Gafka, Savor, and Wilson (2023) - "Sources of return predictability"

Discussion by Ahmed Gueciouer

MFA
8 March 2024

Summary

Comments

Summing up

Context

- Asset prices evolve differently on days with macroeconomic announcements (A-days) than on days without (N-days). On A-days,
  - timeseries: returns are higher, and so much of the equity risk premium is realized on these days (Savor and Wilson 2013; Ai and Bansal 2018; Ernst, Gilbert, and Hrdlicka 2019)
  - cross-section: the CAPM fits better (Savor and Wilson 2014), potentially because A-day information releases decrease disagreement between investors (Andrei, Cujean, and Wilson 2023)
- Gafka, Savor, and Wilson (2023) exploit this A-day vs. N-day dichotomy to tease apart the sources of (timeseries) market return predictability
  - Their idea is to use time variation in predictive relationships (A-days vs. N-days) to reason about what drives market return predictability
  - Great idea! Connects to prior work examining business cycle variations in t.s. market return predictability (Cujean and Hasler 2017; Gómez-Cram 2022), on measuring FOMC risk premia (Liu, Tang, and Zhou 2022), and on highlighting short-lived x.s. stock return predictability (Chinco, Clark-Joseph, and Ye 2019;...)

Context

- Asset prices evolve differently on days with macroeconomic announcements (A-days) than on days without (N-days). On A-days,
  - timeseries: returns are higher, and so much of the equity risk premium is realized on these days (Savor and Wilson 2013; Ai and Bansal 2018; Ernst, Gilbert, and Hrdlicka 2019)
  - cross-section: the CAPM fits better (Savor and Wilson 2014), potentially because A-day information releases decrease disagreement between investors (Andrei, Cujean, and Wilson 2023)
- Gafka, Savor, and Wilson (2023) exploit this A-day vs. N-day dichotomy to tease apart the sources of (timeseries) market return predictability
  - Their idea is to use time variation in predictive relationships (A-days vs. N-days) to reason about what drives market return predictability
  - Great idea! Connects to prior work examining business cycle variations in t.s. market return predictability (Cujean and Hasler 2017; Gómez-Cram 2022), on measuring FOMC risk premia (Liu, Tang, and Zhou 2022), and on highlighting short-lived x.s. stock return predictability (Chinco, Clark-Joseph, and Ye 2019;...)

Setting

Defining A-days & N-days

- A-days: inflation, unemployment & scheduled FOMC interest rate announcements (Savor and Wilson 2013, 2014)
- N-days: complement of A-days

Low-frequency predictions

- Market excess return timeseries are aggregated up from individual log-A-day and log-N-day returns to quarterly-frequency (typically) returns
- 48 predictor variables drawn from prior literature, mostly varying at monthly frequency

Findings

New empirical facts: N-days are important

1. According to univariate timeseries quarterly return predictive regressions,
   - 1.1 there are more variables that predict N-day returns than variables that predict A-day returns
   - 1.2 mostly, variables predict either one or the other
2. The Shiller (1981) excess volatility puzzle is concentrated on N-days
3. N-day returns are significantly associated with investor (I/B/E/S brokerage) disagreement

Proposed explanation: investor disagreement

- A-days: announcements resolve disagreement
- N-days: disagreement matters; must be linked to predictability
- Related to the Andrei, Cujean, and Wilson (2023) explanation for different A-day vs. N-day CAPM fits
- I'll talk more about this mechanism later...

Summary

Comments

Summing up

1. Strengthening the predictability results

- Rather than separate regressions, what happens when you include interaction terms like $\times \mathbb{1}\{\text{Is N-day}\}_t$, plus baseline effects (and cluster)? This would let you directly test for changes in predictive relationships.
- Are out-of-sample (OOS) and in-sample (IS) results similar?
  - Norm to also report OOS return predictability; e.g. Welch and Goyal (2008) & Campbell and Thompson (2008)
  - In-sample tests can be problematic; e.g. Martin and Nagel (2022)
  - You could compare IS vs. OOS R^2 statistics, or some related measure of the degree of predictability
- What happens when you use multivariate regressions?
  - "kitchen-sink" approach (Welch and Goyal 2008)
  - Could use machine learning techniques if concerned about high dimensionality; some even allow for significance testing of coefficients (e.g. ridge regression)
- Why quarterly returns? Robust to using monthly?
- Table 6, IS univariate regressions: do t-stats account for potential heteroskedasticity & serial correlation?

2. Variation within A-days

FOMC vs non-FOMC?

- On FOMC days, the Fed might be revealing its private information about fundamentals (Romer and Romer 2000), but it may simply be responding to news itself (Bauer and Swanson 2023)...
- ... so financial market participants could be learning about the Fed's policy rule (Bauer, Pflueger, and Sunderam 2022)
- If the mechanism is indeed about the release of information about fundamentals, all A-day results should be robust to dropping FOMC days

Variation with uncertainty reduction?

- Savor and Wilson (2013) document that the VIX, an ex-ante measure of uncertainty, decreases from just before to just after announcements
- Wachter and Zhu (2022) argue these VIX decreases show uncertainty is being resolved for investors, rather than the price of risk changing
- How does A-day predictive power vary with A-day ΔVIX?
  - What about more generally, beyond just A-days?

3. Disagreement and predictability (1/2)

- Following Andrei, Cujean, and Wilson (2023), Section 5 of Gafka, Savor, and Wilson (2023) applies the Law of Total Variance to across-investor i disagreement, plus assumptions, to arrive at a stock-specific equality:

$$\mathbb{V}\text{ar}[\tilde{R}^e] = \mathbb{V}\text{ar}_i[\tilde{R}^e] + \mathbb{V}\text{ar}[\underbrace{\overline{\mathbb{E}}[\tilde{R}^e]}_{\text{x.s. mean}}] + \underbrace{\mathbb{V}\text{ar}[\mathbb{E}_i[\tilde{R}^e] - \overline{\mathbb{E}}[\tilde{R}^e]]}_{\text{x.s. dispersion}}. \quad (1)$$

- Aggregating up to market returns, and conditioning on t,

$$\mathbb{V}\text{ar}_t[\tilde{R}^e_{M,t+1}] = \mathbb{V}\text{ar}_{i,t}[\tilde{R}^e_{M,t+1}] + \mathbb{V}\text{ar}_t\left[\overline{\mathbb{E}}[\tilde{R}^e_{M,t+1}]\right] + \mathbb{V}\text{ar}_t\left[\mathbb{E}_i[\tilde{R}^e_{M,t+1}] - \overline{\mathbb{E}}[\tilde{R}^e_{M,t+1}]\right]. \quad (2)$$

- Next, Gafka, Savor, and Wilson (2023):

  1. go on to plug in the conditional CAPM into Eqn. (2) so that $\overline{\mathbb{E}}[\tilde{R}^e_{M,t+1}]$ makes an appearance, and then take a simplified version to the data on A-days and N-days, and
  2. find that disagreement $\mathbb{V}\text{ar}_t\left[\mathbb{E}_i[\tilde{R}^e_{M,t+1}] - \overline{\mathbb{E}}[\tilde{R}^e_{M,t+1}]\right]$ across investors (I/B/E/S brokerages) is associated only with realized N-day returns.

- Let's stay within this framework, and focus on Eqn. (2)...

3. Disagreement and predictability (2/2)

- The following is a simple rearrangement of Eqn. (2):

$$\underbrace{\frac{\mathbb{V}\text{ar}_t[\tilde{R}^e_{M,t+1}] - \mathbb{V}\text{ar}_{i,t}[\tilde{R}^e_{M,t+1}]}{\mathbb{V}\text{ar}_t[\tilde{R}^e_{M,t+1}]}}_{\text{(i) } R^2\text{, frac. of return variance explained}} = \underbrace{\frac{\mathbb{V}\text{ar}_t\left[\overline{\mathbb{E}}[\tilde{R}^e_{M,t+1}]\right]}{\mathbb{V}\text{ar}_t[\tilde{R}^e_{M,t+1}]}}_{\text{(ii) also a frac. of var.}} + \underbrace{\frac{\mathbb{V}\text{ar}_t\left[\mathbb{E}_i[\tilde{R}^e_{M,t+1}] - \overline{\mathbb{E}}[\tilde{R}^e_{M,t+1}]\right]}{\mathbb{V}\text{ar}_t[\tilde{R}^e_{M,t+1}]}}_{\text{(iii) x.s. disagreement, normalized}}$$

- How might we cast the Gafka, Savor, and Wilson (2023) story in terms of the above? Public information releases on A-days reduce cross-investor disagreement, so term (iii) decreases, since private info. is crowded-out.
- Next, how to define return predictability? An investor i who can predict returns can achieve some positive R^2, which is term (i) above.
- N-day results interpreted from this perspective: when disagreement (iii) is higher, return predictability (i) may be higher, if (ii) is held constant.
- But can reasoning about investor disagreement tell us why certain variables predict on one day vs. the other?
- Also, why shouldn't the public signals used in this paper affect the explanatory power of consensus beliefs ($1 - (ii) = R^2$), in addition to (i)?

4. More direct evidence on sources of return predictability

Gafka, Savor, and Wilson (2023): *"While A-day predictors are driven by future fundamentals, N-day predictors seem to be predicting the 'noise' component of stock market movements."*

- Both of the above statements can be tested more directly
- Are A-day predictors actually informative about A-day macro variable announcements, even when controlling for professional economists' forecasts (e.g. SPF, Blue Chip, Bloomberg consensus estimates)?
- Can you rule out if N-day predictors might (partly) reveal information about "fundamentals" that's then (more fully) revealed on A-days?
  - For example, extract the residual "noise" component of returns from your Shiller (1981)-style analysis and try to predict that

Other points

- What do you mean by "reversals"?
- Section 5: update notation to clarify vectors vs. scalars, following Andrei, Cujean, and Wilson (2023)
- Pages 24-25: subscripts $\{M, t+1\}$ for returns inconsistent?
- Empirically reassure that the omitted term $\mathbb{V}\text{ar}_t[\overline{\mathbb{E}}[\tilde{R}^e_{M,t+1}]]$ in eqn. (14) does not enter into the residual when estimating eqn. (15) on the data
- What exactly is the t in the empirical proxy of $\mathbb{V}\text{ar}_t\left[\mathbb{E}_i[\tilde{R}^e_{M,t+1}] - \overline{\mathbb{E}}[\tilde{R}^e_{M,t+1}]\right]$? i.e. at what point(s) t is the cross-sectional dispersion calculated: over some look-back prior to the beginning of the quarter?
- Table 8: why not correct standard errors for heteroskedasticity and serial correlation?

Summary

Comments

Summing up

Conclusion

- Promising approach and thought-provoking findings!
- Additional analyses could help us better understand what is driving both A-day and N-day predictability
- How important actually is investor disagreement as an explanation for the A-day vs. N-day dichotomy?

Good luck!

References I

Ai, Hengjie, and Ravi Bansal. 2018. "Risk preferences and the macroeconomic announcement premium." *Econometrica* 86 (4): 1383–1430.

Andrei, Daniel, Julien Cujean, and Mungo Wilson. 2023. "The lost capital asset pricing model." *Review of Economic Studies* 90 (6): 2703–2762.

Bauer, Michael D, Carolin E Pflueger, and Adi Sunderam. 2022. "Perceptions about monetary policy." *NBER Working Paper 30480.*

Bauer, Michael D, and Eric T Swanson. 2023. "An alternative explanation for the "Fed information effect"." *American Economic Review* 113 (3): 664–700.

Campbell, John Y, and Samuel B Thompson. 2008. "Predicting excess stock returns out of sample: Can anything beat the historical average?" *Review of Financial Studies* 21 (4): 1509–1531.

Chinco, Alex, Adam D Clark-Joseph, and Mao Ye. 2019. "Sparse signals in the cross-section of returns." *Journal of Finance* 74 (1): 449–492.

Cujean, Julien, and Michael Hasler. 2017. "Why does return predictability concentrate in bad times?" *Journal of Finance* 72 (6): 2717–2758.

References II

Ernst, Rory, Thomas Gilbert, and Christopher M Hrdlicka. 2019. "More than 100% of the equity premium: How much is really earned on macroeconomic announcement days?" *Available at SSRN 3469703.*

Gafka, Beata, Pavel Savor, and Mungo Wilson. 2023. "Sources of return predictability."

Gómez-Cram, Roberto. 2022. "Late to recessions: Stocks and the business cycle." *Journal of Finance* 77 (2): 923–966.

Liu, Hong, Xiaoxiao Tang, and Guofu Zhou. 2022. "Recovering the FOMC risk premium." *Journal of Financial Economics* 145 (1): 45–68.

Martin, Ian WR, and Stefan Nagel. 2022. "Market efficiency in the age of big data." *Journal of Financial Economics* 145 (1): 154–177.

Romer, Christina D, and David H Romer. 2000. "Federal Reserve information and the behavior of interest rates." *American Economic Review* 90 (3): 429–457.

Savor, Pavel, and Mungo Wilson. 2014. "Asset pricing: A tale of two days." *Journal of Financial Economics* 113 (2): 171–201.

References III

Savor, Pavel, and Mungo Wilson. 2013. "How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements." *Journal of Financial and Quantitative Analysis* 48 (2): 343–375.

Shiller, Robert J. 1981. "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review* 71 (3): 421–436.

Wachter, Jessica A, and Yicheng Zhu. 2022. "A model of two days: Discrete news and asset prices." *Review of Financial Studies* 35 (5): 2246–2307.

Welch, Ivo, and Amit Goyal. 2008. "A comprehensive look at the empirical performance of equity premium prediction." *Review of Financial Studies* 21 (4): 1455–1508.