

Hourly Oil Price Volatility: The Role of COVID-19

Neluka Devpura, Paresh Kumar Narayan

<https://doi.org/10.46557/001c.13683>

DEPTH	EXPERTISE	VOICE	GENERATED
Overview (10-15 min)	Balanced	Maya	July 15, 2026

On April 20th, 2020, the price of West Texas Intermediate crude oil fell below zero for the first time in history. It didn't just collapse; it became negative. Sellers were paying buyers to take oil off their hands. That single number captures something about 2020 that no economic model had anticipated. However, the price level, as dramatic as it was, isn't actually the most revealing part of the story. The more telling aspect is the intensity of the swings surrounding it — hour by hour, day by day — and whether the pandemic itself, measured in cases and deaths, was driving those swings.

That's the question Neluka Devpura and Paresh Kumar Narayan set out to answer. They sought to determine not whether COVID-19 disrupted oil markets — that much was clear to anyone following the news — but whether the epidemiological data itself, the raw counts of infected people and fatalities, had a measurable, independent effect on oil price volatility, beyond everything else already known to move markets.

Volatility, to be precise, refers to the magnitude of price swings, not their direction. It's about uncertainty. This matters because it ripples into everything downstream — investor portfolios, firm profitability, global industrial production, and risk management. A market that swings wildly is one where planning becomes almost impossible. Thus, the question isn't just academic; it has real consequences for anyone exposed to energy prices, which is to say, almost everyone.

What makes Devpura and Narayan's approach distinctive is the frequency of their data. Most oil market research uses daily data, but they utilized hourly data — seventeen observations per trading day, running from 1:00 am to 5:00 pm, assembled from Refinitiv's Datascope price feed. The dataset covers July 1st, 2019, through June 12th, 2020, yielding 4,250 hourly observations. Hourly data allows research-

ers to see intraday dynamics that daily data simply smooths away. The West Texas Intermediate price is calculated as the average of the opening bid, opening ask, closing bid, and closing ask for each hour. Returns are computed as the logarithm of the current price divided by the previous hour's price, scaled to percentage terms.

To measure volatility, they used three established range-based formulas. The first, VOL1, follows the Parkinson method from 1980 — it uses the high and low prices within each period and assumes prices follow a geometric Brownian motion without drift. The second, VOL2, uses the Rogers and Satchell method from 1991. The third, VOL3, follows Garman and Klass from 1980. All three are constructed from the same hourly high, low, open, and close series, allowing them to capture genuine intraperiod price movement. The COVID-19 variables — global daily cases and deaths — were converted to hourly values by dividing each day's count by seventeen to align them with the seventeen hourly oil observations.

The descriptive statistics alone tell a stark story. Before December 31st, 2019 — the date of the first globally reported COVID-19 case, which the authors use as their sample split — the mean hourly West Texas Intermediate price was sixty-two dollars and twenty-six cents. During the COVID period, it fell to forty-two dollars and seventy-two cents, a decline of about thirty-one percent. However, price levels are one aspect. The volatility measures each rose roughly tenfold in their hourly means. VOL1 went from 0.0016 to 0.017. VOL2 increased from 0.0037 to 0.039. The standard deviation of the price itself rose roughly sixfold. The skewness of hourly returns flipped from positive 8.7 to negative 4.8 — meaning the distribution of returns shifted from having a long right tail to a long left tail, indicating a change from occasional large gains to occasional catastrophic losses.

The regression analysis is where the paper's central finding emerges. Devpura and Narayan estimate ordinary least squares models for each of the three volatility measures, using one-hour lags of COVID-19 cases and deaths as the key explanatory variables. They run two versions of each model: one without controls and one that adds lagged returns, the bid-ask spread, and trading volume as conventional predictors. The preferred results are from the controlled specifications. To guard against the statistical artifacts that affect time-series data — autocorrelation and heteroskedasticity — they use Newey-West corrected standard errors, with lag lengths chosen by the Schwarz information criterion starting from a maximum of twelve.

The results are consistent across every specification. COVID-19 cases and deaths both carry positive, statistically significant coefficients — in most cases significant

at the one percent level — on all three volatility measures, regardless of whether controls are included. The slope coefficients for cases in the controlled models are small in isolation. Yet, the sample standard deviation of global cases is 2,578. Multiply through, and a one-standard-deviation increase in cases raises daily oil price volatility by 8.56 percent for VOL1, 8.42 percent for VOL2, and 17.07 percent for VOL3. The effects from deaths are larger. With a standard deviation of 157 daily deaths, a one-standard-deviation increase raises daily volatility by 11.41 percent, 11.28 percent, and 22.02 percent across the three measures. That's the eight to twenty-two percent range that Devpura and Narayan report as their headline finding.

The pattern deserves a moment of reflection. Deaths drive larger volatility effects than cases across all three measures. VOL3 — the Garman-Klass measure — consistently shows the largest effects, roughly double those of VOL1 and VOL2. The direction never flips, the statistical significance never disappears, and the magnitude never becomes trivial. That consistency is what lends the finding its weight.

The skeptic's version of this story would suggest that COVID-19 occurred simultaneously with a historic oil price war between Russia and Saudi Arabia, a demand collapse as travel and industry shut down globally, and a liquidity crisis in financial markets. Perhaps those other factors are driving the volatility, and COVID case counts are merely a proxy for the broader catastrophe. Devpura and Narayan address this concern directly. The control variables in their models — trading volume, bid-ask spread, lagged returns — exhibit sensible, significant effects. Lagged trading volume is positive and significant, with a t-statistic of 2.9. The bid-ask spread is negative and significant. These findings make sense: higher volume tends to be associated with higher volatility, and a tighter spread suggests a more liquid, less erratic market. The COVID coefficients endure the addition of all these controls, and they are consistent across all three ways of measuring volatility. That's not what a spurious correlation looks like.

What Devpura and Narayan have demonstrated, then, is that pandemic news — the raw epidemiological signal of how many people were getting sick and dying — became a direct input into oil market uncertainty. This occurred not through some long causal chain involving policy responses and demand forecasts, but rather because the cases and deaths themselves, as they accumulated daily, were influencing hourly price swings. Their conclusion puts it plainly: the results imply greater uncertainty in the oil market and the need to rebalance investor portfolios that hold oil as an asset.

There's a broader implication here that the paper conveys rather than simply asserts. Biological systems and financial systems are typically treated as separate domains — the realm of epidemiologists on one side, economists and traders on the other. In 2020, those domains intersected. An epidemiological count, the type of number public health officials monitor to understand disease spread, turned out to be a variable that moved the market in energy trading. The standard deviation of global deaths — 157 per day — raised daily oil volatility by up to twenty-two percent. That's not a trivial figure. That's a signal. And capturing it required viewing the market not by the day, but by the hour.

This methodological choice — the use of hourly data — is quietly one of the paper's most significant contributions. It enabled Devpura and Narayan to see the texture of the pandemic's effect on markets, to differentiate COVID's contribution from the surrounding noise, and to translate regression coefficients into economically meaningful numbers that matter to anyone managing exposure to energy prices. The pandemic caused oil markets to behave in unprecedented ways. This paper quantified precisely how much of that was due to the virus itself.

This lecture was created by ennepō.

Go to <https://ennepo.ai> to Discover, Create and Follow the latest research in your field.

Read when you can. Listen when you want to.

CREDITS

AI-generated lecture by ennepō.ai, adapted from:

- "Hourly Oil Price Volatility: The Role of COVID-19", Neluka Devpura, Paresh Kumar Narayan (Energy RESEARCH LETTERS, 2020). doi:10.46557/001c.13683. CC-BY-SA (<https://creativecommons.org/share-your-work/cclicenses/>)

Restructured, summarized and explained. Not endorsed by the authors.

This lecture (script and audio) is licensed under CC BY-SA 4.0 (<https://creativecommons.org/licenses/by-sa/4.0/>).

<https://ennepo.ai/lecture/92663613-1c62-41be-9b4a-5f0da033bb3a>