

A Model of Multi-Frequency Trade*

Nicolas Crouzet, Ian Dew-Becker, and Charles G. Nathanson

March 5, 2017

Abstract

We develop a noisy rational expectations model of financial trade featuring investors who acquire information and trade at a range of different frequencies. In the model, a restriction on high-frequency trading affects efficiency of prices at high frequencies, but leaves low-frequency efficiency unaffected. In a particular equilibrium of the model, traders specialize into trading at individual frequencies. We show that high- and low-frequency investors coexist, trade with each other, and make money from each other. The model matches numerous basic features of financial markets: investors endogenously specialize into strategies distinguished by frequency; volume is disproportionately driven by high-frequency traders; and the portfolio holdings of informed investors forecast returns at the same frequencies as those at which they trade.

Investors in financial markets follow many different strategies, including value investing, technical analysis, macro strategies, and algorithmic trading. These strategies differ in two salient ways. First, they require investors to learn about different aspects of asset prices; market-makers or algorithmic traders care more about the high-frequency movements of prices, while value investing puts more emphasis on their slow-moving features. These investors all understand that their information sets may not overlap, and yet they trade with each other, presumably making some money in the process. Second, these strategies differ in the frequency at which they require investors to trade, or equivalently the rate at which they turn over their positions.

This paper proposes an equilibrium model in which investors endogenously specialize in acquiring information and trading at different frequencies. There is a single fundamentals process, and a continuum of investors who trade forward contracts on the fundamental. These investors also learn about different aspects of asset dynamics. An example of the fundamentals process is the spot price of oil: investors are able to acquire information that tells them about the future path of oil prices, allowing them to potentially earn profits on the forward contracts. As is common elsewhere, in order to grease the wheels of the market, we assume that investors trade against an exogenous flow of demand for forward contracts that fluctuates stochastically over time.

We show that in such a model, there exists a natural (though not necessarily unique) equilibrium in which individual investors endogenously choose to focus on specific frequencies of the

*Crouzet: Northwestern University. Dew-Becker: Northwestern University and NBER. Nathanson: Northwestern University. We appreciate helpful comments from Stijn van Nieuwerburgh.

fundamentals. Some investors learn about low-frequency aspects of oil prices in the sense that they get a signal about their average path over, say, a period of decades, while others learn about higher-frequency behavior, receiving a signal about how oil prices vary from day to day or month to month. This occurs despite the fact that the learning technology is fully general, and in no way tilts investors towards frequency specialization *ex-ante*.

Given attention allocation – what aspect of fundamentals investors choose to learn about – in equilibrium we show that their positions fluctuate at the frequency at which they receive signals. That is, investors who learn about long-run fundamentals hold positions in forward contracts that fluctuate slowly over time, whereas those who do high-frequency research have positions that vary at high frequencies. So we have a model in which people choose to learn about high- or low-frequency aspects of fundamentals, and that learning causes them to endogenously become high- or low-frequency traders.

While there is other research on investors who trade at different frequencies, that work typically endows investors with investment horizons that differ exogenously.¹ In our setting, all investors have the same objective, maximizing utility over identical horizons. We view this as an important restriction in our setting because it is obviously not the case that people who trade at high frequencies, e.g. turning over their portfolios once per day, really have investment horizons of only 24 hours. Rather, all investors want to maximize the same utility function over wealth, they just go about it in different ways.

What is particularly interesting about the equilibrium that we obtain is that it is not the case that the informed investors trade only with the exogenous demand (i.e. liquidity traders). In fact, high- and low-frequency traders trade with each other. The simple reason is that a high-frequency trader cannot distinguish uninformative demand shocks from the orders of informed low-frequency traders (and vice versa). So in periods when fundamentals are persistently strong, low-frequency traders tend to hold persistently more forward contracts than high-frequency traders and earn profits from them. Similarly, if there is a very transitory increase in fundamentals, the high-frequency traders tend take advantage and earn profits while the low-frequency traders lose, as they ignore the temporary trading opportunity. In that sense, then, everybody is a noise trader sometimes, and they all understand that, but they still participate and make money on average.

The model has a number of predictions for observable features of financial markets. First, as we have already discussed, it predicts that there are traders who can be distinguished by the frequencies at which their asset holdings change over time, and they do research about fundamentals at the same frequencies. So we obtain endogenous high- and low-frequency traders with a specific prediction for how research aligns with trade.

The model also matches salient facts about differences in volume across investors. We can

¹See, e.g., Amihud and Mendelson (1986), who assume that investors are forced to sell after random periods of time; Hopenhayn and Werner (1996), who assume that investors vary in their rates of pure time preference, and Defusco, Nathanson, and Zwick (2016) who assume that there are sets of investors who are exogenously forced to sell at deterministic horizons that vary across groups. Turley (2012), like us, studies a setting in which investors endogenously choose to learn about high- or low-frequency information, though he studies only a two-period case.

very easily show analytically that high-frequency investors account for a fraction of aggregate volume that is out of proportion to their fraction of total asset holdings. An interesting implication of that result is that incorporating trading costs into the model can have substantial effects on optimal information acquisition strategies. Since high-frequency trade requires paying much larger transaction costs than low-frequency trade, any trading costs cause prices to naturally be less efficient and for there to be less liquidity at high than low frequencies.² Moreover, as transaction costs fall, we expect to see a shift towards higher-frequency trade and for prices to become more efficient at high frequencies (see also Turley (2012)).

The model is fundamentally about differences in information across investors. People obtain information in order to make money, and so their asset holdings in general should forecast returns. We see that both in the model and in the data.³ But different investors' holdings do not forecast returns in the same way. The frequency at which an investor's portfolio holdings forecasts returns is the same as the frequency at which they trade: high-frequency investors' positions forecast returns at very short horizons, while buy-and-hold investors' portfolios forecast returns over much longer periods.

The idea that an investors' asset holdings should forecast returns over a period related to how long those assets will be held is perhaps not surprising. Studies of the holdings of mutual funds and other institutional investors typically examine returns over a period of perhaps 3–12 months. At the other extreme, Brogaard, Hendershott, and Riordan (2014) show that the holdings of high-frequency traders (as defined by NASDAQ) forecast returns over periods of 1–5 seconds – a horizon 7 orders of magnitude smaller than a calendar quarter.

To empirically test our model, we provide novel evidence on the relationship between turnover and asset return predictability. Using form 13F data on institutional asset holdings, we first show that asset turnover within funds is highly persistent over time, suggesting that it is a salient feature of investor strategies. Next, after confirming past results that institutional holdings predict returns, we show that the predictive power of the holdings of high-turnover funds decays much more quickly than those of low-frequency funds, consistent with the model.

Finally, we use the model to study the effects of a policy that restricts high-frequency trade. Such a policy has the obvious effect of reducing the informativeness of prices at high frequencies, but it has no effect at low frequencies. The practical implication is that while prices at any single moment contain less information than without the policy, moving averages of prices remain almost equally informative about moving averages of dividends. So to the extent that economic decisions are made based on an average of prices over time, rather than a price at a single moment, the model implies that restricting high-frequency trade will not reduce the information available for those decisions.

²Gârleanu and Pedersen (2013) also discuss how high-frequency information is less valuable in the presence of trading costs, while Dávila and Parlato (2016) study in a related setting how trading costs can affect information acquisition, but without our focus on differences across frequencies.

³See, e.g., the literature on the predictive power of mutual fund and institutional investor asset holdings for future returns, such as Carhart (1997) and Yan and Zhang (2009), among many others.

To summarize, we develop a model that matches a number of major features of trade in financial markets: investors can be distinguished by the frequencies at which they trade; volume is accounted for by high-frequency traders; and the holdings of investors forecast returns at horizons similar to their holding periods. The model can then be used to analyze the effects of restricting trade at specific frequencies.

In general models with dynamic trade are extremely difficult to solve; solutions typically require some kind of restriction, such as to a very narrow class of driving processes (e.g. AR(1) and Ornstein–Uhlenbeck processes studied in Wang (1993, 1994) and He and Wang (1995)), or to only two or three period horizons. We allow for a very long investment horizon and place only technical constraints on the fundamental processes driving the model and obtain fully analytic solutions. The sacrifice that we make is that information sets are fixed on date 0 – investors only obtain signals once. The model should be thought of as essentially a stationary equilibrium: it gives a steady-state description of trade, volume, and returns. It is not well suited to studying how investors and markets respond to shocks to information sets.

The major advantage of our particular information structure is that it allows us to take a long-horizon dynamic model and solve it as a series of parallel scalar problems. In particular, solving our model is only marginally more difficult than solving a standard single-period/single asset noisy rational expectations model – it reduces to a parallel set of such equilibria. The paper thus has useful methodological contributions for analyzing models of trade over time.

This paper builds on a growing recent literature that tries to understand optimal information acquisition in financial markets. The most important building blocks are the models of van Nieuwerburgh and Veldkamp (2010) and Kacperczyk, van Nieuwerburgh, and Veldkamp (2016) in that we use a highly similar information and market structure and build on their results on optimal information acquisition (their reverse water-filling solution, in particular). Those papers themselves build critically on work by Grossman and Stiglitz (1980), Hellwig (1980), Diamond and Verrecchia (1981), and Admati (1985) on rational expectations equilibria. More recently, research has tried to understand the effects on the equilibria developed in those earlier papers of various limits on information gathering ability (e.g. Banerjee and Green (2015) and Dávila and Parlatore (2016)).

There is also a literature on price dynamics in rational expectations equilibria, though it is relatively small given how difficult dynamic models are to solve. In particular, a series of papers by Wang (1993, 1994) and He and Wang (1995) study the implications of dynamic equilibria on prices and volume. Those papers are based around AR(1) or Ornstein–Uhlenbeck-type dynamics to maintain tractability (see also Wachter (2002)), whereas we study a setting in which the various exogenous time series may follow processes with minimally constrained autocorrelations. Furthermore, we focus on how investment strategies differ across investors, whereas those papers focus on symmetric strategies. A number of papers also study overlapping generations models, which can help alleviate some of the difficulties with dynamic trade.⁴

There is also a large literature on disagreement in financial markets. In addition to the above

⁴See Spiegel (1998), Watanabe (2008), and Banerjee (2011), among others.

work, see, e.g., Townsend (1983), Basak (2005), Hong and Stein (2007), and Banerjee and Kremer (2010), who focus, like us, on dynamics. In our setting, disagreement arises not just because agents receive signals that have random errors, but also because their signals have different relationships with fundamentals. High-frequency and low-frequency investors will often disagree about the price path of the asset over time because they learn about different characteristics of fundamentals – the path over the next few minutes, say, versus the path over the next several years.

Our desire to develop a model that can match salient features of the cross-section of investment strategies follows from a large empirical literature that documents the behavior of many different types of investors and how it affects the aggregate behavior of financial markets. Chen, Jegadeesh, and Wermers (2000), Gompers and Metrick (2001), Nagel (2005), Griffin and Xu (2009), Yan and Zhang (2009), and Brogaard, Hendershott, and Riordan (2014), for example, study the behavior of institutional investors and how their holdings relate to asset returns. Turley (2012), Bai, Philippon, and Savov (2016), and Weller (2016) study how price informativeness has changed over time and how it is affected by trading costs and the number of investors who trade at different frequencies.

Finally, our work is related to a small literature that studies the properties of asset returns and portfolio choice in the frequency domain including Bandi and Tamoni (2014), Chincio and Ye (2016), Chaudhuri and Lo (2016), and Dew-Becker and Giglio (2016).

The remainder of the paper is organized as follows. Section 1 describes the basic environment, and we solve for optimal information acquisition in section 2. Section 3 examines the implications of the model for the behavior of individual investors in a setting that features investors who specialize in trade at a particular frequency. Section 4 presents empirical evidence on the behavior of institutions consistent with our model of specialization. Finally, section 5 presents our key results on the effects of restrictions on high-frequency trade on return volatility and price efficiency at different frequencies, and section 6 concludes.

1 Asset market equilibrium

We begin by describing the basic market structure and the asset market equilibrium. This section introduces the description of trading strategies in terms of frequencies and shows how the frequency transformation makes multi-period investment a purely scalar problem. The problem is solved from the perspective of date 0.

1.1 Market structure

Time is denoted by $t \in \{-1, 0, 1, \dots, T\}$, with T even, and we will focus on cases in which T may be treated as large. There is a fundamentals process D_t that investors make bets on with realizations on all dates except -1 and 0 . The time series process is stacked into a vector $D \equiv [D_1, D_2, \dots, D_T]'$ (variables without subscripts denote vectors) and is distributed as

$$D \sim N(0, \Sigma_D). \tag{1}$$

The fundamentals process is assumed to be stationary, meaning that it has constant unconditional autocovariances. Stationarity implies that Σ_D is Toeplitz (all diagonals are constant), and we further assume that the eigenvalues of Σ_D are finite and bounded away from zero.⁵

On date 0, there is a market for forward claims on D_t for all $t > 0$. A unit mass of investors indexed by $i \in [0, 1]$ meets on date 0 and commits to a set of trades of futures contracts maturing on all dates. P_t denotes the price of a claim to the fundamental D_t .

There is an exogenous supply of futures, Z , which is distributed as

$$Z \sim N(0, \Sigma_Z). \quad (2)$$

Z_t may be thought of as either exogenous liquidity demand or noise trading. The time series process for supply is also assumed to be stationary. For markets to clear, the net demand of the investors for the fundamental on date t must equal Z_t ,⁶

$$\forall t : \int_i Q_{i,t} di = Z_t, \quad (3)$$

where $Q_{i,t}$ is the number of date- t forward claims agent i buys.

A concrete example of a potential process D_t is the price of crude oil: oil prices follow some stochastic process and investors trade futures on oil at many maturities. D_t can also be interpreted as the dividend on a stock, in which case P_t is the price of a forward claim on a single dividend.

1.1.1 Modeling equities

While the concept of a futures market on the fundamentals will be a useful analytic tool, we can also obviously price portfolios of futures. We model equity as a claim to the stream of fundamentals over time. To purchase such a claim, one would enter into futures contracts for the fundamental on each date $t + j$. Since futures contracts specify that money only changes hands at maturity, the money that must be set aside on date t for a futures contract that expires at $t + j$ is $P_{t+j}R^{-j}$, where R is the discount rate (which is assumed here to be a constant). The date- t cost of a claim to the entire future stream of fundamentals is therefore

$$P_t^{equity} \equiv \sum_{j=1}^{T-t} R^{-j} P_{t+j} \quad (4)$$

Holding any given combination of futures claims on the fundamental D is therefore also equivalent to holding futures contracts on equity claims (i.e. committing to a trading strategy in equities at prices that are agreed on at date 0). Any desired set of exposures to fundamentals over time

⁵The analysis is similar if a transformation of D_t (e.g. its first difference) is stationary. See appendix section A.

⁶It is also possible to assume that there is an exogenous downward-sloping supply curve of the fundamental that shifts stochastically over time; our results go through similarly. This case is treated as part of the analysis of appendix 6.

can be obtained either through purchases of futures or through suitable trading strategies for the equity claim (assuming prices can be committed to or that they are predetermined, which will be the case in our equilibrium).

Our analysis of pricing will focus on futures as they will give the most direct analog to past work. When we discuss volume and trading costs, though, we will take advantage of the equity-based implementation.

1.2 Information structure

The realization of the time series of fundamentals, $\{D_t\}_{t=1}^T$, can be thought of as a single draw from a multivariate normal distribution. Investors are able to acquire signals about that realization. The signals are a collection $\{Y_{i,t}\}_{t=1}^T$ *observed on date 0* with

$$Y_{i,t} = D_t + \varepsilon_{i,t}, \varepsilon_i \sim N(0, \Sigma_i), \quad (5)$$

Information sets in the model are fixed on date 0. Through $Y_{i,t}$, investors can learn about fundamentals potentially arbitrarily far into the future. $\varepsilon_{i,t}$ is a stationary error process in the sense that $cov(\varepsilon_{i,t}, \varepsilon_{i,t+j})$ depends on j but not t (again, Σ_i is Toeplitz).

The information structure here is obviously stylized. One interpretation is that we are collapsing to date 0 all the realizations of a stationary process. That is, agents have a machine that gives them signals about fundamentals plus an error, and that machine reports all of its output on date 0. The information structure is meant to generate two important features in the model. First, obviously on any particular date agents can choose to learn about fundamentals on more than just a single date in the future – they can potentially get information about fundamentals in many different periods (e.g. next quarter vs. over the next five years). Second, by restricting $\varepsilon_{i,t}$ to be “stationary”, we are forcing agents to choose a fixed policy for information. They build a machine (or a research department) that, rather than yielding information about only a single date, returns information about the entire fundamentals stream over time in a way that places no particular emphasis on any single date.

In choosing Σ_i agents will have two choices to make. First, they will be able to choose how informative their signals are by choosing the variance of ε_i . Second, though, they will be able to choose how accurate the signals are about fundamentals over different horizons. Some choices for Σ_i will yield signals that are informative about transitory variation in fundamentals, while others will yield signals that are more useful for forecasting trends.

1.3 Investment objective

All trading decisions are made on date 0. Investors choose demands $\{Q_{i,t}\}_{t=1}^T$ conditional on their observed signals, $\{Y_{i,t}\}_{t=1}^T$, and the set of futures prices, $\{P_t\}_{t=1}^T$. That is, as in past work, agents submit to a central auctioneer demand curves that condition on prices.

We assume that investors have mean-variance utility over cumulative excess returns. Investor i 's objective is

$$U_{0,i} = \max_{\{Q_{i,t}\}} E_{0,i} \left[T^{-1} \sum_{t=1}^T Q_{i,t} (D_t - P_t) \right] - (\rho T)^{-1} \text{Var}_{0,i} \left[\sum_{t=1}^T Q_{i,t} (D_t - P_t) \right], \quad (6)$$

where $E_{0,i}$ is the expectation operator conditional on agent i 's date-0 information set, $\{P, Y_i\}$. $\text{Var}_{0,i}$ is the variance operator conditional on $\{P, Y_i\}$. ρ is risk-bearing capacity per unit of time.

The assumption that all plans are made on date 0 only restricts information sets in a very specific way: it means that investors are not able to condition demand on the *realized* history of fundamentals. That is, it is not free to condition on the history of D_t , even when that history has already been realized. Instead, what agents must condition on is noisy signals about fundamentals, $Y_{i,t}$.

An important implication of that assumption is that agents have no desire to change their investment choices after date 0 since they receive no further information. Agents' trading strategies can thus be equivalently implemented either through a set of purchases of futures contracts or through a dynamic plan for trading the equity claim, as described above.

We interpret the objective as representing a target that an institutional investor might have. Rather than aiming to maximize the discounted sum of returns, as a person who consumes out of wealth might, the investors we study maximize a measure of their performance. The objective can be thought of as representing CARA or quadratic preferences over the sum of excess returns, so it would appear if a manager were paid on date T a fee proportional to total excess returns up to that time. Bhattacharya and Pfleiderer (1985) and Stoughton (1993) also argue that a quadratic contract (which would induce mean-variance preferences) can appear optimally in delegated investment problems. The important characteristic of (6) is that it yields a stationary problem in the sense that there is no discounting to make returns in some periods more important than others.

Finally, note that all investors have the same investment horizon. We show in appendix F that the investment horizon as defined here by T has no effect on information choices in the model – two investors with different T will be equally likely to be high- or low-frequency investors. The simplest way to confirm that fact is to simply note, when we obtain the equilibrium strategies, that T has no effect on the type of information that investors optimally obtain.

1.4 Equilibrium

Conditional on the information choices of the agents – that is, taking the set of Σ_i (which may differ across agents) as given – we study a standard asset market equilibrium.

Definition 1 *An asset market equilibrium is a set of demand functions, $Q_i(P, Y_i)$, and a vector of prices, P , such that investors maximize utility, $U_{0,i}$, and all markets clear, $\int_i Q_{i,t} di = Z_t \forall t$.*

The equilibrium concept is that Grossman and Stiglitz (1980), Hellwig (1980), Diamond and Verrecchia (1981), and Admati (1985). Investors submit demand curves for each futures contract to a Walrasian auctioneer who selects equilibrium prices to clear all markets.

The structure is in fact mathematically that of Admati (1985), who studies investment across a set of assets that might represent stocks in different companies, and the solution from that paper applies directly here. Here we are considering investment across a set of futures contracts that represent claims on some fundamentals process across different dates. We simply rotate the Admati (1985) structure from a cross-section to a time series.

1.5 Trading frequencies

This paper is fundamentally about the behavior of markets at different frequencies, so we need a rigorous concept of what frequencies are. We use the fact that fluctuations at different frequencies represent an (asymptotic) orthogonal decomposition of any time series.

Define a set of vectors of cosines and sines at the fundamental frequencies $\omega_j = 2\pi j/T$ for $j \in \{0, 1, \dots, T/2\}$

$$c_j \equiv \sqrt{2/T} [\cos(2\pi j(t-1)/T); t = 1, 2, \dots, T]' \quad (7)$$

$$s_{j'} \equiv \sqrt{2/T} [\sin(2\pi j(t-1)/T); t = 1, 2, \dots, T]' \quad (8)$$

A cycle at frequency ω_j has an associated wavelength $2\pi/\omega_j$. $\omega_0 = 0$ thus corresponds to an infinite wavelength, or a permanent shock (a constant vector). ω_1 corresponds to a cycle that lasts as long as the sample – c_1 is a single cycle of a cosine. $\omega_{T/2} = \pi$, the highest frequency, corresponds to a cycle that lasts two periods, so that $c_{T/2}$ oscillates between $\pm\sqrt{2/T}$.

The frequency-domain counterpart to the vector of fundamentals, D , is then

$$d = \Lambda' D \quad (9)$$

$$\text{where } \Lambda \equiv \left[c_0/\sqrt{2}, c_1, \dots, c_{T/2}/\sqrt{2}, s_1, s_2, \dots, s_{T/2-1} \right], \quad (10)$$

we use the notation $d_j = c_j' D$ and $d_{j'} = s_{j'}' D$ to refer to fundamentals at particular frequencies. When the distinction is necessary, we use the notation j to refer to a frequency associated with a cosine transform and j' to refer to one with a sine transform. In what follows, lower-case letters denote frequency-domain objects. Note that Λ is orthonormal with $\Lambda' = \Lambda^{-1}$.

Since d is a linear function of D , it can be thought of as a vector of payoffs on portfolios of futures given by Λ – portfolios with weights on D_t that fluctuate over time as sines and cosines.

For our purposes, the key feature of Λ is that it approximately diagonalizes *all* Toeplitz matrices and thus orthogonalizes stationary time series.⁷

⁷This is a textbook result that appears in many forms, e.g. Shumway and Stoffer (2011). Brillinger (1981) and Shao and Wu (2007) give similar statements under weaker conditions.

Definition 2 f_X is the spectrum of X with elements $f_{X,j}$, defined as

$$f_{X,j} \equiv \sigma_{X,0} + 2 \sum_{s=1}^{T-1} \sigma_{X,j} \cos(\omega_j s) \quad (11)$$

$$f_X \equiv [f_{X,0}, f_{X,1}, \dots, f_{X,T/2}, f_{X,1}, f_{X,2}, \dots, f_{X,T/2-1}]'. \quad (12)$$

Lemma 1 For a stationary time series $X_t \sim N(0, \Sigma_X)$ with autocovariances $\sigma_{X,j} \equiv \text{cov}(X_t, X_{t-j})$,

$$x \equiv \Lambda' X \Rightarrow N(0, \text{diag}(f_X)) \quad (13)$$

where \Rightarrow denotes convergence in the sense that

$$|\Lambda' \Sigma_X \Lambda - \text{diag}(f_X)| \leq c_X T^{-1/2} \quad (14)$$

for a constant c_X and for all T .⁸ $\text{diag}(f_X)$ is a matrix with the vector f_X on its main diagonal and zero elsewhere.

Proof. This is a textbook result (e.g. Brockwell and Davis (1991)). See appendix B for a derivation specific to our case. ■

For any finite horizon, the matrix Λ does not exactly diagonalize the covariance matrix of D . But as T grows, the error induced by ignoring the off-diagonal elements of the covariance matrix $\Lambda' D$ becomes negligible (it is of order $T^{-1/2}$), and x is well approximated as a vector of independent random variables.⁹ The spectrum of X , f_X , measures the variance in X coming from fluctuations at each frequency. It also represents an approximation to the eigenvalues of Σ_X .¹⁰

To see why this lemma is useful, consider the vector of fundamentals in the frequency domain, $d = \Lambda' D$. Given that $D \sim N(0, \Sigma_D)$, where Σ_D is Toeplitz, we have

$$\Lambda' D = d \Rightarrow N(0, \text{diag}(f_D)). \quad (15)$$

Λ thus approximately diagonalizes the matrix Σ_D , meaning that the elements of d – the fluctuations in fundamentals at different frequencies (with both sines and cosines) are jointly asymptotically independent. Moreover, the same matrix Λ asymptotically diagonalizes the covariance matrix of *any* stationary process. That result will allow us to massively simplify the study of investment over

⁸Two technical points may be noted here. First, as a technical matter, the spectrum f_X must be extended as T grows. A simple way to do that is to suppose that there is a true process for X with a spectrum that is a continuous function f_X , and in any finite sample of length T , there is then an associated spectrum $f_{X,T}$ defined in (11). The second point is that the constant c_X is then a function of that true spectrum f_X ; the appendix elaborates on that fact.

⁹For all the stationary processes studied in the paper, we assume that the autocovariances are summable in the sense that $\sum_{r=1}^{\infty} |j\sigma_{X,j}|$ is finite (which holds for finite-order stationary ARMA processes, for example).

¹⁰ f_X represents an approximation to the eigenvalues only in the sense that $\Lambda' \Sigma_X \Lambda \approx \text{diag}(f_X)$. Providing a sense in which f_X is actually close to the true eigenvalues of Σ_X is a subtler problem that we do not address here. The specific result in lemma 1 is all that we actually need for our results.

many horizons. It says that set of orthogonal factors underlying all stationary processes is (nearly) the same.

1.6 Market equilibrium in the frequency domain

The approximate diagonalization induced by Λ allows us to solve the model through a series of parallel scalar problems that can be easily solved by hand. Using the asymptotic approximation that d and z are independent across frequencies (and across sines and cosines), we obtain the following frequency-by-frequency solution to the asset market equilibrium.¹¹

Solution 1 *Under the approximations $d \sim N(0, \text{diag}(f_D))$ and $z \sim N(0, \text{diag}(f_Z))$, the prices of the frequency-specific portfolios, p_j , satisfy, for all j, j'*

$$p_j = a_{1,j}d_j - a_{2,j}z_j \quad (16)$$

$$a_{1,j} \equiv 1 - \frac{f_{D,j}^{-1}}{\left(\rho f_{avg,j}^{-1}\right)^2 f_{Z,j}^{-1} + f_{avg,j}^{-1} + f_{D,j}^{-1}} \quad (17)$$

$$a_{2,j} \equiv \frac{a_{1,j}}{\rho f_{avg,j}^{-1}} \quad (18)$$

$$\text{where } f_{avg,j}^{-1} \equiv \int_i f_{i,j}^{-1} di \quad (19)$$

where p_j , d_j , and z_j represent the frequency- j components of prices, fundamentals, and supply, respectively. f_i is the spectrum of the matrix Σ_i . See appendix C for the derivation.

The price of the frequency- j portfolio depends only on fundamentals and supply at that frequency. As usual, the informativeness of prices, through $a_{1,j}$, is increasing in the precision of the signals that investors obtain, while the impact of supply on prices is decreasing in signal precision and risk tolerance. The frequency domain analog to the usual demand function is

$$q_{i,j} = \rho \frac{E[d_j - p_j \mid y_{i,j}, p_j]}{\text{Var}[d_j - p_j \mid y_{i,j}, p_j]} \quad (20)$$

These solutions for the prices and demands are the standard results for scalar markets. What is novel here is that the choice problem refers to trades over time. p_j is the price of a portfolio whose exposure to fundamentals fluctuates over time at frequency $2\pi j/T$. Both prices and demands at frequency j depend only on signals and supply at frequency j – the problem is completely separable across frequencies.

The appendix shows that the frequency domain solution provides a close approximation to the true solution in the time domain. Specifically, the true time domain solution from Admati (1985)

¹¹A simple way to see where this solution comes from is to note that, under the asymptotic approximation, $\Lambda' A_1 \Lambda$ and $\Lambda' A_2 \Lambda$ from the Admati solution can be written purely in terms of diagonal matrices, for which addition, multiplication, and inversion are simply scalar operations on the main diagonal.

(with no approximations) can be written as

$$P = A_1 D - A_2 Z \tag{21}$$

for a pair of matrices A_1 and A_2 defined in the appendix that are complicated matrix functions of Σ_Z , Σ_D , and the precisions of the signals agents obtain.

Proposition 1 *The difference between calculating the prices directly in the time domain using the Admati (1985) solution in the time domain and rotating the frequency domain solution back into the time domain is small in the sense that*

$$|A_1 - \Lambda \text{diag}(a_1) \Lambda'| \leq c_1 T^{-1/2} \tag{22}$$

$$|A_2 - \Lambda \text{diag}(a_2) \Lambda'| \leq c_2 T^{-1/2} \tag{23}$$

for constants c_1 and c_2 . Furthermore, while prices and demands are stochastic, the time- and frequency-domain solutions are related through an even stronger result

$$e_{\max} [\text{Var}(\Lambda p - P)] \leq c_P T^{-1/2} \tag{24}$$

$$e_{\max} [\text{Var}(\Lambda q_i - Q_i)] \leq c_Q T^{-1/2} \tag{25}$$

where the operator $e_{\max}[\cdot]$ denotes the maximum eigenvalue of a matrix (that is, the operator norm), for constants c_P and c_Q .

In other words, among portfolios whose squared weights sum to 1, the maximum variance of the pricing and demand errors – the difference between the truth from the time domain solution and the frequency-domain approximation that assumes that Λ diagonalizes the covariance matrices – is of order $T^{-1/2}$ (that is, the bound holds for *any* portfolio of futures, not just the frequency- or time-domain claims). We note also that these are not limiting results – they are true for all T .

Result 1 shows that for large T , the standard time-domain solution for stationary time series processes becomes arbitrarily close to a simple set of parallel scalar problems in the frequency domain. The time domain solution is obtained from the frequency domain solution by premultiplying by Λ .

2 Optimal information choice

We now model a constraint on information acquisition and characterize optimal strategies. The objective, constraint, and solution are drawn from van Nieuwerburgh and Veldkamp (2009) and Kacperczyk, van Nieuwerburgh, and Veldkamp (2016; KvNV). Our analysis follows theirs closely, except that we are studying a time-series model and a frequency transformation. Whereas KvNV study a symmetric equilibrium in which all investors follow the same information acquisition strategy, we will subsequently argue for the relevance of a separating equilibrium in our setting.

2.1 Objective

Following KvNV, we assume that investors choose information to maximize the expectation of their mean-variance objective (6) subject to a linear constraint on total precision:

$$\max_{\{f_{i,j}\}} E_{-1} [U_{i,0} | \Sigma_i^{-1}] \text{ such that } T^{-1} \text{tr} (\Sigma_i^{-1}) \leq \bar{f}^{-1} \quad (26)$$

where E_{-1} is the expectation operator on date -1 , i.e. prior to the realization of signals and prices (as distinguished from $E_{i,0}$, which conditions on P and Y_i). The trace function $\text{tr} (\Sigma_i^{-1})$ measures the total cost of acquiring a private signal with precision matrix Σ_i^{-1} and is subject to the bound \bar{f}^{-1} .¹² This cost function is also equal to the sum of the eigenvalues of the precision matrix. Since the eigenvalues represent the precisions of the orthogonalized signals, it can be thought of as measuring the total precision of the independent parts of the signals. Moreover, since the trace operator is invariant under rotations, this measure of information is invariant to the domain of analysis, time or frequency.¹³ That is,

$$T^{-1} \text{tr} (\Sigma_i^{-1}) = T^{-1} \sum_{j,j'} f_{i,j}^{-1} \quad (27)$$

The information constraint is linear in the frequency-specific precisions. Investors also face the constraint that $f_{i,j} = f_{i,j'}$, which ensures that the variance matrix of ε_i is symmetric and Toeplitz.¹⁴

The appendix shows that, given the optimal demands, an agent's expected utility is linear in the precision they obtain at each frequency.

Lemma 2 *Under the frequency domain representation, when informed investors optimize, each investor's expected utility may be written as a function of their own precisions, $f_{i,j}^{-1}$, and the average across other investors, $f_{avg,j}^{-1} \equiv \int_i f_{i,j}^{-1} di$, with*

$$E_{-1} [U_{0,i} | \{f_{i,j}\}] = \frac{1}{2} T^{-1} \sum_{j,j'} \lambda_j \left(f_{avg,j}^{-1} \right) f_{i,j}^{-1} + \text{constants} \quad (28)$$

$\lambda_j(x)$ is a function determining the marginal benefit of information at each frequency with the properties $\lambda_j(x) > 0$ and $\lambda_j'(x) < 0$ for all $x \geq 0$. The fact that $\lambda_j' < 0$ says that the marginal benefit to an investor of allocating attention to frequency j is decreasing in the amount of attention that other investors allocate to that frequency – attention decisions are strategic substitutes. If

¹²Our main analysis considers the case where signals about fundamentals are costly but investors can condition on prices freely. Appendix I considers a case where it is costly to condition expectations on prices and shows that model's predictions results go through similarly with the caveat that investors never choose to become informed about prices, as in Kacperczyk, van Nieuwerburgh, and Veldkamp (2016).

¹³This result relies on the approximation $\Sigma_i \approx \Lambda' \text{diag}(f_i) \Lambda$.

¹⁴KvNV show that the solution of the optimal attention allocation problem (26) are identical if one assumes that the cost of information is measured by the entropy of the investor's signals, which corresponds to the function $\ln |\Sigma_i^{-1}| \approx \sum_{j,j'} \log f_{i,j}^{-1}$. The key feature of the two cost functions is that they are non-convex in precision.

$f_{avg,j}^{-1}$, the average precision of the signals obtained by other agents, is high, then prices are already efficient at frequency j , so there is little benefit to an investor from learning about that frequency.

The frequency-domain transformation is what allows us to write utility as a simple sum across frequencies. An investor's utility depends additively on the amount of information that they obtain at each frequency. In the time domain, utility is a complicated function of matrices.

2.2 Characterizing the optimum

The critical feature of (28) is that expected utility is linear in the set of precisions that agent i chooses, $\{f_{i,j}^{-1}\}$. Since the both the objective (28) and the constraint (27) are linear in the choice variables, it immediately follows that agents either allocate all attention to a single frequency, or that they are indifferent between allocating attention across some subset of the frequencies. We then obtain the following solution for attention allocation.

Solution 2 *Information is allocated so that*

$$f_{avg,j}^{-1} = \begin{cases} \lambda_j^{-1}(\bar{\lambda}) & \text{if } \lambda_j(0) \geq \bar{\lambda} \\ 0 & \text{otherwise} \end{cases} \quad (29)$$

where $\bar{\lambda}$ is obtained as the solution to

$$T^{-1} \sum_{j,j':\lambda_j^{-1}(\bar{\lambda})>0} \lambda_j^{-1}(\bar{\lambda}) = \bar{f}^{-1}. \quad (30)$$

This is the reverse water-filling solution from KvNV. While it may appear mathematically complicated, the intuition is simple: investors allocate attention to signals in such a way that the marginal benefit is equalized to the extent possible across frequencies. It is impossible to allocate negative attention, though, so if the marginal benefit of paying attention to a particular frequency, $\lambda_j(0)$, is below the cutoff $\bar{\lambda}$, then $f_{i,j}^{-1} = 0$ there for all investors.

The intuition is easiest to develop graphically. Figure 1 plots the functions $\lambda_j(0)$ and $\lambda_j(f_{avg,j}^{-1})$ across frequencies ω_j , where

$$\lambda_j(0) = f_{D,j} (1 + \rho^{-2} f_{D,j} f_{Z,j}). \quad (31)$$

The initial marginal benefit of allocating attention is increasing in the amount of fundamental information and the volatility of supply.

The details of the calibration are reported in appendix J. What is important here is simply that $\lambda_j(0)$ has peaks at low, middle, and high frequencies. Those are the frequencies at which D_t or Z_t is more volatile, so there is more information to potentially be gathered and a larger reward for doing so. For a given value of \bar{f}^{-1} , $\lambda_j(f_{avg,j}^{-1})$ is a flat line for all j such that $\lambda_j(0) \geq \bar{\lambda}$. Those are the frequencies that investors learn about. The term ‘‘reverse water filling’’ refers here to the idea that the curve $\lambda_j(0)$ is inverted and one pours water into it. $\bar{\lambda}$ is then the level of the water's

surface.¹⁵ As the information constraint is relaxed, $\bar{\lambda}$ falls and potentially more frequencies receive attention.

Given the calibration, we see that there are investors acquiring information in three disconnected ranges of frequencies. At the places where $\lambda_j(0)$ is farther above $\bar{\lambda}$, there is more information acquisition, whereas the locations where $\lambda_j(0) = \bar{\lambda}$ are marginal in the sense that they are the next to receive attention if $\bar{\lambda}$ falls.

Another way to interpret the results is to observe the following

Result 1 *The return at frequency j has variance*

$$\text{Var}[r_j] = \lambda_j \left(f_{avg,j}^{-1} \right) \tag{32}$$

$$\text{where } r_j \equiv d_j - p_j. \tag{33}$$

The marginal benefit of acquiring information at a particular frequency is exactly equal to the unconditional variance of returns at that frequency. When returns have high variance, there are potentially large profits to be earned from acquiring information. When returns have zero variance, on the other hand, prices are already perfectly informative, so there is no reason to study fundamentals at such a frequency. So agents desire to learn at the frequencies where returns are most volatile.

The solution derived here characterizes aggregate information acquisition – the sum of the precisions obtained by all the agents at each frequency – but it does not describe exactly what strategy each agent follows; and in fact there are infinitely many strategies for individual investors consistent with the aggregate solution. We now examine one particular solution that leads to the existence of traders who can be characterized by their trading frequencies.

3 Specialization

3.1 The separating equilibrium

Given the assumptions we have made so far, the only restrictions on information allocation are those that ensure that the information allocation condition (29) holds. There are numerous equilibria with that characteristic, though. KvNV focus on the symmetric equilibrium in which all investors allocate their attention in proportion to $f_{avg,j}^{-1}$ at each frequency. There are also many asymmetric and mixed-strategy equilibria.

Since one of our goals is to understand the potential existence and behavior of high-and low-frequency traders, we now focus on equilibria in which all investors learn about only a single frequency. Specifically, we assume that for every agent i , there is a frequency j_i^* such that $f_{i,j_i^*}^{-1} =$

¹⁵Again, each frequency (except 0 and π) has an associated sine and cosine. The same amount of precision is required to be allocated to both the sine and cosine at each frequency.

$\bar{f}^{-1}/2$, and $f_{i,j}^{-1} = 0$ for all other j

$$f_{i,j} = f_{i,j'} = \begin{cases} \bar{f}^{-1}/2 & \text{if } j = j_i^* \text{ and } j_i^* \notin \{0, T/2\} \\ \bar{f}^{-1} & \text{if } j = j_i^* \text{ and } j_i^* \in \{0, T/2\} \\ 0 & \text{otherwise} \end{cases} \quad (34)$$

(\bar{f}^{-1} is divided by 2 in the first case because the agent pays attention to both the sine and the cosine at frequency j_i^*). Specialization here means that agents obtain information at a single frequency and are uninformed at all other frequencies.

We offer two potential explanations for why specialization would be natural. First, it could be the case that people must pay a fixed cost for each frequency that they learn about. That is, just starting to learn about some aspect of fundamentals might be costly. In that case people would naturally choose to learn about only a single frequency, since all frequencies pay the same marginal benefit, so learning about more than one frequency requires an extra payment of the fixed cost with zero benefit.

Another motivation for specialization is that people might simply have different natural aptitudes or desires for learning about fluctuations at different frequencies. The appendix develops a simple example of such a case that can generate specialization. The basic idea is that if the preference for learning about particular frequencies is sufficiently small (i.e. it is second-order or in a sense lexicographic) then the equilibrium described in the previous section still holds, but with each investor focusing their attention on only a single frequency.

Now obviously in reality nobody learns about just a single aspect of the world. It is also not the case, though, that everybody learns about everything. We focus here on the case with specialization as it is consistent with the evidence discussed above and with new results presented below on the wide divergences in behavior and research across investors.

3.2 Specialization model predictions

We now examine the implications of the model with specialization for the behavior of individual investors, obtaining the following results:

1. Investors can be distinguished by the frequencies at which their portfolio positions fluctuate, and those fluctuations match the frequencies at which they obtain information.
2. The average volume accounted for by an investor is proportional to the frequency at which they trade. In the presence of quadratic trading costs, costs can be linearly decomposed across frequencies and are quadratic in frequency.
3. Investors' positions are correlated with returns most strongly at the frequency they learn about.
4. Investors earn money from liquidity provision, they earn money from trading at the frequency at which they are informed, and they *lose* money to other investors from trading at frequencies at which they are uninformed.

3.2.1 Fluctuations in positions

Result 2 *Investor i 's demand at frequency j is*

$$q_{i,j} = z_j + \rho \left[\left(f_{i,j}^{-1} - f_{avg,j}^{-1} \right) r_j + f_{i,j}^{-1} \tilde{\varepsilon}_{i,j} \right] \quad (35)$$

where $\tilde{\varepsilon}_{i,j}$ is equal to the j th column of Λ multiplied by ε_i , i.e. the noise in investor i 's signal at frequency j , and r_j is the realized return on the j th frequency portfolio.

Investor i 's demand depends on three terms. z_j is the stochastic supply at frequency j . Each investor is equally willing to absorb supply, so they all take equal fractions, giving them a common component z_j .

The second term, $\rho \left(f_{i,j}^{-1} - f_{avg,j}^{-1} \right) r_j$ reflects investor i 's information. At the frequency that investor i pays attention to, $f_{i,j}^{-1} - f_{avg,j}^{-1}$ is positive, so investor i 's demand covaries positively with returns at that frequency. That is, investors who learn about low-frequency dynamics hold portfolios that are long when returns are high over long periods, while high-frequency investors hold portfolios that covary positively with transitory fluctuations in returns. At the other frequencies, where investor i does not pay attention, $f_{i,j}^{-1} = 0$, so the investor's demand actually covaries slightly negatively with returns, holding z_j fixed.

The third term, $\rho f_{i,j}^{-1} \tilde{\varepsilon}_{i,j}$ is the idiosyncratic part of demand that is due to the random error in the signal that agent i receives. Note that the standard deviation of $f_{i,j}^{-1} \tilde{\varepsilon}_{i,j}$ is equal to $f_{i,j}^{-1/2}$, so these errors are equal to zero at the frequencies that the investor ignores (i.e. all but one).

When the number of active frequencies (i.e. with $f_{avg,j}^{-1} > 0$) is large, $f_{avg,j}^{-1}$ becomes small relative to \bar{f}^{-1} . That means that the term $\left(f_{i,j}^{-1} - f_{avg,j}^{-1} \right)$ is close to zero at all frequencies except for the one that the agent pays attention to, j_i^* . Since $\left(f_{i,j}^{-1} - f_{avg,j}^{-1} \right) \approx 0$ for all other frequencies, we have

$$Q_{i,t} \approx Z_t + \cos(\omega_{j_i^*} t/T) (\bar{f}^{-1} r_{j_i^*} + \tilde{\varepsilon}_j) + \sin(\omega_{j_i^*} t/T) (\bar{f}^{-1} r_{j_i^{*'}} + \tilde{\varepsilon}_{j'}) \quad (36)$$

Investor i 's demand on date t thus is approximately equal to supply on that date plus a multiple the part of returns depending on frequency $\omega_{j_i^*}$, $r_{j_i^*}$ and $r_{j_i^{*'}}$, plus an error. The second line shows that what is really going on is that investor i 's information can be thought of as a signal about returns interacted with a cosine and a sine.

The important feature of equations (35) and (36) is that they show that each agent's position is equal to Z_t plus fluctuations that come primarily at the frequency that they pay attention to.¹⁶ That is, if some agent allocates all attention to frequency $\omega_{j_i^*}$, then their relative position, $Q_{i,t} - Z_t$,

¹⁶More formally, the variance of $Q_{i,t} - Z_t$ can be decomposed as $Var(Q_{i,t} - Z_t) = \sum_j \left(\rho^2 (f_{i,j}^{-1} - f_{avg,j}^{-1})^2 f_{R,j} + f_{i,j}^{-1} \right)$. Now consider a simple case where there are N frequencies that receive equal allocations of information. Furthermore, denote the spectrum of returns as $f_{R,j}$. Then we have

$$\lim_{N \rightarrow \infty} Var(Q_{i,t} - Z_t) = \rho^2 \bar{f}^{-2} f_{R,j_i^*} + f_{i,j}^{-1} \quad (37)$$

which shows that $Q_{i,t} - Z_t$ is driven by fluctuations at a single frequency.

fluctuates over time at frequency $\omega_{j_i^*}$. This can be seen by noting that the sum of a sine and a cosine at frequency $\omega_{j_i^*}$, even with different coefficients, remains a cosine that fluctuates at frequency $\omega_{j_i^*}$, just shifted by a constant. Specifically,

$$Q_{i,t} \approx Z_t + \sqrt{\frac{(\bar{f}^{-1}r_{j_i^*} + \tilde{\varepsilon}_{j_i^*})^2}{+(\bar{f}^{-1}r_{j_i^{*'}} + \tilde{\varepsilon}_{j_i^{*'}})^2}} \cos(\omega_{j_i^*}t/T + C_{i,j}) \quad (38)$$

where $C_{i,j}$ is a function of $(\bar{f}^{-1}r_j + \tilde{\varepsilon}_j)$ and $(\bar{f}^{-1}r_{j'} + \tilde{\varepsilon}_{j'})$. So agent i 's excess demand is approximately a cosine with a random translation and amplitude.

As a numerical example, figure 2 plots a hypothetical history for a particular agent's position $Q_{i,t}$ in the same calibration that we studied above. We see that $Q_{i,t}$ looks like a sinusoid with noise added; the noise is from the Z_t term in (38). The noise in the agent's signal, $\tilde{\varepsilon}_{j_i^*}$ and $\tilde{\varepsilon}_{j_i^{*'}}$, simply changes the amplitude and translation of the cosine in (38).

So equations (35) and (36) deliver our first two basic results for the behavior of individual specialized investors: the investors can be distinguished by the frequencies at which their asset holdings fluctuate, and those frequencies are linked to the type of information that they acquire. The first result, that there are traders at different frequencies, is essentially obtained by design: it follows from the assumption that agents specialize across frequencies. Nevertheless, the finding is interesting for its novelty in a theoretical setting.

The fact that the frequency of trading is related to information acquisition, while not surprising, is certainly not obtained by assumption. In past work, different trading behavior has sometimes been obtained by simply assuming that different agents have different exogenously specified trading horizons. In our case, any investor can potentially trade at any frequency. That choice is entirely endogenous – investors are not forced to trade any particular frequency by assumption (the assumption is that they gather information at a single frequency).

The reason that buy-and-hold investors in our model buy and hold is that they have persistent low-frequency information about fundamentals – they have signals that fundamentals will be strong or weak over long time spans. Similarly, high-frequency investors have transitory high-frequency information. So the model provides a testable prediction that we should observe investors doing research about asset return dynamics that aligns in terms of frequency or time horizon with their average holding periods.

3.2.2 How do investors earn money?

Investors earn returns in the model through two basic mechanisms: providing liquidity and trading on private signals. We can see from the results on demand above that the liquidity function is spread equally across investors. The effects of private information are more interesting.

Result 3 *Investor i 's expected profits (which are also equal to the covariance of their positions with*

returns) are

$$E [Q'_i R] = \sum_j E [q_{i,j} r_j] \quad (39)$$

$$= \sum_j E [z_j r_j] + \rho \left(\bar{f}^{-1} - f_{avg,j_i^*}^{-1} \right) f_{R,j_i^*} - \sum_{j \neq j_i^*} \rho f_{avg,j}^{-1} f_{R,j} \quad (40)$$

where the spectrum of returns is (from (32))

$$f_{R,j} = \max(\bar{\lambda}, \lambda_j(0)) \quad (41)$$

The first term on the right-hand side is the contribution from each investor's liquidity provision. The second term is the positive covariance of the investor's holdings with returns at the frequency they are informed about. Informed investors have demands that covary positively with returns at a particular frequency, then the investors who are uninformed about that frequency must have demands that covary *negatively* with returns (after accounting for $E [z_j r_j]$). That is the third term above: there is a negative contribution to the correlation of the investor's demands with returns from the frequencies they do not pay attention to.

It is not the case that trading from frequencies $j \neq j_i^*$ is unprofitable. Investors still earn profits from liquidity provision. It is just the case that some of their profits at those frequencies are taken by investors who are more informed. In some sense, this result is inevitable. The total profits that the informed investors earn as a group come from trading with liquidity demand. If an investor earns more money by becoming informed at some frequency, that must come at the cost of other investors.

Now since the information allocation we find is an optimum, obviously investors must be in some sense comfortable with the losses we see here. Intuitively, the slight trading losses they bear at frequencies other than j_i^* are offset by their gains at j_i^* . But obviously any trading that informed investors do that is not related to exogenous supply must ultimately come at the cost of other informed investors.

So the model has the feature that high-frequency investors earn money at high frequencies, but they lose money at lower frequencies relative to other investors. Low-frequency investors might know that oil prices are on a long-term downward trend. In such a situation, the high-frequency investors can still earn profits by betting on day-to-day movements in oil prices, but they will lose money to those who understand that prices are generally drifting down. Similarly, low-frequency investors will tend to lose out at high frequencies by, for example, failing to trade at precisely the right time, buying slightly too high and selling slightly too low compared to where they would if they had high-frequency information.

3.2.3 Volume and trading costs

We study volume in the representation of the model in terms of equity holdings. Recall that equity is modeled as a discounted claim to dividends on all future dates. An investor's position $Q_{i,t}$ can be acquired either by holding $Q_{i,t}$ units of forwards or $Q_{i,t}$ units of equity. In modeling volume, we consider trading in equity. Using equity to measure volume ensures that a person who has position that does not change between dates t and $t+1$ ($Q_{i,t} = Q_{i,t+1}$) induces no trade volume, whereas if we assumed that every forward position required volume, then each investor's contribution would be $|Q_{i,t}|$ each date, meaning, unrealistically, that buy-and-hold investors would contribute constantly to volume.

The equity volume contributed by investor i is

$$V_{i,t} = |\Delta Q_{i,t}| \quad (42)$$

$$\text{where } \Delta Q_{i,t} \equiv Q_{i,t} - Q_{i,t-1} \quad (43)$$

Recall that investors' positions can be written as functions of cosines and sines. The appendix derives the following result for volume for each investor.

Result 4 *The volume induced by investor i , $|\Delta Q_{i,t}|$, may be approximated as*

$$|\Delta Q_{i,t}| \approx |\Delta Z_t| + \omega_{j_i^*} \bar{f}^{-1} \rho \left| \begin{array}{l} \sin(\omega_{j_i^*} t) (r_{j_i^*} + \tilde{\varepsilon}_{i,j_i^*}) \\ + \cos(\omega_{j_i^*} t) (r_{j_i^{*'}} + \tilde{\varepsilon}_{i,j_i^{*'}}) \end{array} \right| \quad (44)$$

and has expectation

$$E[|\Delta Q_{i,t}|] - E[|\Delta Z_t|] \approx \omega_{j_i^*} \sqrt{\frac{2}{\pi}} \rho (\bar{f}^{-1} \bar{\lambda} + 1/2). \quad (45)$$

The approximations converge to true equalities as $T \rightarrow \infty$.

So we find that agent i 's contribution to volume depends on the volume induced by exogenous supply and also the magnitude of returns at frequency $\omega_{j_i^*}$ ($\bar{\lambda}$).

Agent i 's contribution to aggregate volume is also exactly proportional to the frequency they allocate attention to, $\omega_{j_i^*}$. High-frequency investors contribute relatively more to aggregate volume because they have portfolios that change most rapidly. An investor at the very lowest frequency, $\omega = 0$, contributes zero volume beyond that induced by exogenous supply, since their position is nearly constant. Investors at $\omega = \pi$, on the other hand, contribute the maximum possible volume as they approximately turn over their entire portfolios in each period.

Not surprisingly, it is also straightforward to show that high-frequency investors typically will face larger trading costs. As an example, consider quadratic trading costs proportional to $\sum_{t=2}^T (Q_{i,t} - Q_{i,t-1})^2$. The appendix shows that trading costs can, just like volume, be decomposed across frequencies.

Result 5 *The quadratic variation in an investor’s position can be approximated (with convergence as $T \rightarrow \infty$) by*

$$\sum_{t=2}^T (Q_{i,t} - Q_{i,t-1})^2 = \sum_{j,j'} (2\pi j)^2 T^{-1} q_{i,j}^2. \quad (46)$$

The quadratic trading costs associated with a given demand vector Q_i can be written as a simple sum across frequencies. Trading costs are proportional to the frequency squared. It is thus immediately apparent from our frequency-domain analysis that changes in the magnitude of trading costs have the largest effects on the highest frequencies.

4 Institutional portfolio turnover and return forecasting

In this section, we demonstrate empirically that investment funds specialize in the frequency at which they trade, and we show that the portfolio holdings of high turnover funds forecast returns at relatively shorter horizons than those of lower turnover funds.

4.1 Data

We obtain data on institutional asset holdings from SEC form 13F. These forms list the identities and quantities of securities held by each institution at the end of the filing quarter.¹⁷ The data cover the period 1980–2015. Data on monthly stock returns is taken from CRSP and is aggregated to a quarterly frequency with delisting returns included. We obtain data on the risk-free rate, market return, and Fama–French (1993) factors from Kenneth French’s website.

4.2 Fund specialization

Yan and Zhang (2009) define the churn rate $c_{i,t}$ of institution i in quarter t as

$$c_{i,t} \equiv \frac{\min \left(\sum_{s|\Delta S_{i,s,t} > 0} P_{s,t} \Delta S_{i,s,t}, \sum_{s|\Delta S_{i,s,t} \leq 0} P_{s,t} |\Delta S_{i,s,t}| \right)}{\frac{1}{2} \sum_s P_{s,t-1} S_{i,s,t-1} + \frac{1}{2} \sum_s P_{s,t} S_{i,s,t}}, \quad (47)$$

where $P_{s,t}$ is the price and $S_{i,s,t}$ is the number of shares of stock S held by institution i at the end of quarter t . The churn rate is equal to the minimum of net purchases and sales during quarter t as a fraction of the institution’s average value over the two quarters, and it is used to measure the turnover of each institution’s portfolio. Due to the presence of the minimum operator, institutions

¹⁷Institutions are required to report only their holdings of 13(f) securities, a category defined by the SEC that includes exchange-traded equities and some securities that can be converted to equity. Only institutions holding more than \$100,000,000 in 13(f) securities at the end of the quarter must file form 13F, and each institution is required to report only securities for which its holdings exceed \$200,000 or 10,000 shares. Gompers and Metrick (2001) provide more information on these filings. We use Thompson Reuters’s database of these filings, which includes the price of each security at the filing date.

must both buy and sell large fractions of their portfolios to register a high churn rate. The mean churn rate is 0.12 and the standard deviation is 0.14, indicating a high degree of right skewness as the minimum churn rate is zero.

If institutions specialize in the rate at which they trade, then the churn rate should be persistent over time within institutions. Figure 5 plots the sample autocorrelations $\text{corr}(c_{i,t}, c_{i,t-\Delta t})$ for $\Delta t \geq 2$. The churn rate strongly persists over time, with an autocorrelation of 0.51 over 10 years and 0.21 over 30 years.¹⁸ We also find that institution fixed effects (δ_i) account for 65 percent of the variance in the churn rate in the regression $c_{i,t} = \delta_i + \varepsilon_{i,t}$, where $\varepsilon_{i,t}$ a residual.

4.3 Fund performance

To separate institutions according to their trading frequency, at each quarter t we divide institutions into deciles, denoted d , based on the mean of $c_{i,t}$ over the previous five years.¹⁹ For simplicity, we restrict attention to top and bottom deciles ($d = 10$ and $d = 1$). Table 1 lists several institutions in these extreme deciles during the most recent quarter in our data. The top decile contains several well-known quantitative and high-frequency trading firms, whereas the bottom contains endowments and insurance companies.

Table 1: Institutions in the top and bottom deciles of churn rate in the fourth quarter of 2015

Top decile	Bottom decile
Arrowstreet Capital	Berkshire Hathaway
Citadel	Bill & Melinda Gates Foundation
Dynamic Capital Management	Lilly Endowment
Ellington Management Group	Longview Asset Management
Quantlab	MetLife
Renaissance Technologies	New York State Teachers' Retirement System
Soros Fund Management	University of Notre Dame
Virtu Financial	University of Chicago

At the beginning of quarter t , we average the portfolio holdings of all the funds in each decile d at the end of quarter $t - 1$ (with equal weight on each fund) and then track the returns on that aggregate decile-level portfolio over subsequent quarters, reinvesting proceeds from any delistings in the remaining stocks in the portfolio according to their value weights at that time. The return during quarter t of the decile d portfolio formed in quarter $t - k$ is denoted $r_{d,k,t}$.

We measure the performance of each portfolio by its alpha,

$$r_{d,k,t} - r_t^f = \alpha_{d,k} + \beta_{d,k} F_t + \varepsilon_{d,k,t}, \quad (48)$$

¹⁸Institution identifiers can be reassigned over time in the 13F data, leading to measurement error that biases the longer-term autocorrelations towards zero.

¹⁹We restrict attention to institutions for which the $t + 1$ return on some of their holdings appears in CRSP, as these are the only institutions that can be analyzed in our return regressions.

F_t is a vector of market risk factors; $F_t = r_t^m - r_t^f$ in the CAPM specification and $F_t = (r_t^m - r_t^f \quad r_t^{smb} \quad r_t^{hml})'$ in the Fama-French specification (r_t^{smb} and r_t^{hml} are returns on the SMB and HML portfolios, respectively). We focus on returns over the first two years after portfolio formation by estimating (48) only for $k \leq 8$.

Figure 6 displays alphas in the two specifications. For simplicity, we compare the alphas in the first quarter ($\alpha_{d,1}$) to those in the following seven quarters, $\alpha_{d,>1} \equiv \frac{1}{7} \sum_{j=2}^8 \alpha_{d,j}$. The holdings of high-turnover funds out-perform more during the first quarter, while those of low-turnover funds out-perform more during subsequent quarters. The difference in differences $(\alpha_{10,1} - \alpha_{10,>1}) - (\alpha_{1,1} - \alpha_{1,>1})$, which measures the relative out-performance of high churn holdings versus low churn holdings at short versus long horizons, is equal to 0.005 in both specifications and is significant at the 5 percent level.²⁰ So, consistent with the model, high-turnover funds hold stocks that outperform relatively more in the short-run, while low-turnover funds hold assets that display more persistent outperformance.

5 The effects of eliminating high-frequency trade

Recently there has been interest in policies that might restrict high-frequency trading. Some of those policies are aimed at investors who trade at the very highest frequencies (such as the CFTC's recently proposed Regulation AT; see CFTC (2016)). But there are also proposals to discourage even portfolio turnover at the monthly or annual level.²¹ There are two ways to interpret such policies. One would be that regulators might impose a tax on trading, which would simply represent a transaction cost. Such a regulation would obviously have the strongest effects on high-frequency traders (that result can be derived in an extension to the present setting with trading costs), but it would ultimately affect all trading strategies. A more targeted policy would be one that simply outlawed following a trading strategy in which positions fluctuate at frequencies above some bound. Such a policy is straightforward to analyze in our framework.

We show in this section that a policy that restricts high-frequency trading by professional investors (as opposed to liquidity traders) reduces liquidity and price informativeness and increases return volatility at high frequencies. At the frequencies not targeted by the policy, though, price informativeness is, if anything, increased.

5.1 The policy

If elimination of high-frequency trading means that investors cannot hold portfolios with components that fluctuate rapidly, it means that those investors are restricted to setting $q_{i,j} = 0$ for ω_j in

²⁰Yan and Zhang (2009) similarly find that the fraction of the outstanding shares of a stock held by high-churn institutions predicts subsequent returns, while the fraction held by low-churn institutions does not.

²¹The US tax code, for example, encourages holding assets for at least a year through the higher tax rates on short-term capital gains. There have been recent proposals to further expand such policies (a plan to create a schedule of capital gains tax rates that declines over a period of six years was attributed to Hillary Clinton during the 2016 US Presidential election; see Auxier et al. (2016)).

the relevant frequency range. But when $q_{i,j}$ must equal zero at some set of frequencies, obviously no trader will allocate attention to those frequencies.

In cases where the sophisticated investors must set $q_j = 0$, there can obviously be no equilibrium since liquidity demand is perfectly inelastic. We therefore consider a simple extension to the baseline model where the exogenous supply curve is elastic,

$$Z_t = \tilde{Z}_t + kP_t \quad (49)$$

$$z_j = \tilde{z}_j + kp_j \quad (50)$$

where \tilde{Z}_t is the exogenous supply process and k is the response of supply to prices.

The appendix solves this extended version of the model. We obtain the same water filling equilibrium. For this section, what is most important is the volatility of returns.

For the frequencies that are restricted

$$\text{for } q_{i,j} = 0: p_j^{\text{restricted}} = -k^{-1}\tilde{z}_j. \quad (51)$$

That is, prices at the restricted frequencies are now completely uninformative, depending only on supply, with no relationship with fundamentals. Moreover, the market is completely illiquid in the sense that when exogenous supply increases, there is no change in trade – prices just move so that trade remains at zero. In other words, prices equilibrate the market instead of quantities.

5.2 Return volatility

Result 6 *Given an information policy $f_{\text{avg},j}^{-1}$, the variance of returns at frequency j , when trade is unrestricted (i.e. in the benchmark model from above), is*

$$f_R(\omega_j) \equiv \text{Var}(r_j) \quad (52)$$

$$= \lambda_j \left(f_{\text{avg},j}^{-1} \right) \quad (53)$$

$$= \min(\bar{\lambda}, \lambda_j(0)). \quad (54)$$

where

$$\lambda_j(0) = f_{D,j} + \frac{f_{\tilde{Z},j}}{\left(k + \rho f_{D,j}^{-1}\right)^2} \quad (55)$$

Recall that f_R is also the spectrum of returns, $R_t \equiv D_t - P_t$.

So the spectrum of returns inherits exactly the water-filling property of the marginal benefits of information. In the context of our benchmark calibration, the spectrum of returns is exactly the $\lambda_j \left(f_{\text{avg},j}^{-1} \right)$ curve plotted in figure 1.

That result does not apply when the sophisticated investors are restricted from trading, though.

Result 7 *The variance of returns at any restricted frequency, where $q_{i,j}$ must equal zero, is*

$$f_{R,j}^{restricted} = f_{D,j} + \frac{f_{\tilde{Z},j}}{k^2} \quad (56)$$

and

$$f_{R,j}^{restricted} > \lambda_j(0). \quad (57)$$

The volatility of returns at a restricted frequency is higher than it would be if the sophisticated investors were allowed to trade, even if they gathered no information. Intuitively, when the active investors have risk-bearing capacity ($\rho > 0$), they can absorb some of the exogenous supply. The greater is the risk-bearing capacity, the smaller is the effect of supply volatility on return volatility.

We examine the quantitative implications of restricting trade in the context of the calibration used above. The top panel of figure 3 plots $f_{avg,j}^{-1}$ in the restricted and unrestricted scenarios. The restriction is that investors are not allowed to trade at frequencies above $\omega = 3$ (cycle lengths shorter than 2.1 periods). We see, then, that no information is acquired at those frequencies. That means, though, that investors can allocate their attention elsewhere, so we observe more information acquisition at other frequencies.

The bottom panel of figure 3 plots return volatility in the two regimes and also when investors can trade at all frequencies, but they are just restricted from gathering information at high frequencies. At high frequencies, we see that the restrictive policy has two separate effects that both strongly affect return volatility. First, when investors can trade but do not gather information, there is a jump in return volatility at the frequencies where $f_{avg,j}^{-1}$ is constrained to zero (up to $\lambda_j(0)$). But under the full restriction, where they cannot trade at all, we see that the effect on return volatility is much larger, due to the reduced risk bearing capacity. At the unrestricted frequencies, return volatility actually weakly declines, again due to the fact that more attention is allocated to those frequencies.

Restricting sophisticated investors (such as dealers or proprietary trading firms) from trading at high frequencies in this model can thus substantially raise asset return volatility at high frequencies – it can lead to, for example, large minute-to-minute fluctuations in prices and returns. Sophisticated traders typically play a role of smoothing prices over time, essentially intermediating between excess inelastic demand in one minute and excess inelastic supply in the next. When they are restricted from holding positions that fluctuate from minute to minute, they can no longer provide that intermediation service. Such behavior has no impact on low-frequency volatility in prices, though. Even when there is no high-frequency trading, changes in average prices between one year and the next are essentially unaffected.

5.3 Price informativeness and efficiency

The fact that the sophisticated investors choose to allocate no attention to high frequencies under the trading restriction obviously has implications for price efficiency there. To see precisely how, we

measure price informativeness as the precision of a person's prediction of fundamentals conditional on observing prices only, $Var(D_t | P)$. In the frequency domain we have

$$\bar{\tau}_j \equiv Var(d_j | p)^{-1} \quad (58)$$

$$= \left(\rho f_{avg,j}^{-1}\right)^2 f_{Z,j}^{-1} + f_{D,j}^{-1} \quad (59)$$

price-based precision, $\bar{\tau}_j$ is higher at frequencies where there is less fundamental uncertainty ($f_{D,j}^{-1}$ is lower), where there is less variation in liquidity demand ($f_{Z,j}^{-1}$ is lower) or where investors acquire more information ($f_{avg,j}^{-1}$ is higher). So when trading strategies are restricted and $f_{avg,j}^{-1}$ endogenously falls to zero at the restricted frequencies, price informativeness clearly falls. In fact, when $f_{avg,j}^{-1} = 0$, the ex-post precision at each frequency is exactly $f_{D,j}^{-1}$, which is the prior precision; prices contain no information. The decline in informativeness happens, though, only at the restricted frequencies.²²

Result 8 *If trade is restricted at some set of frequencies, prices become (weakly) less informative at those frequencies ($\bar{\tau}_j$ falls) but informativeness is unaffected or increased at all other frequencies.*

5.3.1 Informativeness for moving averages of D_t

If a person is making decisions based on estimates of fundamentals from prices and they are worried that prices are contaminated by high-frequency noise, a natural response would be to examine an average of fundamentals and prices over time. For averages of fundamentals, we have the following:

Result 9 *The variance of an estimate of the average of fundamentals over dates t to $t + n - 1$ conditional on observing the vector of prices, P , is*

$$Var\left(\frac{1}{n} \sum_{m=0}^{n-1} D_{t+n} | P\right) = \frac{1}{nT} \sum_{j,j'} F_n(\omega_j) \bar{\tau}_j^{-1} \quad (60)$$

$$\text{where } F_n(\omega_j) \equiv \frac{1}{n} \frac{1 - \cos(n\omega_j)}{1 - \cos(\omega_j)} \quad (61)$$

and $\bar{\tau}_j$ is defined in (58).

F_n is the Fejér kernel. $F_1 = 1$, and as n rises, the mass of the Fejér kernel migrates towards the origin. That is, it places progressively less mass on high frequencies and more on low frequencies (it always integrates to 1). Specifically,

$$\frac{1}{T} \sum_{j,j'} F_n(\omega_j) = 1 \quad (62)$$

$$\lim_{n \rightarrow \infty} F_n(\omega) = 0 \text{ for all } \omega \neq 0 \quad (63)$$

²²To see that result in the time domain, the appendix shows that $Var(D_t | P) = \frac{1}{T} \sum_{j,j'} \bar{\tau}_j^{-1}$. The variance of an estimate of fundamentals conditional on prices at a particular date is equal to the average of the variances across all frequencies. So when uncertainty, $\bar{\tau}_j^{-1}$, rises at some set of frequencies, the informativeness of prices for fundamentals on every date falls by an equal amount.

The total weight allocated across the frequencies always sums to 1, and as n rises, the mass becomes allocated eventually purely to frequencies local to zero.

This result shows that the informativeness of prices for moving averages of fundamentals places relatively more weight on low- than high-frequency informativeness. So even if prices have little or no information at high frequencies – $\bar{\tau}_j$ is small for large j , there need not be any degradation of information about averages of fundamentals over multiple periods, as they depend primarily on precision at lower frequencies (smaller values of j).

The top panel of figure 4 plots the Fejér kernel, F_n , for a range of values of n . One can see that even with $n = 2$, the weight allocated to frequencies above the cutoff of $\omega = 3$ that we use in the example in figure 3 is close to zero. As n rises higher, the weight falls towards zero at a progressively wider range of frequencies. Equation (60) therefore shows that while a reduction in precision at high frequencies due to trading restrictions will reduce the informativeness of prices about fundamentals on any single date, it has quantitatively small effects on the informativeness of prices for fundamentals over longer periods.

Moving averages of fundamentals depend less on the precise high-frequency details of the world, so when high-frequency information is reduced, we would not expect to see a reduction in the informativeness of prices for moving averages. More concretely, going back to our example of oil futures, when high-frequency trade is not allowed, prices become noisier, making it more difficult to obtain an accurate forecast of the spot price of oil at some specific moment in the future. If one is interested in the average of spot oil prices over a year, on the other hand, then we would expect futures prices to remain informative, even when high-frequency trade is restricted.

5.3.2 How should one forecast D_t conditional on prices?

As an alternative to estimating an average of fundamentals over some number of periods, one's goal might alternatively be to specifically forecast fundamentals on just one date. In that case, we see that the effect of high-frequency trade restrictions is to cause investors to focus on averages of prices over multiple periods. That is, to forecast the spot oil price in one particular month, one might use an average of futures prices over neighboring months.

In general, the expectation of the full time series of fundamentals, D , is

$$E[D | P] = \Lambda \text{diag}(\phi_j(f_{avg,j})) \Lambda' P \quad (64)$$

ϕ_j is a function of only $f_{avg,j}$ and the behavior of fundamentals and liquidity demand at frequency j with $\phi_j(0) = 0$. Intuitively, this equation says that dividends are obtained from prices using a filter: prices are transformed into the frequency domain ($\Lambda' P$), a filter is applied that depends on the informativeness of prices at each frequency ($\phi_j(f_{avg,j})$), and then D is obtained by transforming back into the time domain.

The effect of eliminating information – setting $f_{avg,j}^{-1} = 0$ – at high frequencies is then simple to analyze. The frequency domain step sets to zero the weight on any frequencies at which there

is no information. That is, in obtaining the expectation of dividends, one first applies a low-pass filter to prices (e.g. Christiano and Fitzgerald (2003)).

Result 10 *When information acquisition is set to zero for frequencies above a cutoff \bar{j} , so that $f_{avg,j}^{restricted} = f_{avg,j} 1\{j \leq \bar{j}\}$ (where $1\{\cdot\}$ is the indicator function) the expectation of fundamentals conditional on prices is*

$$E[D | P] = \Lambda \text{diag}(\phi_j(f_{avg,j})) \text{diag}(1\{j \leq \bar{j}\}) \Lambda' P \quad (65)$$

Specifically, $E[D_t | P]$ is equal to the t 'th row of $\Lambda \text{diag}(\phi_j(f_{avg,j})) \text{diag}(1\{j \leq \bar{j}\}) \Lambda'$ multiplied by the vector of prices, P .

Under the restriction on information acquisition (which, as we saw above, happens when investors may not trade at frequencies above \bar{j}), expectations are now calculated by first applying a filter to prices that eliminates high-frequency fluctuations (that is, it sets to zero all components of the price vector P that are spanned by high-frequency ($j \geq \bar{j}$) sines and cosines). That filter is implemented by multiplying prices by $\text{diag}(1\{j \leq \bar{j}\}) \Lambda'$, which eliminates the high-frequency components. Intuitively, since those components are completely devoid of information, they should be filtered out before estimating fundamentals.

A simple benchmark case is where the variance of fundamentals and liquidity demand is constant across frequencies, so that $f_{avg,j}$ is also constant across frequencies in the absence of trading restrictions and $\text{diag}(\phi_j(f_{avg,j}))$ is a multiple of the identity matrix. The question, then, is how investors estimate fundamentals on some date t based on the history of prices.

When there is no restriction on prices and $\phi_j(f_{avg,j})$ is constant, we see immediately that $E[D | P] \propto P$, so that the expectation of fundamentals on date t depends only on the price on date t (for any t).

When trade is restricted, we have $E[D | P] \propto \Lambda \text{diag}(1\{j \leq \bar{j}\}) \Lambda' P$. The bottom panel of figure 4 therefore plots a representative row of $\Lambda \text{diag}(1\{j \leq \bar{j}\}) \Lambda'$ for different values of \bar{j} (the interior rows are all highly similar; the boundaries induce some differences). We see that with the trading restriction, the estimate of fundamentals on date t now depends on prices on t and its neighbors. Moreover, the smaller is \bar{j} – the more frequencies that are restricted – the wider is the set of weights applied to prices. Intuitively, then, figure 4 confirms the natural result that when prices are less informative at high frequencies, the simple response is to estimate fundamentals based on a moving average of prices.

5.3.3 Summary

In the end, this section shows that the model has two key predictions for the effects of restrictions on high-frequency trade. First, at the frequencies at which trade is restricted, price informativeness falls and return volatility rises (due to both information effects and liquidity effects). Second, though, price informativeness at low frequencies is, if anything, *improved* by the policy. So if a

manager is making investment decisions based on fundamentals only at a particular moment, then that decision will be hindered by the policy since prices now have more noise. But if decisions are made based on averages of fundamentals over longer periods, e.g. over a year, then the model predicts that there need not be adverse consequences. If anything, low-frequency price informativeness may increase as investors reallocate attention to lower frequencies.

6 Conclusion

This paper develops a model in which there are many different investors who all trade at different frequencies. Investors in real-world markets follow countless strategies that are associated with rates of turnover that differ by multiple orders of magnitude. We show that in fact it is entirely natural that investors would differentiate along the dimension of investment frequency.

It has been standard in the literature for decades to focus on factors or principal components when studying the cross-section of asset returns. For stationary time series, the analog to factors or principal components is the set of fluctuations at different frequencies. So just as it seems natural for investors to focus on particular factors in the cross-section of returns, e.g. value stocks, a particular industry, or a particular commodity, it is also natural for investors to focus on fluctuations in fundamentals at a particular frequency, like quarters, business cycles, or decades.

Such an attention allocation problem can be solved using a combination of standard tools from time series econometrics and the literature on equilibria in financial markets. We show that the model fits a wide range of basic stylized facts about financial markets: investors can be distinguished by turnover rates; trading frequencies align with research frequencies; volume is driven primarily by high-frequency traders; and the positions of informed traders forecast returns at a horizon similar to their holding period.

Since the model has a rigorous concept of what being a high- or low-frequency entails, it is particularly useful for studying the effects of regulatory policies that would restrict trade at certain frequencies, whether by outlawing it or by simply making it more costly. We find that not only do such policies reduce the informativeness of prices at those frequencies, they also reduce liquidity and increase return volatility. In fact, return volatility will in general be raised even above where it would be in the complete absence of information, since eliminating active traders from the market removes their risk-bearing capacity.

At this point, the primary drawback of the model in our view is that it is not fully dynamic. In a certain sense we have to assume that investors do not update information sets over time. While that simplification does not interfere with the model's ability to match a wider range of basic facts about financial markets, a simple desire for realism suggests that incorporating dynamic learning is an obvious next step.

References

- Admati, Anat R.**, “A Noisy Rational Expectations Equilibrium for Multi-Asset Securities Markets,” *Econometrica*, 1985, pp. 629–657.
- Amihud, Yakov and Haim Mendelson**, “Asset Pricing and the Bid-Ask Spread,” *Journal of Financial Economics*, 1986, *17* (2), 223–249.
- Auxier, Richard, Len Burman, Jim Nunns, Ben Page, and Jeff Rohaly**, “An Updated Analysis of Hillary Clinton’s Tax Proposals,” October 2016.
- Bandi, Federico and Andrea Tamoni**, “Business-cycle consumption risk and asset prices,” *Working Paper*, 2014.
- Banerjee, Snehal**, “Learning from Prices and the Dispersion in Beliefs,” *Review of Financial Studies*, 2011, *24* (9), 3025–3068.
- and **Brett Green**, “Signal or noise? Uncertainty and learning about whether other traders are informed,” *Journal of Financial Economics*, 2015, *117* (2), 398–423.
- and **Ilan Kremer**, “Disagreement and learning: Dynamic patterns of trade,” *The Journal of Finance*, 2010, *65* (4), 1269–1302.
- Basak, Suleyman**, “Asset pricing with heterogeneous beliefs,” *Journal of Banking & Finance*, 2005, *29* (11), 2849–2881.
- Bhattacharya, Sudipto and Paul Pfleiderer**, “Delegated Portfolio Management,” *Journal of Economic Theory*, 1985, *36* (1), 1–25.
- Brillinger, David R.**, *Time Series: Data Analysis and Theory*, McGraw Hill, 1981.
- Brockwell, Peter J. and Richard A. Davis**, *Time Series: Theory and Methods*, Springer, 1991.
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan**, “High-Frequency Trading and Price Discovery,” *Review of Financial Studies*, 2014, *27* (8), 2267–2306.
- Carhart, Mark M.**, “On Persistence in Mutual Fund Performance,” *Journal of Finance*, 1997, *52* (1), 57–82.
- Chaudhuri, Shomesh E. and Andrew W. Lo**, “Spectral Portfolio Theory,” 2016. Working paper.
- Chen, Hsiu-Lang, Narasimhan Jegadeesh, and Russ Wermers**, “The value of active mutual fund management: An examination of the stockholdings and trades of fund managers,” *Journal of Financial and quantitative Analysis*, 2000, *35* (03), 343–368.

- Chinco, Alexander and Mao Ye**, “Investment-Horizon Spillovers: Evidence From Decomposing Trading-Volume Variance,” 2016. Working paper.
- Christiano, Lawrence J. and Terry J. Fitzgerald**, “The Band Pass Filter,” *International Economic Review*, 2003, *44*(2), 435–465.
- Commission, Commodity Futures Trading**, “Regulation Automated Trading; Supplemental notice of proposed rulemaking,” in “Federal Register” 2016.
- Dávila, Eduardo and Cecilia Parlatore**, “Trading Costs and Informational Efficiency,” 2016. Working paper.
- DeFusco, Anthony A., Charles G. Nathanson, and Eric Zwick**, “Speculative Dynamics of Prices and Volume,” 2016. Working paper.
- Dew-Becker, Ian and Stefano Giglio**, “Asset Pricing in the Frequency Domain: Theory and Empirics,” *Review of Financial Studies*, 2016. Forthcoming.
- Diamond, Douglas W and Robert E Verrecchia**, “Information aggregation in a noisy rational expectations economy,” *Journal of Financial Economics*, 1981, *9* (3), 221–235.
- Fama, Eugene F. and Kenneth R. French**, “Common risk factors in the returns on stocks and bonds,” *Journal of Financial Economics*, 1993, *33* (1), 3–56.
- Gârleanu, Nicolae and Lasse Heje Pedersen**, “Dynamic Trading with Predictable Returns and Transaction Costs,” *The Journal of Finance*, 2013, *68* (6), 2309–2340.
- Gompers, Paul A and Andrew Metrick**, “Institutional Investors and Equity Prices,” *Quarterly Journal of Economics*, 2001, pp. 229–259.
- Griffin, John M and Jin Xu**, “How smart are the smart guys? A unique view from hedge fund stock holdings,” *Review of Financial Studies*, 2009, *22* (7), 2531–2570.
- Grossman, Sanford J. and Joseph Stiglitz**, “On the Impossibility of InfoInformation Efficient Markets,” *American Economic Review*, 1980, *70*, 393–408.
- He, Hua and Jiang Wang**, “Differential information and dynamic behavior of stock trading volume,” *Review of Financial Studies*, 1995, *8* (4), 919–972.
- Hellwig, Martin F**, “On the aggregation of information in competitive markets,” *Journal of Economic Theory*, 1980, *22* (3), 477–498.
- Hong, Harrison and Jeremy C Stein**, “Disagreement and the stock market,” *The Journal of Economic Perspectives*, 2007, *21* (2), 109–128.

- Hopenhayn, Hugo A and Ingrid M Werner**, “Information, Liquidity, and Asset Trading in a Random Matching Game,” *Journal of Economic Theory*, 1996, *68* (2), 349–379.
- Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp**, “Rational attention allocation over the business cycle,” Technical Report, National Bureau of Economic Research 2016.
- Nagel, Stefan**, “Short sales, institutional investors and the cross-section of stock returns,” *Journal of Financial Economics*, 2005, *78* (2), 277–309.
- Shao, Xiaofeng and Wei Biao Wu**, “Asymptotic Spectral Theory for Nonlinear Time Series,” *The Annals of Statistics*, 2007, *35* (4), 1773–1801.
- Shumway, Robert H. and David T. Stoffer**, *Time Series Analysis and Its Applications*, New York: Springer, 2011.
- Spiegel, Matthew**, “Stock price volatility in a multiple security overlapping generations model,” *Review of Financial Studies*, 1998, *11* (2), 419–447.
- Stoughton, Neal M**, “Moral Hazard and the Portfolio Management Problem,” *The Journal of Finance*, 1993, *48* (5), 2009–2028.
- Townsend, Robert M**, “Forecasting the forecasts of others,” *The Journal of Political Economy*, 1983, pp. 546–588.
- Turley, Robert**, “Informative Prices and the Cost of Capital Markets,” 2012. Working paper.
- Wachter, Jessica A**, “Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets,” *Journal of financial and quantitative analysis*, 2002, *37* (01), 63–91.
- Wang, Jiang**, “A model of intertemporal asset prices under asymmetric information,” *The Review of Economic Studies*, 1993, *60* (2), 249–282.
- , “A model of competitive stock trading volume,” *Journal of political Economy*, 1994, pp. 127–168.
- Watanabe, Masahiro**, “Price volatility and investor behavior in an overlapping generations model with information asymmetry,” *The Journal of Finance*, 2008, *63* (1), 229–272.
- Weller, Brian**, “Efficient Prices at Any Cost: Does Algorithmic Trading Deter Information Acquisition?,” 2016. Working paper.
- Yan, Xuemin Sterling and Zhe Zhang**, “Institutional Investors and Equity Returns: Are Short-Term Institutions Better Informed?,” *Review of Financial Studies*, 2009, *22* (2), 893–924.

A Non-stationary fundamentals

If fundamentals are non-stationary, e.g. if D_t has a unit root, then Σ_D is no longer Toeplitz and our results do not hold. In that case, we assume that D_0 is known by all agents and that the distribution of $\Delta D_t \equiv D_t - D_{t-1}$ is known, with covariance matrix $\Sigma_{\Delta D}$. Then the entire problem can simply be rescaled by defining $\tilde{P}_t \equiv P_t - D_{t-1}$, so that

$$R_t = D_t - P_t \tag{66}$$

$$= \Delta D_t - \tilde{P}_t \tag{67}$$

Our analysis then applies to \tilde{P}_t and ΔD_t , with $Q_{i,t}$ continuing to represent the number of forward contracts on D_t that agent i buys. That is, we are allowing agents to condition demand $Q_{i,t}$ not just on signals and prices, but also the level of D_{t-1} , simply through differencing.

B Proof of lemma 1

Gray (2006) shows that for any circulant matrix (a matrix where row n is equal to row $n - 1$ circularly shifted right by one column, and thus one that is uniquely determined by its top row), the discrete Fourier basis, $u_j = [\exp(i\omega_j t), t = 0, \dots, T - 1]'$ for $j \in \{0, \dots, T - 1\}$, is the set of eigenvectors.

Let Σ be a symmetric Toeplitz matrix with top row $[\sigma_0, \sigma_1, \dots, \sigma_{T-1}]$. Define the function $circ(x)$ to be a circulant matrix with σ_{circ} as its top row. Define a vector σ

$$\sigma \equiv [\sigma_0, \sigma_1 + \sigma_{T-1}, \sigma_2 + \sigma_{T-2}, \dots, \sigma_{T-2} + \sigma_2, \sigma_{T-1} + \sigma_1]'$$
(68)

Following Rao (2016), we “approximate” Σ by the circulant matrix $\Sigma_{circ} \equiv circ(\sigma_{circ})$. Since Σ_{circ} is symmetrical, one may observe that its eigenvalues repeat in the sense that $u'_j \Sigma_{circ} = u'_{T-j} \Sigma_{circ}$ for $0 < j < T$. Since pairs of eigenvectors with matched eigenvalues can be linearly combined to yield alternative eigenvectors, it immediately follows that the matrix Λ from the main text contains a full set of eigenvectors for Σ_{circ} . The associated eigenvalues are

$$f_{\Sigma_{circ}}(\omega_j) = \sigma_0 + 2 \sum_{t=1}^{T-1} \sigma_t \cos(\omega_j t)$$
(69)

We can write this relationship more compactly as:

$$\begin{aligned} \Sigma_{circ} \Lambda &= \Lambda f_{\Sigma} \\ \Lambda' \Sigma_{circ} \Lambda &= f_{\Sigma} \end{aligned}$$

where the $T \times T$ diagonal matrix f_Σ is given by:

$$f_\Sigma = \text{diag} \left(f_\Sigma(\omega_0), f_\Sigma(\omega_1), \dots, f_\Sigma\left(\omega_{\frac{T}{2}}\right), f_\Sigma(\omega_1), f_\Sigma(\omega_2), \dots, f_\Sigma\left(\omega_{\frac{T}{2}-1}\right) \right)'.$$

The approximate diagonalization of the matrix Σ consists in writing:

$$\begin{aligned} \Lambda' \Sigma \Lambda &= f_\Sigma + R_\Sigma \\ \text{where } R_\Sigma &\equiv \Lambda' (\Sigma - \Sigma_{\text{circ}}) \Lambda \end{aligned}$$

By direct inspection of the elements of $\Sigma - \Sigma_{\text{circ}}$, one may see that the m, n element of R_Σ , denoted $R_\Sigma^{m,n}$ satisfies (defining λ_m to be the m th column of Λ and $\lambda_{m,n}$ to be its m, n element)

$$R_\Sigma^{m,n} \equiv \lambda_m' (\Sigma - \Sigma_{\text{circ}}) \lambda_n \quad (70)$$

$$= \sum_{i=1}^T \sum_{j=1}^T \lambda_{m,i} \lambda_{n,j} (\Sigma - \Sigma_{\text{circ}})^{m,n} \quad (71)$$

$$\leq \sum_{i=1}^T \sum_{j=1}^T \frac{2}{T} (\Sigma - \Sigma_{\text{circ}})^{m,n} \quad (72)$$

$$\leq \frac{4}{T} \sum_{j=1}^{T-1} j |\sigma_j| \quad (73)$$

where $(\Sigma - \Sigma_{\text{circ}})^{m,n}$ is the m, n element of $(\Sigma - \Sigma_{\text{circ}})$. So R_Σ is bounded elementwise by a term of order T^{-1} . One may show that the weak norm satisfies $|\cdot| \leq \sqrt{T} |\cdot|_{\max}$, where $|\cdot|_{\max}$ denotes the elementwise max norm, which thus yields the result that $|\Lambda \Sigma \Lambda' - \text{diag}(f_\Sigma)| \leq bT^{-1/2}$ for some b .

B.1 Convergence in distribution and \bar{O} bounds

Define the notation \Rightarrow to mean that $\Lambda X \Rightarrow N(0, \hat{\Sigma}_X)$ if $\Lambda X \sim N(0, \Sigma_X)$ and $|\hat{\Sigma}_X - \Sigma_X| = \bar{O}(T^{-1/2})$.

The notation \bar{O} indicates

$$|A - B| = \bar{O}(T^{-1/2}) \iff |A - B| \leq bT^{-1/2} \quad (74)$$

for some constant b and for all T . This is a stronger statement than typical big- O notation in that it holds for all T , as opposed to holding only for some sufficiently large T .

Trigonometric transforms of stationary time series converge in distribution under more general conditions. See Shumway and Stoffer (2011), Brillinger (1981), and Shao and Wu (2007).

C Derivation of solution 1

Since the optimization is entirely separable across frequencies (confirmed below), we can solve everything in scalar terms. To save notation, we suppress the j subscripts indicating frequencies in this section when they are not necessary for clarity. So in this section f_D , for example, is a scalar representing the spectral density of fundamentals at some arbitrary frequency.

C.1 Statistical inference

We guess that prices take the form

$$p = a_1 d - a_2 z \quad (75)$$

The joint distribution of fundamentals, signals, and prices is then

$$\begin{bmatrix} d \\ y_i \\ p \end{bmatrix} \sim N \left(0, \begin{bmatrix} f_D & f_D & a_1 f_D \\ f_D & f_D + f_i & a_1 f_D \\ a_1 f_D & a_1 f_D & a_1^2 f_D + a_2^2 f_Z \end{bmatrix} \right) \quad (76)$$

The expectation of fundamentals conditional on the signal and price is

$$E[d | y_i, p] = \begin{bmatrix} f_D & a_1 f_D \end{bmatrix} \begin{bmatrix} f_D + f_i & a_1 f_D \\ a_1 f_D & a_1^2 f_D + a_2^2 f_Z \end{bmatrix}^{-1} \begin{bmatrix} y_i \\ p \end{bmatrix} \quad (77)$$

$$= [1, a_1] \begin{bmatrix} 1 + f_i f_D^{-1} & a_1 \\ a_1 & a_1^2 + a_2^2 f_Z f_D^{-1} \end{bmatrix}^{-1} \begin{bmatrix} y_i \\ p \end{bmatrix} \quad (78)$$

and the variance satisfies

$$\tau_i \equiv \text{Var}[d | y_i, p]^{-1} = f_D^{-1} \left(1 - [1 \ a_1] \begin{bmatrix} 1 + f_i f_D^{-1} & a_1 \\ a_1 & a_1^2 + a_2^2 f_Z f_D^{-1} \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ a_1 \end{bmatrix} \right)^{-1} \quad (79)$$

$$= \frac{a_1^2}{a_2^2} f_Z^{-1} + f_i^{-1} + f_D^{-1} \quad (80)$$

We use the notation τ to denote a posterior precision, while f^{-1} denotes a prior precision of one of the basic variables of the model. The above then implies that

$$E[d | y_i, p] = \tau_i^{-1} \left(f_i^{-1} y_i + \frac{a_1}{a_2^2} f_Z^{-1} p \right) \quad (81)$$

C.2 Demand and equilibrium

The agent's utility function is (where variables without subscripts here indicate vectors),

$$U_i = \max_{\{Q_{i,t}\}} \rho^{-1} E_{0,i} [T^{-1} Q'_i (D - P)] - \rho^{-2} \text{Var}_{0,i} [T^{-1/2} Q'_i (D - P)] \quad (82)$$

$$= \max_{\{Q_{i,t}\}} \rho^{-1} E_{0,i} [T^{-1} q'_i (d - p)] - \rho^{-2} \text{Var}_{0,i} [T^{-1/2} q'_i (d - p)] \quad (83)$$

$$= \max_{\{Q_{i,t}\}} \rho^{-1} T^{-1} \sum_{j=0}^{T-1} q_{i,j} E_{0,i} [(d_j - p_j)] - \rho^{-2} T^{-1} \sum_{j=0}^{T-1} q_{i,j}^2 \text{Var}_{0,i} [d_j - p_j] \quad (84)$$

where the last line follows by imposing the asymptotic independence of d across frequencies (we analyze the error induced by that approximation below). The utility function is thus entirely separable across frequencies, with the optimization problem for each $q_{i,j}$ independent from all others.

Taking the first-order condition associated with the last line above for a single frequency, we obtain

$$\begin{aligned} q_i &= \rho \tau_i E [d - p \mid y_i, p] \\ &= \rho_i \left(f_i^{-1} y_i + \left(\frac{a_1}{a_2} f_Z^{-1} - \tau_i \right) p \right) \end{aligned}$$

Summing up all demands and inserting the guess for the price yields

$$z = \int_i \rho \left(f_i^{-1} y_i + \left(\frac{a_1}{a_2} f_Z^{-1} - \tau_i \right) (a_1 d - a_2 z) \right) di \quad (85)$$

$$= \int_i \rho \left(f_i^{-1} d + \left(\frac{a_1}{a_2} f_Z^{-1} - \tau_i \right) (a_1 d - a_2 z) \right) di \quad (86)$$

Where the second line uses the law of large numbers. Matching coefficients then yields

$$\int_i \rho \left(\frac{a_1}{a_2} f_Z^{-1} - \tau_i \right) di = -a_2^{-1} \quad (87)$$

$$\int_i \rho f_i^{-1} + \rho \left(\frac{a_1}{a_2} f_Z^{-1} - \tau_i \right) a_1 di = 0 \quad (88)$$

and therefore

$$\rho \int_i f_i^{-1} di = \frac{a_1}{a_2} \quad (89)$$

Inserting the expression for τ_i into (87) yields

$$a_1 = \frac{\frac{a_1}{a_2} + \rho \left(\frac{a_1}{a_2} \right)^2 f_Z^{-1}}{\rho \left(\int_i f_i^{-1} di + f_D^{-1} + \left(\frac{a_1}{a_2} \right)^2 f_Z^{-1} \right)} \quad (90)$$

Now define aggregate precision to be

$$f_{avg}^{-1} \equiv \int_i f_i^{-1} di \quad (91)$$

We then have

$$\tau_i = \frac{a_1^2}{a_2^2} f_Z^{-1} + f_i^{-1} + f_D^{-1} \quad (92)$$

$$\tau_{avg} \equiv \int \tau_i di = (\rho f_{avg}^{-1})^2 f_Z^{-1} + f_{avg}^{-1} + f_D^{-1} \quad (93)$$

$$a_1 = \tau_{avg}^{-1} \left(f_{avg}^{-1} + (\rho f_{avg}^{-1})^2 f_Z^{-1} \right) = 1 - \frac{f_D^{-1}}{\tau_{avg}} = \frac{\tau_{avg} - f_D^{-1}}{\tau_{avg}} \quad (94)$$

$$a_2 = \frac{a_1}{\rho f_{avg}^{-1}} \quad (95)$$

C.3 Proof of proposition 1

In the time domain, the solution from Admati (1985) is

$$P = A_1 D - A_2 Z \quad (96)$$

$$A_1 \equiv I - S_{avg}^{-1} \Sigma_D^{-1} \quad (97)$$

$$A_2 \equiv \rho^{-1} A_1 \Sigma_{avg} \quad (98)$$

Standard properties of norms yield the following result. If $|A - B| = \bar{O}(T^{-1/2})$ and $|C - D| = \bar{O}(T^{-1/2})$, then

$$|cA - cB| = \bar{O}(T^{-1/2}) \quad (99)$$

$$|A^{-1} - B^{-1}| = \bar{O}(T^{-1/2}) \quad (100)$$

$$|(A + C) - (B + D)| = \bar{O}(T^{-1/2}) \quad (101)$$

$$|AC - BD| = \bar{O}(T^{-1/2}) \quad (102)$$

In other words, convergence in weak norm carries through under addition, multiplication, and inversion. Since A_1 is a function of Toeplitz matrices using those operations, it follows that $|\Lambda' A_1 \Lambda - \text{diag}(a_1)| = \bar{O}(T^{-1/2})$, and the same holds for A_2 .

For the variance of prices, we define

$$R_1 = A_1 - \Lambda \text{diag}(a_1) \Lambda' \quad (103)$$

$$R_2 = A_2 - \Lambda \text{diag}(a_2) \Lambda' \quad (104)$$

$$|Var [P - \Lambda p]| \leq |R_1 \Sigma_D R_1'| + |R_2 \Sigma_Z R_2'| \quad (105)$$

$$\leq |R_1 \Sigma_D| |R_1| + |R_2 \Sigma_Z| |R_2| \quad (106)$$

$$\leq \|\Sigma_D\| |R_1|^2 + \|\Sigma_Z\| |R_2|^2 \quad (107)$$

$$\leq K \left(|R_1|^2 + |R_2|^2 \right) \quad (108)$$

The first line follows from the triangle inequality; the second line comes from the sub-multiplicativity of the weak norm; the third line uses the fact that, as indicated by Gray (2006), for any two square matrices G, H , $\|GH\|_2 \leq \|G\| \|H\|$; and the last line follows from the assumption that the eigenvalues of Σ_D and Σ_Z are bounded by some K .

Since the weak norm is invariant under unitary transformations,

$$|R_1| = |\Lambda' R_1 \Lambda| = |\Lambda' A_i \Lambda - \text{diag}(a_i)| \quad , \quad i = 1, 2.$$

Therefore,

$$|Var [P - \Lambda P]| \leq K \left(|\Lambda' A_1 \Lambda - \text{diag}(a_1)|^2 + |\Lambda' A_2 \Lambda - \text{diag}(a_2)|^2 \right) \quad (109)$$

$$= \bar{O} \left(\frac{1}{T} \right) \quad (110)$$

Since $\|\cdot\| \leq \sqrt{T} \|\cdot\|$, $\|Var [P^c - P]\| = \bar{O} (T^{-1/2})$.

D Proof of lemma 2

Inserting the optimal value of $q_{i,j}$ into the utility function, we obtain

$$E_{-1} [U_{i,0}] \equiv \frac{1}{2} E \left[T^{-1} \sum_{j=0}^{T-1} \tau_{i,j} E [d_j - p_j | y_{i,j}, p_j]^2 \right] \quad (111)$$

$U_{i,0}$ is utility conditional on an observed set of signals and prices. $E_{-1} [U_{i,0}]$ is then the expectation taken over the distributions of prices and signals.

$Var [E [d_j - p_j | y_{i,j}, p_j]]$ is the variance of the part of the return on portfolio j explained by $y_{i,j}$ and p_j , while $\tau_{i,j}^{-1}$ is the residual variance. The law of total variance says

$$Var [d_j - p_j] = Var [E [d_j - p_j | y_{i,j}, p_j]] + E [Var [d_j - p_j | y_{i,j}, p_j]] \quad (112)$$

where the second term on the right-hand side is just $\tau_{i,j}^{-1}$ and the first term is $E [E [d_j - p_j | y_{i,j}, p_j]^2]$ since everything has zero mean. The unconditional variance of returns is

$$Var [d_j - p_j] = (1 - a_{1,j})^2 f_{D,j} + \frac{a_{1,j}^2}{\left(\rho f_{avg,j}^{-1} \right)^2} f_{Z,j} \quad (113)$$

So then

$$E_{-1}[U_{i,0}] = \frac{1}{2}T^{-1} \sum_{j=0}^{T-1} \left((1 - a_{1,j})^2 f_{D,j} + \frac{a_{1,j}^2}{(\rho f_{avg}^{-1})^2} f_{Z,j} \right) \tau_{i,j} - \frac{1}{2} \quad (114)$$

We thus obtain the result that agent i 's expected utility is linear in the precision of the signals that they receive (since $\tau_{i,j}$ is linear in $f_{i,j}^{-1}$).

Furthermore,

$$(1 - a_{1,j})^2 f_{D,j} + \frac{a_{1,j}^2}{(\rho f_{avg,j}^{-1})^2} f_{Z,j} = \tau_{avg,j}^{-2} f_{D,j}^{-1} + \rho^{-2} f_{avg,j}^2 \tau_{avg,j}^{-2} \left(\tau_{avg,j} - f_{D,j}^{-1} \right)^2 f_{Z,j} \quad (115)$$

$$= \tau_{avg,j}^{-2} f_{D,j}^{-1} + \rho^{-2} \tau_{avg,j}^{-2} \left(\rho^2 f_{avg}^{-1} f_Z^{-1} + 1 \right)^2 f_{Z,j} \quad (116)$$

$$= \tau_{avg,j}^{-2} \left(f_{D,j}^{-1} + \rho^{-2} \left(\rho^2 f_{avg}^{-1} f_Z^{-1} + 1 \right)^2 f_{Z,j} \right) \quad (117)$$

So

$$E_{-1}[U_{i,0}] = \frac{1}{2}T^{-1} \sum_{j=0}^{T-1} \tau_{avg,j}^{-2} \left(f_{D,j}^{-1} + \rho^{-2} \left(\rho^2 f_{avg}^{-1} f_Z^{-1} + 1 \right)^2 f_{Z,j} \right) \left(\left(\rho f_{avg,j}^{-1} \right)^2 f_{Z,j}^{-1} + f_{i,j}^{-1} + f_{D,j}^{-1} \right) - \frac{1}{2} \quad (118)$$

E Derivation of solution 2

Investors allocate attention, $f_{i,j}^{-1}$, to maximize $E_{-1}[U_{i,0}]$ subject to the constraint

$$\sum_{j,j'} f_{i,j}^{-1} \leq \bar{f}^{-1} \quad (119)$$

and that $f_{i,j}^{-1} = f_{i,j'}^{-1}$. Since the investors are maximizing a linear objective subject to a linear constraint, the optimal policy is clearly to allocate attention $f_{i,j}^{-1}$ only to the frequencies j at which the marginal benefit is equal to the maximum available marginal benefit.

Define the function λ_j

$$\lambda_j(x) \equiv \left((\rho x)^2 f_{Z,j}^{-1} + x + f_{D,j}^{-1} \right)^{-2} \left(f_{D,j}^{-1} + \rho^{-2} \left(\rho^2 x f_{Z,j}^{-1} + 1 \right)^2 f_{Z,j} \right)$$

then $\lambda_j \left(f_{avg,j}^{-1} \right)$ is the marginal benefit from attention to frequency j . Note that $d\lambda_j(x)/dx < 0$. In equilibrium, then, there is a number $\bar{\lambda}$ such that

$$\lambda_j \left(f_{avg,j}^{-1} \right) \leq \bar{\lambda} \text{ for all } j \quad (120)$$

Now define $\mathcal{J}(\bar{\lambda})$ to be the set of frequencies j such that $\lambda_j^{-1}(\bar{\lambda}) > 0$.²³ That is the set of

²³Technically, it is the set of frequencies for which $\lambda_j^{-1}(\min(\bar{\lambda}, \lambda_j(0))) > 0$.

frequencies for which there is positive attention.

For any frequency that investors allocate attention to,

$$f_{avg,j}^{-1} = \lambda_j^{-1}(\bar{\lambda}) \quad (121)$$

$$f_{avg,j}^{-1} = \int f_{i,j}^{-1} di \quad (122)$$

Now

$$\sum_{j,j' \in \mathcal{J}} \int f_{i,j}^{-1} di = \int \sum_{j,j' \in \mathcal{J}} f_{i,j}^{-1} di \quad (123)$$

$$= \int \bar{f}^{-1} di = \bar{f}^{-1} \quad (124)$$

So then

$$\sum_{j,j' \in \mathcal{J}(\bar{\lambda})} \lambda_j^{-1}(\bar{\lambda}) = \sum_{j,j' \in \mathcal{J}(\bar{\lambda})} f_{avg,j}^{-1} = \bar{f}^{-1} \quad (125)$$

So $\bar{\lambda}$ is obtained by solving $\sum_{j,j' \in \mathcal{J}(\bar{\lambda})} \lambda_j^{-1}(\bar{\lambda}) = \bar{f}^{-1}$.

F Time horizon and investment

At first glance, the assumption of mean-variance utility over cumulative returns over a long period of time ($T \rightarrow \infty$) may appear to give investors an incentive to primarily worry about long-horizon performance, whereas a small value of T would make investors more concerned about short-term performance. In the present setting, that intuition is not correct – the $T \rightarrow \infty$ limit determines how detailed investment strategies may be, rather than incentivizing certain types of strategies.

The easiest way to see why the time horizon controls only the detail of the investment strategies is to consider settings in which T is a power of 2. If $T = 2^k$, then the set of fundamental frequencies is

$$\left\{ 2\pi j / 2^k \right\}_{j=0}^{2^{k-1}} \quad (126)$$

For $T = 2^{k-1}$, the set of frequencies is

$$\left\{ 2\pi j / 2^{k-1} \right\}_{j=0}^{2^{k-2}} = \left\{ 2\pi (2j) / 2^k \right\}_{j=0}^{2^{k-2}} \quad (127)$$

That is, when T falls from 2^k to 2^{k-1} , the effect is to simply eliminate alternate frequencies. Changing T does not change the lowest or highest available frequencies (which are always 0 and π , respectively). It just discretizes the $[0, \pi]$ interval more coarsely; or, equivalently, it means that the matrix Λ is constructed from a smaller set of basis vectors.

When T is smaller – there are fewer available basis functions – Q and its frequency domain analog $q \equiv \Lambda'Q$ have fewer degrees of freedom and hence must be less detailed. So the effect

of a small value of T is to make it more difficult for an investor to *isolate* particularly high- or low-frequency fluctuations in fundamentals (or any other narrow frequency range). But in no way does T cause the investor's portfolio to depend more on one set of frequencies than another. While we take $T \rightarrow \infty$, we will see that the model's separating equilibrium features investors who trade at both arbitrarily low and high frequencies, and T has no effect on the distribution of investors across frequencies.

G Proofs of specialization model predictions

G.1 Results 2 and 3

$$q_i = \rho \left(f_i^{-1} y_i + \left(\frac{a_1}{a_2} f_Z^{-1} - \tau_i \right) p \right)$$

The coefficient on $\tilde{\varepsilon}_i$ is f_i^{-1} . Straightforward but tedious algebra confirms that the coefficient on d is

$$\rho (f_{avg}^{-1} - f_i^{-1}) (a_1 - 1)$$

The coefficient on z is

$$1 + \rho (f_i^{-1} - f_{avg}^{-1}) a_2$$

We thus have

$$q_i = \rho (f_{avg}^{-1} - f_i^{-1}) (a_1 - 1) d + (1 + \rho (f_i^{-1} - f_{avg}^{-1})) a_2 z \quad (128)$$

Now note that

$$r = (1 - a_1) d + d_2 z \quad (129)$$

So then

$$q_i = \rho (f_i^{-1} - f_{avg}^{-1}) r + \rho \tilde{\varepsilon}_i + z \quad (130)$$

The result on the covariance then follows trivially.

G.2 Result 4

Approximating first differences with derivatives, we obtain

$$\Delta Q_{i,t} - \Delta Z_t \approx - \sum_{j=0}^{T/2} \frac{2\pi j}{T} \left[\begin{array}{l} \sin(2\pi j t/T) \left(\rho \left[\left(f_{i,j}^{-1} - f_{avg,j}^{-1} \right) r_j + f_{i,j}^{-1} \tilde{\varepsilon}_{i,j} \right] \right) \\ + \cos(2\pi j t/T) \left(\rho \left[\left(f_{i,j}^{-1} - f_{avg,j}^{-1} \right) r_{j'} + f_{i,j}^{-1} \tilde{\varepsilon}_{i,j'} \right] \right) \end{array} \right] \quad (131)$$

where the approximation becomes a true equality as $T \rightarrow \infty$. Now if we furthermore use the approximations $f_{i,j_i^*}^{-1} - f_{avg,j_i^*}^{-1} \approx \bar{f}^{-1}/2$ and suppose that the exogenous supply process is small enough that it rarely causes a trader's demand to change signs, then we have

$$|\Delta Q_{i,t}| \approx |\Delta Z_t| + \omega_{j_i^*} \bar{f}^{-1} \rho \left| \begin{array}{l} \sin(\omega_{j_i^*} t) (r_{j_i^*} + \tilde{\varepsilon}_{i,j_i^*}) \\ + \cos(\omega_{j_i^*} t) (r_{j_i^{*'}} + \tilde{\varepsilon}_{i,j_i^{*'}}) \end{array} \right|. \quad (132)$$

G.3 Result 5

$$\begin{aligned}
QV\{q_j\} &\equiv \sum_{t=2}^T (Q_{i,t} - Q_{i,t-1})^2 \approx \sum_{t=2}^T \left(\sum_j \frac{2\pi j}{T} \begin{bmatrix} q_j \sin(2\pi jt/T) \\ +q_{j'} \cos(2\pi jt/T) \end{bmatrix} \right)^2 & (133) \\
&= \sum_{t=2}^T \sum_{j,k} \left(\frac{2\pi}{T} \right)^2 jk \begin{bmatrix} q_j \sin(2\pi jt/T) q_k \sin(2\pi kt/T) + q_{j'} \cos(2\pi jt/T) q_k \sin(2\pi kt/T) \\ q_j \sin(2\pi jt/T) q_{k'} \cos(2\pi kt/T) + q_{j'} \cos(2\pi jt/T) q_{k'} \cos(2\pi kt/T) \end{bmatrix} & (134) \\
&\approx \sum_{j,j'} (2\pi j)^2 T^{-1} q_j^2 & (135)
\end{aligned}$$

where the equality in the first line is approximate in assuming that $\cos(2\pi jt/T) - \cos(2\pi j(t-1)/T) \approx \frac{2\pi j}{T} \sin(2\pi jt/T)$ and the same for the differences in the sines. The third line uses the fact that sines of unequal frequencies are orthogonal (it is approximate because $t=1$ is not included in the sum inserts the integral for \sin^2 and \cos^2 , rather than the exact finite sums. All the approximations here are accurate for large T .

H Proofs of trading restriction results

H.1 Results 6 and 7

If trade by the investors is not allowed at certain frequencies, then obviously markets cannot clear at those frequencies when supply is inelastic. In this section we therefore first solve the model for the case with an upward sloping supply curve and then analyze the effect of eliminating trade on asset prices and returns.

H.1.1 Equilibrium with elastic supply

We now assume that there is exogenous supply on each date of

$$Z_t = \tilde{Z}_t + kP_t \quad (136)$$

where k is a constant determining the slope of the supply curve. One could imagine allowing k to differ across frequencies, which would be equivalent to allowing supply to depend on prices on multiple dates (intuitively, maybe supply increases by more when prices have been persistently high than when they are just temporarily high). Here, though, we simply leave k constant across frequencies. Multiplying by Λ' yields

$$z_j = \tilde{z}_j + kp_j \quad (137)$$

Solving the inference problem as before, we obtain

$$\tau_i \equiv \text{Var} [d | y_i, p]^{-1} \quad (138)$$

$$= \frac{a_1^2}{a_2^2} f_{\tilde{Z}}^{-1} + f_i^{-1} + f_D^{-1} \quad (139)$$

and

$$E [d | y_i, p] = \tau_i^{-1} \left(f_i^{-1} y_i + \frac{a_1}{a_2} f_{\tilde{Z}}^{-1} p \right) \quad (140)$$

H.1.2 Demand and equilibrium

The investors' demand curves are again

$$q_i = \rho_i \left(f_i^{-1} y_i + \left(\frac{a_1}{a_2} f_{\tilde{Z}}^{-1} - \tau_i \right) p \right)$$

Summing up all demands and inserting the guess for the price process yields

$$\tilde{z} + k (a_1 d - a_2 \tilde{z}) = \int_i \rho \left(f_i^{-1} d + \left(\frac{a_1}{a_2} f_{\tilde{Z}}^{-1} - \tau_i \right) (a_1 d - a_2 \tilde{z}) \right) di \quad (141)$$

Matching coefficients yields

$$\int_i \rho \left(\frac{a_1}{a_2} f_{\tilde{Z}}^{-1} - \tau_i \right) di = -a_2^{-1} (1 - ka_2) \quad (142)$$

$$\int_i \rho f_i^{-1} + \rho \left(\frac{a_1}{a_2} f_{\tilde{Z}}^{-1} - \tau_i \right) a_1 di = ka_1 \quad (143)$$

Combining those two equations, we have

$$\rho f_{avg}^{-1} = a_1 (k + a_2^{-1} (1 - ka_2)) \quad (144)$$

$$= \frac{a_1}{a_2} \quad (145)$$

$$a_1 = \frac{f_{avg}^{-1} + (\rho f_{avg}^{-1})^2 f_{\tilde{Z}}^{-1}}{\tau_{avg} + \rho^{-1} k} \quad (146)$$

$$= \frac{\tau_{avg} - f_D^{-1}}{\tau_{avg} + \rho^{-1} k} \quad (147)$$

$$a_2 = \frac{a_1}{\rho f_{avg}^{-1}} \quad (148)$$

H.1.3 Utility

As before, the contribution to optimized utility from frequency j is

$$\left((1 - a_{1,j})^2 f_{D,j} + \frac{a_{1,j}^2}{(\rho f_{avg}^{-1})^2} f_{Z,j} \right) \tau_{i,j} \quad (149)$$

Furthermore,

$$\begin{aligned} (1 - a_{1,j})^2 f_{D,j} + \frac{a_{1,j}^2}{(\rho f_{avg,j}^{-1})^2} f_{\tilde{Z},j} &= \left(\frac{\rho^{-1}k + f_D^{-1}}{\tau_{avg} + \rho^{-1}k} \right)^2 f_{D,j} + \rho^{-2} f_{avg,j}^2 \left(\frac{\tau_{avg} - f_D^{-1}}{\tau_{avg} + \rho^{-1}k} \right)^2 f_{\tilde{Z},j} \\ &= (\tau_{avg} + \rho^{-1}k)^{-2} \left((\rho^{-1}k + f_D^{-1})^2 f_{D,j} + \rho^{-2} f_{avg,j}^2 \left((\rho f_{avg}^{-1})^2 f_{\tilde{Z}}^{-1} + f_{avg}^{-1} \right)^2 f_{\tilde{Z},j} \right) \\ &= (\tau_{avg} + \rho^{-1}k)^{-2} \left((\rho^{-1}k + f_D^{-1})^2 f_{D,j} + \rho^{-2} \left(\rho^2 f_{avg,j}^{-1} f_{\tilde{Z}}^{-1} + 1 \right)^2 f_{\tilde{Z},j} \right) \end{aligned}$$

So

$$E_{-1}[U_{i,0}] = \frac{1}{2} T^{-1} \sum_{j=0}^{T-1} \left[\begin{aligned} &(\tau_{avg,j} + \rho^{-1}k)^{-2} \left((\rho^{-1}k + f_{D,j}^{-1})^2 f_{D,j} + \rho^{-2} \left(\rho^2 f_{avg,j}^{-1} f_{\tilde{Z},j}^{-1} + 1 \right)^2 f_{\tilde{Z},j} \right) \\ &\times \left((\rho f_{avg,j}^{-1})^2 f_{Z,j}^{-1} + f_{i,j}^{-1} + f_{D,j}^{-1} \right) \end{aligned} \right] - \frac{1}{2} \quad (150)$$

When there are no active investors and just exogenous supply, we have

$$0 = \tilde{z}_j + k p_j \quad (151)$$

$$p_j = -k^{-1} \tilde{z}_j \quad (152)$$

$$r_j = d_j - k^{-1} \tilde{z}_j \quad (153)$$

We then have

$$f_R = f_D + \frac{f_Z}{k^2} \quad (154)$$

$$f_{R,0} = f_{D,j} + \frac{f_{\tilde{Z},j}}{(k + \rho f_{D,j}^{-1})^2} \quad (155)$$

H.2 Result 9

We have

$$D | Y, P \sim N(\bar{D}, \Lambda \text{diag}(\tau_0^{-1}) \Lambda') \quad (156)$$

where τ_0 is a vector of frequency-specific precisions conditional on prices. Now consider some average over D , $F'D$, where F is a column vector. Then

$$\text{Var}(D_t) = \mathbf{1}'_t \Lambda \text{diag}(\tau_0^{-1}) \Lambda' \mathbf{1}_t \quad (157)$$

$$= (\Lambda' \mathbf{1}_t)' \text{diag}(\tau_0^{-1}) (\Lambda' \mathbf{1}_t) \quad (158)$$

$$= \sum_{j,j'} \lambda_{t,j}^2 \tau_{0,j}^{-1} \quad (159)$$

$$= \lambda_{t,0}^2 \tau_{0,0}^{-1} + \lambda_{t,T/2}^2 \tau_{0,0}^{-1} + \sum_{n=1}^{T/2-1} (\lambda_{t,n}^2 + \lambda_{t,n'}^2) \tau_{0,n}^{-1} \quad (160)$$

where $\mathbf{1}_t$ is a vector equal to 1 in its t th element and zero elsewhere and $\lambda_{t,j}$ is the j th trigonometric transform evaluated at t , with

$$\lambda_{t,j} = \sqrt{2/T} \cos(2\pi j(t-1)/T) \quad (161)$$

$$\lambda_{t,j'} = \sqrt{2/T} \sin(2\pi j(t-1)/T) \quad (162)$$

$$\lambda_{t,0} = \sqrt{1/T} \quad (163)$$

$$\lambda_{t,T/2} = \sqrt{1/T} \cos(\pi(t-1)) = \sqrt{1/T} (-1)^{t-1} \quad (164)$$

More generally, then

$$\text{Var}\left(\frac{1}{s} \sum_{m=0}^{s-1} D_{t+m}\right) = \frac{1}{s^2} \left(\sum_{m=0}^{s-1} \mathbf{1}_{t+m}\right)' \Lambda \text{diag}(\tau_0^{-1}) \Lambda' \left(\sum_{m=0}^{s-1} \mathbf{1}_{t+m}\right) \quad (165)$$

$$= \frac{1}{s^2} \left(\sum_{m=0}^{s-1} \lambda_{t+m,0}\right)^2 \tau_{0,0}^{-1} + \frac{1}{s^2} \left(\sum_{m=0}^{s-1} \lambda_{t+m,T/2}\right)^2 \tau_{0,T/2}^{-1} \quad (166)$$

$$+ \frac{1}{s^2} \sum_{n=1}^{T/2-1} \left[\left(\sum_{m=0}^{s-1} \lambda_{t+m,n}\right)^2 + \left(\sum_{m=0}^{s-1} \lambda_{t+m,n'}\right)^2 \right] \tau_{0,n}^{-1} \quad (167)$$

where $\tau_{0,n}$ is the frequency- n element of τ_0 . For $0 < n < T/2$

$$\left(\sum_{m=0}^{s-1} \lambda_{t+m,n}\right)^2 + \left(\sum_{m=0}^{s-1} \lambda_{t+m,n'}\right)^2 = \sum_{m=0}^{s-1} \sum_{k=0}^{s-1} \frac{2}{T} \left[\begin{array}{l} \cos(2\pi n(t+m-1)/T) \cos(2\pi n(t+k-1)/T) \\ + \sin(2\pi n(t+m-1)/T) \sin(2\pi n(t+k-1)/T) \end{array} \right] \quad (168)$$

Now note that

$$2 \cos(x) \cos(y) + 2 \sin(x) \sin(y) = 2 \cos(x-y) \quad (169)$$

So we have

$$\left(\sum_{m=0}^{s-1} \lambda_{t+m,n}\right)^2 + \left(\sum_{m=0}^{s-1} \lambda_{t+m,n}\right)^2 = \frac{2}{T} \sum_{m=0}^{s-1} \sum_{k=0}^{s-1} \cos\left(\frac{2\pi n}{T}(m-k)\right) \quad (170)$$

$$= 2\frac{s}{T} \sum_{m=-(s-1)}^{s-1} \frac{s-|m|}{s} \cos\left(\frac{2\pi n}{T}m\right) \quad (171)$$

$$= 2\frac{s}{T} F_s\left(\frac{2\pi n}{T}\right) \quad (172)$$

$$= \frac{2}{T} \frac{1 - \cos\left(s\frac{2\pi n}{T}\right)}{1 - \cos\left(\frac{2\pi n}{T}\right)} \quad (173)$$

where F_s denotes the s th-order Fejér kernel. Note that when $s = T$, the above immediately reduces to zero, since $\cos(2\pi n) = 0$. That is the desired result, as an average over all dates should be unaffected by fluctuations at any frequency except zero.

For $n = 0$,

$$\left(\sum_{m=0}^{s-1} f_{t+m,0}\right)^2 = \left(\sum_{m=0}^{s-1} \sqrt{1/T}\right)^2 \quad (174)$$

$$= \left(s\frac{1}{T^{1/2}}\right)^2 \quad (175)$$

$$= \frac{s}{T} F_s(0) \quad (176)$$

Since $F_s(0) = s$ (technically, this holds as a limit: $\lim_{x \rightarrow 0} F_s(x) = s$).

For $n = T/2$,

$$\left(\sum_{m=0}^{s-1} f_{t+m,T/2}\right)^2 = \frac{1}{T} \left(\sum_{m=1}^s (-1)^m\right)^2 = \begin{cases} \frac{1}{T} & \text{for odd } s \\ 0 & \text{otherwise} \end{cases} \quad (177)$$

$$= \frac{s}{T} \frac{1}{s} \left(\frac{\sin(s\pi/2)}{\sin(\pi/2)}\right)^2 = \frac{s}{T} F_s(\pi) \quad (178)$$

So we finally have that

$$\text{Var}\left(\frac{1}{s} \sum_{m=0}^{s-1} D_{t+m}\right) = \frac{1}{sT} \sum_{j,j'} F_s(\omega_j) \tau_{0,j}^{-1} \quad (179)$$

I Costly learning about prices

I.1 Generic result: no learning from prices

Lemma 3 *Assume that learning from prices is costly. At that at time -1 , if agent i decides to infer information from prices, then their capacity constraint is:*

$$\text{Tr}(f_i^{-1} + f_P^{-1}) \leq \bar{f}^{-1},$$

where f_P^{-1} is inverse of the variance-covariance matrix of signals contained in prices, and f_i^{-1} is the variance-covariance of the private signals of agent i . On the other hand, if agent i decides not to infer information from prices, then his capacity constraint is:

$$\text{Tr}(f_i^{-1}) \leq \bar{f}^{-1}.$$

Then, agents always prefer not to learn from prices.

Proof. If agent i has decided not to learn from prices, then at time 0, their posterior distribution over d is:

$$\begin{aligned} d|y_i &\sim N(\mu(y_i), \tau_i^{-1}) \\ \tau_i^{NP} &= f_D^{-1} + f_i^{-1} \\ \mu(y_i) &= (\tau_i^{NP})^{-1} f_i^{-1} y_i \end{aligned} \tag{180}$$

Agent i still observes prices; their first-order condition leads to the demand schedule:

$$q_i = \rho \tau_i^{NP} (\mu(y_i) - p).$$

His time-0 utility is:

$$U_{0,i}^{NP}(y_i; p) = \frac{1}{2T} (\mu(y_i) - p)' \tau_i^{NP} (\mu(y_i) - p). \tag{181}$$

Since τ_i^{NP} is symmetric, this implies:

$$E_{-1,i} [U_{0,i}^{NP}] = \frac{1}{2T} \text{tr}(\tau_i^{NP} V_i^{NP}) + \frac{1}{2T} (\mu_i^{NP})' \tau_i^{NP} \mu_i^{NP}, \tag{182}$$

where as before:

$$\begin{aligned} \mu_i^{NP} &= E_{-1,i} [\mu(y_i) - p] \\ V_i^{NP} &= \text{Var}_{-1,i} [\mu(y_i) - p] \end{aligned} \tag{183}$$

As before, because all fundamentals are mean 0, $\mu_i = 0$. Moreover, by the law of total variance:

$$V_i = \underbrace{\text{Var}_{-1} [d - p]}_{\equiv V_{-1}} - (\tau_i^{NP})^{-1}$$

Therefore,

$$\begin{aligned}
E_{-1,i} \left[U_{0,i}^{NP} \right] &= \frac{1}{2T} \text{tr}(\tau_i^{NP} V_i) \\
&= \frac{1}{2T} \text{tr}(\tau_i^{NP} V_{-1}) - \frac{1}{2T} \text{tr}(I) \\
&= \frac{1}{2T} \text{tr}(f_D^{-1} V_{-1}) - \frac{1}{2T} \text{tr}(I) + \frac{1}{2T} \text{tr}(f_i^{-1} V_{-1})
\end{aligned} \tag{184}$$

The time-(-1) attention allocation problem of such an agent is therefore:

$$\begin{aligned}
U_{-1,i}^{NP}(f_{avg}^{-1}) &= -\frac{1}{2} + \frac{1}{2T} \text{tr}(f_D^{-1} V_{-1}) + \frac{1}{2T} \max_{f_i^{-1}} \text{tr}(f_i^{-1} V_{-1}) \\
\text{s.t.} \quad f_{i,j}^{-1} &\geq 0 \quad \forall j \in [0, \dots, T-1] \\
\text{tr}(f_i^{-1}) &\leq \bar{f}^{-1}
\end{aligned} \tag{185}$$

For an agent who does learn from prices (but shares the other agent's ex-ante distribution over p and d , summarized by V_{-1}), the attention allocation problem has already been derived; it is given by:

$$\begin{aligned}
U_{-1,i}(f_{avg}^{-1}) &= -\frac{1}{2} + \frac{1}{2T} \text{tr}((f_D^{-1} + f_P^{-1}) V_{-1}) + \frac{1}{2T} \max_{f_i^{-1}} \text{tr}(f_i^{-1} V_{-1}) \\
\text{s.t.} \quad f_{i,j}^{-1} &\geq 0 \quad \forall j \in [0, \dots, T-1] \\
\text{tr}(f_i^{-1} + f_P^{-1}) &\leq \bar{f}^{-1}
\end{aligned} \tag{186}$$

Since f_i^{-1} is diagonal, $f_i^{-1} \rightarrow \text{tr}(f_i^{-1} V_{-1})$ can be thought of as a linear map on R^T . By the Riesz representation theorem, there is $\lambda \in R^T$ such that $\forall f_i^{-1}, \text{tr}(f_i^{-1} V_{-1}) = \sum_{j=0}^{T-1} f_{i,j}^{-1} \lambda_j$. Let $\tilde{\lambda}$ denote the element-wise maximum of λ . Note, in particular, that:

$$\text{tr}(f_P^{-1} V_{-1}) = \sum_{j=0}^{T-1} f_{P,j}^{-1} \lambda_j.$$

Moreover, after optimization, not learning through prices yields utility:

$$U_{-1,i}^{NP}(f_{avg}^{-1}) = -\frac{1}{2} + \frac{1}{2T} \text{tr}(f_D^{-1} V_{-1}) + \frac{1}{2T} \tilde{\lambda} \bar{f}^{-1}.$$

Learning through prices yields utility:

$$U_{-1,i}(f_{avg}^{-1}) = -\frac{1}{2} + \frac{1}{2T} \text{tr}((f_D^{-1} + f_P^{-1}) V_{-1}) + \frac{1}{2T} \tilde{\lambda} (\bar{f}^{-1} - \text{tr}(f_P^{-1}))$$

The difference between the two is:

$$\begin{aligned}
U_{-1,i}^{NP}(f_{avg}^{-1}) - U_{-1,i}(f_{avg}^{-1}) &= \frac{1}{2T} \tilde{\lambda} \text{tr}(f_P^{-1}) - \frac{1}{2T} \text{tr}(f_P^{-1} V_{-1}) \\
&= \frac{1}{2T} \tilde{\lambda} \text{tr}(f_P^{-1}) - \frac{1}{2T} \sum_{j=0}^{T-1} f_{P,j}^{-1} \lambda_j \\
&\geq 0
\end{aligned} \tag{187}$$

Therefore, the agent always prefer not to learn from prices. ■

I.2 The equilibrium when agents do not learn about prices

Guess:

$$p = a_3 d - a_4 z$$

with a_3, a_4 diagonal matrices of size $T \times T$. Straightforward derivations lead to:

$$\begin{aligned}
a_3 &= I - (\tau_{avg} + kI)^{-1} (f_D^{-1} + kI) \\
&= (\tau_{avg} + kI)^{-1} f_{avg}^{-1} \\
a_4 &= \frac{1}{\rho} a_3 f_{avg} \\
&= \frac{1}{\rho} (\tau_{avg} + kI)^{-1} \\
\tau_{avg} &= f_{avg}^{-1} + f_D^{-1} \\
\tau_i &= f_i^{-1} + f_D^{-1}
\end{aligned} \tag{188}$$

Moreover, expected utility is given by:

$$\begin{aligned}
E_{-1,i} [U_{0,i}^{NP}] &= C^{NP} + \frac{1}{2T} \text{tr}(V_{-1}^{NP} f_i^{-1}) \\
V_{-1}^{NP} &= f_D \left((I + k f_D)^2 + \frac{f_Z f_D}{\rho^2} \right) (I + k f_D + f_D f_{avg}^{-1})^{-2} \\
C^{NP} &= \frac{1}{2T} \text{tr}(f_D^{-1} V_{-1}^{NP}) - \frac{1}{2}
\end{aligned} \tag{189}$$

J Calibration

$$\bar{f}^{-1} = 0.01$$

$$T = 1000$$

$$f_D(\omega) = \frac{1}{4} \left| 1 - \frac{1}{2} e^{i\omega} \right|^{-2} + 1 - .55 \cos(2\omega) + \frac{7}{16} \left| 1 + \frac{1}{2} e^{i\omega} \right|^{-2}$$

$$f_Z(\omega) = \frac{1}{10} \left| 1 - \frac{1}{2} e^{i\omega} \right|^{-2}$$

$$\rho = 1$$

Figure 1: Optimal information acquisition and waterfilling

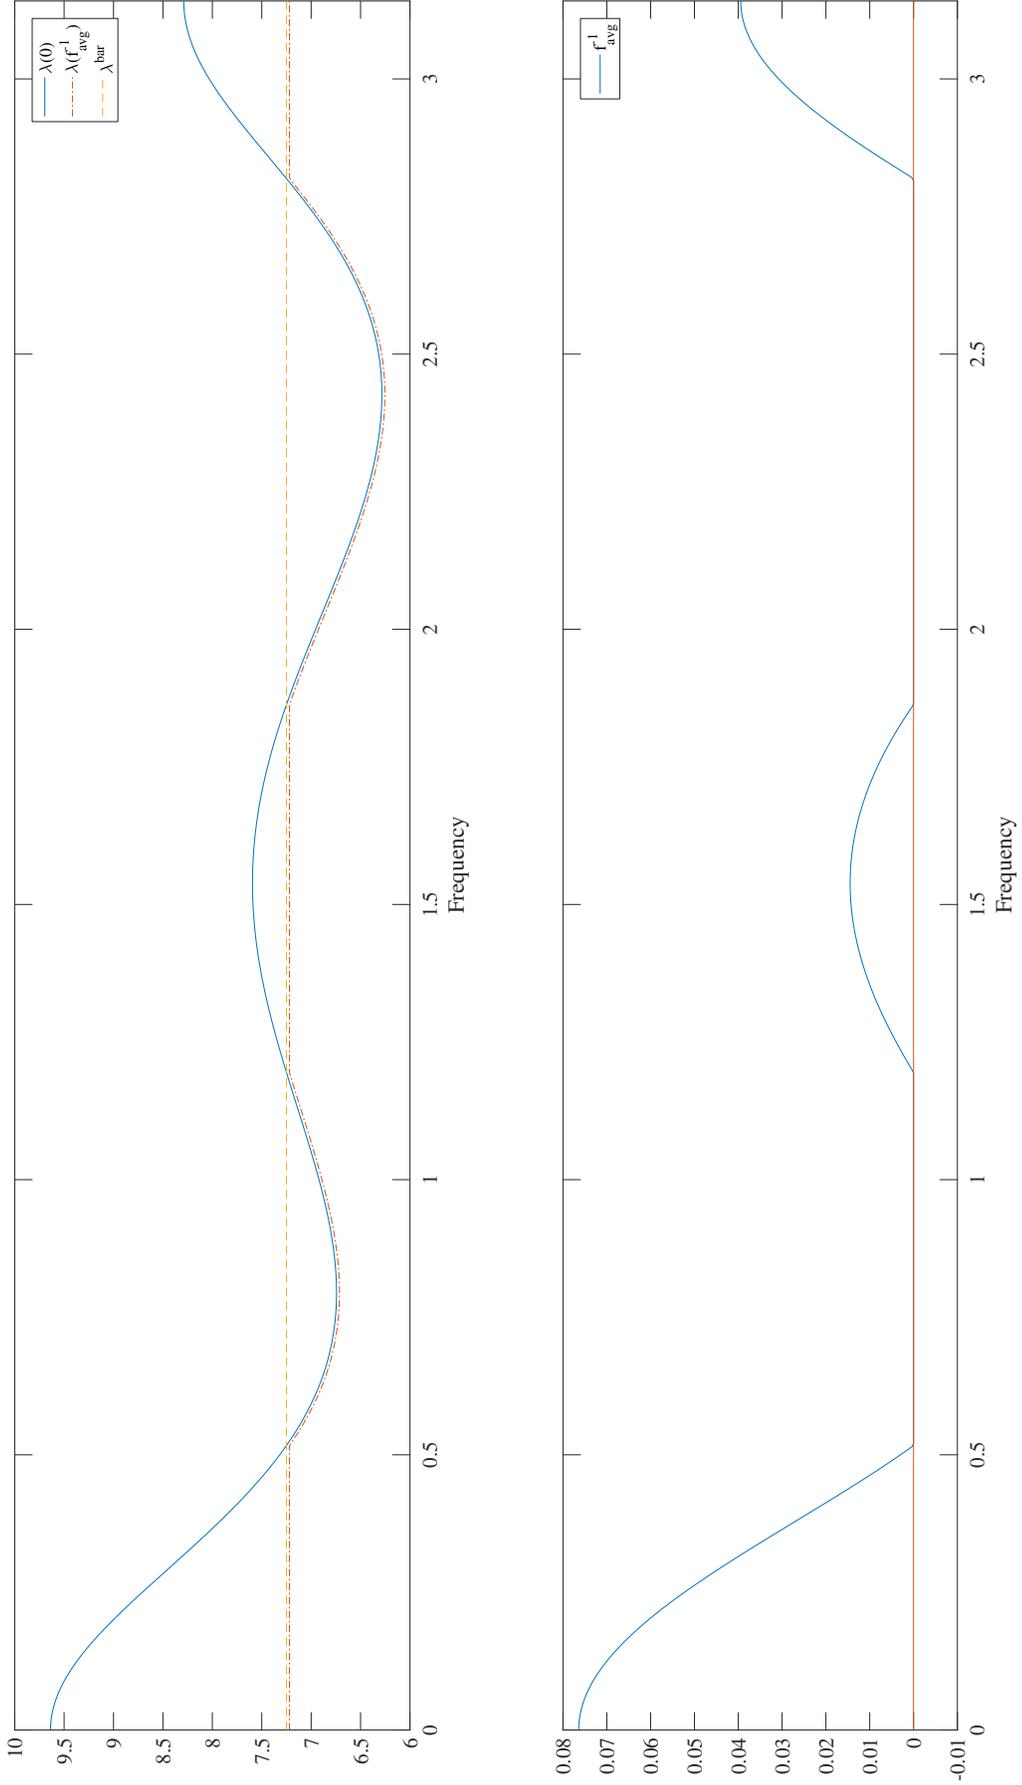


Figure 2: Example investor's demand

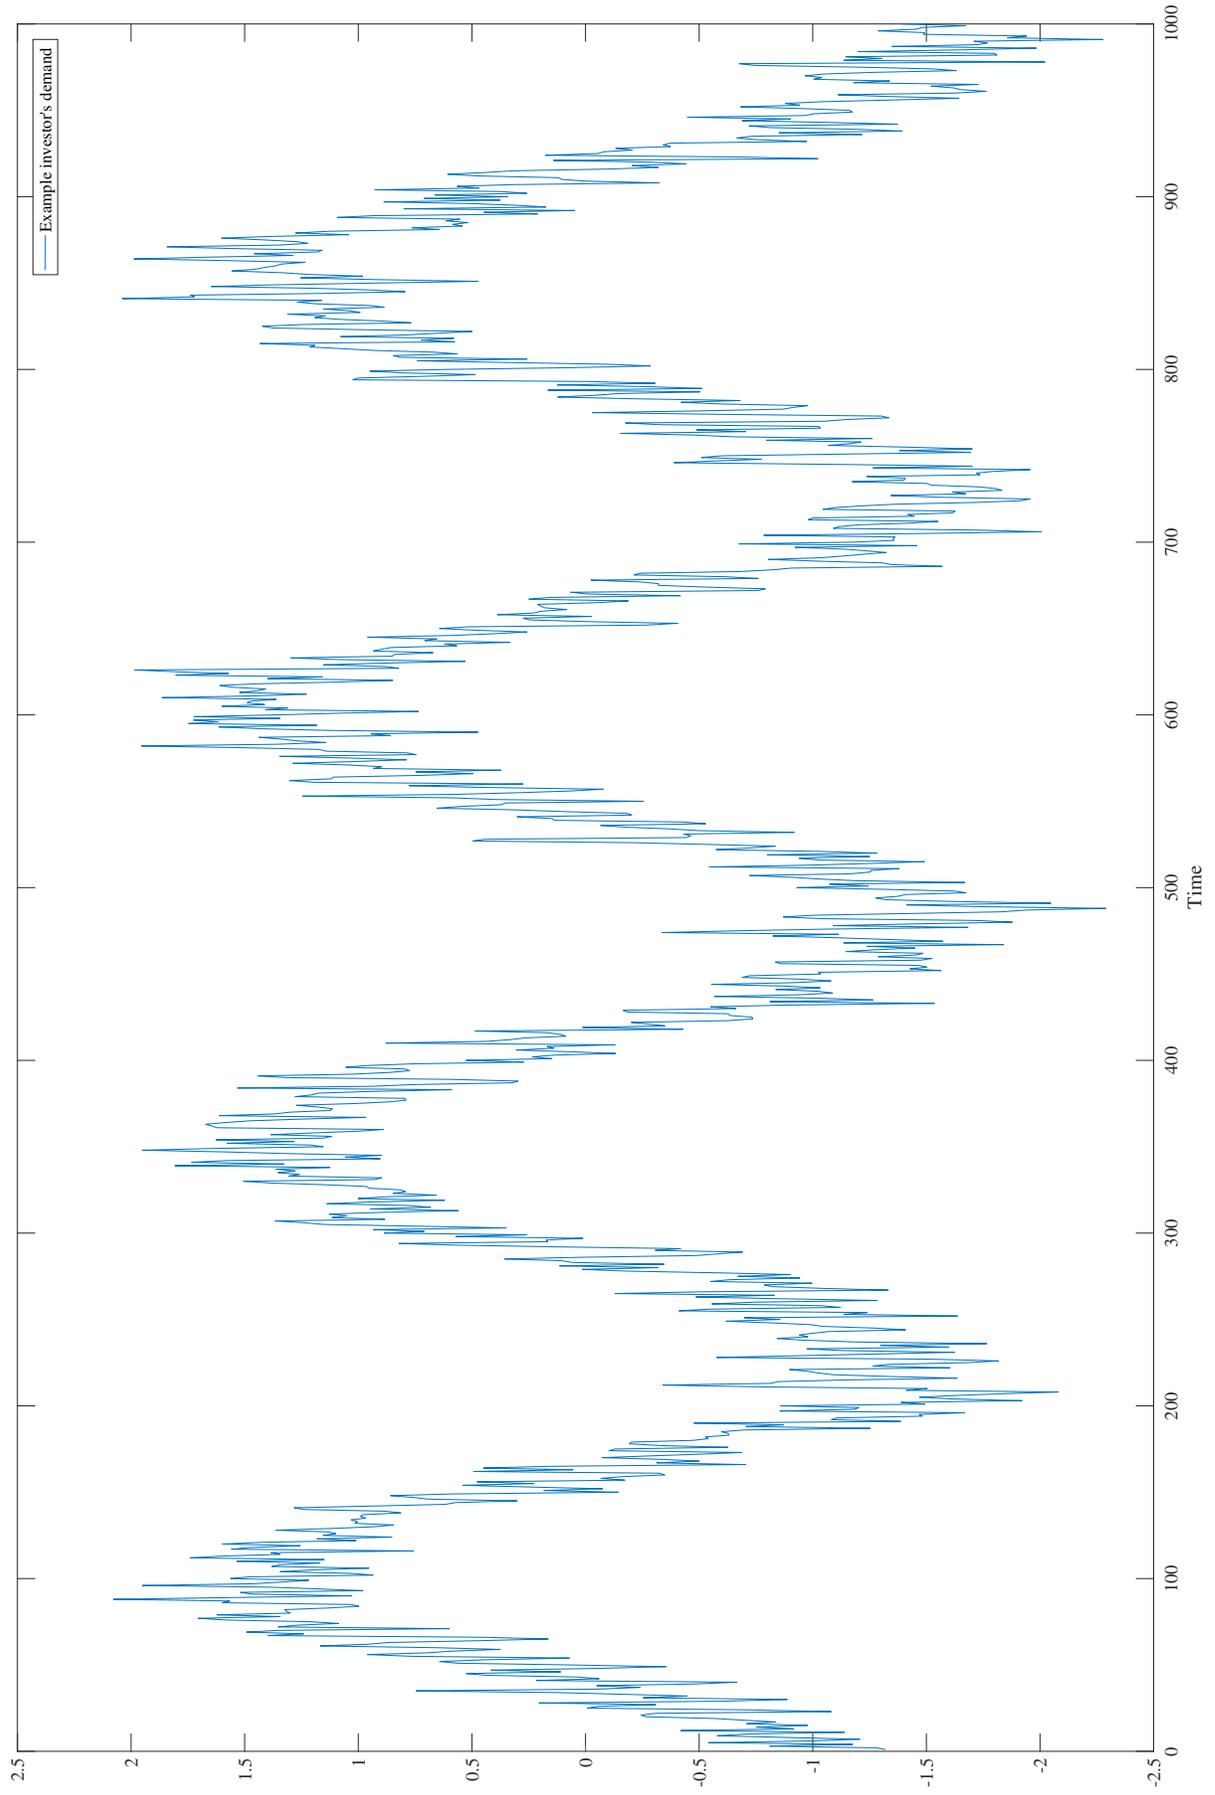


Figure 3: Effects of restricting high-frequency trade

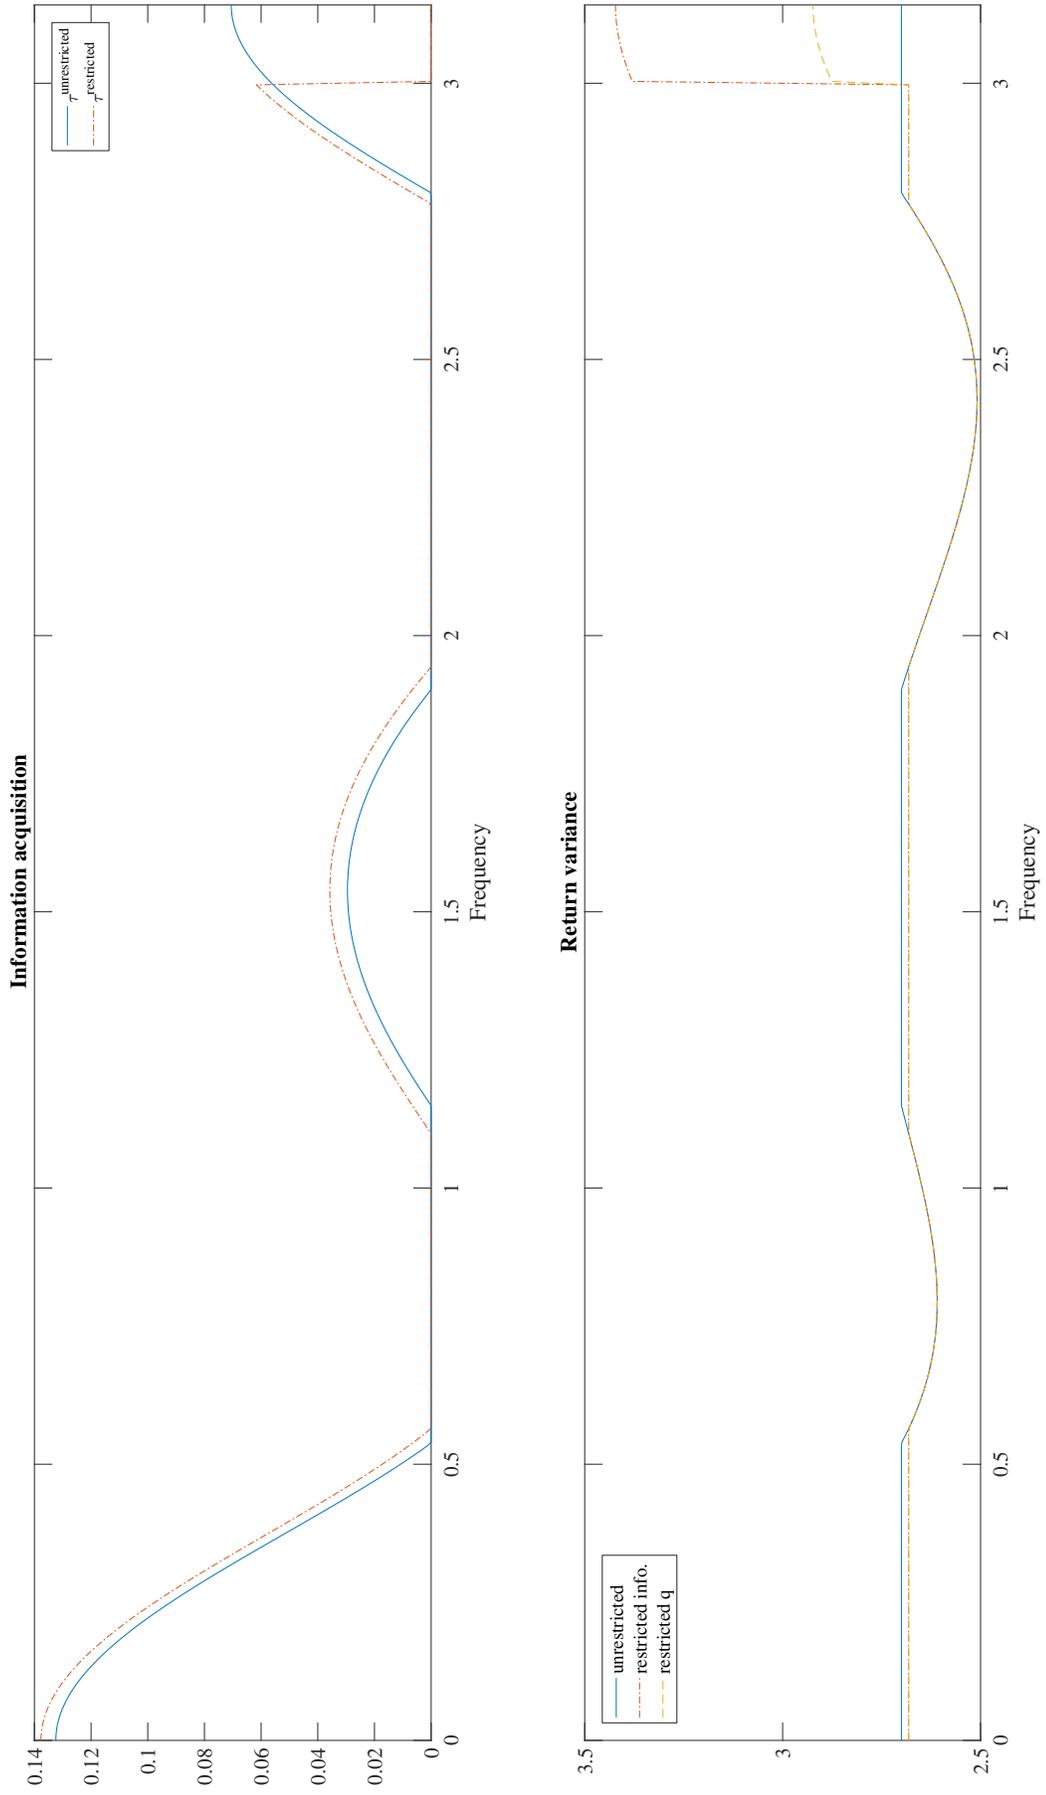


Figure 4: Weights across frequencies and time

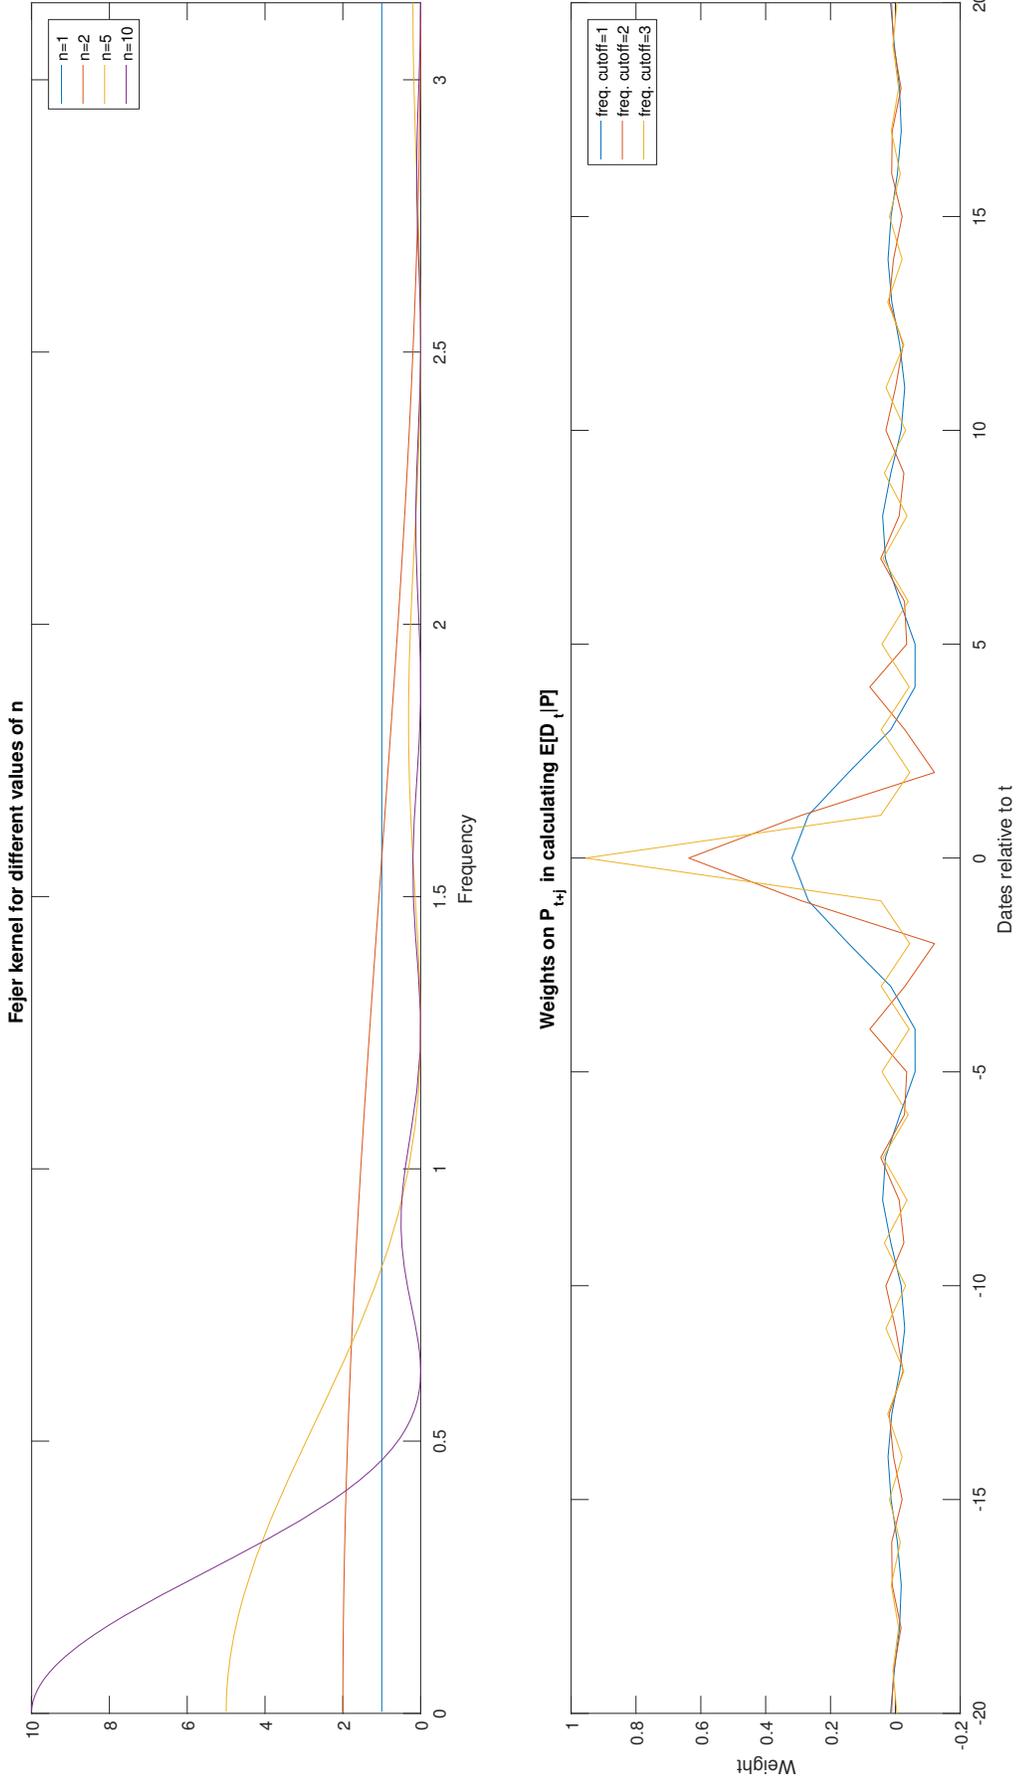


Figure 5: Persistence of the churn rate over time

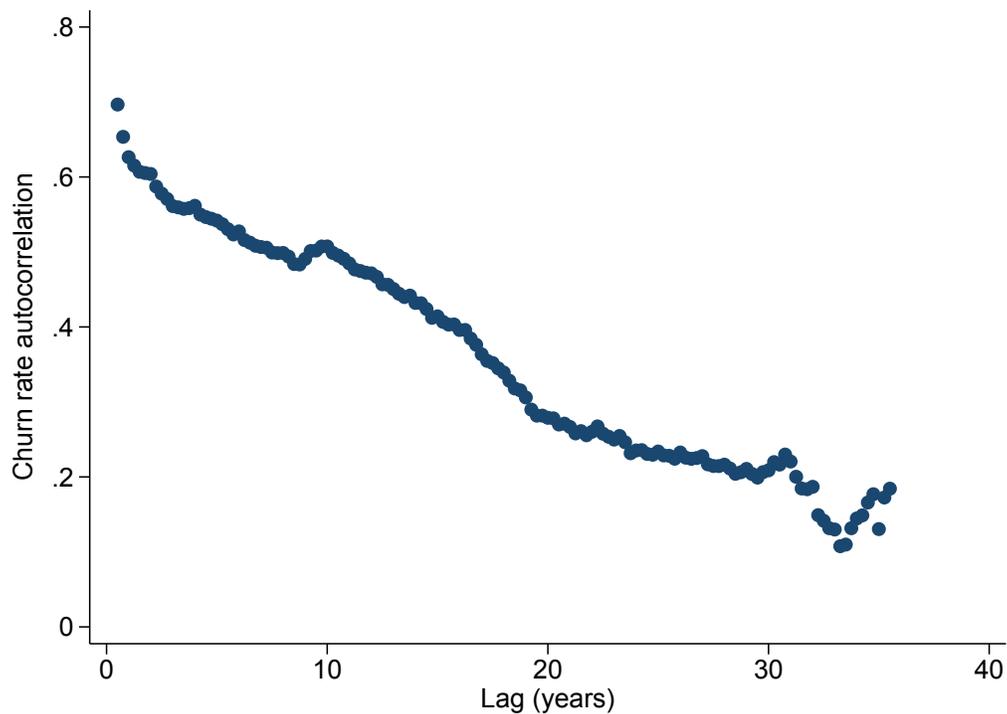


Figure 6: Out-performance of institution holdings at different horizons

