Do central bank cycles drive stock returns? New evidence from the US, UK, and Japan

Ali Uppal*

23 June 2022

Abstract

Cieslak et al. (2019) show that the equity premium in the US since 1994 is earned entirely in even weeks of the Federal Open Market Committee meeting cycle and that these same even weeks also drive international stock returns. Updating their data, I find that their US result does not hold out-of-sample and show that with an extended sample, the result loses its robustness as early as 2004. I construct central bank cycles for the Bank of England and the Bank of Japan and show, when accounting for potential pre-announcement effects, their international result also no longer holds.

Keywords: Central Banks, Monetary Policy, Stock Returns, Federal Reserve

JEL Codes: E50 and G12

*Department of Economics, UC San Diego, email: asuppal@ucsd.edu. I would like to thank Giampaolo Bonomi, Jim Hamilton, Molly Shatto, Allan Timmermann, Annette Vissing-Jørgensen, Johannes Wieland, Steven Yee and seminar participants at the UCSD Graduate Student Research Seminar for useful comments and suggestions.

1 Introduction

If you invested $100 in the stock market at the start of 1994, you would have $768 in 2016. However, if you only held stocks in even weeks of the Federal Open Markets Committee (FOMC) cycle[1] over the same time period, your $100 would become a striking $1,522. This 'even-week' result, which I will refer to as the 'odd' result, was meticulously derived by Cieslak et al. (2019) (hereafter CMVJ), inspired by the earlier work of Lucca and Moench (2015). In this paper, I explore the odd result in greater detail, evaluating its robustness to different time periods, different definitions of FOMC cycle weeks, and different countries, before exploring some potential explanations for its existence.

The question of whether the Federal Reserve (the Fed) has a substantial impact on the stock market, as well as how it does, are important questions and have meaningful implications for our understanding of monetary policy and the functioning of central banks. However, the question is fraught with challenges given the difficulty in identifying monetary policy shocks (Ramey (2016)). Many papers seek identification by exploring the impact of monetary policy around the time of the announcement (see for example Lucca and Moench (2015) and Brusa et al. (2020)).

CMVJ use neither esoteric measures of monetary policy shocks nor the high-frequency approach around announcements. Instead, they focus on the evolution of stock returns over the full cycle of days between scheduled FOMC meetings and find that the equity premium is earned entirely in even weeks of the FOMC cycle, claiming that the FOMC cycle in stock returns appears to be a general phenomenon, one that is strengthening over time. By ruling out alternative explanations methodically, they claim that the odd result is causally driven by systematic informal communication around biweekly board meetings at the Fed. CMVJ's findings are significant and novel. Indeed, not only have they driven momentum into this area of the literature (see for example Laarits (2020) and Hu et al. (2021)), their findings have

[1] Even weeks refer to weeks 0, 2, 4, and 6 of the FOMC cycle. The FOMC cycle captures all the days from one FOMC announcement to the next. The cycle starts the week before the announcement and ends approximately six weeks after the announcement. See Table 1 for a full mapping cycle days to cycle weeks.

been widely reported in the media, with both The Economist[2] and The Wall Street Journal[3] featuring articles discussing the odd result. Furthermore, CMVJ find that the FOMC cycle, in addition to driving US stock returns, also drives international stock returns. Brusa et al. (2020) add that high equity returns around monetary policy announcements are unique to the Fed.

In the context of the findings above, this paper has two key contributions. First, I find that the odd result does not hold out-of-sample, casting doubt on the claim that it is a general phenomenon. Furthermore, I show that when using data up to the end of 2021, the odd result loses statistical significance as early as 2004 and the coefficient trends down following the financial crisis, before becoming negative from 2016 onwards. Second, in the context of FOMC cycle driving international stock returns, I show that neither UK nor Japanese stock returns are driven by FOMC cycle when one controls for a pre-announcement effect.

The remainder of this paper is set as follows. Section 2 describes the data used. Section 3 evaluates the robustness of the odd result in the US. Section 4 considers the odd result in the context of the UK and Japan. Section 5 provides a conclusion.

2 Data

A key contribution of this paper is the collection of a variety of datasets, transformed in a way to ensure consistency and comparability across central bank cycles of varying lengths. This section is split into data relating to central bank meetings, financial data, and Fed Chair calendar data. While much of my data goes up to 2016, in line with that of CMVJ, the extended data for the US goes up to the end of 2021.

2.1 Central Bank Meeting Data

I collect data on central bank meetings for the US, the UK, and Japan. I then compute days in terms of the central bank's cycle time. I use both the definition of central bank cycle time

[2] https://www.economist.com/finance-and-economics/2016/09/03/the-long-arm-of-the-fed
[3] https://on.wsj.com/3li9ypk

of CMVJ and an alternative definition of weeks as a robustness check (see Section 3.2).

The cycle time captures the rate-setting process of the central bank. CMVJ define the FOMC cycle by having week 0 of the FOMC cycle start the day before a scheduled FOMC announcement. The announcement day is considered day 0, therefore week 0 starts on day -1. Their rationale for having week 0 start the day before the announcement rather than the day of the announcement is to capture the pre-announcement effect documented by Lucca and Moench (2015) (i.e., large excess stock returns in advance of FOMC meetings). One could also try an alternative definition of the FOMC cycle where week 0 starts on day 0. Both CMVJ's definition and my alternative definition exclude weekends and measure the cycle from the week before the announcement (week -1) until 6 weeks after the announcement. The two definitions can be seen in the table below.

Table 1: Defining weeks of central bank cycles

Week of the cycle	Days counted under Cieslak et al. (2019)	Days counted under alternative definition
-1	$-6, \ldots, -2$	$-5, \ldots, -1$
0	$-1, \ldots, 3$	$0, \ldots, 4$
1	$4, \ldots, 8$	$5, \ldots, 9$
2	$9, \ldots, 13$	$10, \ldots, 14$
3	$14, \ldots, 18$	$15, \ldots, 19$
4	$19, \ldots, 23$	$20, \ldots, 24$
5	$24, \ldots, 28$	$25, \ldots, 29$
6	$29, \ldots, 33$	$30, \ldots, 34$

2.1.1 FOMC Meetings

I start with the data compiled by CMVJ which collects FOMC meeting dates from 1982 to 2016. I then update their data to include FOMC meetings up to the end of 2020, using the Federal Reserve website. Figure 1 below shows the days on which the meetings took place for the post-1994 period as this will be the period relevant for analysis.[4] There are 216 FOMC meetings in total and as can be seen by the peaks in Figure 1 below, the FOMC meets eight times per year.

[4] It was only since 1994 that the Fed publicly announced its decision following a scheduled FOMC meeting and so interpretations before this date are less meaningful.

Figure 1: FOMC Meeting Frequency and Timing

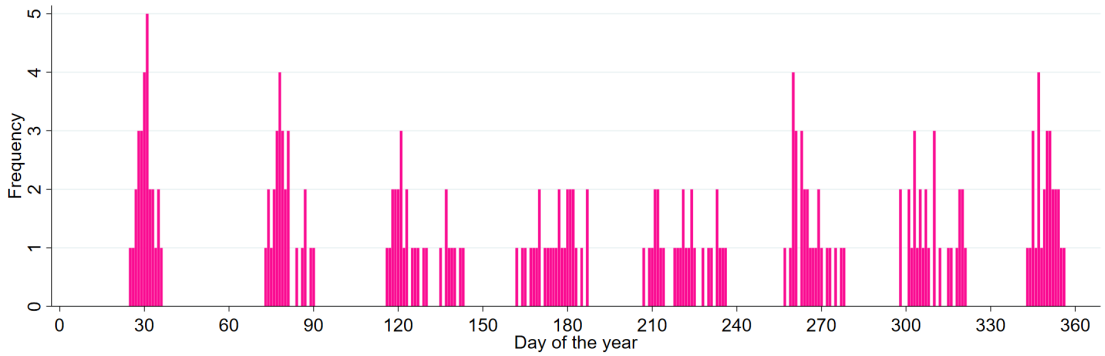

2.1.2 Bank of England (BoE) Meetings

Using the BoE website, I collect information on meetings of the Monetary Policy Committee (I will refer to these as BoE meetings). My sample starts in July 1997 (the BoE gained operational independence in June 1997). Unlike the FOMC, the BoE meets monthly. However, in late 2016, it changed its meeting schedule to an eight-meetings-a-year schedule. Therefore, for consistency, I focus on the period from 1997 to 2016. There are 232 BoE meetings in total and we can see the 12 monthly meetings in Figure 2 below.

Figure 2: BoE Meeting Frequency and Timing

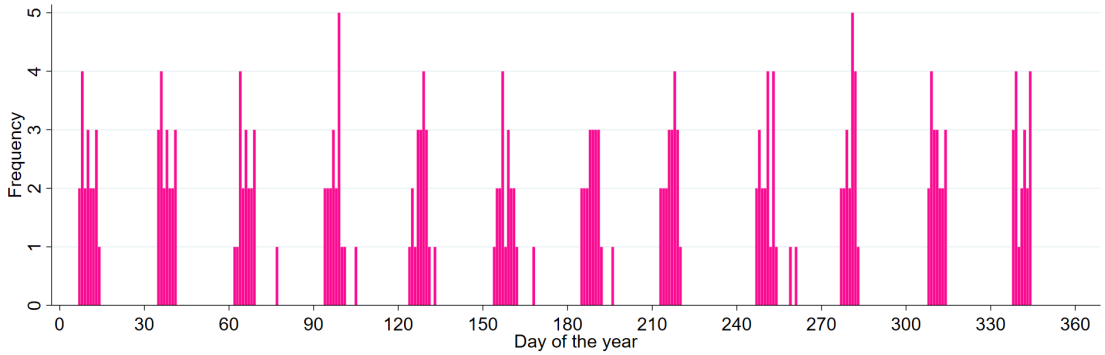

2.1.3 Bank of Japan (BoJ) Meetings

I collect information on BoJ monetary policy meetings from the BoJ website. Given the BoJ gained independence in April 1998, I start my sample in May 1998. The BoJ's meeting schedule was much less regular than that of either the FOMC or BoE (see Figure 3). In line with CMVJ, I exclude unscheduled meetings. As highlighted by Brusa et al. (2020), up until 2005, the BoJ progressively decreased the number of meetings from 20 per year, before settling on 14 in 2006. Like the BoE, the BoJ decided to shift to an eight-meeting schedule

in 2016. I focus on the period from 1998 to 2016 which most closely matches the BoE sample and the FOMC sample of CMVJ. This sample consists of 286 BoJ meetings.

Figure 3: BoJ Meeting Frequency and Timing

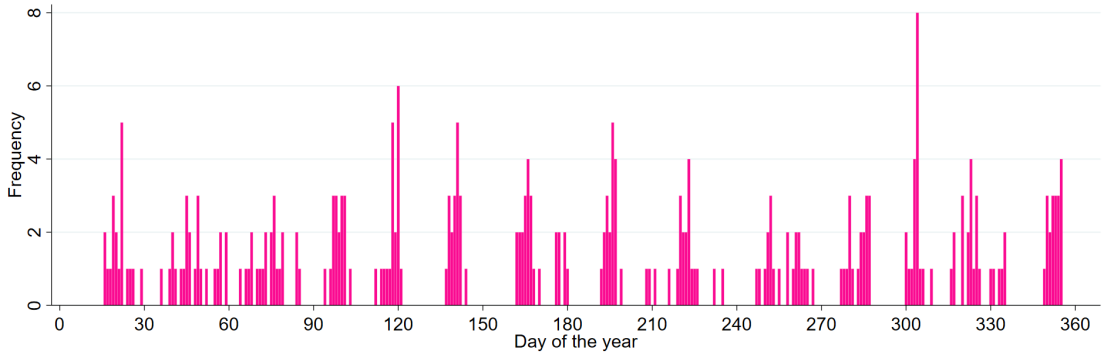

2.2 Financial Data

US and Japanese daily excess return data comes from Kenneth French's website while the equivalent for the UK comes from Gregory et al. (2013) with updates from the Xfi Centre for Finance and Investment at the University of Exeter Business School. The VIX data is from CBOE via WRDS. Table 2 below shows the key summary statistics for these variables.

Table 2: Financial Data Summary Statistics

Variable	Obs	Mean	Std. Dev.	First-Order AC
US Daily Excess Return	7,259	0.04	1.17	-0.08
UK Daily Excess Return	4,864	0.02	1.15	-0.01
Japan Daily Excess Return	4,957	0.02	1.42	-0.06
Daily VIX	7,049	19.78	8.31	0.98

3 Robustness of the Odd Result in the US

This section explores the robustness of the odd result documented by CMVJ. First, I replicate CMVJ's main results and test their robustness out-of-sample. I then test whether the odd result is robust to a different definition of weeks. Next, I briefly explore CMVJ's results on a day-of-the-cycle basis and then conclude the section with an analysis of the potential role of uncertainty.

3.1 Out-of-Sample Robustness

CMVJ show that their results are robust out-of-sample. Specifically, they split their sample into two periods: 1994-2013 and 2014-2016. The latter three-year period constitutes an out-of-sample test for CMVJ as the initial draft of their paper only employed data from 1994 to 2013. Not only do they show that their results hold in their out-of-sample period, leading them to conclude that the FOMC cycle in stock returns is a general phenomenon, they further claim that the odd result is strengthening over time. I replicate CMVJ's results and test whether the odd result remains robust to a more recent out-of-sample period. My out-of-sample period is the five-year period from 2017 to the end of 2021. Further, I combine this with CMVJ's out-of-sample period so that I can also test whether the result holds from 2014-2021. Here I obtain my first contribution in this paper: the FOMC cycle in stock returns no longer holds out-of-sample in either my out-of-sample or the combined out-of-sample.

Table 3 reports regressions of daily excess US stock returns on FOMC cycle dummies. t-statistics are calculated on the basis of robust standard errors as is the case in the rest of this paper. Columns 3 and 4 of Table 3 replicate the results reported by CMVJ in Table I Panel B of their paper, in columns 1 and 2, respectively. The interpretation of the results is that, between 1994 and 2013, the average excess return per day is 13.6 bps higher on days that fall in week 0 in FOMC cycle time and 9.9 bps higher on days that fall in week 2, 4, or 6 compared to days that fall in odd weeks in FOMC cycle time. The coefficient on the week 2, 4, or 6 dummy strengthens substantially in the 2014-2016 period. However, column 2 shows that the result does not hold in the 2017-2021 period.[5] Specifically, the week 0 dummy, although significant at the 5% level, is now negative rather than positive while the dummy for even weeks 2, 4, or 6 is no longer significant at all. Note that the number of observations in CMVJ's out-of-sample period (783) is much smaller than in mine (1259). Finally, in column 1, I combine the two out-of-sample periods so that there is one larger out-of-sample test with eight years of data from the start of 2014 to the end of 2021. I find that none of the coefficients are statistically significant.

[5] My results remain consistent when using the out-of-sample period 2017-2019, i.e., excluding 2020 and 2021 given potential concerns about the global pandemic.

Table 3: Regressions of Daily Excess U.S. Stock Returns on FOMC Cycle Dummies

	Combined OoS 2014 to 2021	New OoS 2017 to 2021	CMVJ OoS 2014 to 2016	CMVJ Main 1994 to 2013
Dummy=1 in Week 0	-0.0651	-0.211**	0.174*	0.136***
	(-0.90)	(-2.07)	(1.92)	(2.76)
Dummy=1 in Week 2,4,6	0.0564	-0.0180	0.176***	0.0993***
	(1.08)	(-0.24)	(2.67)	(2.65)
N	2042	1259	783	5214

t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

While Table 3 replicates and extends CMVJ's out-of-sample test, it does not test the significance of each individual even-week dummy (i.e., dummy=1 in week 0, 2, 4, 6). Table 4 therefore repeats the exercise in Table 3 but with individual even-week dummies. A few points are worth noting. First, the odd result appears less significant in CMVJ's out-of-sample period (2014-2016) when using individual even-week dummies. In my out-of-sample period (2017-2021) and the combined out-of-sample period, the results are consistent with those in Table 3: the FOMC cycle in stock returns does not hold out-of-sample.

Table 4: Regressions of Daily Excess U.S. Stock Returns on FOMC Cycle Dummies

	Combined OoS 2014 to 2021	New OoS 2017 to 2021	CMVJ OoS 2014 to 2016	CMVJ Main 1994 to 2013
Dummy=1 in Week 0	-0.0651	-0.211**	0.174*	0.136***
	(-0.90)	(-2.07)	(1.92)	(2.76)
Dummy=1 in Week 2	0.0859	0.0488	0.146*	0.0811*
	(1.29)	(0.51)	(1.82)	(1.70)
Dummy=1 in Week 4	0.0268	-0.0847	0.200**	0.107**
	(0.42)	(-0.93)	(2.41)	(1.99)
Dummy=1 in Week 6	0.0151	-0.143	0.325	0.177**
	(0.07)	(-1.03)	(0.52)	(1.98)
N	2042	1259	783	5214

t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

The lack of robustness of the odd result naturally leads to the question of when it stopped being robust. One puzzle when reading the final version of CMVJ's paper (published in 2019) is that given the initial paper was available in 2014, why had the odd result not been arbitraged away? Indeed, looking at column 1 of Table 3 and Table 4, it would appear the result is no longer significant after the initial draft was made available.[6] This would be consistent with the findings of McLean and Pontiff (2016) who find that investors learn about mispricing from academic publications.

In order to understand when the odd result began to lose significance, I estimate a regression of the daily excess stock return on the even-week dummy (i.e., dummy=1 in week 0, 2, 4, or 6) using an expanding window going backwards. First, I do the analysis using CMVJ's time horizon (i.e., 1994-2016) as shown in Figure 4. Then, I use my updated time horizon with data going up until the end of 2021 as shown in Figure 5.[7] Each point in the figure represents the even-week regression coefficient estimated for the time horizon along the horizontal axis. For example, the last marker in Figure 4 represents the excess return when estimating the even-week coefficient using data from the start of 2015 to the end of 2016.

Figure 4 clear supports the CMVJ interpretation. Specifically, that the even-week effect is robust in their sample period (1994-2016) and it appears to be strengthening over time. CMVJ argue that this "pattern fits a 'Fed put' interpretation, meaning that the Fed reacts to low stock returns by providing (a promise of) monetary policy accommodation, with this accommodation being unexpectedly strong in our sample" (p.2202-2203). Indeed, Cieslak and Vissing-Jørgensen (2020) find that since the mid-1990s, negative stock returns predict policy accommodation. However, Figure 5 tells a surprisingly different story. First, the result stops being statistically significant as early as 2004. Second, the even-week effect goes from being positive (as in CMVJ) to negative, with the decline starting after the global financial crisis. Together, these two points suggest that CMVJ's result no longer holds and is not as robust when using a longer time horizon. Moreover, if one believes the Fed Put interpretation put forth by CMVJ, then by the same argument, Figure 5 would suggest that

[6] The first draft of the CMVJ paper was available on 23 April 2014. See https://faculty.haas.berkeley.edu/morse/research/papers/cycle_paper_cieslak_morse_vissingjorgensen.pdf

[7] The results are similar when excluding the global pandemic, i.e., ending the sample at end-2019.

the Fed is no longer being as accomodative in response to negative stock returns.

Figure 4: Expanding Window Regression, 1994-2016

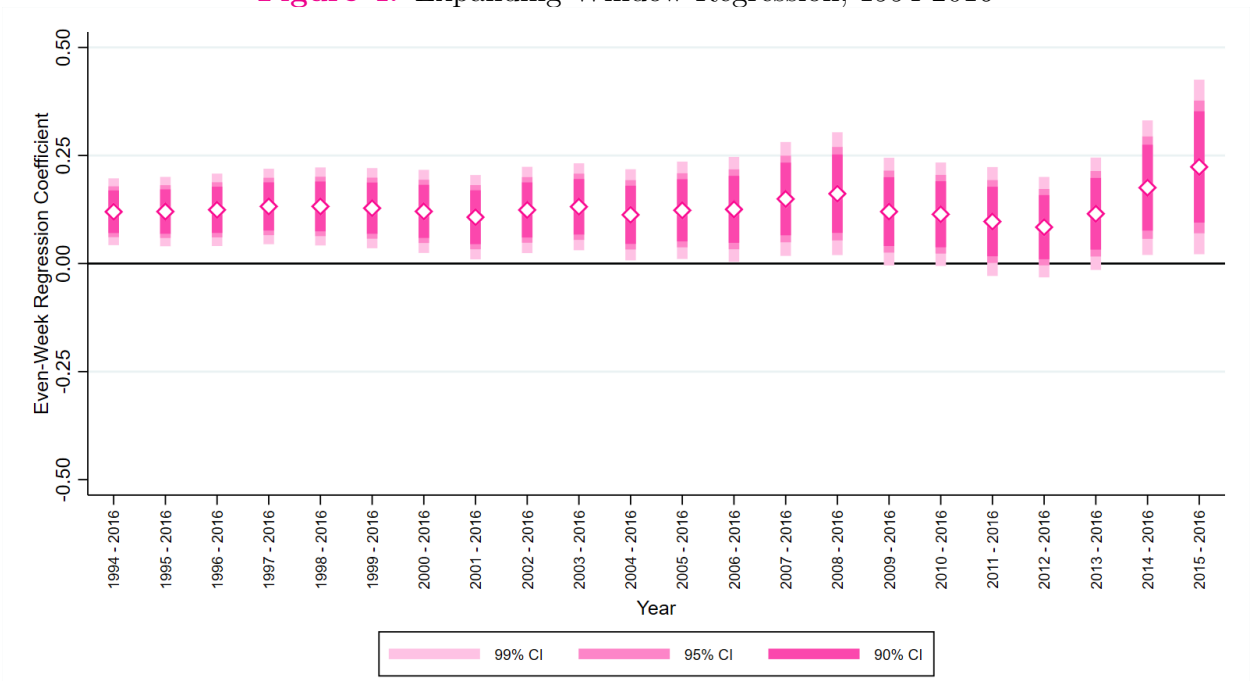

Figure 5: Expanding Window Regression, 1994-2021

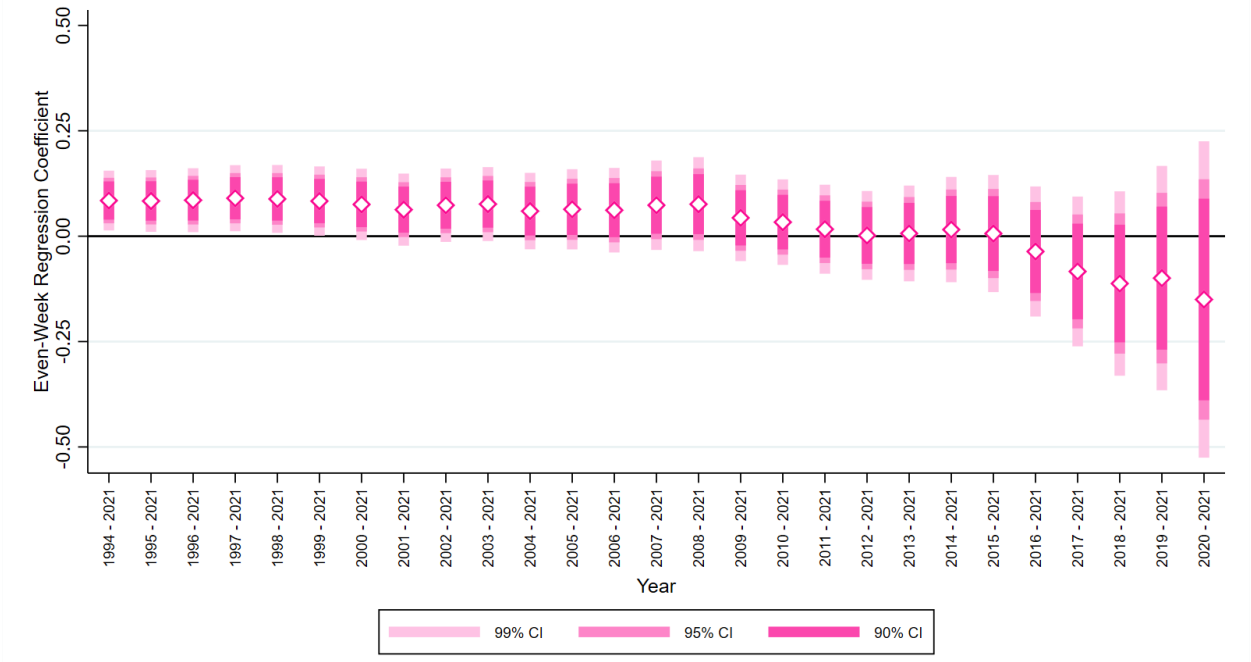

3.2 Week Definition Robustness

While CMVJ are thorough in exploring the robustness of the odd result, they do not test the robustness of their finding to different definitions of FOMC cycle weeks. In a sense, they are correct not to be overly concerned with this definition as it is somewhat arbitrary; an investor will not care how an even week is defined so long as it is defined consistently and known in advance in order to enable them to generate excess returns. However, one could argue that there is a conceptual distinction between day -1 (the day before the announcement) and day 0 as the latter involves a substantive public information release. As such, combining both days into week 0 may capture two different effects: the pre-FOMC announcement drift documented by Lucca and Moench (2015) and the even-week effect that CMVJ document. Therefore, to test the robustness of CMVJ's findings to different week definitions, I replicate their regressions under their definition and my alternative definition as shown in Table 1, where week 0 begins with day 0 (i.e., the week starts with the FOMC announcement) rather than day -1 as in CMVJ. To keep this exercise focused on the implication of different week definitions only, I use the same sample as CMVJ (i.e., 1994-2016).

Table 5: Regressions of Daily Excess U.S. Stock Returns, Different Week Definitions

	CMVJ Week		Alternative Week	
	(1)	(2)	(3)	(4)
Dummy=1 in Week 0	0.141***	0.141***	0.101**	0.101**
	(3.17)	(3.17)	(2.33)	(2.33)
Dummy=1 in Week 2,4,6	0.109***		0.0935***	
	(3.24)		(2.81)	
Dummy=1 in Week 2		0.0896**		0.0768*
		(2.10)		(1.83)
Dummy=1 in Week 4		0.120**		0.112**
		(2.52)		(2.48)
Dummy=1 in Week 6		0.187**		0.0905
		(2.07)		(0.98)
N	5997	5997	5997	5997

t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

Table 5 shows that, while the results remain significant regardless of week definition, the magnitude of each coefficient and its associated t-statistic both fall when moving from CMVJ's week definition to mine (regardless of whether one uses individual week dummies). Therefore, it seems prudent to take an approach that is agnostic as to the definition of the week.

3.3 Day-of-the-Cycle Analysis

While CMVJ focus on the weeks in their analysis, the previous section has highlighted some potential pitfalls of doing so. Assessing individual days over the FOMC cycle is less susceptible to such critiques. Therefore, I run a separate regression for each day of the cycle where I regress daily excess returns on a dummy for that *day* only. Figure 6 plots the coefficients of each of the regressions as well as their confidence intervals. Each CMVJ-defined week is represented by a different colour. As in Section 3.2, I use the CMVJ sample.

Figure 6: Regression Coefficient from Regressing Daily Excess Return on Day of the Cycle

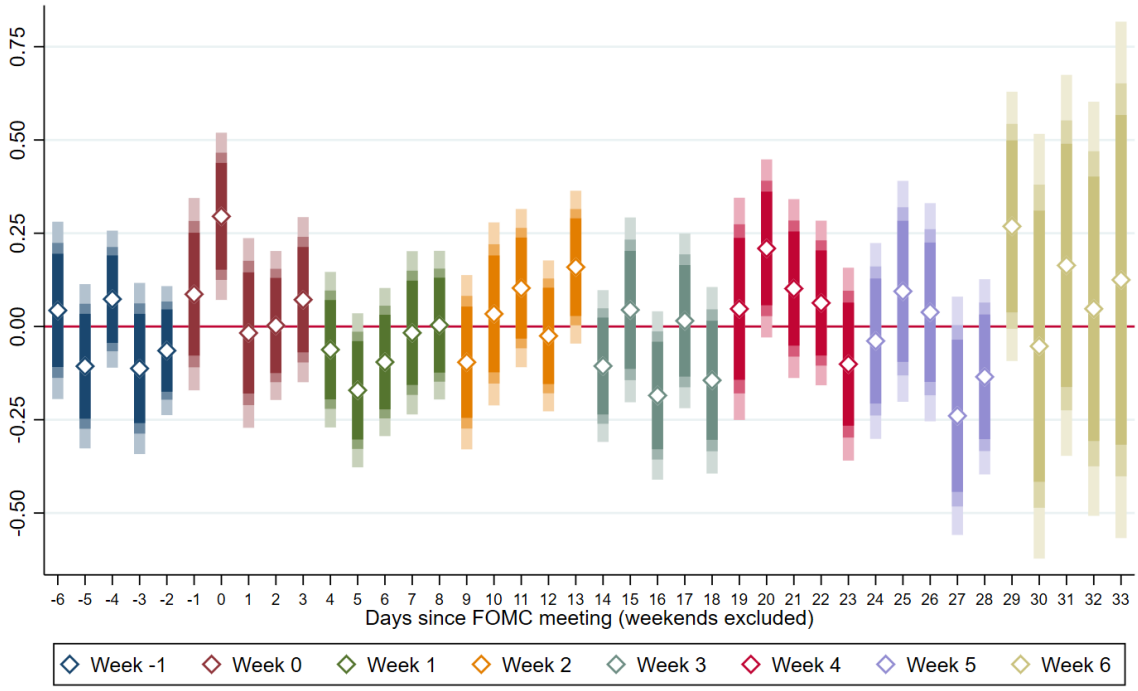

Notes: The 40 regressions underpinning this chart have excess returns as the dependent variable. The only regressors are a constant and the dummy for the specific day of the cycle. The diamonds reflect the coefficient on the day dummy of each regression. The confidence intervals are shown by coloured bars. The darkest shade represents the 99% confidence interval, one shade lighter represents the 95% confidence interval, and the lightest colour represents the 90% confidence interval.

Figure 6 shows that certain days appear to be more important than specific weeks. The only days that are statistically significant and positive are Day 0 (daily excess return of 30 bps, significant at the 1% level), Day 13 (daily excess return of 16 bps, significant at the 5% level), Day 20 (daily excess return of 21 bps, significant at the 5% level), and Day 29 (daily excess return of 27 bps, significant at the 10% level). Given that these results essentially show a spike every 10 days (i.e., every two weeks when excluding weekends), they appear consistent with the even-week effect in CMVJ and also highlight potentially narrower windows for future research to examine.

While CMVJ do not do a daily analysis, they specifically discuss the role of informal communication by the Fed in even weeks. Therefore, it is possible that the specific spikes in Figure 6 are informative about these instances of informal communication. Several papers discuss potential informal communication by the Fed. For example, Bradley et al. (2020) use a particularly creative method to document potential leaks. They use taxi ridership between the New York Fed and large financial institutions in New York to show that late-night meetings and lunchtime interactions increase around FOMC meetings. They claim that this suggests increased opportunities for Fed information to flow to markets along informal or discreet channels. Indeed, CMVJ themselves provide compelling anecdotal evidence of leaks.[8] More systematically, Morse and Vissing-Jørgensen (2020) analyze nearly thirty thousand entries in Federal Reserve governors' calendars from 2007-2018 and document the role that the media plays in information transmission on even-week days with policy-maker interactions. Moreover, they highlight the importance of the policy-maker interactions taking place during stressed times, as measured by CBOE Volatility Index (VIX). Hence, given uncertainty appears to play an important role in their mechanism, I explore it in greater detail in the next section.

[8] CMVJ provide a powerful quote from Richard Fisher, former president of the Dallas Fed: "There is one former Governor who recently visited my Bank [...] who told the staff [...] that this individual—I'll let you guess who it is—was, in essence, the 18th or 19th member, depending on how many we have, of the FOMC, and the equivalent of a voting member. He makes money off of us when he talks and sells."

3.4 Potential Role of Uncertainty

CMVJ, together with Morse and Vissing-Jørgensen (2020), provide compelling evidence of informal systematic communication (and leaks) by the Fed. Specifically, CMVJ argue that "the Fed can use informal communication to reduce uncertainty about the policy rule and guide policy expectations" (p.2243). Therefore, I examine whether uncertainty rises in advance of even weeks to test this motivation both within the CMVJ sample (i.e., 1994-2016) and my five-year out-of-sample period (2017-2021).

If it is indeed the case that the Fed is seeking to reduce uncertainty, then we would expect uncertainty to increase in the days leading up a spike in returns. Figure 7 replicates Figure 1 of CMVJ and overlays VIX movements over the preceding 5 days to visualize this question. For a given day in the FOMC cycle, the chart shows the change over the previous week in the VIX and the 5-day cumulative return in the following week. As can be seen, the even-week increase in returns tend to be preceded by a rise in uncertainty, providing suggestive, though not conclusive, evidence of the FOMC responding to uncertainty.

Figure 7: Stock Returns and VIX changes over the FOMC Cycle, 1994-2016

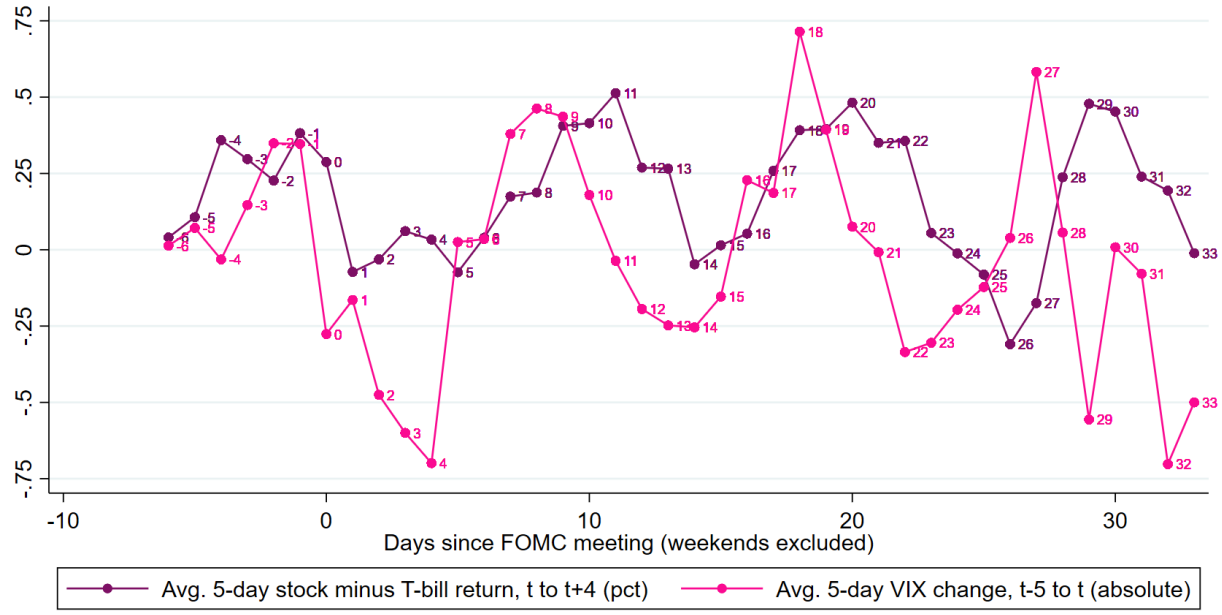

In a similar fashion to Table 4, I run regressions of the daily change in the VIX and find that the VIX is rising in odd weeks (confirming the pattern in Figure 7). Therefore, it is plausible that FOMC is reacting to this uncertainty which leads to stock returns rising in even weeks. However, it is unclear why the VIX would consistently rise in odd weeks (a truly odd result!). Indeed, without an identification assumption, we cannot separate out the effect of the VIX on excess returns from that of excess returns on the VIX. Columns 3 and 4 in Table 6 show that this result, much like the excess stock return result, does not hold out-of-sample. This strengthens the case that the results of CMVJ are sample-specific.

Table 6: Regressions of Daily VIX Change on FOMC Cycle Dummies

	1994-2016		2017-2021	
	(1)	(2)	(3)	(4)
Dummy=1 in Week -1,1,3,5	0.161***		-0.0689	
	(3.83)		(-0.55)	
Dummy=1 in Week -1		0.178***		-0.0226
		(2.86)		(-0.10)
Dummy=1 in Week 1		0.174***		-0.212
		(2.71)		(-1.45)
Dummy=1 in Week 3		0.143**		-0.0139
		(2.17)		(-0.09)
Dummy=1 in Week 5		0.134*		0.0483
		(1.89)		(0.16)
N	5583	5583	1258	1258

t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

4 The Odd Result in an International Context

CMVJ show that international stock market returns also follow the FOMC cycle, though they do not test the robustness of this result in the same way as they do for US stock returns, nor do they explore central bank cycles of other countries. While there is work that explores

stock returns around announcements of the domestic central bank for several countries (e.g., Brusa et al. (2020)), I am not aware of any work that explores the central bank cycle in other countries, as CMVJ have done for the US. Therefore, this section has two questions. First, is CMVJ's finding that international stock returns are driven by the FOMC cycle robust? Second, do central bank cycles, like the FOMC cycle, exist in other countries? In the case of the UK and Japan, I find that the answer to both questions is no.

4.1 BoE Cycle

Figure 8 provides preliminary answers to the two questions posed above. The chart on the left shows 5-day forward cumulative UK excess stock returns over the BoE cycle. There is no odd result in the BoE cycle. However, it appears there is a type of pre-announcement drift.[9] The chart on the right shows UK stock returns over the FOMC cycle and appears to confirm CMVJ's finding that international stock returns are driven by the FOMC cycle.

Figure 8: UK Stock Returns over the BoE and FOMC Cycles, 1997-2016

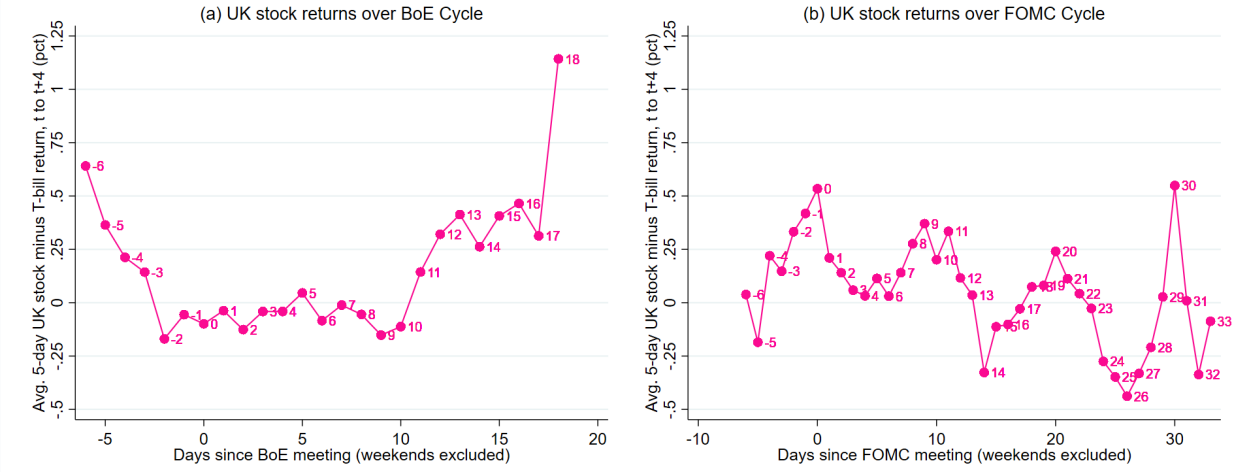

While the above visual evidence is useful, I run regressions to answer the two questions more rigorously. Table 7 shows the results of regressing daily UK excess stock returns on a combination of BoE and FOMC cycle dummies. I control for a type of pre-announcement drift given the evidence in the left chart of Figure 8. The FOMC week dummy is the standard

[9] This would not be same as the pre-announcement drift documented by Lucca and Moench (2015) which would appear in week 0 when using CMVJ's week definition.

CMVJ even-week dummy. Column 1 shows that only week -1 is significant (i.e., stocks rise in the week leading up to BoE meetings). Column 2 shows UK stock returns seem to spike in even weeks of the FOMC cycle, though the result is marginally significant. While this is consistent with CMVJ's findings, I find that after including the BoE week -1 dummy (i.e., controlling for the pre-announcement effect), the FOMC cycle is no longer significant in predicting UK stock returns (column 3). If CMVJ's hypothesis that the central bank cycle is a result of informal communication and leaks, then this suggests that BoE is perhaps stricter about communication flows.

Table 7: Regressions of Daily Excess UK Stock Returns on BoE/FOMC Cycle Dummies

	(1)	(2)	(3)
Dummy=1 in BoE Week -1	0.143***		0.138***
	(3.74)		(3.59)
Dummy=1 in FOMC Week 0,2,4,6		0.0628*	0.0532
		(1.90)	(1.60)
N	4861	4861	4861

t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

4.2 BoJ Cycle

I repeat the above analysis, but this time in the case of Japan. Given the BoJ had much less consistency in the number of meetings per year, any interpretation should be treated with caution. The chart on the left of Figure 9 shows 5-day forward cumulative excess stock returns over the BoJ cycle. Like the BoE, there is no odd result for the BoJ. The chart on the right shows Japanese stock returns over the FOMC cycle and while they do not follow the FOMC cycle as closely as UK stock returns, they still appear to have even-week spikes.

Figure 9: Japan Stock Returns over the BoJ and FOMC Cycles, 1998-2016

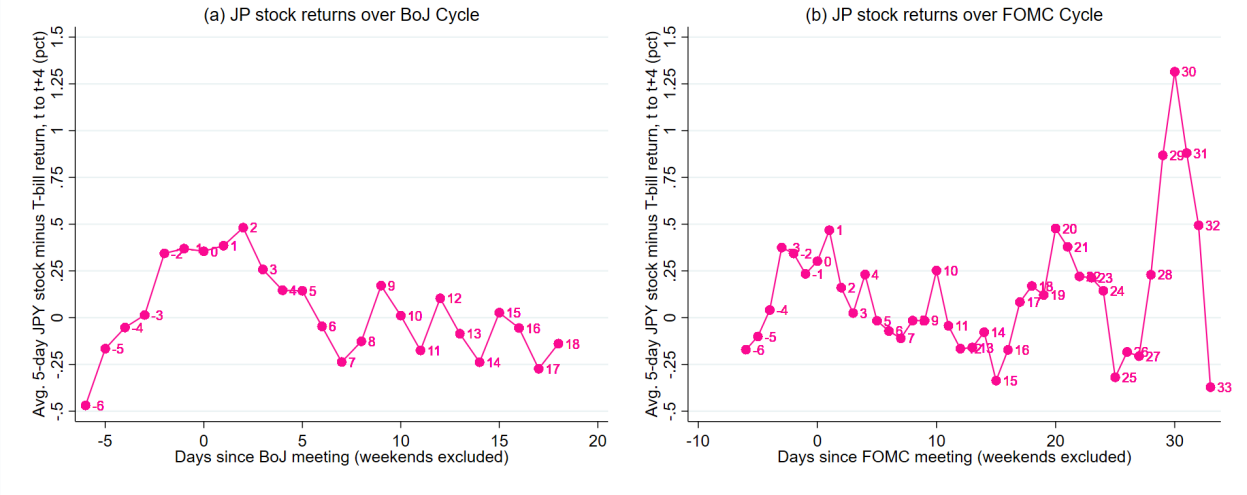

As before, I run regressions to answer the two questions more rigorously. Table 8 shows the results of regressing daily Japanese excess stock returns on a combination of BoJ and FOMC cycle dummies. I test for the pre-announcement effect (i.e., a dummy in BoJ week -1). The FOMC week dummy is the standard CMVJ even-week dummy. Column 1 shows that while week -1 is significant, it is negative. Column 2 shows that even without any controls, Japanese stock returns are not driven by the FOMC cycle. Again, if CMVJ's argument is correct, then the lack of a BoJ cycle suggests that like the BoE, and unlike the FOMC, the BoJ is potentially stricter about information flow.

Table 8: Regressions of Daily Excess JP Stock Returns on BoJ/FOMC Cycle Dummies

	(1)	(2)	(3)
Dummy=1 in BoJ Week -1	-0.156***		-0.156***
	(-3.45)		(-3.45)
Dummy=1 in FOMC Week 0,2,4,6		0.0188	0.0167
		(0.46)	(0.41)
N	4954	4954	4954

t statistics in parentheses

* p<0.10, ** p<0.05, *** p<0.01

5 Conclusion

CMVJ's odd result, both in the US and internationally (i.e., that even weeks of the FOMC cycle predict excess stock returns), is very important in raising necessary questions about the conduct of central banks. Therefore, in this paper, I have evaluated some of CMVJ's key findings and make two key contributions. First, I find that the FOMC cycle does not drive US stock returns after 2016. Specifically, I find that the odd result is only robust when one concludes the sample for analysis in 2016. When expanding the sample to end-2021, the result stops being robust as early as 2004. My second contribution relates to evaluating the robustness of CMVJ's odd result internationally. I find that the FOMC cycle does not drive UK excess stock returns when one controls for potential pre-announcement effects in the UK and that FOMC cycle does not drive Japanese stock returns regardless of whether one includes controls. I further show that neither the BoE or the BoJ, unlike the FOMC, have especially unusual patterns in their cycle.

I have also briefly explored one of CMVJ's interpretations of their odd result (i.e., the Fed Put). While I do find that uncertainty appears to be rising in the weeks preceding even weeks, this only appears to hold in CMVJ's main sample from 1994 to 2016. When conducting this analysis up to end-2021, I again conclude that the result no longer appears to hold. It is unclear whether this is because there is less informal communication by the Fed or whether the Fed is being less accomodative. However, regardless of whether there is informal communication, this area of research raises an important question: does informal communication (or targeted leaks) improve the effectiveness of monetary policy? CMVJ convincingly highlight some of the negative welfare consequences associated with informal communication. The lack of robustness of CMVJ's result may reflect less reliance on informal communication in recent years. This could reduce some of the aforementioned negative welfare effects and might indeed be a direct result of CMVJ's careful empirical analysis, highlighting the need for such important work to continue.

References

Bradley, D., Finer, D. A., Gustafson, M., and Williams, J. (2020). When Bankers Go to Hail: Insights into Fed-Bank Interactions from Taxi Data. *Mimeo*.

Brusa, F., Savor, P., and Wilson, M. (2020). One Central Bank to Rule Them All. *Review of Finance*, 24(2):263–304.

Cieslak, A., Morse, A., and Vissing-Jørgensen, A. (2019). Stock returns over the FOMC cycle. *The Journal of Finance*, 74(5):2201–2248.

Cieslak, A. and Vissing-Jørgensen, A. (2020). The Economics of the Fed Put. *The Review of Financial Studies*, 34(9):4045–4089.

Gregory, A., Tharyan, R., and Christidis, A. (2013). Constructing and testing alternative versions of the Fama–French and Carhart models in the UK. *Journal of Business Finance & Accounting*, 40(1-2):172–214.

Hu, G. X., Pan, J., Wang, J., and Zhu, H. (2021). Premium for Heightened Uncertainty: Explaining Pre-Announcement Market Returns. *National Bureau of Economic Research Working Paper Series*, No. 25817.

Laarits, T. (2020). Pre-Announcement Risk. *NYU Stern School of Business*.

Lucca, D. O. and Moench, E. (2015). The Pre-FOMC announcement drift. *Journal of Finance*, 70(1):329–371.

McLean, R. D. and Pontiff, J. (2016). Does academic research destroy stock return predictability? *The Journal of Finance*, 71(1):5–32.

Morse, A. and Vissing-Jørgensen, A. (2020). Information transmission from the federal reserve to the stock market: Evidence from governors' calendars. *Unpublished working paper. University of California at Berkeley*.

Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. *Handbook of Macroeconomics*, 2:71–162.