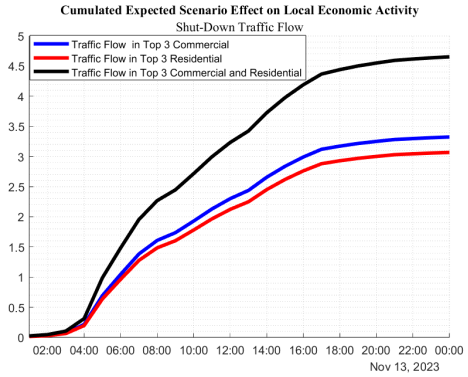
Turning to point (ii), the economic impact of the restrictions varies significantly across hours of the day, with the largest effect occurring during the central part of the day (6 a.m.—6 p.m.) and a much smaller effect in the late evening and early morning hours. This finding is consistent with the empirical features of diurnal variation in local economic activity.

Finally, for point (iii) we find that the largest losses in economic output are observed when the shutdown affects commercial zip codes. This result supports the economic interpretation of our local economic activity factor which reflects business activity patterns and is therefore expected to be more pronounced in high-intensity commercial areas than in residential zones.

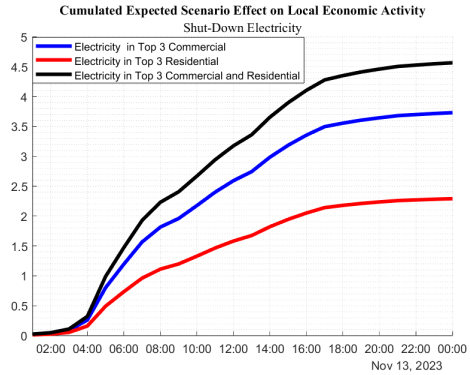
[52]Our results are computed using $\Sigma^{\dagger}_{y|\mathcal{Y}} = \Sigma_{y|\mathcal{Y}} + \lambda I$ where $(\lambda = \left(\frac{1}{S}\sum_s \mu_{S,s}\right)^{-1})$ is the inverse of the grand mean of the observation vector. Appendix D.4 reports robustness checks for different values of λ and for alternative regularization methods.



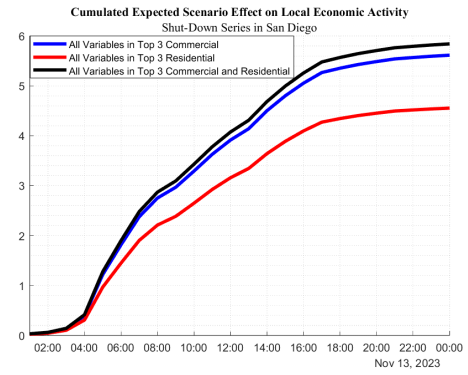
(a) **Foot Traffic.**



(b) **Traffic Flow.**



(c) **Electricity.**



(d) **All.**

Figure 12: **Spatio-Temporal Scenario Analysis.** This figure plots the scenario analysis under four different conditioning events. In all cases, two sets of zip codes within the San Diego area are affected: the top three residential and top three commercial zip codes. The remaining zip codes are left unrestricted. Panel 12a shows the cumulative effect of shutting down foot traffic for a full 24-hour period on November 13, 2023. Panel 12b reports the cumulative impact of shutting down traffic flow under the same conditions, while Panel 12c illustrates the cumulative effect of shutting down electricity consumption. Finally, Panel 12d considers the simultaneous shutdown of all variables within those zip codes over the same 24-hour period. In all panels, the lines track the cumulative expected scenario effect expressed as a percentage of the daily San Diego GDP.

6.3 Event Detections

We now show how our GEM approach can be used to detect events or disruptions across series and locations. For each zip code i and series s, consider again the GEM-implied residuals $\widehat{\epsilon}_{s,i,t_h} = y_{s,i,t_h} - \widehat{y}_{s,i,t_h}$. Let $W = [t_h^{start}, t_h^{end}]$ denote a window of interest and by W^c its complement. We estimate the mean and standard deviations of the residuals in

zip code i and series s over the complement window:

$$\bar{\widehat{\epsilon}}_{s,i}(W) = \frac{1}{T_{W^c}} \sum_{t_h \in W^c} \widehat{\epsilon}_{s,i,t_h},$$

$$\widehat{\sigma}_{s,i}(W) = \left(\frac{1}{T_W} \sum_{t_h \in W^c} (\widehat{\epsilon}_{s,i,t_h} - \bar{\widehat{\epsilon}}_{s,i}(W^c))^2, \right)^{1/2},$$

where T_W is the length of the window over which the sum is taken.

Next, we consider the following two event detection statistics. The first criterion is based on the deviations of the residuals from their baseline standard deviation:

$$|\widehat{\epsilon}_{s,i,t_h}| \geq \kappa \, \widehat{\sigma}_{s,i}(W), \; t_h \in W, \tag{28}$$

where κ can be chosen depending on the application.[53] The second criterion is based on the standardized residuals,

$$\widehat{z}_{s,i,t_h} := \frac{\widehat{\epsilon}_{s,i,t_h} - \bar{\widehat{\epsilon}}_{s,i,t_h}}{\widehat{\sigma}_{s,i}(W)}, \tag{29}$$

and the following monitoring statistic

$$Q_{s,i}(W) = \frac{1}{T_W} \sum_{t_h \in W} |\widehat{z}_{s,i,t_h}|. \tag{30}$$

Using this approach we are able to identify periods characterized by large hour-zip code z-scores in (29). For example, using a search window that spans from August to September 2021, the largest negative scores are almost all associated with the hours (12-5 p.m.) on 08/24/2021 when a plane landed on interstate 5, a major freeway in San Diego. The incident caused a severe traffic disruption that lasted for several hours.[54] Although this event was short-lived, and only affected a small set of zip codes connected through the impacted interstate freeway, using the scenario analysis, we estimate the cost of this disruption to be about $45 million in losses to that day's San Diego GDP. To validate this estimate, we compute an approximation following the methodology used in the traffic congestion literature, obtaining an estimate of about $39.2 million.[55]

As a second example, we detect large positive z-scores during the heat wave affecting San Diego from August 31 through September 9, 2022.[56] The spike in the z-score from the LEA was caused by a sharp increase in electricity consumption, which lead the California Independent System Operator to issue Flex Alerts for those days. Spikes in electricity consumption were particularly large in the late afternoon and early evening hours on each of those days.

[53] In our implementation we use a value of κ equal to 3.

[54] See this San Diego Union Tribune's article for details.

[55] Specifically, we multiply the average hourly traffic flow in the five affected zip codes—computed over the month of August and for the time interval affected by the disruption (12 pm–5 pm)—by the duration of the disruption and by the average hourly earnings in California for August 2021.

[56] Details on this event can be found at this Public Office's report.

7 Economic Validation

In this section, we validate the economic interpretation of the LEA factor by relating it to key economic measures. We begin by establishing a direct link between the LEA factor and economic activity using local payroll growth. Next, we corroborate this interpretation by documenting a strong negative relationship with weekly unemployment claims, a series that has not been used to construct our LEA estimates.

7.1 Links to Payroll Growth

Verifying that our economic activity factor is closely linked to economic activity is difficult at the highest levels of granularity because data on economic activity is not observed at that level (e.g., hourly by zip code). As a second-best alternative, instead, we start by examining the correlation between our LEA factor and a more aggregate measure of economic activity in San Diego, namely monthly payroll growth. Specifically, we compare our growth factor with the quarterly growth rate of payroll in San Diego, $\Delta^3 PAYROLL_{t_m} = PAYROLL_{t_m} - PAYROLL_{t_m-3}$. We use the quarterly difference to get a more reliable measure of economic activity than the month-on-month variation, which tends to be noisier and affected by measurement errors.[57] We compare this growth rate to the monthly change in our LEA factor aggregated to the monthly level, $\Delta f^m_{t_m} = f^m_{t_m} - f^m_{t_m-1}$, where the monthly factor $f^m_{t_m}$ aggregates approximately 720 (24×30) hourly observations.

Figure 13a plots the time series of the two growth rates over the sample January 2022 to December 2023. The two series move closely together, with a correlation of 0.53. This finding aligns with our objective of recovering a factor that captures fundamental forces in the local economy.

We next show how our LEA factor can be used in real time to forecast the growth rate of payroll data. A serious challenge in forecasting this series is that data on this variable only become available with a lengthy delay.[58] Simple autoregressive models such as an AR(1) are therefore infeasible to use to forecast local payroll growth.

The literature on empirical features of the business cycle has documented that variables related to the labor market, such as the unemployment rate and payroll data, are lagging indicators of economic activity (see, e.g., Stock and Watson (1999) and Stock and Watson (2002)), with an average lag of around four to six months. Taking a conservative stance, we consider an autoregressive model with a six-month lag:

$$\Delta^{(3)} PAYROLL_{t_m} = a + \beta \Delta PAYROLL_{t_m-6} + \epsilon_{t_m}.$$

[57] Reinforcing this point, the month-on-month growth rate allows an $MA(1)$ representation with a negative first-order autocorrelation of -0.45.

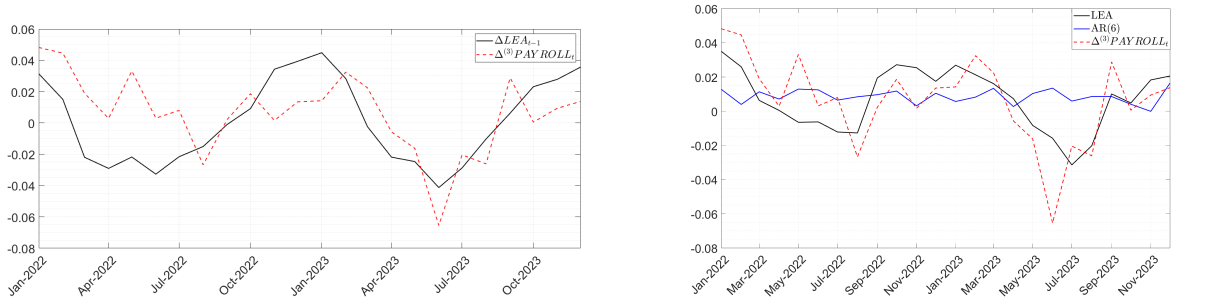
[58] The payroll survey is conducted monthly, concluding on the 12th of each and the estimates start being released the month after and are updated on a rolling basis until March when tax data gets released and the CAEDD produces the final revisions (up till one year later).

We compare forecasts from this model to forecasts from a model based on our real-time LEA forecasts:

$$\Delta^{(3)} PAYROLL_{t_m} = a + \gamma \Delta LEA_{t_m - 6} + \epsilon_{t_m}.$$

Estimates of this equation show that the six-month lagged value of LEA obtains a highly significant positive coefficient with a t-statistic of around 3.5.

Figure 13b illustrates the relationship. We see a clear positive relation between next-month payroll growth and the six-month lagged change in our LEA factor.



(a) **Time Series.** (b) **Forecasted Payroll vs Actual Payroll.**

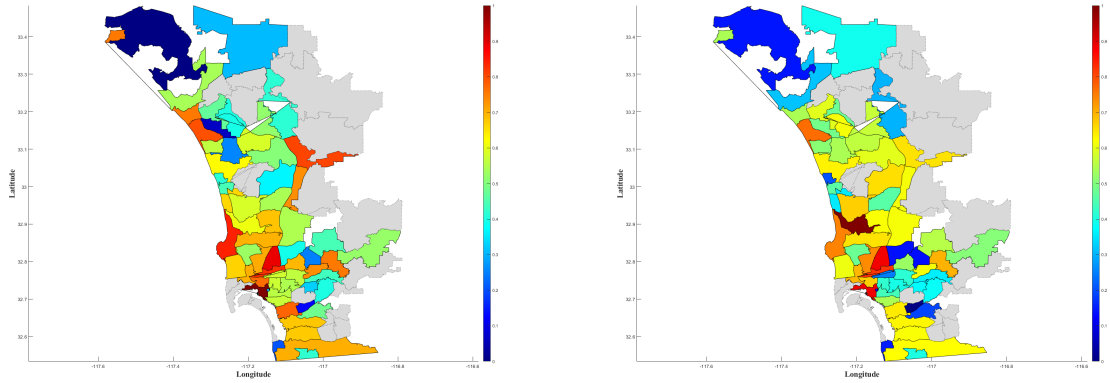
Figure 13: **LEA and Payroll Growth.** The left figure plots the growth rate of the monthly San Diego LEA factor (solid line) along with the growth rate of San Diego payroll (dashed line). The right figure plots the forecasts of payroll growth (solid lines) and the actual values (dashed line). The data frequency is monthly and the sample spans January 2022 - December 2023.

The spatio-temporal radar of economic activity developed in Section 5.4 offers an additional validation tool. If our local economic activity factor accurately captures economic patterns, it should give us a representation of economic activity that can be related to that offered by an analogous payroll-based representation.

In Figure 14, we compare the STAR constructed using our LEA factor, based on foot traffic, traffic flow, electricity usage, and spending data (Figure 14a) to an alternative version constructed based exclusively on payroll data (Figure 14b), which uses the ratio of local $PAYROLL_{zip,t}$ to total $PAYROLL_{SD,t}$ to determine the economic importance of a given zip code, and obtain payroll economic activity importance $PEAI_{s,t} = PAYROLL_{s,t}/PAYROLL_{SD,t}$.[59]

The rank correlation between the LEA- and payroll-based STARs equals 0.75, consistent with a strong alignment in the spatial information based on our approach and that implied by more traditional economic variables such as payroll.

[59]In the appendix (Figure D.3) we also report the STAR constructed using the individual series to assess the contribution of each variable across the zip codes.



(a) **STAR Based on the LEAI.** (b) **STAR Based on the PEAI.**

Figure 14: **Spatio-Temporal Economic Radars.** The left figure plots the STAR based on the LEA factor estimated using the GEM approach. The right panel plots the STAR based exclusively on the payroll data. Grey areas are zip codes within San Diego county excluded due to incomplete data.

7.2 Links to Weekly Unemployment Claims

We finally corroborate the economic interpretation of our local economic activity factor by establishing a strong (negative) relationship with weekly unemployment claims. Weekly unemployment claims serve as a timely labor market indicator. Given the strong empirical link between aggregate measures of economic activity and labor market conditions, establishing a link between our LEA factor and local labor markets therefore offers an important economic validation step.

Let UNC_t^W denote the log of the weekly unemployment claims measured at the aggregate San Diego level for the week ending at time t. To smooth short-term noise related to data measurement and reporting, we consider the following four-week moving average of weekly unemployment claims:

$$Z_t^{UNC} = \frac{\sum_{j=0}^{3} UNC_{t-j}^W}{4}.$$

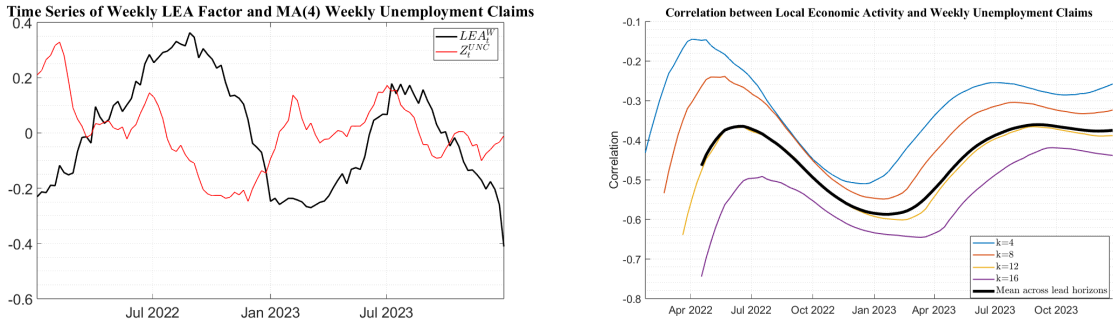
Figure 15a displays the time series of Z_t^{UNC} and our weekly LEA factor. The figure shows periods of highly negative correlation between the LEA and Z_t^{UNC}. For example, the decline in economic activity occurring towards the end of 2022 and early 2023 precedes a rise in weekly unemployment claims in the subsequent months.

We begin by quantifying this negative relationship computing lead-lag correlations between the local economic activity factor and Z_t^{UNC}.[60] Specifically, we compute expanding-window correlations between the weekly LEA and leads of the four-week moving-average of unemployment claims, $Corr(LEA_t^w, Z_{t+k}^{UNC})$, for $k = 4, 8, 12, 16$.[61]

[60] By construction Z_t^{UNC} lags with respect to both UNC_t and the LEA factor.

[61] We consider at least 4 lags because Z_{t+k}^{UNC} is centered between $t-2$ and $t-1$ and we expect our weekly factor to be leading the weekly unemployment claims of at least one week.

Figure 15b plots the recursive correlations computed using an expanding window approach and a warm-up period of one year of weekly data (from January 2021 to December 2021). We find a strongly negative correlation between our LEA factor and weekly unemployment claims. The average correlation, across leads and time, is -0.42, establishing a robust negative relationship between our local economic activity factor and local labor market conditions.



(a) **Time Series.** (b) **Correlations.**

Figure 15: **LEA and Weekly Unemployment Claims.** The left figure plots the (demeaned) weekly LEA factor and the (demeaned) MA(4) representation for weekly unemployment claims (Z_t^{UNC}) over the sample from January 2022 to December 2023. The right figure plots the correlation between the weekly LEA factor and the leads of the MA(4) transformation of weekly unemployment claims, Z_{t+k}^{UNC}, for $k = 4, 8, 12, 16$. The sample spans from January 2021 to December 2023. The black thick line is the mean across all lead horizons.

8 Conclusion

Measuring, modeling, and forecasting hourly economic activity at the local (zip code) level would, until recently, have been an infeasible task. It can now be achieved, we argue, through an approach that combines highly granular spatio-temporal data on variables such as foot traffic, traffic flows, electricity usage and credit card spending with more aggregate data on payroll and income through bridging equations and aggregation constraints. We show how these series share both common diurnal patterns (night versus day, week days versus weekends, season of the year) that are largely common across different locations and idiosyncratic terms that vary by series and location.

Building on this observation, we propose a dynamic factor approach that extracts a local economic activity factor while accounting for both repeated diurnal patterns as well as differences across variables and locations. We then develop an estimation methodology that nests aggregation constraints and economic restrictions within the recursions and produces estimates of the local economic activity factor. We show that this factor is highly correlated with external measures of local economic activity such as monthly payroll data or weekly unemployment claims, justifying its interpretation as a local economic activity factor. In addition, we show that our model can be used to generate accurate forecasts of the variables underlying economic activity at the hourly, daily, and weekly horizons.

By pushing the boundaries of economic modeling at finer levels of granularity, we aim to uncover new insights into the structure and dynamics of economic activity. Among several novel stylized facts of the economic activity process at the hourly zip-code level, our local economic activity factor shows a very wide intra-day range for economic activity: at its nadir around 2 a.m., economic activity declines by two-thirds while, conversely, at its peak, activity is double its average daily level. Our scenario analysis also shows that local shocks to individual variables such as foot traffic or electricity usage can have economically large spillover effects on the regional economy.

Our modeling approach is very general and can easily be extended to even finer levels of spatial granularity such as geo-fenced data (e.g., blocks of buildings or even individual buildings). Similarly, data such as foot traffic, electricity usage, and traffic flows are measured nearly continuously, so monitoring of minute-by-minute economic activity is also feasible in principle. Plant or store-level activity can also be exploited in a panel context to obtain better estimates of the underlying economy. We hope that our contribution will provide new avenues for understanding economic activity.

References

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., and Wolf, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, 83(6):2127–2189.

Altissimo, F., Cristadoro, R., Forni, M., Lippi, M., and Veronese, G. (2010). New eurocoin: Tracking economic growth in real time. *The Review of Economics and Statistics*, 92(4):1024–1034.

Arias, M. A., Gascon, C. S., and Rapach, D. E. (2016). Metro business cycles. *Journal of Urban Economics*, 94:90–108.

Aruoba, S. B. and Diebold, F. X. (2010). Real-time macroeconomic monitoring: Real activity, inflation, and interactions. *American Economic Review*, 100(2):20–24.

Aruoba, S. B., Diebold, F. X., and Scotti, C. (2009). Real-time measurement of business conditions. *Journal of Business & Economic Statistics*, 27(4):417–427.

Bańbura, M., Giannone, D., Modugno, M., and Reichlin, L. (2013). Now-casting and the real-time data flow. In *Handbook of Economic Forecasting*, volume 2, pages 195–237. Elsevier.

Bassi, A., Colacito, R., and Fulghieri, P. (2013). 'o sole mio: An experimental analysis of weather and risk attitudes in financial decisions. *The Review of Financial Studies*, 26(7):1824–1852.

Batch, A., Bridgman, B., Dunn, A., and Gholizadeh, M. (2025). Consumption zones. *Journal of Economic Geography*, 25(2):191–213.

Baumeister, C., Leiva-León, D., and Sims, E. (2024). Tracking weekly state-level economic conditions. *Review of Economics and Statistics*, 106(2):483–504.

Bok, B., Caratelli, D., Giannone, D., Sbordone, A. M., and Tambalotti, A. (2018). Macroeconomic nowcasting and forecasting with big data. *Annual Review of Economics*, 10(1):615–643.

Buda, G., Carvalho, V. M., Corsetti, G., Duarte, J. B., Hansen, S., Moura, A. S., Ortiz, Á., Rodrigo, T., Rodríguez Mora, J., and Alves da Silva, G. (2025). The short lags of monetary policy.

Camacho, M. and Perez-Quiros, G. (2010). Introducing the euro-sting: Short-term indicator of euro area growth. *Journal of Applied Econometrics*, 25(4):663–694.

Carriero, A., Clark, T. E., and Marcellino, M. (2015). Realtime nowcasting with a bayesian mixed frequency model with stochastic volatility. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 178(4):837–862.

Chen, J. C., Dunn, A., Hood, K., Driessen, A., and Batch, A. (2019). Off to the races: A comparison of machine learning and alternative data for predicting economic indicators. In *Big data for 21st century economic statistics*. University of Chicago Press.

Chetty, R., Friedman, J. N., and Stepner, M. (2024). The economic impacts of covid-19: Evidence from a new public database built using private sector data. *The Quarterly Journal of Economics*, 139(2):829–889.

Colacito, R., Hoffmann, B., and Phan, T. (2019). Temperature and growth: A panel analysis of the united states. *Journal of Money, Credit and Banking*, 51(2-3):313–368.

De Polis, A., Galvão, A. B., and Petrella, I. (2025). Tracking weekly activity using new data sources.

Demmel, J. W. (1997). *Applied Numerical Linear Algebra*. SIAM.

Desmet, K., Nagy, D. K., and Rossi-Hansberg, E. (2018). The Geography of Development. *Journal of Political Economy*, 126(3):903–983.

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 20(1):134–144.

D'Amuri, F. and Marcucci, J. (2017). The predictive power of google searches in forecasting us unemployment. *International Journal of Forecasting*, 33(4):801–816.

Ghysels, E., Sinko, A., and Valkanov, R. (2007). Midas regressions: Further results and new directions. *Econometric Reviews*, 26(1):53–90.

Giannone, D., Reichlin, L., and Small, D. (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55(4):665–676.

Hamilton, J. D. (1994). *Time Series Analysis*. Princeton University Press.

Harvey, A. (2006). Forecasting with unobserved components time series models. *Handbook of Economic Forecasting*, 1:327–412.

Harvey, A. C. (1990). Forecasting, structural time series models and the kalman filter.

Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069.

Kleinman, B., Liu, E., and Redding, S. J. (2023). Dynamic Spatial General Equilibrium. *Econometrica*, 91(2):385–424.

Klopack, B. and Luco, F. (2025). Jue insight: Measuring local consumption with payment cards and cell phone pings. *Journal of Urban Economics*.

Koop, G., McIntyre, S., Mitchell, J., Poon, A., and Wu, P. (2025a). Measuring subregional economic activity: missing frequencies and missing data. In *Recent Developments in Bayesian Econometrics and Their Applications: Festschrift in Honour of Sune Karlsson*, pages 47–65. Springer.

Koop, G., McIntyre, S., Mitchell, J., and Raftapostolos, A. (2025b). Monthly gdp growth estimates for the us states. *arXiv preprint arXiv:2501.04607*.

Koopman, S. J. and Durbin, J. (2000). Fast filtering and smoothing for multivariate state space models. *Journal of Time Series Analysis*, 21(3):281–296.

Kose, M. A., Otrok, C., and Whiteman, C. H. (2003). International business cycles: World, region, and country-specific factors. *American Economic Review*, 93(4):1216–1239.

Lewis, D. J., Mertens, K., Stock, J. H., and Trivedi, M. (2021). High-frequency data and a weekly economic index during the pandemic. In *AEA Papers and Proceedings*, volume 111, pages 326–330. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Lewis, D. J., Mertens, K., Stock, J. H., and Trivedi, M. (2022). Measuring real activity using a weekly economic index. *Journal of Applied Econometrics*, 37(4):667–687.

Marcellino, M., Porqueddu, M., and Venditti, F. (2016). Short-term gdp forecasting with a mixed-frequency dynamic factor model with stochastic volatility. *Journal of Business & Economic Statistics*, 34(1):118–127.

Marcellino, M., Stock, J. H., and Watson, M. W. (2006). A comparison of direct and iterated multistep ar methods for forecasting macroeconomic time series. *Journal of Econometrics*, 135(1-2):499–526.

Mariano, R. S. and Murasawa, Y. (2003). A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics*, 18(4):427–443.

Moench, E., Ng, S., and Potter, S. (2013). Dynamic hierarchical factor models. *Review of Economics and Statistics*, 95(5):1811–1817.

Owyang, M. T., Rapach, D. E., and Wall, H. J. (2009). States and the business cycle. *Journal of Urban Economics*, 65(2):181–194.

Redding, S. J. and Rossi-Hansberg, E. (2017). Quantitative Spatial Economics. *Annual Review of Economics*, 9(Volume 9, 2017):21–58.

Sims, C. A. and Zha, T. (1998). Bayesian methods for dynamic multivariate models. *International Economic Review*, pages 949–968.

Stock, J. H. and Watson, M. W. (1989). New indexes of coincident and leading economic indicators. *NBER Macroeconomics Annual*, 4:351–394.

Stock, J. H. and Watson, M. W. (1999). Business cycle fluctuations in us macroeconomic time series. *Handbook of Macroeconomics*, 1:3–64.

Stock, J. H. and Watson, M. W. (2002). Macroeconomic forecasting using diffusion indexes. *Journal of Business & Economic Statistics*, 20(2):147–162.

Stock, J. H. and Watson, M. W. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In *Handbook of Macroeconomics*, volume 2, pages 415–525. Elsevier.

Waggoner, D. F. and Zha, T. (1999). Conditional forecasts in dynamic multivariate models. *Review of Economics and Statistics*, 81(4):639–651.

Wallis, K. F. (1974). Seasonal adjustment and relations between variables. *Journal of the American Statistical Association*, 69(345):18–31.

A Data

In this appendix we report data sources and details of the variables used to construct our spatio-temporal granular dataset. We begin by reporting a description of each data series and any pre-processing step that we apply. We then present additional empirical features of the data series.

A.1 Data Description and Pre-Processing Operations

We organize our data discussion starting with the most granular variables and progressing toward the most aggregated ones.

Data on foot traffic and spending patterns are sourced from Dewey Data: `https://www.deweydata.io/`. Dewey is a data service provider for academics whose subscription guarantees access to several datasets.

Foot Traffic Data. The foot traffic data is sourced from Dewey, which in turn obtains its data from Advan Research.[62] This dataset provides aggregated counts of visits to points of interest (POIs) from a panel of mobile devices at an hourly frequency across U.S. states. Advan uses geolocation data from mobile phones to detect when a device enters the boundaries of a given POI and records such instances as visits. The data are available from 2018 onward. The identification of POIs follows the same classification used by SafeGraph. We map each POI to its corresponding zip code based on geometry and sum the number of visits for each zip code at the hourly level. The time series reflect daylight saving time adjustments, resulting in either 25-hour or 23-hour days, depending on the direction of the clock change. To align all dates to a consistent 24-hour format, we proceed as follows. On the first Sunday in November, when clocks move back one hour, we average the visit counts across the two 1:00 a.m. hours. Conversely, during the spring transition, when the 2:00 a.m. hour is skipped, we impute the missing hour by using the visit count from the corresponding hour of the previous week.

Traffic Flow Data. The traffic flow data is retrieved from the Performance Measurement System (PeMS) on Caltrans. Traffic flow is defined as the number of vehicles passing by a given location within a given amount of time. The data is available from 2014 and collected in real time at the hourly frequency from individual detectors on highways across all major metropolitan areas of California. We map the location of detectors to zip codes and sum the traffic flow for each zip code at hourly intervals. We apply the same approach to daylight saving time adjustments as used in the processing of the foot traffic data.

Electricity Usage Data. The electricity usage data is obtained from the San Diego Gas and Electric (SDG&E) company, a major electricity supplier in San Diego county. SDG&E furnished us with zip-code level and hourly interval readings of total electricity usage, disaggregated by revenue class: agricultural, commercial, industrial, and

[62] `www.advanresearch.com`.

residential, available from 2019 to 2023. The data sampling follows an anonymization procedure known as the 5% rule which ensures that no individual data point represents more than 5% of the sample to protect user privacy. We aggregate all four categories and calculate the whole electricity volume for each zip code. Daylight saving time adjustments are handled in the same manner as applied to the visit data.[63]

Weather Data. We obtain weather data from the National Centers for Environmental Information (NCEI), which provides weather variables recorded at numerous monitoring stations.[64] The weather variables include variables such as temperature, humidity, rainfall, and emissions and are available at the hourly frequency. We choose to include only temperature in our granular dataset, as we expect it to provide a good overall summary of the weather component most correlated with high-frequency economic activity. Temperature is available in two forms: dry bulb and wet bulb. Both are included in the NCEI dataset. We focus on dry bulb temperature, as it represents the standard air temperature measured by a dry thermometer. It is generally considered the closest approximation to human thermal perception and is therefore more likely to be correlated with economic activity. Although we do not include other weather variables such as air pollution at this stage they represent a promising extension to our dataset and are left for future research.

Spending Data. We obtain our spending data from DEWEY, which sources it from SafeGraph, one of the most accurate providers of point-of-interest (POI) data.[65] The dataset contains aggregated and anonymized credit and debit card transactions associated with specific POIs, recorded daily across U.S. locations. SafeGraph provides transaction data at the POI level, including various characteristics of each transaction, such as whether it occurred online or in person.[66] The data is available at a daily frequency starting in 2019 and covers more than 20 million POI locations. We incorporate spending data by aggregating all POIs within each zip code and record the spending data for day d on the first hour of day $d+1$. This ensures that the full economic activity occurring on day d is elapsed before the data are used. The underlying panel is unbalanced and time-varying, with changes in coverage over time. The data provider matches the transaction data to reflect the actual date of purchase but documents evidence of a systematic lag only on Saturdays and Sundays which results in higher reported spending on Mondays.

We use the following data pre-processing steps to ensure robustness with respect to the documented measurement lag and the changing size of the POI panel. First, to correctly attribute weekend spending data, we redistribute the excess spending recorded on Mondays to Saturdays and Sundays. Specifically, we allocate 40% of Monday's spending measurement across the two weekend days.[67] Second, due to changing coverage, SafeGraph documents a declining panel size over time. To ensure that the spending data reflect reliable patterns, we de-trend the raw data using the trend documented by Safe-

[63]The data provider informed us of documented issues in their measurement system, which result in substantial mis-measurement for March 2021. To address this, we impute the missing or unreliable data at the zip-code/hour level using the average of the corresponding observations from February and April.

[64]`https://www.ncei.noaa.gov/access/search/data-search/local-climatological-data`

[65]`www.safegraph.com`

[66]Purchases made online and shipped to a residence are not included, as they are not associated with a physical POI. However, pick-up and delivery orders are included.

[67]The 40% share is a conservative, data-informed heuristic that balances underreporting on weekends with potential overstatement on Mondays, improving the temporal alignment of spending signals.

Graph, and we correct this trend using data from Bloomberg Second Measure.[68] As an additional robustness check, we apply a similar methodology using the 'large bank consumer credit card balances' from the Federal Reserve Bank of St. Louis (FRED).[69]

Weekly Unemployment Insurance Claims. We obtain weekly unemployment insurance claims from the California Employment Development Department (CAEDD). The CAEDD computes new, additional and transitional claims at the weekly-county level. We use data on new claims, where a new claim is defined as the first claim for a benefit year period. The measurement process of the CAEDD ensures that each individual can have at most one single new claim during a benefit year period.

Payroll Data. We obtain data on earnings from the California Employment Development Department (CAEDD) which collaborates with the U.S. Bureau of Labor Statistics to conduct an employer survey collecting information on employment and payroll. Participation in this survey is mandatory for all firms with at least 1,000 employees and a sample of smaller firms is also included. The California EDD reports a long-standing consistency in survey participation. The data reflect individuals who were employed during, or received compensation for, any portion of the pay period that includes the 12th day of the month. Employers report this information based on payroll records maintained for tax and accounting purposes. Following the standard definition, we define local payroll as the total labor compensation paid within a given location. We construct this measure at the county/MSA level by combining the monthly number of employed individuals with the average weekly wage. We use the civilian employment category from the Official Estimates of Industry Employment as our measure of the employed labor force. To compute monthly payroll, we first convert the average weekly wage to a monthly equivalent by multiplying it by the average number of weeks in a month. To account for the varying number of days and weeks across different months we use as scaling factor of 4.4, which is the average number of weeks within each month. We then multiply this monthly wage by the number of employed individuals. The resulting estimate is recorded at 8:00 a.m. on the first day of the following month to align it with other time-stamped data in our dataset. From the County Business Patterns (CBP) dataset we obtain employment data at the zip-code level for the annual frequency. We use these data to construct a panel of zip-code monthly payroll data in the following way. For each year in the dataset, we compute the proportion of payroll in each zip code. Next, we assume that this proportion does not vary within the year, and we use it to compute the monthly zip-code employees. We then multiply the monthly zip-code employees by the average monthly wage to obtain the panel of monthly zip-code payroll.

GDP. Gross Domestic Product (GDP) data are obtained from the U.S. Bureau of Economic Analysis (BEA) which provides both county-level and state-level estimates. While state-level GDP is available at a quarterly frequency, county-level GDP is only reported annually. In our dataset, we assign the county GDP for year $yyyy$ to the first day of the following year, i.e., January 1st, $yyyy + 1$, at 1:00 a.m.[70] This timing ensures that the GDP measurement reflects the full elapse of the underlying year's economic activity. We use GDP in current dollar terms, thereby maintaining consistency across all nominal

[68] `secondmeasure.com`

[69] `https://fred.stlouisfed.org/series/RCCCBPURCHASETOT`

[70] For example, GDP for 2023 is recorded at 1:00 a.m. on January 1st, 2024.

variables.

Variable Transformations Let y^{raw}_{s,i,t_h} denote the raw observation for series y_i at location s and time t_h. In our empirical analysis, we apply the following transformation:

$$y_{s,i,t_h} = \log\left(1 + y^{\text{raw}}_{s,i,t_h}\right).$$

This transformation, also known as the *log1p* transformation, is widely used in the data science and machine learning literature. It is particularly useful when raw observations can take large values but may also include zeros, as is often the case with alternative data sources.

Nominal Variables. All variables expressed in dollar terms such as GDP, payroll, and spending, are considered in nominal values.

Data Series Summary

Series	Spatial Level	Temporal Level	Sample	Source
Foot Traffic	zip code	hourly	2018 - 2024	Advan
Traffic	zip code	hourly	2014 - 2024	Caltrans
Electricity	zip code	hourly	2018 - 2023	SDG&E
Weather	County	hourly	2001 - 2024	NCEI
Spend	zip code	daily	2019 - 2024	SafeGraph
Weekly Unemployment Claims	County/MSA	weekly	2019 - 2025	CAEDD
Average Weekly Earnings	County/MSA	monthly	2010 - 2024	CAEDD
Employment	County/MSA	monthly	2010 - 2024	CAEDD
Payroll and Employment	zip code	yearly	1994 - 2024	CBP
GDP - San Diego	County/MSA	yearly	2001 - 2023	BEA

Table A.1: **Data Series.** This table reports a summary of the data series used to assemble the spatio-temporal granular dataset. For each series, we provide the spatial and temporal dimension, the availability of the sample, and the data source. Data on foot traffic (Advan) and spending (SafeGraph) are obtained from Dewey. SDG&E stands for San Diego Gas & Electric.

B Individual Payroll Regressions

Here we provide additional details on the univariate zip-code level payroll regressions reported in Table 1 of Section 2.3. These payroll regressions take the form:

$$\Delta PAY_{s,m} = b_0 + b_{TF}\Delta TF_{s,m} + \phi\Delta PAY_{s,m-1} + \nu_{s,m} + e_{s,m},$$
$$\Delta PAY_{s,m} = b_0 + b_{FT}\Delta FT_{s,m} + \phi\Delta PAY_{s,m-1} + \nu_{s,m} + e_{s,m},$$
$$\Delta PAY_{s,m} = b_0 + b_{ElU}\Delta ElU_{s,m} + \phi\Delta PAY_{s,m-1} + \nu_{s,m} + e_{s,m},$$
$$\Delta PAY_{s,m} = b_0 + b_{Temp}\Delta Temp_{s,m} + \phi\Delta PAY_{s,m-1} + \nu_{s,m} + e_{s,m},$$
$$\Delta PAY_{s,m} = b_0 + b_{SPEND}\Delta SPEND_{s,m} + \phi\Delta PAY_{s,m-1} + \nu_{s,m} + e_{s,m}.$$

FG: I modified these equations based on the new payroll regression specification. We can keep Figure B.1 or we can remove it, as with the new approach we do not have the individual coefficients.

Figure B.1 reports histograms of the estimates of the coefficients of the univariate payroll regressions for traffic flow, foot traffic, electricity usage and spending across the 62 zip codes in our sample. None of the coefficients result in a negative point estimate, providing further evidence of the positive correlation between the alternative series (foot traffic, traffic flows, electricity usage, and spending) and economic aggregates.

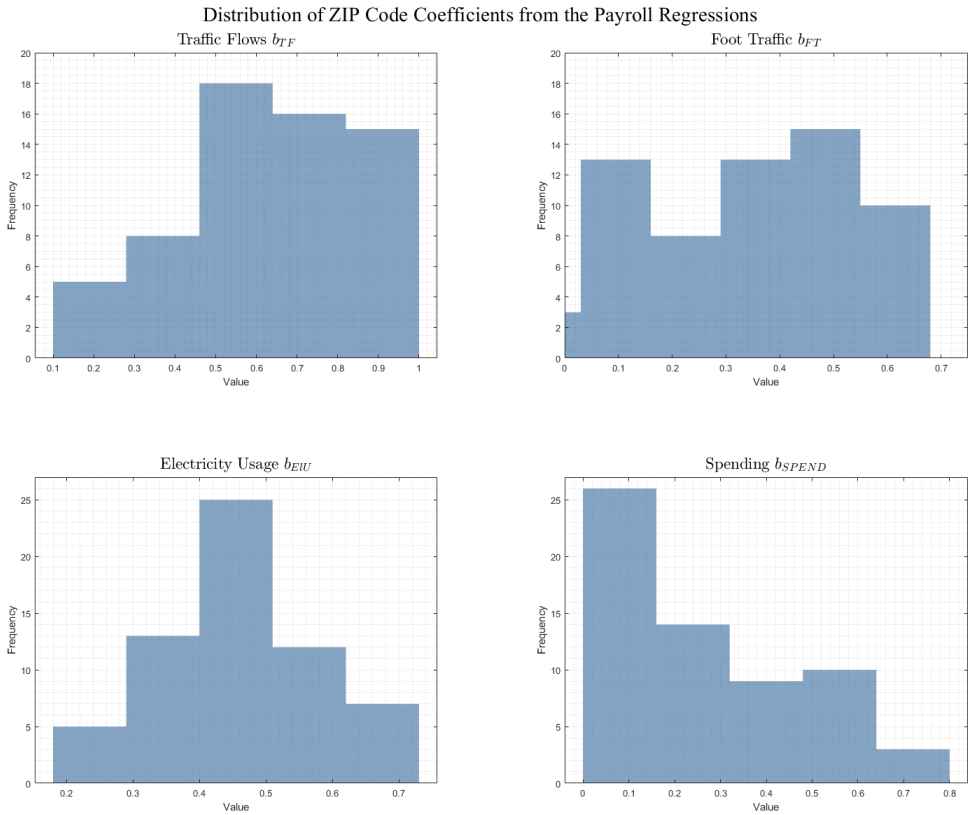


Figure B.1: **Histogram of Coefficients Individual Payroll Regression.**

C Economic Identification

C.1 Identifying Economic Activity Components

Consider a set of granular variables $y^h_{s,t_h},\ s \in S_G$, and an aggregate $y^h_{S_A,t_h}$. Following our model in Section 3.1, the process describing the evolution in the granular variables is

$$y^h_{s,t_h} = \varrho_s + \beta_s f^h_{t_h} + \epsilon^h_{s,t_h}, \quad s \in S_G,$$

while that for the aggregate variable is

$$y^h_{S_A,t_h} = \varrho_{S_A} + \beta_{S_A} f^h_{t_h} + \epsilon^h_{S_A,t_h},$$

where $f^h_{t_h}$ is the unobserved common factor.

The observation equation that stacks all variables is given by

$$y^h_{t_h} = \begin{bmatrix} y^h_{S_A,t_h} \\ y^h_{s_1,t_h} \\ \vdots \\ y^h_{S_G,t_h} \end{bmatrix}.$$

Its variance-covariance matrix is

$$\Sigma_y = B\Sigma_f B' + \Sigma_\epsilon,$$

where

$$B = \begin{bmatrix} \beta_{S_A} \\ \beta_{s_1} \\ \vdots \\ \beta_{S_G} \end{bmatrix}, \quad \Sigma_f = Var(f^h_{t_h}), \quad \Sigma_\epsilon = Var(\epsilon^h_{t_h}).$$

The variance-covariance matrix of the factor is given by

$$\Sigma_f = \begin{bmatrix} \sigma^2_\mu & 0 \\ 0 & 0 \end{bmatrix}.$$

For identification purposes, we impose that the loading of the slow-moving mean on the aggregate variable is equal to 1, i.e., $\beta_{\mu,S_A} = 1$.[71]

To see the effects of these identification restrictions, consider the case $\beta_{\mu,S_A} \neq 1$. By

[71] To simplify the argument, we examine the case with scalar factor and aggregate variable.

rescaling the factor loadings in B by $1/\beta^2_{\mu,S_A}$, we obtain:

$$B\Sigma_f B' = \sigma^2_\mu \begin{bmatrix} 1 \\ \beta_{\mu,s_1}/\beta_{\mu,S_A} \\ \vdots \\ \beta_{\mu,S_G}/\beta_{\mu,S_A} \end{bmatrix} \begin{bmatrix} 1 & \beta_{\mu,s_1}/\beta_{\mu,S_A} & \cdots & \beta_{\mu,S_G}/\beta_{\mu,S_A} \end{bmatrix},$$

$$= \sigma^2_\mu \begin{bmatrix} 1 & \beta_{\mu,s_1}/\beta_{\mu,S_A} & \cdots & \beta_{\mu,S_G}/\beta_{\mu,S_A} \\ \beta_{\mu,s_1}/\beta_{\mu,S_A} & \left(\frac{\beta_{\mu,s_1}}{\beta_{\mu,S_A}}\right)^2 & \cdots & \left(\frac{\beta_{\mu,s_1}\beta_{\mu,S_G}}{\beta^2_{\mu,S_A}}\right) \\ \vdots & \vdots & \cdots & \vdots \\ \beta_{\mu,S_G}/\beta_{\mu,S_A} & \left(\frac{\beta_{\mu,s_1}\beta_{\mu,S_G}}{\beta^2_{\mu,S_A}}\right) & \cdots & \left(\frac{\beta_{\mu,S_G}}{\beta^2_{\mu,\mu,S_A}}\right)^2 \end{bmatrix}.$$

This shows that it would be possible to increase the factor loading of GDP without changing the distribution of the observables. The restriction $\beta_{\mu,S_A} = 1$ anchors the factor to the aggregate variable.

C.2 The $GDP-$Payroll Relationship

In this section, we show how we obtain a measure of the relationship between $\log(GDP_t)$ and $\log(PAY_t)$, which we use to impose an additional identification restriction in the estimation of the LEA factor.

We start from the linear regression

$$\Delta \log(GDP_t) = a + b\Delta \log(PAY_t) + e_t,$$

where GDP_t is annual San Diego GDP and PAY_t is the total pay variable for San Diego in year t, constructed from the monthly (m) payroll data as $PAY_t = \sum_{m \in year\ t} PAY_{month}$ and $PAY_m = PAYROLL_m \times WAGE_m$.[72]

To obtain the relationship between $\log(GDP_t)$ and $\log(PAY_t)$ we write

$$\Delta \log(GDP_t) = a + b\Delta \log(PAY_t) + e_t,$$

from which we obtain

$$\log(GDP_t) = a + \log(GDP_{t-1}) + b\log(PAY_t) - b\log(PAY_{t-1}) + e_t.$$

We re-write the last equation as

$$\log(GDP_t) = \widetilde{a}_{t-1} + b\log(PAY_t) + e_t, \tag{31}$$

where

$$\widetilde{a}_{t-1} \equiv a + \log(GDP_{t-1}) - b\log(PAY_{t-1}).$$

[72] This step is necessary to match the annual frequency of GDP_t.

We use the relation in equation (31) to obtain an estimate of b, and use this coefficient as *a priori* information to impose a restriction on the loading of the LEA factor on PAY_t in the filtering algorithm.

Specifically, we estimate (31) using San Diego data at the annual frequency over the period 2001-2023 and we obtain a point estimate of b equal to 0.719.

To link this back to the economic activity factor, we have

$$\begin{aligned}
\underbrace{\log(GDP_t)}_{f_t^{yr}+\epsilon_{GDP,t}^{yr}} &= \widetilde{a}_{t-1} + b\log(PAY_t) + e_t, \\
b\log(PAY_t) &= \log(GDP_t) - \widetilde{a}_{t-1} - e_t, \\
b\log(PAY_t) &= f_t^{yr} + \epsilon_{GDP,t}^{yr} - \widetilde{a}_{t-1} - e_t, \\
\log(PAY_t) &= \underbrace{\frac{1}{b}}_{\gamma} f_t^{yr} + \frac{1}{b}(\epsilon_{GDP,t}^{yr} - \widetilde{a}_{t-1} - e_t).
\end{aligned}$$

Defining

$$\epsilon_{PAY,t}^{yr} := \frac{1}{b}(\epsilon_{GDP,t}^{yr} - \widetilde{a}_{t-1} - e_t),$$

we obtain the economic-anchoring equation:

$$\log(PAY_t) = \gamma f_t^{yr} + \epsilon_{PAY,t}^{yr},$$

where $\gamma = \frac{1}{b} = \frac{1}{0.719} = 1.39$, the economic anchoring restriction discussed in Section 3.

C.3 Incremental Value of New Variables

To help better understand how adding information from alternative data series would affect our estimates of granular economic activity, we provide details of how our approach uses information contained in the observable data series to estimate local economic activity.

Consider an initial information set $\Omega_{t_h}^1$. Suppose we augment this with another series to obtain $\Omega_{t_h}^2 = \{\Omega_{t_h}^1, y_{S_{\mathcal{G}},2,t_h}^h\}$. We assume that $y_{S_{\mathcal{G}},2,t_h}^h$ contains information on $S_{\mathcal{G}}$ granular locations and so is $S_{\mathcal{G}}-$dimensional.

Updates of the latent state estimate in an unobserved component model takes the form:[73]

$$E(f_{t_h}|\Omega_{t_h}^2) = E(f_{t_h}|\Omega_{t_h}^1) + \omega_{2,t_h} \times \left(y_{S_{\mathcal{G}},2,t_h}^h - E(y_{S_{\mathcal{G}},2,t_h}^h|\Omega_{t_h}^1)\right), \tag{32}$$

where ω_{2,t_h} is an $S_{\mathcal{G}}-$dimensional vector of weights that determines the importance of observations $y_{S_{\mathcal{G}},2,t_h}^h$ in updates of the state, and $y_{S_{\mathcal{G}},2,t_h}^h - E(y_{S_{\mathcal{G}},2,t_h}^h|\Omega_{t_h}^1)$ is a news term. The weights can be expressed as (see Hamilton (1994) p.390)

$$\omega_{2,t_h} = (\beta_2 Var(f_{t_h}|\Omega_{t_h}^1)\beta_2' + \Sigma_2)^{-1} Var(f_{t_h}|\Omega_{t_h}^1)\beta_2', \tag{33}$$

[73] See Harvey (1990) and Hamilton (1994).

where β_2 is the $S_{\mathcal{G}}-$dimensional vector of factor loadings on the observables $y^h_{S_{\mathcal{G},2,t_h}}$, and Σ_2 is the variance-covariance matrix of $y^h_{S_{\mathcal{G},2,t_h}}$.

Equation (33) shows that the contribution of the series $y^h_{S_{\mathcal{G},2,t_h}}$ to the filtering of the unobserved state depends on the strength of the relationship between f_{t_h} and $y^h_{S_{\mathcal{G},2,t_h}}$ given by the loadings β_2, and the noise of the series, Σ_2. Series that are poor proxies for local economic activity have either sparse (granular) loadings or a low signal-to-noise ratio.

Adding series that are uncorrelated with local economic activity or have a low signal-to-noise ratio should therefore not introduce a sizable bias in the estimates of the latent state, particularly in large samples where ω_{2,t_h} will be close to zero (see equation (32)) and the state filtered conditional on using the information in $y^h_{S_{\mathcal{G},2,t_h}}$ is nearly identical to that obtained using only $\Omega^1_{t_h}$: $E(f_{t_h}|\Omega^2_{t_h}) = E(f_{t_h}|\Omega^1_{t_h})$.

Adding redundant series to our model therefore does not seriously impair our estimates of the economic activity factor obtained with a smaller information set $\Omega^1_{t_h}$.

D Granular Economic Modeling: Further Results

D.1 Likelihood Decomposition

This appendix reports details of how we construct the likelihood function for our granular modeling approach.

We begin by noting that the hierarchical structure we use to model the variables observed at different time frequencies induces a natural subsetting of the parameter space. We denote the parameter vector containing the parameters of the τ block by θ^τ while the parameters that do not depend on the time frequency are denoted by $\theta^\dagger$. Finally, we denote the parameter vector that contains *all* parameters by $\widetilde{\theta}$:

$$\widetilde{\theta} = (\theta^{yr}, \ldots, \theta^\tau, \ldots, \theta^h, \theta^\dagger).$$

The joint log likelihood of our granular model introduced in equation (10) is given by

$$\log L = \text{constant} - \frac{1}{2}\sum_{t_h=1}^{T} \log|F_{t_h}| - \frac{1}{2}\sum_{t_h=1}^{T} v'_{t_h} F^{-1}_{t_h} v_{t_h},$$

where F_{t_h} is the conditional variance of the observation vector at time t_h, $Var(y_{t_h}|\mathcal{Y}_{t_h-1})$, and v_{t_h} is the prediction error.

The observation vector, $y_{t_h} = [y^{yr}_{t_h}, \ldots, y^\tau_{t_h}, \ldots, y^h_{t_h}]'$, contains data sampled at different

frequencies. Building on this hierarchical structure, we rewrite the log likelihood as

$$\log L = \text{constant} - \frac{1}{2}\sum_{\tau=h}^{yr}\left(\sum_{t_h=1}^{T}\log|F_{t_h}^{\tau}| - \frac{1}{2}\sum_{t_h=1}^{T} v_{t_h}^{\tau\prime} F_{t_h}^{\tau\,-1} v_{t_h}^{\tau}\right)$$

$$\log L = \sum_{\tau=h}^{yr}\log L^{\tau}.$$

The $\tau-$likelihood depends on $(\theta^{\tau}, \theta^{\dagger})$, where the dependence on the common parameters, $\theta^{\dagger}$, is derived from the spatio-temporal aggregation restrictions of Section 3.3. Specifically, because the $\tau-$frequency variables depend on $\alpha_{t_h}^{\tau}$, which is the temporally aggregated counterpart of $\alpha_{t_h}^{h}$, the underlying parameters of the distribution that governs the state vector are those of $\alpha_{t_h}^{h}$.

The likelihood function can be decomposed further by exploiting the filtering approach proposed by Koopman and Durbin (2000). In our case, we have:

$$\log L = \sum_{\tau=h}^{yr}\log L^{\tau} = \sum_{\tau=h}^{yr}\sum_{i=1}^{N_{\tau}}\log L_i^{\tau},$$

where i denotes the $i-$th series of the $\tau-$block, and N_{τ} denotes the total number of series in the same block. This likelihood function is the one that we optimize in our estimation algorithm.

D.2 Further Empirical Results

Figure D.1 plots the weekly average of the hourly LEA factor and its slow-moving mean component in deviations from the mean hourly local economic activity expressed in percentage unit.

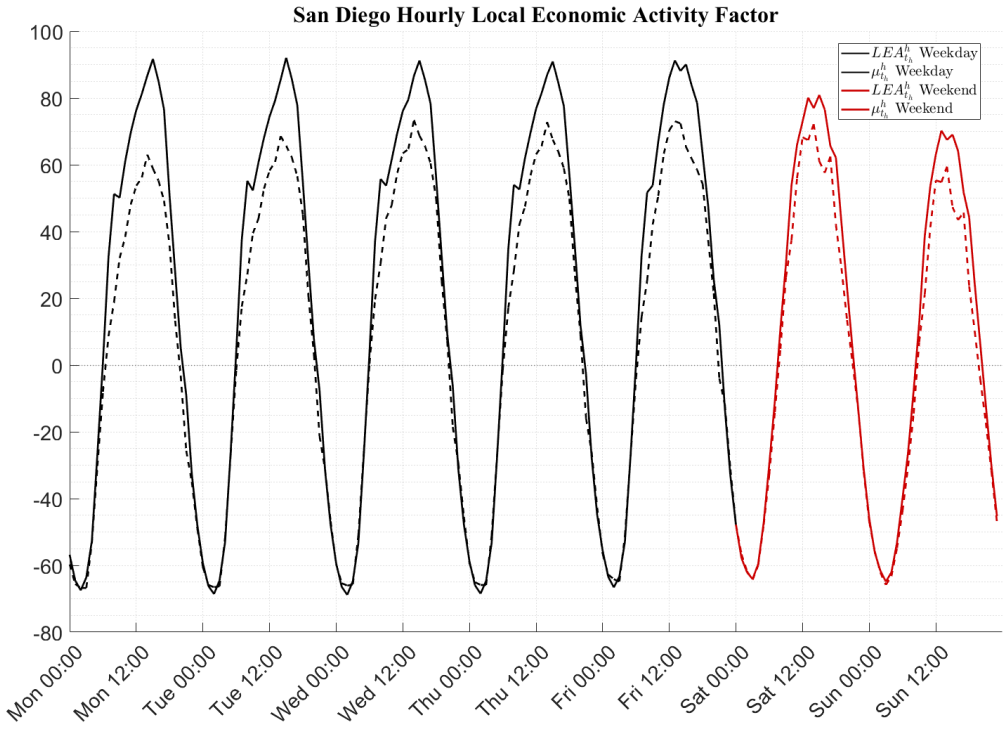


Figure D.1: **Hourly Local Economic Activity Factor** This figure plots the average hourly San Diego local economic activity factor expressed in percentage deviations from its hourly mean. The solid line tracks $LEA^h_{t_h}(\Delta\%)$, while the broken line plots $\mu^h_{t_h}(\Delta\%)$.

Figure D.2 plots the distributions of the granular loadings of the calendar components on the series for traffic flow, foot traffic and electricity consumption.

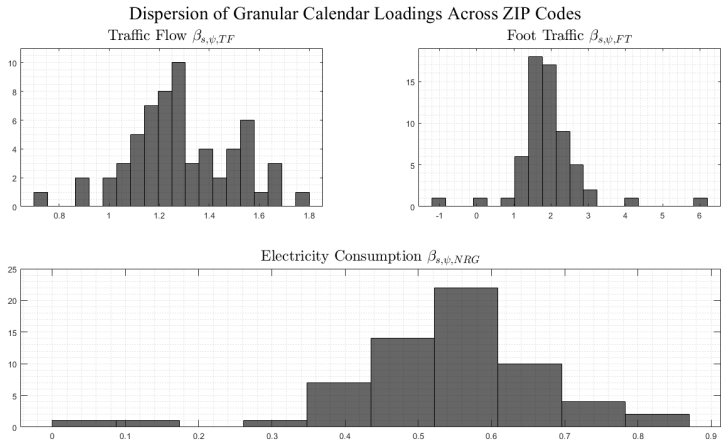
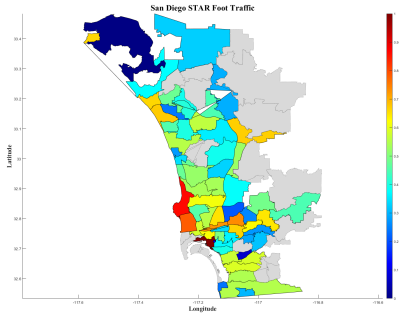


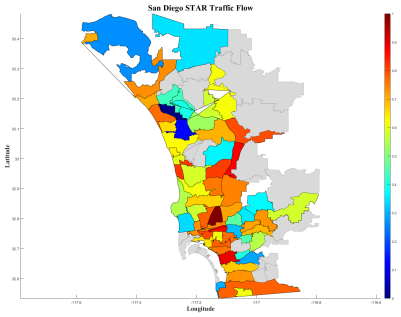
Figure D.2: **Granular Calendar Factor Loadings.** This figure plots the histograms of the granular calendar betas of the hourly LEA factor, $\beta_{s,\psi,i}$, for hourly traffic flow (upper left plot), hourly foot traffic (upper right panel), and hourly electricity consumption (bottom panel) for each of the 62 San Diego zip codes in our dataset.

D.2.1 STAR: Individual Series

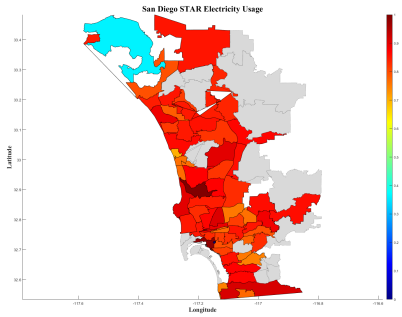
We next explore the role of the individual series in the spatio temporal activity radar introduced in Section 5.4. Specifically, in Figure D.3 we construct the STAR using only one series at a time. The figure shows that the STARs from the individual series have different dynamics, highlighting a role for each series in accounting for different spatio-temporal patterns. For example, we see that there is little heterogeneity across space in electricity usage, consistent with electricity having a rigid spatial demand. Conversely, the STARs of foot traffic, traffic flow, and spending show substantial spatial variation across zip codes.



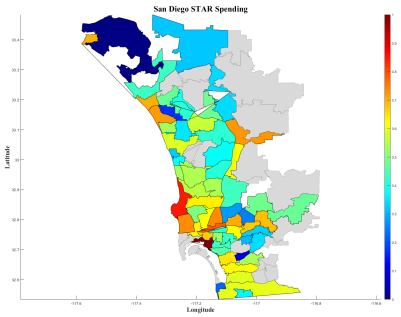
(a) **STAR Foot Traffic.**



(b) **STAR Traffic Flow.**



(c) **STAR Electricity**



(d) **STAR Spending.**

Figure D.3: **STAR plots constructed from individual series**. This figure reports the STAR computed as in Section 5.4 using the individual series on hourly foot traffic (top-left panel), hourly traffic flow (top-right panel), hourly electricity usage (bottom-left panel), and daily spending (bottom-right panel). Grey areas are zip codes within San Diego county that are excluded due to incomplete data.

D.3 Forecasting

D.3.1 Hour-Ahead Forecasts

The models used for the short-term forecasts are summarized by:

$$\begin{aligned}
\text{benchmark aggregate level:} \quad & \widehat{y}^{h}_{SD,i,t_h+1|t_h,bmk_{SD}} = y^{h}_{SD,i,t_h}, \\
\text{benchmark zip-code level:} \quad & \widehat{y}^{h}_{zip,i,t_h+1|t_h,bmk_{zip}} = \widehat{a}_{zip,i} + \widehat{b}_{zip,i}\overline{y}^{h}_{\overline{zip},i,t_h}, \\
\text{surrogate-proxy forecasts:} \quad & \widehat{y}^{h}_{zip,i,t_h+1|t_h} = \widehat{a}_{zip,i} + \widehat{b}_{zip,i}\ \widetilde{LEA}^{h}_{SUR,t_h+1|t_h}, \\
\text{granular economic forecasts:} \quad & \widehat{y}^{h}_{zip,i,t_h+1|t_h} = \widehat{\varrho}_{zip,i} + \widehat{\beta}_{zip,i}\widehat{f}^{h}_{t_h+1|t_h} + \widehat{\phi}\widehat{\epsilon}^{h}_{zip,i,t_h|t_h}.
\end{aligned}$$

Figure D.4 plots hourly one-step-ahead forecasts of local economic activity. The solid and dashed lines plot $\widehat{LEA}^{h}_{t_h+1}$ and $\widehat{\mu}^{h}_{t_h+1}$, *respectively*.

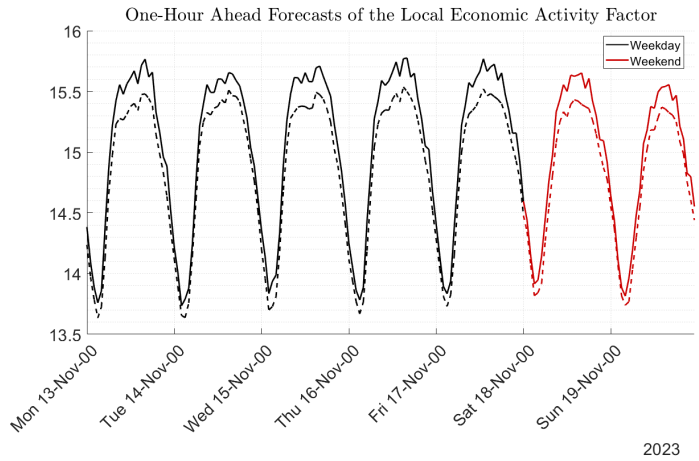


Figure D.4: **One-step-ahead forecasts of hourly Local Economic Activity.** This figure illustrates the one-hour-ahead forecasts of San Diego local economic activity dynamics. The solid and dashed lines plot $\widehat{LEA}^{h}_{t_h+1}$ and $\widehat{\mu}^{h}_{t_h+1}$, respectively. The plot goes from 00 a.m., Monday, November 13^{th} 2023 to 00 a.m. Monday, November 20^{th} 2023.

D.4 Robustness Checks for the Scenario Analysis

We next conduct some robustness checks for the scenario analysis reported in Section 6.2.

D.4.1 Regularization Schemes

We begin by assessing the robustness of the scenario analysis to the regularization of the variance-covariance matrix of the observation vector that we use use to generate the results reported in the main text.

I. Ridge-type Regularization

We invert the variance-covariance matrix $\Sigma_{y|Y}$ using the following Ridge-like regularization:

$$\left(\Sigma_{y|\mathcal{Y}} + \lambda I\right)^{-1}, \tag{34}$$

where λ is a penalty parameter and I is a conformable identity matrix.

Letting μ_S be the vector of means of the granular variables, we study the robustness of this method for four different choices of the penalty parameter:

$$\begin{aligned}
\lambda_0 &= \left(\frac{1}{S}\sum_s \mu_{S,s}\right)^{-1}, \\
\lambda_1 &= \min_s \frac{1}{\mu_{S,s}}, \\
\lambda_2 &= \max_s \frac{1}{\mu_{S,s}}, \\
\lambda_3 &= diag(\mu_S^{-1}).
\end{aligned}$$

II. Low-rank Approximation

As a second regularization method, we consider a low-rank approximation of the variance-covariance matrix that discards near-zero eigenvalues.[74] Let Λ be the diagonal matrix that contains the eigenvalues of $\Sigma_{y|\mathcal{Y}}$ on its main diagonal. If an eigenvalue is smaller than the threshold $10^{-6}\max\{eig(\Sigma_{y|\mathcal{Y}})\}$, we truncate it. This yields a low-rank approximation to $\Sigma_{y|\mathcal{Y}}$, $\Sigma^{LR}_{y|\mathcal{Y}}$ whose inverse, $(\Sigma^{LR}_{y|\mathcal{Y}})^{-1}$, we use in our scenario calculations.

D.4.2 Multivariate vs Univariate Approach

We can also compare our estimates of scenario effects under different approximations of $\Sigma_{y|\mathcal{Y}}$ to those from an approximation that assumes it to be a diagonal matrix. We refer to the first approach as multivariate, and the second as univariate.

Under the multivariate approach, we compute the scenario effects using the expression:

$$E[f^h_{t_h+1}|\mathcal{E}_{t_h+1}] = E[f^h_{t_h+1}|\mathcal{Y}_{t_h}] + \Sigma_{f,y|\mathcal{Y}}\Sigma^{-1}_{y|\mathcal{Y}}\left(y_{t_h+1}(\mathcal{E}_{t_h+1}) - E[y_{t_h+1}|\mathcal{Y}_{t_h}]\right),$$

with weights

$$\Sigma_{f,y|\mathcal{Y}}\Sigma^{-1}_{y|\mathcal{Y}},$$

where $\Sigma^{-1}_{y|\mathcal{Y}}$ is the 'full' variance-covariance matrix estimated using either the Ridge-type regularization (equation (34)) or the low-rank approximation.

In the univariate case, we instead consider

$$E[f^h_{t_h+1}|\mathcal{E}_{t_h+1}] = E[f^h_{t_h+1}|\mathcal{Y}_{t_h}] + \Sigma_{f,y|\mathcal{Y}}D^{-1}_{y|\mathcal{Y}}(y_{t_h+1}(\mathcal{E}_{t_h+1}) - E[y_{t_h+1}|\mathcal{Y}_{t_h}]),\ t_h \in T_h, \tag{35}$$

[74] See Demmel (1997).

where $D_{y|\mathcal{Y}}$ is a diagonal matrix that contains the variances of the y_{t_h} granular units on its main diagonal. This model leads to the following expression for the weights:

$$(\Sigma_{f,y|\mathcal{Y}} D_{y|\mathcal{Y}}^{-1})_{s,i} = \frac{\sigma_{f,y_{s,i}}}{\sigma^2_{y_{s,i}}} \equiv \beta_{s,i} \times \frac{Var(f)}{Var(y_{s,i})}.$$

Figure D.5 reports the results of our robustness analysis. The plot shows that the maximum value of the scenario effect is obtained for the smallest value of the penalty parameter in the Ridge-type regularization ($\lambda_1 = 0.086$). Moderate values of the penalty ($\lambda_0 = 0.12$ and λ_3 which imposes different penalties on each term) result in intermediate scenario effects. For the highest penalty ($\lambda_2 = 0.29$), the effect is similar to the effect from the low-rank approximation. Unsurprisingly, the smallest scenario effect is observed for the univariate approach which neglects the covariance terms in $\Sigma_{y|\mathcal{Y}}$ and, in a sense, provides a lower-bound approximation for the total effect.

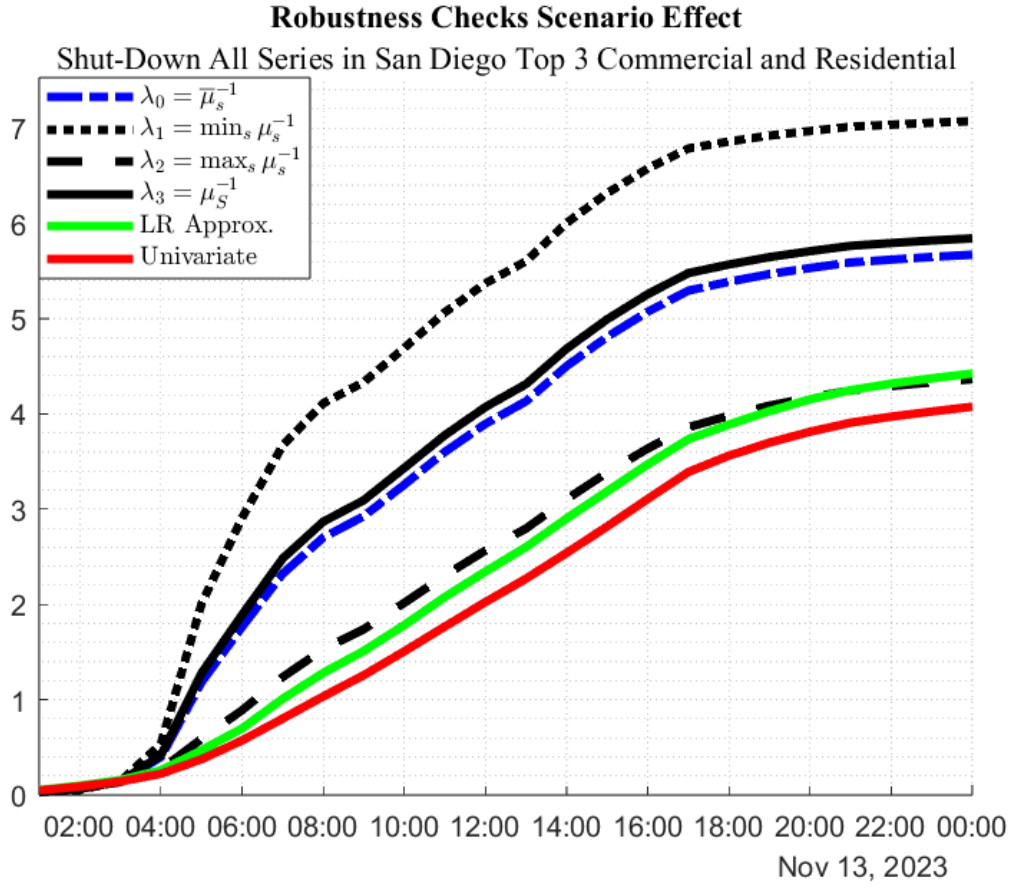


Figure D.5: **Scenario Analysis: Robustness Checks** This figure plots the scenario effects under different regularization methods. We consider the same exercise as in Figure (12d) and refer to the main text and the Figure's caption for a detailed description of the scenario analysis. The cumulative expected scenario effect is expressed as a percentage of the daily San Diego GDP.