A Note on Trading the Term Structure of VIX Futures

Anusar Farooqui*

March 18, 2020

Abstract

The term structure of VIX futures contains a very strong signal of dealer risk appetite. Unlike balance sheet quantities, this feature is available at very high frequencies. Here we exhibit two systematic strategies to mine the attendant risk premium from the term structure of expected volatility. We optimize our two hyperparameters by OOS cross-validation. We compare our strategies to holding the S&P 500, selling short-term vol unhedged, and a portfolio that sells short-term vol and hedges by going long on medium-term vol. We find that our strategies allow us to harvest a considerable portion of the risk premium associated with the balance sheet management of market-based intermediaries. Both in-sample and OOS, the risk-adjusted returns on our strategies are at least twice as high as the three benchmarks.

1 Introduction

The marginal investor in many asset classes is a market-based intermediary; not a retail investor. The marginal value of wealth to these investors prices most asset classes. More precisely, fluctuations in the risk-bearing capacity of US securities broker-dealers drive

*anusar.farooqui@outlook.com

fluctuations in the risk premia embedded in asset prices.[1] Dealer risk-appetite is priced into the cross-section of excess returns on US stocks, exchange rates, commodities, and fixed-income securities. These are the main insights of the extant literature on intermediary asset pricing. There is a problem, however. Most proxies of risk-appetite are based on balance sheet quantities that are only available at the quarterly frequency. This makes the insights of intermediary asset pricing practically useless to market-based intermediaries themselves since they operate at a much higher frequency. Indeed, dealers care more about daily fluctuations, or even intra-day fluctuations, than fluctuations at the quarterly frequency.

The problem is not insurmountable. We can, in fact, proxy dealer risk-appetite with its dual. Specifically, we can isolate the signal of risk-appetite in real-time by asset prices that are most responsive to dealer balance sheet management. We know that dealers calculate their value-at-risk as a product of their leverage and daily volatility. Given prevailing levels of volatility, dealers choose their capital ratios to target a constant probability of survival. This means that the price of volatility is a highly sensitive barometer of dealer risk-appetite. Indeed, He and Krishnamurthy (2013) find that virtually all variation in the risk premia embedded in option prices is explained by dealer risk-appetite.

The VIX is an aggregate measure of expected volatility. In earlier work, we have shown that innovations in the VIX are priced into the cross-section of excess returns on US stocks. That allows us to harvest some portion of the intermediary risk premium. In order to harvest a greater portion of the intermediary risk premium we must invest in variance assets such as VIX futures and variance swaps. In particular, we can harvest a greater portion of the intermediary risk premium by holding the term structure of expected volatility. But as we shall see, we can do much better.

In what follows, we document the performance of a feature derived from the term

[1] Etula (2013), Danielsson et al. (2011), Adrian et al. (2014), He and Krishnamurthy (2013), Farooqui (2017).

structure of expected volatility. We first exploit the predictive information contained in our feature to predict the probability of risk-off on the following day. Then we construct portfolios of variance assets that are either long or short the term structure of volatility. We evaluate the performance of our strategies both in-sample and out-of-sample (OOS). The results speak for themselves.

2 The Probability of a Risk-Off Tomorrow

We cannot reveal the specific feature that allows us to predict a risk-off tomorrow due to the commercial possibilities.[2] We can, however, reveal that our feature is computed from VIX futures. We obtain VIX futures data from the CBOE's website, data on volatility ETFs from the *Financial Times* website, and the S&P 500 Index from *FRED*. We use the secondary market rate on the 3-month bill as the risk-free rate; also obtained from *FRED*. Our dataset begins on Jan 2, 2013 and ends on March 16, 2020. We use the first 252 days to initially train out predictive model. All tests that follow are based on data from Jan 1, 2014 to March 16, 2020.

We sort all VIX futures on maturity and place them into three buckets, Low, Mid, and High, by maturity quantile. We use High *minus* Low as our VIX futures instrument for what we call our *wholesale strategy*. We use ProShares's ETFs for short-term and mid-term volatility, VIXY and VIXM respectively, as our instruments for what we call our *retail strategy*. The wholesale and retail versions of our systematic strategy are especially interesting to compare since they may reveal the existence of premium earned by institutional investors over and above retail investors. Conversely, the premium may turn out to be negative, implying that institutional investors have more than bid away the wholesale premium in their hunt for yield.

[2]Please get in touch with the author if you'd like to discuss commercial arrangements.

Our predictive model relies on a hyperparameter that we cross-validate out-of-sample using 5-fold CV. We use Prado's *log loss* as our loss function since our regime switching strategy is especially exposed to bad predictions with high confidence.[3] Figure 1 displays the OOS log loss as a function of the prediction model hyperparameter α. Luckily for us, log loss turns out to have a unique global minimum, considerably simplifying our prediction problem.

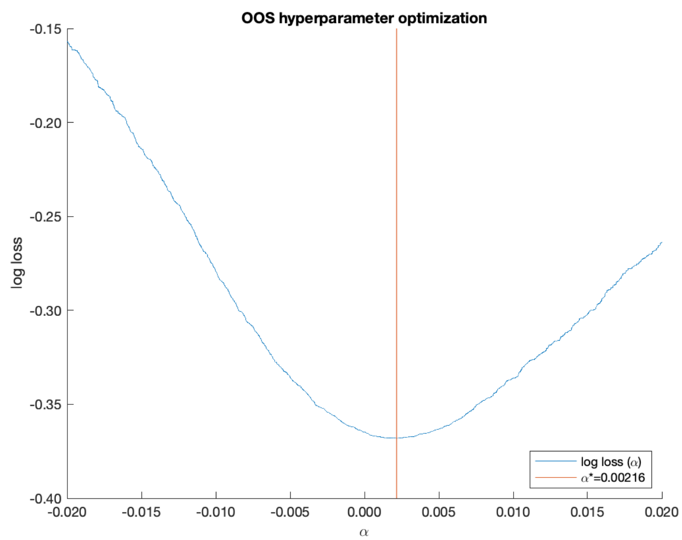

Figure 1: OOS crossvalidation of prediction model hyperparameter.

Figure 2 displays the predicted probability of risk-off on the next day. The major spikes correspond to known events. The first big spike corresponds to the "China panic" on August 24, 2015. The second one, corresponds to the dramatic return of volatility on February 5, 2018. Our predictive model called both of them. It also called the doldrums towards the end of 2018, and, of course, the dramatic revival of systematic volatility associated with the Coronavirus pandemic that got going on February 24, 2020 and is still underway.

On all of these and a few others, the predicted probability of risk-off exceeded 50

[3] Marcos López de Prado. *Advances in Financial Machine Learning*, 2018, p. 134.

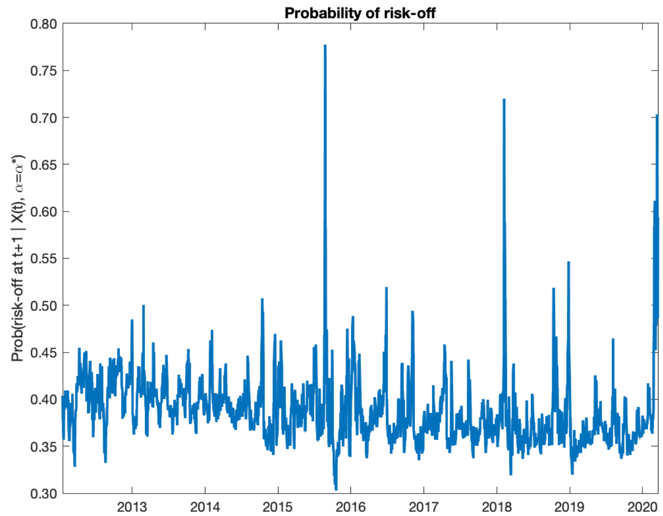

Figure 2: Predicted probability of risk-off on the next day.

percent. But what is the appropriate threshold to call a risk-off? When should we move from being short the term spread, ie selling vol while hedging with medium-term vol, to being long? The probability cutoff is location hyperparameter than we tune, again with 5-fold CV, to achieve the highest Sharpe ratio OOS. Figure 3 and Figure 4 show that the same parameter is optimal for both the retail and the wholesale strategies. This is striking confirmation that the globally optimal hyperparameter is independent our choice of instruments. Ie, no matter which instrument we use, the OOS results suggest that $\theta^* = 0.439$.

The OOS Sharpe ratio is slightly higher for the retail strategy. This suggests that institutional investors bid away whatever wholesale premium existed. Figure 5 and Figure 6 display the returns on our retail and wholesale strategies along with the days on which we predicted a risk-off. Figure 7 displays the Sharpe ratios of the two strategies along with those of the benchmarks over the whole sample. Figure 8 displays the OOS Sharpe ratios for the same. Figure 9 displays the max drawdowns, and Figure 10 displays the return on

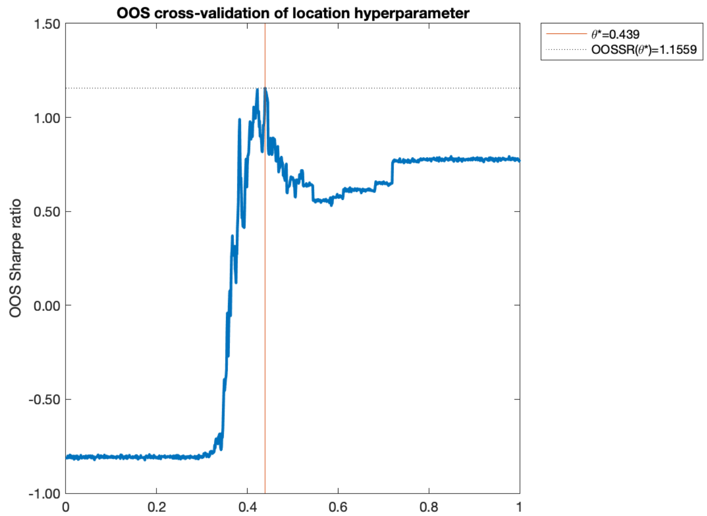

Figure 3: OOS crossvalidation of location hyperparameter for wholesale strategy.

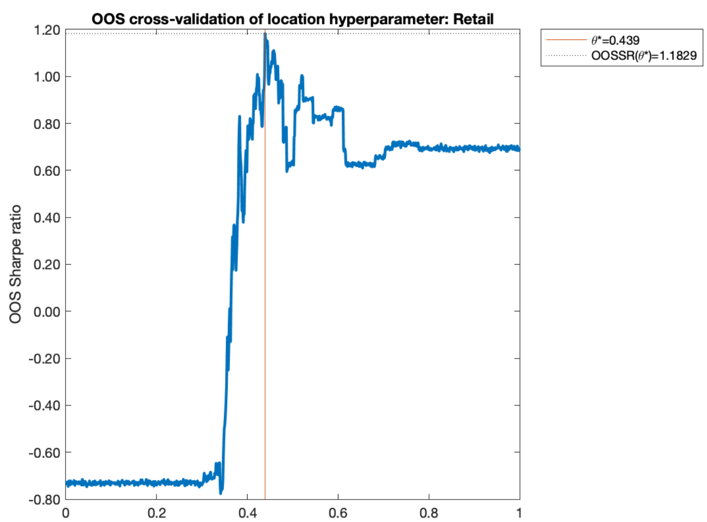

Figure 4: OOS crossvalidation of location hyperparameter for retail strategy.

max drawdown. Figure 11 displays the cumulative returns on the five strategies.

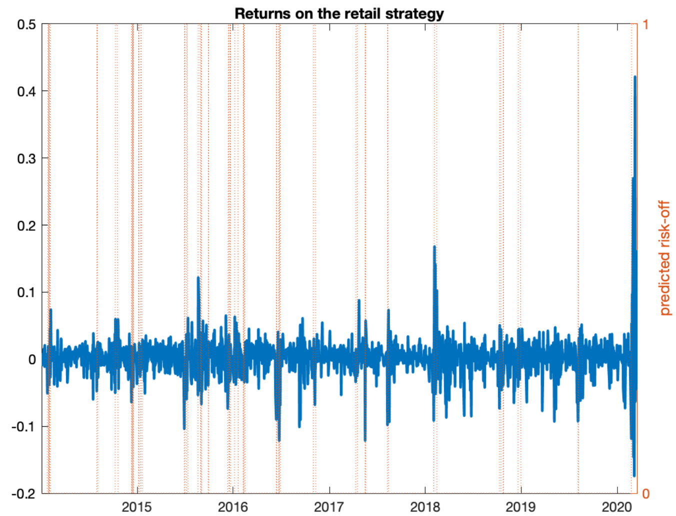

Figure 5: Return on the retail strategy.

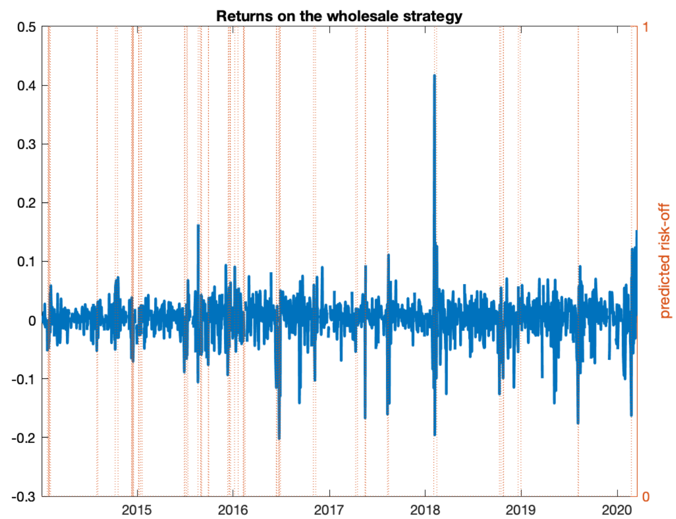

Figure 6: Return on the wholesale strategy.

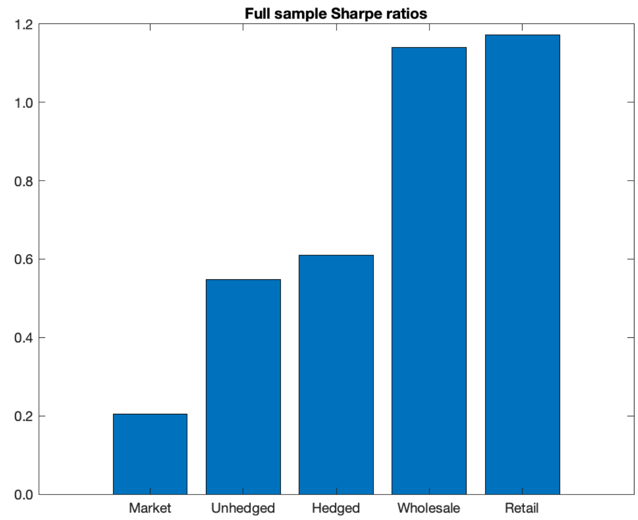

Figure 7: Full sample Sharpe ratios.

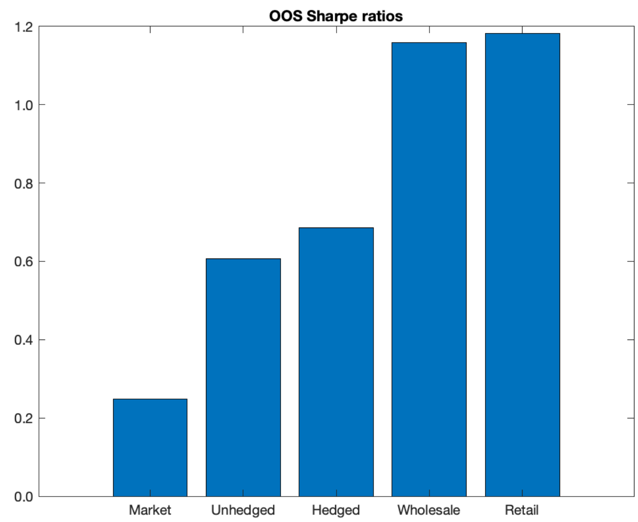

Figure 8: OOS Sharpe ratios.

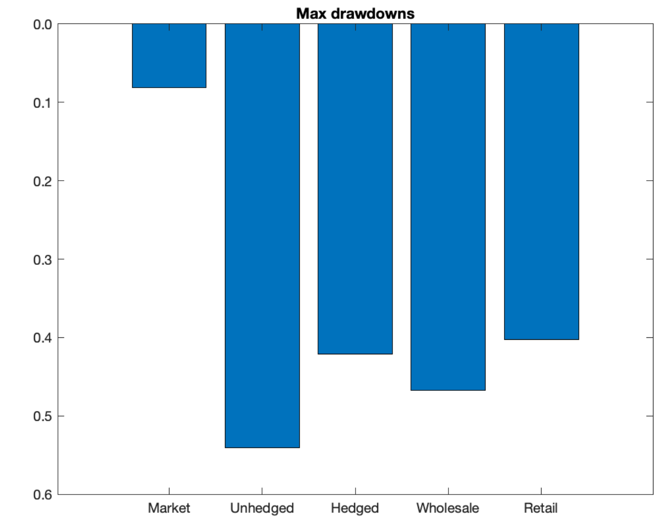

Figure 9: Max drawdowns.

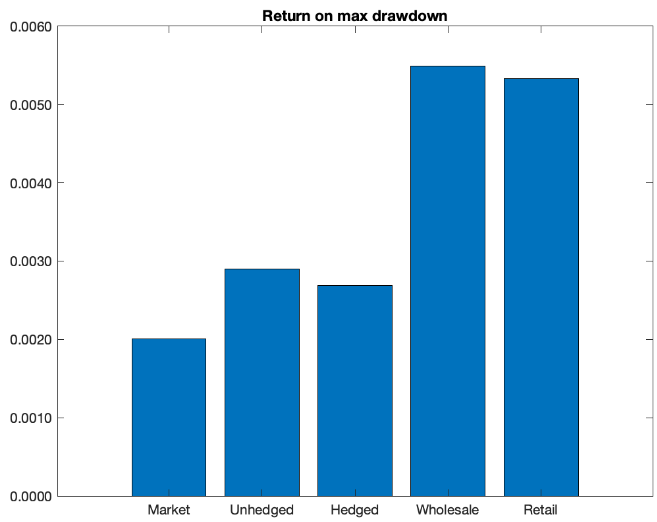

Figure 10: Return on max drawdown.

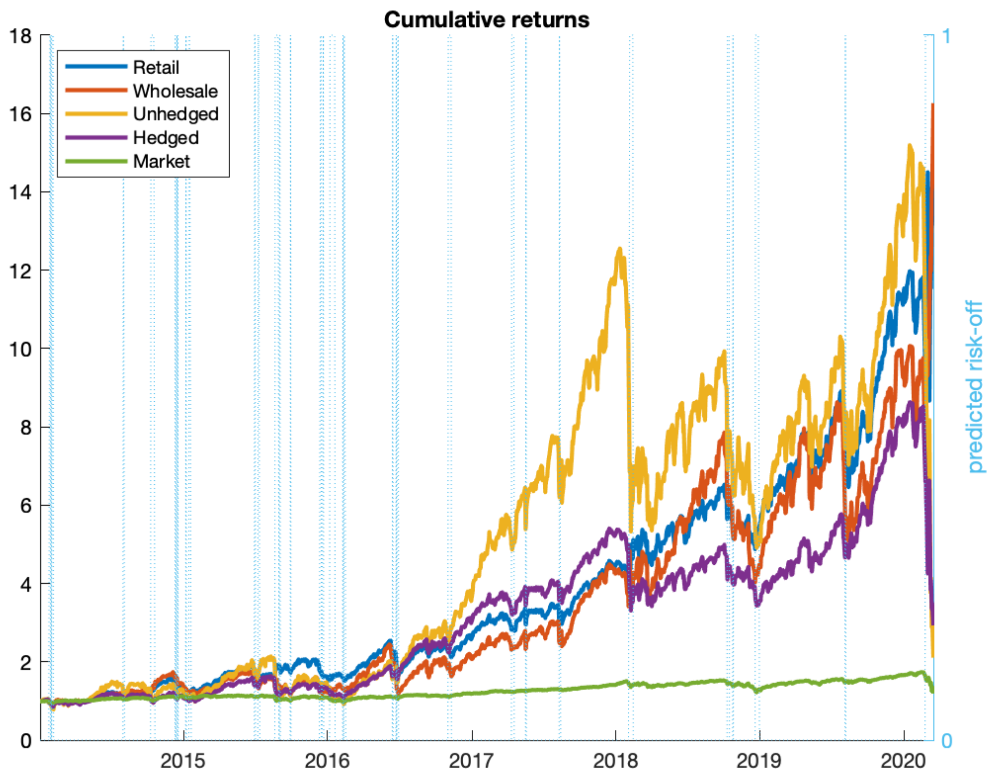

Figure 11: Cumulative returns.

3 Summary

Table 1: Return statistics.

	Market	Unhedged	Hedged	Wholesale	Retail
Mean daily return	0.0002	0.0016	0.0011	0.0026	0.0021
Volatility	0.0102	0.0445	0.0287	0.0352	0.0286
Sharpe ratio	0.2559	0.5596	0.6282	1.1555	1.1910
Sharpe ratio (OOS)	0.2479	0.6061	0.6859	1.1587	1.1818
Skewness	-2.0625	-2.4015	-3.2425	0.2903	2.1231
Return on MaxDD	0.0020	0.0029	0.0027	0.0055	0.0053
Annualized risk-adj. return[4]	0.0334	0.0923	0.1032	0.2016	0.2079

Table 1 displays summary statistics of our two systematic strategies as well as the three benchmarks. What is particularly interesting is that our tactical strategies to trade the vol term structure exhibit positive skew. This is probably driven by last month's outsized returns on our strategies. As the market has cratered, our predictive algorithm has called most risk-off days with uncanny prescience. This gives us more reason, not less, to trust our feature.

The bottom line is that the risk-adjusted return on our systematic strategies is twice as high as benchmark portfolios both in- and out-of-sample. This is the strongest possible evidence that we have truly isolated the signal for intermediary risk-appetite at arbitrary frequencies.