


Nighttime Light Dynamics as a Geographic Risk Signal for Municipal Bonds

Abstract: This study examines whether satellite-derived nighttime-light measures can serve as an alternative indicator of municipal bond risk that can be an indication to the future trading. Focusing on transportation-related municipal bonds with similar geographic exposure, we analyze whether changes in nighttime-light intensity and spatial clustering are associated with bond-price movements over a ten-trading-day horizon. Because municipal bond prices are strongly influenced by interest-rate movements, the analysis first isolates the portion of price variation related to Treasury rates and then evaluates whether night while considering the nearby toll-gate revenue to debt maturity profiles:

The bond that is covered is CUSIP **89602HAB3**, a municipal bond that covers 9 bridges of **MTA in new York**. The code for the following analysis can be found at:

https://github.com/sup21072/bridge_alternative



=== Average MoM % by facility (over study period) ===

facility

Bronx-Whitestone Bridge	9.13
Hugh L. Carey Tunnel	4.24
Queens-Midtown Tunnel	2.57
Cross Bay Veterans Mem. Br.	1.32
Marine Parkway-Gil Hodges Br.	0.76
Henry Hudson Bridge	0.42
RFK (Triborough) Bridge	0.28
Throgs Neck Bridge	0.26
Verrazzano-Narrows Bridge	0.15

Brightening most on average: Bronx-Whitestone Bridge (+9.13% per month)

Dimming most on average: Verrazzano-Narrows Bridge (+0.15% per month)

=== MoM volatility (std dev) by facility ===

facility

Bronx-Whitestone Bridge	50.95
Hugh L. Carey Tunnel	31.11
Queens-Midtown Tunnel	26.88
Henry Hudson Bridge	15.98
Cross Bay Veterans Mem. Br.	15.54
Marine Parkway-Gil Hodges Br.	15.44
RFK (Triborough) Bridge	15.27
...	
Verrazzano-Narrows Bridge	12.01

Most volatile: Bronx-Whitestone Bridge ($\pm 51.0\%$ MoM)

Most stable: Verrazzano-Narrows Bridge ($\pm 12.0\%$ MoM)

```

# =====
# STEP 3: Load the 9 MTA Bridges & Tunnels coordinates
# =====
# 7 bridges + 2 tunnels operated by MTA Bridges & Tunnels (TBTA).
# Coordinates are approximate midpoints, sufficient for a 2 km buffer.
crossings = [
    # (name, lat, lon, type)
    ("RFK (Triborough) Bridge",      40.7889, -73.9293, "bridge"),
    ("Bronx-Whitestone Bridge",      40.8009, -73.8347, "bridge"),
    ("Throgs Neck Bridge",           40.8023, -73.7948, "bridge"),
    ("Verrazzano-Narrows Bridge",     40.6066, -74.0447, "bridge"),
    ("Henry Hudson Bridge",          40.8779, -73.9220, "bridge"),
    ("Marine Parkway-Gil Hodges Br.", 40.5719, -73.8853, "bridge"),
    ("Cross Bay Veterans Mem. Br.",   40.5878, -73.8189, "bridge"),
    ("Queens-Midtown Tunnel",        40.7430, -73.9700, "tunnel"),
    ("Hugh L. Carey Tunnel",         40.6985, -74.0150, "tunnel"),
]
bt = pd.DataFrame(crossings, columns=["name", "lat", "lon", "type"])
bt_pts = gpd.GeoDataFrame(
    bt, geometry=gpd.points_from_xy(bt.lon, bt.lat), crs="EPSG:4326"
)

# Build buffers in a metric projection (NAD83 / NY Long Island EPSG:32118 is good for NYC)
bt_metric = bt_pts.to_crs("EPSG:32118")
bt_metric["buffer"] = bt_metric.geometry.buffer(BUFFER_KM * 1000)

# Union all 9 buffers (some will overlap, e.g. RFK / Queens-Midtown / Henry Hudson)
buffer_union = unary_union(bt_metric["buffer"].values)
buffer_union_gdf = gpd.GeoDataFrame(
    {"name": ["BT_buffer_union"], "geometry": [buffer_union]},
    crs="EPSG:32118",
).to_crs("EPSG:4326")

# Per-facility buffers in lat/lon (for plotting & per-bridge breakdown)
# Drop the original point geometry column first to avoid name collision,
# then promote 'buffer' to be the active geometry, rename it, and reproject.
bt_buffers = (
    bt_metric
    .drop(columns="geometry")      # remove original point geometry column
    .set_geometry("buffer")        # promote 'buffer' to active geometry
    .rename_geometry("geometry")   # safe now - no collision
    .to_crs("EPSG:4326")
    [{"name", "type", "geometry"}]
)

print(f"Built {len(bt_buffers)} individual {BUFFER_KM} km buffers.")
print(f"Union buffer area: {buffer_union_gdf.to_crs('EPSG:32118').geometry.area.iloc[0]/1e6:.1f} km²")

# Save for reuse
bt_buffers.to_file(f"{OUT_PREFIX}buffers.geojson", driver="GeoJSON")
buffer_union_gdf.to_file(f"{OUT_PREFIX}buffer_union.geojson", driver="GeoJSON")


```

[7]

```


... Built 9 individual 2.0 km buffers.
Union buffer area: 112.0 km²

```





The initial analysis focuses on nine MTA Bridges and Tunnels facilities and constructs a monthly nighttime-light index from average VIIRS radiance within two-kilometer buffers surrounding each facility. The analysis is subsequently expanded to twenty transportation crossings across the broader New York metropolitan region. Rather than relying solely on average brightness, the study develops multiple dimensions of the satellite signal, including changes in mean brightness, within-buffer brightness dispersion, the geographic dispersion of illuminated pixels, and principal components summarizing cross-sectional lighting patterns.

time-light variables helps to explain the residual of traded price or the nearby toll gates.



Several adjustments are introduced to separate economic information from non-economic variation in satellite imagery. NOAA weather observations are matched to individual facilities to control for snowfall, which can mechanically increase observed nighttime radiance. Municipal bond prices are also decomposed to account for movements associated with broader interest-rate conditions, allowing the analysis to examine whether geospatial signals are related specifically to the residual component of bond-price variation.

The traffic data provide an intermediate validation test for the satellite signal. If nighttime activity around a bridge or tunnel genuinely reflects changes in transportation intensity, increases in satellite-observed activity should be associated with increases in actual vehicle crossings at the corresponding facility. This allows the study to test whether the satellite signal reflects real economic activity rather than unrelated changes in nighttime illumination. The proposed economic relationship can therefore be interpreted as a sequence. First, changes in nighttime activity may capture changes in local transportation activity. Second, increased transportation activity may lead to higher bridge and tunnel usage. Third, greater usage may contribute to higher toll revenue. Finally, stronger toll-related fundamentals may affect the valuation or credit spread of MTA transportation bonds.



This creates a more direct economic explanation for the relationship between satellite imagery and municipal bond prices. Rather than relying only on a statistical correlation between nighttime lights and bond prices, the analysis tests whether the satellite signal is connected to an observable source of economic activity that supports the transportation system.

Toll-rate changes must also be considered when analyzing toll revenue. Toll revenue can increase either because more vehicles use a facility or because the toll charged per vehicle has increased. A rise in revenue following a scheduled toll increase therefore does not necessarily

indicate stronger transportation demand. The analysis should distinguish changes in traffic volume from changes in toll rates and incorporate major toll-policy changes as controls.



The traffic analysis can subsequently be combined with reported toll revenue from MTA financial or engineering reports. Comparing vehicle crossings, satellite activity, and toll revenue would provide a stronger test of the underlying economic mechanism. If increases in nighttime activity are followed by increases in traffic and toll revenue, the satellite signal would have a clearer interpretation as an alternative indicator of transportation-related economic activity.

The final stage of the analysis examines whether these changes in transportation fundamentals are reflected in MTA bond prices or credit spreads. An important question is whether the municipal bond market incorporates changes in traffic and toll conditions immediately or whether there is a measurable delay.

```
Merged: 39 overlapping months

=== Cluster-shape MoM signals vs bond ===

      target lag (mo) signal n      r      p
spread (level) 0 mean brightness MoM % (baseline) (lag 0) 39 +0.058 +0.726
spread MoM %    0 mean brightness MoM % (baseline) (lag 0) 38 -0.031 +0.851
spread (level) 1 mean brightness MoM % (baseline) (lag 1) 38 +0.112 +0.503
spread MoM %    1 mean brightness MoM % (baseline) (lag 1) 38 +0.029 +0.863
spread (level) 2 mean brightness MoM % (baseline) (lag 2) 37 -0.046 +0.787
spread MoM %    2 mean brightness MoM % (baseline) (lag 2) 37 -0.085 +0.618
spread (level) 3 mean brightness MoM % (baseline) (lag 3) 36 -0.001 +0.996
spread MoM %    3 mean brightness MoM % (baseline) (lag 3) 36 +0.054 +0.754
spread (level) 0 within-buffer std MoM % (lag 0) 39 +0.087 +0.597
spread MoM %    0 within-buffer std MoM % (lag 0) 38 -0.029 +0.864
spread (level) 1 within-buffer std MoM % (lag 1) 38 +0.156 +0.349
spread MoM %    1 within-buffer std MoM % (lag 1) 38 +0.056 +0.737
spread (level) 2 within-buffer std MoM % (lag 2) 37 -0.052 +0.758
spread MoM %    2 within-buffer std MoM % (lag 2) 37 -0.185 +0.272
spread (level) 3 within-buffer std MoM % (lag 3) 36 +0.034 +0.845
spread MoM %    3 within-buffer std MoM % (lag 3) 36 +0.147 +0.393
spread (level) 0 within-buffer tightness MoM % (lag 0) 39 +0.000 +0.998
spread MoM %    0 within-buffer tightness MoM % (lag 0) 38 +0.075 +0.656
spread (level) 1 within-buffer tightness MoM % (lag 1) 38 +0.105 +0.532
spread MoM %    1 within-buffer tightness MoM % (lag 1) 38 +0.060 +0.719
...
mean brightness MoM % (baseline) r=-0.085 (n=37, lag=2m)
within-buffer std MoM % r=-0.185 (n=37, lag=2m)
within-buffer tightness MoM % r=-0.284 (n=37, lag=2m)
cross-facility dispersion MoM % r=-0.148 (n=37, lag=2m)
```



This study examines whether satellite-derived nighttime-light measures contain information about municipal bond risk beyond conventional interest-rate variables. Using geographically mapped nighttime-light intensity and spatial clustering measures, we investigate whether changes in local light conditions are associated with subsequent bond-price movements over a ten-trading-day horizon. The analysis first removes the component of bond-price variation associated with Treasury-rate movements and then tests whether geographic light variables explain the remaining variation. To distinguish bond-specific effects from regional effects, the methodology is evaluated across multiple bonds with similar geographic and economic exposure. The study contributes to the growing use of alternative geospatial data in asset pricing by examining whether satellite-observed economic activity can provide an additional risk-state indicator in relatively opaque municipal bond markets.

That is actually a very useful distinction. The research claim becomes much more interesting: satellite imagery may tell us something about the local economic/risk environment that traditional bond-market variables do not fully capture.




If you develop it into a larger panel of bonds and show that the relationship survives Treasury controls, issuer controls, maturity controls, weather adjustments, and out-of-region testing, you would have a much more credible empirical paper rather than simply a profitable backrest.

That is actually a very useful distinction. The research claim becomes much more interesting: satellite imagery may tell us something about the local economic/risk environment that traditional bond-market variables do not fully capture.

If you develop it into a larger panel of bonds and show that the relationship survives Treasury controls, issuer controls, maturity controls, weather adjustments, and out-of-region testing, you would have a much more credible empirical paper rather than simply a profitable backtest.

The traffic data were collected from the MTA bridge and tunnel traffic dataset available through NYC Open Data. The raw dataset was downloaded in manageable chunks because of its large size, after which facility names, dates, and vehicle-crossing counts were cleaned and standardized. Observations were then restricted to facilities with usable coverage in the sample Henry Hudson Bridge, Hugh L. Carey Tunnel, and Throgs Neck Bridge and individual traffic records were aggregated to the monthly level by summing vehicle crossings for each facility. Month-over-month (MoM) traffic growth was calculated as the percentage change in monthly crossings relative to the preceding month. To construct the series shown in the figure, the facility-level MoM growth rates were averaged equally across the covered facilities for each month, and the visualization was restricted to January 2024 onward. Thus, the plotted series represents the average monthly change in traffic activity across the three MTA facilities for which consistent observations were available.

Key results:



Relationship

Pearson

Approx. p-

R²

Interpretation

r

value

Lights vs bond

-0.365

0.022


13.3%


Moderate negative

spread

relationship

Monthly MTA Bridge and Tunnel Traffic







Different light source strategies:

The analysis considers three complementary satellite-light strategies to capture different dimensions of transportation-related activity around MTA bridge and tunnel facilities. The first strategy uses the month-to-month change in mean nighttime brightness, which measures whether overall illuminated activity across the monitored locations is increasing or decreasing. The second strategy focuses on changes in the spatial shape of the light distribution within each facility buffer, capturing whether illuminated clusters become more concentrated or more dispersed over time. The third strategy measures cross-facility dispersion, which evaluates whether the lighting patterns of individual bridges and tunnels are becoming more similar or increasingly differentiated from one another. These measures are compared with monthly

changes in the bond spread to determine whether different forms of geospatial activity contain information relevant to municipal credit conditions. Rather than relying on a single brightness measure, the three strategies distinguish between the overall intensity of activity, the local spatial organization of that activity, and the degree of divergence across facilities. The figure suggests that these light-based indicators can experience substantial changes in specific months, while the bond spread also exhibits periods of pronounced widening and tightening; however, the relationship is not uniformly one-to-one, indicating that each light measure may capture a different aspect of the underlying economic environment rather than functioning as an identical trading signal.

MTA TRBTRN PMT Refunding Series A-1 (CUSIP 89602HAB3)
Month-to-month change in cluster shape vs bond spread MoM (n=39)






The trading scenario evaluates whether satellite-derived changes in nighttime activity around MTA bridge and tunnel facilities can be used to time investments in the TBTA General Revenue Refunding 2017B bond. The backtest begins with a hypothetical \$100,000 portfolio and constructs trading signals from unusually weak changes in nighttime brightness and in the

geographic concentration of illuminated activity. A “dimming” signal is triggered when the change in brightness falls into the lower quartile of the historical distribution, while a “tightening” signal is defined similarly using changes in the spatial dispersion of nighttime lights. A combined strategy requires both conditions to occur at the same time. After allowing for an assumed publication delay in the satellite data, the strategy purchases the bond at the next available daily VWAP and holds the position for ten trading days. Performance is then compared with a benchmark strategy that enters the bond every month regardless of the satellite signal. The purpose of the exercise is to test whether selectively trading only when the alternative-data conditions are present can improve portfolio outcomes relative to mechanical monthly exposure,

while also examining bond-price movements in the context of changing Treasury yields and rate-adjusted bond-price residuals.



Strategy	Ending value	Approx. return
A: Dim → Buy	\$100,744	+0.74%
B: Tighten → Buy	\$98,945	−1.06%
C: Dim AND Tighten	\$100,010	~0.01%

Benchmark: enter every month	\$94,371	−5.63%
-------------------------------------	-----------------	---------------

The most interesting result is therefore not that the satellite strategy produced a huge profit. It is that selective entry avoided much of the loss experienced by the benchmark. The dimming strategy finished slightly positive while the mechanical monthly strategy lost about 5.6%, a difference of roughly **6.4 percentage points**. Even the tightening and combined rules substantially preserved capital relative to the benchmark.

Starting from a hypothetical \$100,000 portfolio, the dimming-based strategy ended at approximately \$100,744, compared with approximately \$94,371 for a strategy that entered the bond every month. The tightening-only and combined strategies ended near \$98,945 and \$100,010, respectively. The results therefore suggest that the primary benefit of the alternative-data signal may be its ability to avoid unfavorable periods rather than to generate large absolute returns. Because Treasury-rate movements explain a substantial portion of the underlying bond-price variation, the evidence is more consistent with a selective exposure mechanism than with a pure directional forecasting strategy.