On the effects of restricting high-frequency investment*

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

March 8, 2018

Abstract

This paper studies the effects of restrictions on high-frequency investment for price informativeness and the profits and utility of low-frequency investors. We examine a variant of the standard noisy rational expectations framework in which both the exposures of investors and their information about future fundamentals endogenously vary across future dates. In this environment, precluding investors from holding portfolios with exposures to fundamentals that change at high frequency has zero effect on the information embedded in prices about lower-frequency variation in fundamentals. Furthermore, while the entry of high-frequency *investors* reduces the profits of low-frequency investors, restricting high-frequency *investment* in response only makes the problem worse.

The goal of this paper is to understand the effects of restrictions on investment strategies at particular frequencies. There are frequent discussions of whether the rise of high-frequency trade has hurt other investors, and policymakers often cite the goal of discouraging "short-termism" in favor of buy-and-hold investors. There have therefore been various proposals to directly limit trade at the highest frequencies (e.g. the sub-second batch auction mechanism of Budish, Cramton, and Shim (2015)), to tax transactions (i.e. the Tobin (1978) tax), and also to discourage portfolio turnover at frequencies on the order of a year, e.g. through capital gains taxation or changes in corporate voting structure (such as the Long-Term Stock Exchange[1]). While there is some recent work on the consequences of various limits on information gathering ability,[2] and there have been empirical analyses of high-frequency traders,[3] we are not aware of any other work that directly

*Crouzet: Northwestern University. Dew-Becker: Northwestern University and NBER. Nathanson: Northwestern University. An earlier version of this paper was circulated under the title "Multi-frequency trade". We appreciate helpful comments from Bradyn Breon-Drish, Alex Chinco, Stijn Van Nieuwerburgh, Ioanid Rosu, Laura Veldkamp, and seminar participants at Northwestern, the Adam Smith conference, the Finance Cavalcade, and the Red Rock conference.

[1] See LTSE.org and Osipovich and Berman (2017).

[2] See Banerjee and Green (2015), Goldstein and Yang (2015), Dávila and Parlatore (2016), and Farboodi and Veldkamp (2017). Farboodi and Veldkamp (2017) study the effect of changes in information acquisition technologies over time. Biais, Foucault, and Moinas (2014) examine a setting with limits on traders' ability to search for prices across exchanges and to acquire fundamental information.

[3] E.g. Brogaard, Hendershott, and Riordan (2014), Hendershott and Riordan (2013) and Zhang (2010).

studies the effects of restrictions on high- and low-frequency strategies on price informativeness and investor profits.[4]

The question of what effects investment restrictions have strikes us as being of clear objective importance given the interest of investors and policy makers in such restrictions. Moreover, the answer is not obvious. One view is that there might be some sort of separation across frequencies, so that restrictions in one realm do not affect outcomes in another. On the other hand, investors obviously interact (they trade with each other), so it would be surprising if policies targeting a particular type of investor did not act to benefit others. What we find is a mix of the two: market characteristics at high frequencies can affect the profits and utility of low-frequency investors, but they do not affect low-frequency price informativeness.

In order to analyze the effects of frequency-based restrictions on investment strategies, we require a model with two key characteristics: investors must have a meaningful choice about types of information to acquire, and they must also be able to choose among investment strategies with exposures that differ across horizons or frequencies.[5] The paper develops a noisy rational expectations equilibrium model that has precisely those characteristics and remains highly tractable.

Investors trade, on a single date, claims on future values of a dividend process. These can be thought of as equity or dividend futures.[6] Variation in a portfolio's weights across maturities represents variation in the exposure that the portfolio has to fundamentals at different horizons. Investors are also able to acquire information about the future realizations of fundamentals across horizons. All trade happens on a single date, so the model is not fully dynamic, but it has the two features that we desire: exposures and information acquisition choices that may both vary across horizons.

We use the futures market equilibrium to study the effect of restrictions on investment policies on price informativeness and investor profits.[7] A natural constraint that might be imposed on a portfolio manager, either by their investors or by a regulator, is a restriction on how rapidly their exposures can vary across dates. At one extreme are index funds, which are forced to have essentially fixed exposures. Towards the other extreme are trading desks, which are sometimes required to have risk exposures of zero at the end of each trading day, but may still have risk exposure during the day (e.g. Brock and Kleidon (1992) and Menkveld (2013)). More concretely related to our setting with futures markets, some managers are restricted from holding, for example, exposure to "level"

[4]Cartea and Penalva (2012) are perhaps closest. They study a model in which there are exogenously given high- and low-frequency traders and examine how the high-frequency traders affect the prices received by the others.

[5]Shleifer and Vishny (1990) provide an early analysis of a choice by traders to focus on short- versus long-term projects, while Goldstein and Yang (2015) provide a recent study with a choice of different types of signals to learn about. Farboodi and Veldkamp (2017) study a model in which investors can learn whether to learn about fundamentals or demand.

[6]While the futures market is used as a theoretical abstraction, we note that there are growing markets for such claims; see Binsbergen, Brandt, and Koijen (2012) and Binsbergen and Koijen (2017)

[7]A technical contribution of the paper is to show that the type of futures market we study has a general solution that applies when fundamentals follow arbitrary ARMA processes and that can be obtained by hand. For other related work on frequency transformations, see Bandi and Tamoni (2014), Bernhardt, Seiler, and Taub (2010), Chinco and Ye (2017), Chaudhuri and Lo (2016), Dew-Becker and Giglio (2016), and Kasa, Walker, and Whiteman (2013).

or "slope" factors. In the context of our stylized model, the specific investment restrictions say that investors may not hold portfolios of futures whose weights vary across maturities at specified frequencies. If they cannot have fluctuations at higher then daily frequency, then they can effectively only choose portfolio allocations once per day. The restriction thus can also be interpreted as similar to the imposition of an infrequent batch auction mechanism as in Budish, Cramton, and Shim (2015).

Our first main result is that price informativeness and liquidity are reduced at frequencies targeted by such a policy, but not at any others. The model features endogenous information acquisition, as in Goldstein and Yang (2015) and Kacperczyk, Van Nieuwerburgh, and Veldkamp (2016), and at the restricted frequencies, investors have no incentive to acquire information, making prices completely uninformative. But investors continue to obtain information at unrestricted frequencies, meaning that those prices remain equally informative as without the policy.

In the time domain, the consequence of any restriction on trade is to make prices less informative about fundamentals on each individual date. The effects differ, though, for sums of fundamentals over time. As a specific example, consider a policy that discourages investors from holding portfolios with exposures that change within a day. We show that such a policy increases mean-reverting intraday noise in prices. Inference for moving averages of prices, though, such as the average over a day or a week, is inhibited less since the transitory noise in prices averages out.

We view the model in this paper as a neoclassical benchmark. It does not feature frictions in trade such as limit order books, so it should not be interpreted as addressing the type of high-frequency trade that takes advantage of such frictions including arbitrage across exchanges. The model includes a limited form of irrationality (on the part of noise traders), but further irrationality would add more distortions to prices, as could differences in technology across investors. What our analysis shows is that in order for restrictions on investment at some range of frequencies to have effects at on prices others, one must add to the model another friction, or argue that the dynamics that are not modeled here cause the results to change.

All of that said, even without those frictions, the imposition of restrictions on investment policies can still have an impact on investor profits and utility. Our second key finding has to do with how low-frequency investors are affected by the presence of high-frequency investors and by policies restricting high-frequency investment. While the entrance of high-frequency investors reduces the profits and utility of low-frequency investors, restricting high-frequency investment *does not reverse that decline*; in fact, it exacerbates it.

In the model, investors whose positions are driven primarily by low-frequency fluctuations – i.e. buy-and-hold type investors – are those who have low-frequency information about fundamentals. However, even those investors have some amount of transitory variation in their portfolios. An investor who has good information about the long-term value of a stock should be willing to provide liquidity in the short-run, meaning that they will typically have positions that vary continuously, according to variation in exogenous demand (e.g. due to liquidity needs or sentiment shocks). When high-frequency investors enter the market, they are relatively better at providing such liquidity since they have information about the high-frequency features of fundamentals, so they reduce the profits of low-frequency investors, which is the first part of the result.

But the liquidity provision of low-frequency investors is not reduced to zero, which is the source of the second part of the result. When high-frequency investment is outlawed, it does not help the low-frequency investors; in fact it hurts them by completely eliminating their ability to profit from short-term liquidity provision. The only policy that can restore their lost profits is to eliminate the high-frequency *investors*, rather than high-frequency investment itself. Taxes on changes in positions over time (the analog in the model to taxes on trading) have similar effects.

Finally, the fact that both the entry of high-frequency investors and restrictions on high-frequency trade reduce the profits of low-frequency investors makes it all the more surprising that neither of those events reduces the informativeness of prices at low frequencies. That fact again follows from the basic separability of the model across frequencies: while changes in markets at high-frequencies will affect potentially all investors, they do not change incentives for information acquisition at low frequencies, thus leaving price informativeness unaffected.

Our work is broadly related to a recent literature studying the impact of high-frequency trade. Baldauf and Mollner (2017) provide a model in which the entrance of high-frequency traders can hurt informed investors, which is a result similar to what we find, though through a different channel. Whereas in their model high-frequency investors are fundamentally different from other investors, we argue that a nice feature of our setting is that all investors are ex-ante identical, and simply choose to specialize. Biais, Foucault, and Moinas (2015)**?** study a model where traders specialize endogenously, but with a rather different concept of high-frequency trade from what we use. In their setting, high-frequency traders equilibrate prices across trading venues, whereas here they essentially equilibrate prices over time. Similar to us, they examine implications of high-frequency traders for welfare and price informativeness. However, restricting high-frequency trade in their setting almost inevitably reduces price informativeness as the high-frequency traders are also assumed to perfectly observe the true fundamental value of the asset. We do not make that assumption.

Budish, Cramton, and Shim (2015) examine the effect of restricting high-frequency trade, while Biais, Foucault, and Moinas study a tax on speed technology. Both papers argue that such policies can reduce the impact of high-frequency traders, but we suggest a path through which such policies could also actually hurt lower-frequency investors. Hendershott, Jones, and Menkveld (2011) find empirically that the entrance of high-frequency traders can improve price informativeness, which is consistent with our model (though Weller (2017) provides evidence running in the opposite direction).

The most important drawback of the model is that it is not fully dynamic, meaning that we cannot directly study restrictions on trade, but rather study analogous restrictions on investment policies. There is a literature on dynamic trade, but the extant models do not generate our two desired choices for investors of information acquisition and investment exposures that can differ across horizons. The two main difficulties in solving models of dynamic trade are the dynamic portfolio choice problem (which typically does not have closed-form solutions; see Chacko and Viceira (2005) for a discussion) and the infinite-regress problem of Townsend (1983). There is work that has made substantial progress in solving the infinite regress problem, but those models assume that investors have only single-period objectives and they do not allow for a choice of information across horizons.[8] Recent work also examines dynamic models with strategic trade (with similar restrictions regarding horizons), whereas here we study a fully competitive setting in which all investors are price takers.[9] Relative to the highly sophisticated work on dynamic trade available in the literature, our goal is to take a well understood and highly tractable framework that is used in leading models of information acquisition (e.g. Kacperczyk, Van Nieuwerburgh, and Veldkamp (2016)) and study its implications for the effects of investment restrictions.

1 The model

1.1 Market structure

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

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

For our benchmark results, we focus on the case where fundamentals are stationary. Appendix G shows that the results extend naturally to a case in which fundamentals are stationary in their growth rate, rather than their level. We discuss that case further below. Stationarity implies that Σ_D is constant along its diagonals, and we further assume that the eigenvalues of Σ_D are finite and bounded away from zero.

The biggest restriction imposed by the stationarity assumption (whether in levels or differences) is that we are assuming that the distribution of fundamentals is determined entirely by the matrix Σ_D . The model thus does not allow for stochastic volatility or more general changes in the higher moments of D_t over time (though it could handle deterministic changes), nor does it allow for nonlinearities in the time series dependence of D. The fact that we study the level (or change) in fundamentals, rather than their log, is also a restriction, though one that is generally shared by CARA–normal specifications (e.g. Grossman and Stiglitz (1980)).

There is a set of futures claims on realizations of the fundamental. When we say that the model features a choice of investment across dates, we mean that investors will choose portfolio allocations across the futures contracts, which then yield exposures to the realization of fundamentals on

[8]See Makarov and Rytchkov (2012), Kasa, Walker, and Whiteman (2013), and Rondina and Walker (2017).

[9]For models of dynamic strategic trade, see Vayanos (1999, 2001), Ostrovsky (2012), Banerjee and Breon-Drisch (2016), Foucault, Hombert, and Rosu (2016), Du and Zhu (2017), and Dugast and Foucault (2017).

different dates in the future.

A concrete example of a 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 could also be the dividend on a stock, in which case the futures would be claims on dividends on individual dates. The analysis of futures is an abstraction for the sake of the theory, though we note that dividend futures are in fact traded (see Binsbergen and Koijen (2017)). While the concept of a futures market on the fundamentals will be a useful analytic tool, we can also obviously price portfolios of futures. Equity, for example, is a claim to the stream of fundamentals over time. Holding any given combination of futures claims on the fundamental is equivalent to holding futures contracts on equity claims.

1.2 Information structure

There is a unit mass of "sophisticated" investors indexed by $i \in [0,1]$. 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

$$Y_{i,t} = D_t + \varepsilon_{i,t},\ \varepsilon_i \sim N\left(0, \Sigma_i\right), \tag{2}$$

where Σ_i^{-1} is investor i's signal precision matrix (which will be chosen endogenously below). Through $Y_{i,t}$, investors can learn about fundamentals on all dates between 1 and T. $\varepsilon_{i,t}$ is a stationary error process in the sense that $cov\left(\varepsilon_{i,t}, \varepsilon_{i,t+j}\right)$ depends on j but not t. That also implies that $var\left(\varepsilon_{i,t}\right)$ is the same for all t, so all dates are equally difficult to learn about. The stationarity assumption is imposed so that no particular date is given special prominence in the model. Investors must choose an information policy that treats all dates symmetrically, and they are not allowed to choose to learn about a single date.

The signal structure generates one of our desired model features, which is that investors can choose to learn about fundamentals across different dates in the future. When the errors are positively correlated across dates, the signals are relatively less useful for forecasting trends in fundamentals since the errors also have persistent trends. Conversely, when errors are negatively correlated across dates, the signals are less useful for forecasting transitory variation and provide more accurate information about moving averages.

1.3 Investment objective

On date 0, there is a market for forward claims on fundamentals on all dates in the future. Investor i's demand for a date-t forward conditional on the set of prices and signals is denoted $Q_{i,t}$. Investors have mean-variance utility over the net present value of excess returns:

$$U_{0,i} = \max_{\{Q_{i,t}\}} E_{0,i}\left[T^{-1}\beta^t Q_{i,t}\left(D_t - P_t\right)\right] - \frac{1}{2}\left(\rho T\right)^{-1} Var_{0,i}\left[\sum_{t=1}^{T}\beta^t Q_{i,t}\left(D_t - P_t\right)\right], \tag{3}$$

where $0 < \beta \leq 1$ is the discount factor, $E_{0,i}$ and $Var_{0,i}$ are the expectation and variance operators conditional on agent i's date-0 information set, $\{P, Y_i\}$, and ρ is risk-bearing capacity per unit of time. We treat all investors as having identical horizons, T. Appendix B shows that the horizon has no effect on information choices in the model.

The key restriction here (beyond those implicit in the mean-variance assumption) is that signals are acquired and trade occurs on date 0. In general settings there is no known closed-form solution to even the partial-equilibrium dynamic portfolio choice problem, let alone to the full market equilibrium.[10] Moreover, allowing agents to obtain signals repeatedly yields a highly nontrivial updating problem. We therefore use a relatively minimal static model. The model nevertheless has the two characteristics that we stated we desire in the introduction: it allows for investment strategies that place different weight on fundamentals on different dates in the future, and it allows investors to make a choice about how precise their signals are for different types of fluctuations in fundamentals.

The time discounting in (3) has the effect of making dates farther in the future less important in the objective of the investors. We therefore define

$$\tilde{Q}_{i,t} \equiv \beta^t Q_{i,t} \tag{4}$$

to be agent i's discounted demand. In what follows, the $\tilde{Q}_{i,t}$ will be stationary processes. That means that $Q_{i,t} = \beta^{-t}\tilde{Q}_{i,t}$ will generally grow in magnitude with maturity t, though only to a relatively small extent for typical values of β and horizons on the order of 10–20 years.

1.4 Noise trader demand

In order to keep prices from being fully revealing, we assume there is uninformed demand from a set of noise traders. The noise traders are investors with the same objective as the sophisticates, but whose expectations are formed differently. Specifically, their expectations of fundamentals depend only on an exogenous prior for fundamentals and a signal, which we denote Z_t (their expectations are unaffected by prices). The signal Z_t is in reality uncorrelated with fundamentals, so it can be viewed as a type of sentiment shock.

Appendix A shows that when the noise traders maximize an objective of the form of (3) but with their incorrect expectations, then their demand, denoted N_t, can be written as

$$\tilde{N}_t = Z_t - kP_t, \tag{5}$$

$$\text{where } \tilde{N}_t \equiv \beta^t N_t. \tag{6}$$

Z_t depends on the signals the noise traders receive (which are assumed to be common across them)

[10] Frequency-domain solutions to the infinite regress problem, such as Kasa, Walker, and Whiteman (2013) and Makarov and Rytchkov (2012), restrict preferences to depend on wealth one period ahead in order to avoid the dynamic portfolio problem.

and k is a coefficient determining the sensitivity of noise trader demand to prices, which depends on their risk aversion and how precise they believe their signals to be. In principle, N_t can depend on prices on all dates (depending on the structure of priors and signals), but we restrict attention to the case where N_t depends only on P_t for the sake of simplicity.

In the benchmark case where D_t is stationary in levels, we assume that Z_t is also stationary in levels – the noise traders have a signal technology with the same stationarity properties as that of the sophisticates – which yields a useful symmetry between fundamentals, supply, and the signals, in that they are all assumed to be stationary processes. That symmetry is why we use this particular formalization of noise trader demand.

1.5 Asset market equilibrium in the time domain

We begin by solving for the market equilibrium on date 0 that takes the agents' signal precisions, Σ_i^{-1}, as given. The Σ_i^{-1} are chosen on date -1, and that optimization is discussed below.

Definition 1 *For any given set of individual precisions $\{\Sigma_i\}_{i\in[0,1]}$, a date-0 asset market equilibrium is a set of demand functions, $\{Q_i(P, Y_i)\}_{i\in[0,1]}$, and a price vector P, such that investors maximize utility and all markets clear: $\int_i Q_{i,t}di + N_t = 0$ for all $t \geq 1$.*

Investors submit demand curves for each futures contract to a Walrasian auctioneer who selects equilibrium prices to clear all markets.

The structure of the time-0 equilibrium is mathematically that of Admati (1985), who studies investment in a cross-section of assets, and the solution from that paper applies directly here (with the minor difference that supply is also a function of prices). Here we are considering investment across a set of futures contracts that represent claims on some fundamentals process across different dates. The Admati (1985) solution is:

$$P = A_1 D + A_2 Z, \tag{7}$$

$$A_1 \equiv I - \left(\rho^2 \Sigma_{avg}^{-1} \Sigma_Z^{-1} \Sigma_{avg}^{-1} + \Sigma_{avg}^{-1} + \Sigma_D^{-1} + \rho^{-1} k\right)^{-1} \left(\rho^{-1} k + \Sigma_D^{-1}\right), \tag{8}$$

$$A_2 \equiv \rho^{-1} A_1 \Sigma_{avg}^{-1}, \tag{9}$$

$$\text{where } \Sigma_{avg}^{-1} \equiv \int_i \Sigma_i^{-1} di. \tag{10}$$

As Admati (1985) discusses, this equilibrium is not particularly illuminating since standard intuitions, including the idea that increases in demand or decreases in supply should raise prices, do not hold. Prices of futures maturing on any particular date depend on fundamentals and demand for all other maturities except in knife-edge cases. Interpreting the equilibrium requires interpreting complicated products of matrix inverses. The following section shows that the equilibrium can be solved nearly exactly by hand when it is rewritten in terms of frequencies.

2 Frequency domain interpretation

2.1 Frequency portfolios

The basic feature of the model that makes it difficult to interpret is that fundamentals, noise trader demand, and signal errors are all correlated across dates. For any one of those three processes, we could always use a standard orthogonal (eigen-) decomposition to yield a set of independent components. But in general there is no reason to expect that three time series with different correlation properties across dates would have the same orthogonal decomposition (in general they do not). A central result from time series analysis, though, is that a particular frequency transform asymptotically orthogonalizes all standard stationary time series processes.

The way to think about the transformation is that it involves simply analyzing the prices of particular portfolios of futures instead of the futures themselves. The first requirement is that the transformation should be full rank, in the sense that the set of portfolios allows an investor to obtain the same payoffs as the futures themselves. Second, the transformed portfolios should be independent of each other. And third, since we are studying trade at different frequencies, it would be nice if the portfolios also had a frequency interpretation.

Obviously there are many different conceptions of fluctuations at different frequencies. One might imagine step functions switching between +1 and -1 at different rates. For reasons we will see below, it turns out that using sines and cosines will be most natural in our setting. So the portfolios that we study – representing investor exposures – vary smoothly over time in the form $\cos(\omega t)$ and $\sin(\omega t)$.

Formally, the portfolio weights are represented as vectors of the form

$$c_h \equiv \sqrt{\frac{2}{T}}\Big(\cos\left(\omega_h\left(t-1\right)\right)\Big)_{t=1}^{T}, \tag{11}$$

$$s_h \equiv \sqrt{\frac{2}{T}}\Big(\sin\left(\omega_h\left(t-1\right)\right)\Big)_{t=1}^{T}, \tag{12}$$

$$\text{where} \quad \omega_h \equiv 2\pi h/T, \tag{13}$$

for different values of the integer $h \in \{0, 1, ..., T/2\}$. c_0 is the lowest frequency portfolio, with the same weight on all dates, while $c_{\frac{T}{2}}$ is the highest frequency, with weights switching each period between $+/-1$.

Figure 1 plots the weights for a pair of those portfolios. The x-axis represents dates and the y-axis is the weight of the portfolio on each date. The weights vary smoothly over time, with the rate at which they change signs depending on the frequency ω.

Economically, the basic idea is to think about the investment problem as being one of choosing exposure to different types of fluctuations in fundamentals. A long-term investor can be thought of as one whose exposure to fundamentals changes little over time. A high-frequency investor, on the other hand, holds a portfolio whose weights change more frequently and by larger amounts. Our claim is that studying the frequency portfolios is more natural than studying individual futures claims. Investors do not typically acquire exposure to fundamentals on only a single date. Rather, they have exposures on multiple dates, and the portfolios we study are one way to express that. While investors will also obviously not hold a portfolio that takes the exact form of a cosine, any portfolio can be expressed as a sum of cyclical components. An investor whose portfolio loadings change frequently will have a portfolio whose weights are relatively larger on the high-frequency components, which figure 1 shows generate rapid changes in loadings.

2.2 Properties of the frequency transformation

The portfolio weights can be combined into a matrix, Λ, which implements the frequency transformation.

$$\Lambda \equiv \left[\frac{1}{\sqrt{2}} c_0, c_1, s_1, c_2, s_2, ..., c_{\frac{T}{2}-1}, s_{\frac{T}{2}-1}, \frac{1}{\sqrt{2}} c_{\frac{T}{2}} \right] \tag{14}$$

(s_0 and $s_{T/2}$ do not appear since they are identically equal to zero; the $1/\sqrt{2}$ scaling for c_0 and $c_{T/2}$ gives them the same norms as the other vectors).

We use lower-case letters to denote frequency-domain objects. So whereas $\tilde{Q}_t$ is investor i's vector of discounted allocations to the various futures, $\tilde{q}_i$ is their vector of discounted allocations to the frequency portfolios, with

$$\tilde{Q}_i = \Lambda \tilde{q}_i. \tag{15}$$

Rows of Λ represent portfolio weights on different dates and columns represent different frequency portfolios. q_i is then the vector of investor i's allocations to the various frequency portfolios.

In what follows, we use the index $j = 1, ..., T$ to identify each column of Λ, or equivalently, each frequency-domain vector. The jth column of Λ contains a vector that fluctuates at frequency $\omega_{\lfloor \frac{j}{2} \rfloor} = 2\pi \left\lfloor \frac{j}{2} \right\rfloor / T$, where $\lfloor \cdot \rfloor$ is the integer floor operator.[11] So there are two vectors, a sine and a cosine, for each characteristic frequency, with the exceptions of $j = 1$ (frequency 0, the lowest possible) and $j = T$ (frequency $\frac{T}{2}$, the highest possible).

Note also that Λ has the property that $\Lambda^{-1} = \Lambda'$, so that frequency-domain vectors can be obtained through

$$\tilde{q}_i = \Lambda' \tilde{Q}_i. \tag{16}$$

In the same way that q_i represents weights on frequency-specific portfolios, $d \equiv \Lambda' D$ is a representation of the realization of fundamentals written in terms of frequencies instead of time. The first element of d, for example, is proportional to the realized sample mean of D. Equivalently, d is the set of regression coefficients of D on the columns of Λ (which generate an R^2 of 1).

As a simple example, consider the case with $T = 2$. The low-frequency component of dividends is then $d_0 = (D_1 + D_2)/\sqrt{2}$ and the high-frequency component of is $d_1 = (D_1 - D_2)/\sqrt{2}$. Investors trade the low-frequency component d_0 by buying an equal amount of the claims on D_1 and D_2.

[11] $\lfloor x \rfloor$ is the largest integer that is less than or equal to x.

Conversely, investors trade the high-frequency component d_1 by buying offsetting amounts of the claims on D_1 and D_2.

The most important feature of the frequency transformation is that it approximately diagonalizes the variance matrices. We now formalize that idea.

Definition 2 *For an $n \times n$ matrix A with elements $a_{l,m}$, the* ***weak matrix norm*** *is*

$$|A| \equiv \left(\frac{1}{n} \sum_{l=1}^{n} \sum_{m=1}^{n} a_{l,m}^2 \right)^{1/2}. \tag{17}$$

If $|A - B|$ is small, then the elements of A and B are close in mean square.

The frequency transform will lead us to study the spectral densities of the various time series:

Definition 3 *The* ***spectrum*** *at frequency ω of a stationary time series X_t is*

$$f_X(\omega) \equiv \sigma_{X,0} + 2 \sum_{t=1}^{\infty} \cos(\omega t)\, \sigma_{X,t}, \tag{18}$$

$$\text{where } \sigma_{X,t} = cov(X_s, X_{s-t}). \tag{19}$$

The spectrum, or spectral density, is used widely in time series analysis. The usual interpretation is that it represents a variance decomposition. $f_X(\omega)$ measures the part of the variance of X_t associated with fluctuations at frequency ω, which is formalized as follows.

Lemma 1 *For any stationary time series $\{X_t\}_{t=1}^{T}$, with frequency representation $x \equiv \Lambda' X$, the elements of the vector x are approximately uncorrelated in the sense that the covariance matrix of x, $\Sigma_x \equiv \Lambda' \Sigma_X \Lambda$, is nearly diagonal,*

$$|\Sigma_x - diag(f_X)| \leq b T^{-1/2}, \tag{20}$$

and x converges in distribution to

$$x \to_d N(0, diag(f_X)), \tag{21}$$

where b is a constant that depends on the autocorrelations of X,[12] *and $diag(f_X)$ denotes a matrix with the vector $\left\{ f_X\left(\omega_{\lfloor j/2 \rfloor}\right)\right\}_{j=1}^{T}$ on the main diagonal and zeros elsewhere.*[13]

Proof. These are textbook results (e.g. Brockwell and Davis (1991) and Gray (2006)). Appendix C.1 provides a derivation of the inequality (20) specific to our case. The convergence in distribution follows from Brillinger (1981), theorem 4.4.1. ■

[12] Specifically, $b = 4\left(\sum_{j=1}^{\infty} |j \sigma_{X,j}|\right)$.

[13] A requirement of this lemma, which we impose for all the stationary processes studied in the paper, is that the autocovariances are summable in the sense that $\sum_{j=1}^{\infty} |j\sigma_{X,j}|$ is finite (which holds for finite-order stationary ARMA processes, for example). Trigonometric transforms of stationary time series converge in distribution under more general conditions, though. See Shumway and Stoffer (2011), Brillinger (1981), and Shao and Wu (2007).

Lemma 1 says that Λ approximately diagonalizes *all* stationary covariance matrices. So the frequency-specific components of fundamentals, prices, and noise trader demand are all (approximately) independent when analyzed in terms of frequencies. That is, $d = \Lambda' D$, $y_i = \Lambda' Y_i$, and $z = \Lambda' Z$ all have asymptotically diagonal variance matrices. That independence will substantially simplify our analysis, and it is a special property of the sines and cosines, as opposed to other conceptions of frequencies.[14]

2.3 Market equilibrium in the frequency domain

2.3.1 Approximate diagonalization

Instead of solving jointly for the prices of all futures, the approximate diagonalization result allows us to solve a series of parallel scalar problems, one for each frequency. Intuitively, since the frequency-specific portfolios have returns that are uncorrelated with each other, the investors' utility can be written as a sum of mean-variance optimizations

$$U_{0,i} \approx \max_{\{q_{i,j}\}} T^{-1} \sum_{j=1}^{T} \left\{ E_{0,i}\left[\tilde{q}_{i,j}\left(d_j - p_j\right)\right] - \frac{1}{2}\rho^{-1} Var_{0,i}\left[\tilde{q}_{i,j}\left(d_j - p_j\right)\right] \right\}. \tag{22}$$

In what follows, we solve the model using the approximation for $U_{0,i}$, and then show that it converges to the true solution from Admati (1985). When utility is completely separable across frequencies, there is an equilibrium frequency-by-frequency:

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

$$p_j = a_{1,j} d_j + a_{2,j} z_j \tag{23}$$

$$a_{1,j} \equiv 1 - \frac{\rho^{-1} k + f_{D,j}^{-1}}{\left(\rho f_{avg,j}^{-1}\right)^2 f_{Z,j}^{-1} + f_{avg,j}^{-1} + f_{D,j}^{-1} + \rho^{-1} k} \tag{24}$$

$$a_{2,j} \equiv \frac{a_{1,j}}{\rho f_{avg,j}^{-1}} \tag{25}$$

where $f_{avg,j}^{-1} \equiv \int_i f_{i,j}^{-1} di$ is the average precision of the agents' signals at frequency j.

Proof. See appendix C.2. ■

The price of the frequency-j portfolio depends only on fundamentals and supply at that frequency due to the independence across frequencies. As usual, the informativeness of prices, through

[14] Finally, it is should be noted that infill asymptotics, where T grows by making the length of a time period shorter, are not sufficient for lemma 1 to hold. What is important is essentially that T is large relative to the range of autocorrelation of the process X. So, for example, if fundamentals have nontrivial autocorrelations over a horizon of a year, then it is important that T be substantially larger than a year. Van Binsbergen and Koijen (2017), for example, examine data on dividend futures with maturities as long as 16 years.

$a_{1,j}$, is increasing in the precision of the signals that investors obtain, while the impact of noise trader demand on prices is decreasing in signal precision and risk tolerance.

These solutions for the prices are the standard results for scalar markets. What is different here is simply that the agents chose exposures across frequencies, rather than across dates; p_j is the price of a portfolio whose exposure to fundamentals fluctuates over time at frequency $2\pi \lfloor j/2 \rfloor /T$. Both prices and demands at frequency j depend only on signals and supply at frequency j – the problem is completely separable across frequencies.

In what follows, we assume that k is sufficiently small that $ka_{2,j} < 1$ for all j, which simply ensures that z represents a positive demand shock in equilibrium (though most of the results hold without that assumption). The restriction is that noise trader demand not be too sensitive to prices; in the literature k is usually equal to zero.

2.3.2 Quality of the approximation

While solution 1 is an approximation, its error can be bounded. The time domain solution is obtained from the frequency domain solution by premultiplying by Λ (from equation (15)), and we have,

Proposition 1 *The difference between solution 1 and the exact Admati (1985) solution is small in the sense that*

$$\left|A_1 - \Lambda diag\left(a_1\right)\Lambda'\right| \leq c_1 T^{-1/2} \tag{26}$$

$$\left|A_2 - \Lambda diag\left(a_2\right)\Lambda'\right| \leq c_2 T^{-1/2} \tag{27}$$

for constants c_1 and c_2. Furthermore, the variances of the approximation error for prices and quantities are bounded by:

$$\left|Var\left(\Lambda p - P\right)\right| \leq c_P T^{-1/2} \tag{28}$$

$$\left|Var\left(\Lambda \tilde{q}_i - \tilde{Q}_i\right)\right| \leq c_Q T^{-1/2} \tag{29}$$

for some constants c_P and c_Q.

Proof. See appendix C.3. ■

Proposition 1 shows that the frequency domain solution to the market equilibrium provides a close approximation to the true solution, in the sense that the solution in (23), once it is rotated back to the time domain, converges to equations (7-9). Moreover, Λp is stochastically close to P in the sense that the variance of the pricing errors is of order $T^{-1/2}$. So 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 for large T.

2.4 Optimal information choice in the frequency domain

The analysis so far takes the precision of the signals as fixed. Following Van Nieuwerburgh and Veldkamp (2009) and Kacperczyk, Van Nieuwerburgh, and Veldkamp (2016; KVNV), we now allow investors to choose their signal precisions, Σ_i^{-1} to maximize the expectation of their mean-variance objective (3) subject to an information cost,

$$\max_{\{f_{i,j}\}} E_{-1}\left[U_{i,0} \mid \Sigma_i^{-1}\right] - \frac{\psi}{2T} tr\left(\Sigma_i^{-1}\right), \tag{30}$$

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), and ψ is the per-period cost of information. Total information here is measured by the trace operator $tr\left(\Sigma_i^{-1}\right)$.[15]

Given the optimal demands, an agent's expected utility is linear in the precision they obtain at each frequency.

Lemma 2 *Each informed investor's expected utility at time -1 may be written as a function of their own signal precisions, $f_{i,j}^{-1}$, and the average across other investors, $f_{avg,j}^{-1} \equiv \int_i f_{i,j}^{-1} di$, with*

$$E_{-1}\left[U_{0,i} \mid \{f_{i,j}\}\right] = \frac{1}{2T}\sum_{j=1}^{T} \lambda_j\left(f_{avg,j}^{-1}\right) f_{i,j}^{-1} + constant, \tag{31}$$

where the constant does not depend on investor i's precision.

Proof. See appendix C.4. ■

$\lambda_j(x)$ 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 $x \geq 0$. It is possible to show that $\lambda_j\left(f_{avg,j}^{-1}\right) = Var\left[d_j - p_j\right]$.[16]

Since expected utility and the information cost are both linear in the set of precisions that agent i chooses, $\left\{f_{i,j}^{-1}\right\}$, it immediately follows that agents purchase signals at whatever subset of frequencies has $\lambda_j\left(f_{avg,j}^{-1}\right) \geq \psi$.

Solution 2 *Information is allocated so that*

$$f_{avg,j}^{-1} = \begin{cases} \lambda_j^{-1}(\psi) & \textit{if } \lambda_j(0) \geq \psi, \\ 0 & \textit{otherwise.} \end{cases} \tag{32}$$

Because attention cannot be negative, when $\lambda_j(0) \leq \psi$, no attention is allocated to frequency j. Otherwise, attention is allocated so that its marginal benefit and its marginal cost are equated.

[15] KVNV show that the results here are robust to various perturbations of the assumptions: (1) rather than using the trace operator, information can be measured through the entropy of the signals; (2) investors can be given a fixed budget of information rather than a fixed cost; (3) it can be made costly for investors to pay attention to prices in addition to their signals.

[16] These results are established as part of the proof of lemma 2, in appendix C.4.

Note, though, that this result does not pin down precisely how any specific investor's attention is allocated; this class of models, with a non-convex information cost, only determines the aggregate allocation of attention across frequencies. For the purposes of studying price informativeness, though, characterizing this aggregate allocation is all that is necessary.

3 The consequences of restricting investment frequencies for prices

This section focuses on the effects on prices of restrictions on the frequencies at which investment strategies can operate. There are many real-world examples of such restrictions. Some institutional investors face constraints on the speed at which they can change their portfolio weights. For example, a pension fund or endowment might have a policy portfolio that it targets, the weights of which are only updated on an annual or quarterly basis at board meetings. Other investors have restrictions that keep them from holding positions for too long. Market makers and trading desks may have policies restricting their positions to net to zero at the end of each day (e.g. Brock and Kleidon (1992) and Menkveld (2013)).

Those constraints on portfolio managers are in a sense imposed by their own investors. Regulators may also impose restrictions on the types of strategies that investors may undertake. 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 portfolio turnover at the monthly or annual level. The US tax code, for example, encourages holding assets for at least a year through the higher tax rates on short-term capital gains.[17]

Section 5 also shows that a tax on changes in positions over time (approximately, trading) affects most strongly high-frequency investment strategies. So the restrictions in this section are also similar to imposing a quadratic trading tax.

3.1 Restricting investment frequencies

The assumption in this section is that investors are restricted to setting $\tilde{q}_{i,j} = 0$ for $j \in \mathcal{R}$. We leave the noise traders unconstrained, assuming they are perhaps less regulated (like retail investors, in many ways), or that their demand is induced by a rotating set of people, with no individual necessarily trading at high frequency.

Intuitively, if an investor is restricted from exposures at frequencies higher than a day, then they can effectively only choose exposures once per day. Rather than forcing the investor to literally only trade once a day, though, the restriction in our case corresponds to a portfolio that varies smoothly between days. So (approximately) if the investor can choose daily exposures, then their actual exposures, minute-by-minute, might be represented by a spline that smooths between the daily exposures.

[17] 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)).

More formally, a restriction on trading frequencies reduces the degrees of freedom that an investor has in making choices. Suppose we had a model where each time period is an hour, and T is a year, or 1625 trading hours. A restriction that investors cannot invest at a frequency higher than a day (6.5 hours) would mean that they would go from a strategy with 1625 degrees of freedom to one with only 250. A pension that sets a portfolio once a quarter would have only four degrees of freedom. In that sense, then, the restrictions we analyze in this section are similar to a shift from a continuous market to one with infrequent batch auctions, as in Budish, Cramton, and Shim (2015). While that paper proposes holding the auctions still very frequently (i.e. more than once per second), a more aggressive restriction could have auctions only once per day, or once per hour.

Appendix G examines the version of the model in which fundamentals are stationary in differences instead of levels (i.e. they have a unit root). In that case, the analysis in fact goes through nearly identically – frequency restrictions still represent decreases in the degrees of freedom available to investors – but with a single small change: the lowest frequency portfolio, rather than being one that puts equal weight on fundamentals on all dates, puts weight on fundamentals only on the final date, T. Intuitively, an investor who wants to take a position in long-run growth rates does that by buying a claim just to the level on date T. On the other hand, an investor who holds a portfolio that loads on rapid changes in the growth rate of fundamentals will have a portfolio with weights on the level of fundamentals that also change quickly. So in that case, the example of restricting investment at frequencies higher than a day continues to impose the same limit on the set of strategies investors can choose from.

Derivations of the results in the remainder of this section can be found in appendix D.

3.2 Results

We begin by describing price informativeness at different frequencies to demonstrate our key separation result. We then show what happens to prices of standard claims in the time domain.

3.2.1 Price informativeness across frequencies

In terms of frequencies, we obtain a complete separation: prices become uninformative at restricted frequencies, while remaining unaffected at unrestricted frequencies.

Result 1 *When trade by sophisticated investors is restricted at a set of frequencies $\mathcal{R}$, prices satisfy*

$$p_j = \left\{ \begin{array}{c} k^{-1} z_j \text{ for } j \in \mathcal{R} \\ a_{1,j} d_j + a_{2,j} z_j \text{ otherwise} \end{array} \right\}, \tag{33}$$

where a_1 and a_2 are the same as those defined in solution 1.

Intuitively, when sophisticated traders are restricted, prices depend only on sentiment, since the people with information cannot express their opinions. Moreover, the market becomes illiquid, and it is cleared purely through prices rather than quantities.

Since the solution for information acquisition at a frequency j does not depend on anything about any other frequency, the information acquired at a frequency $j \notin \mathcal{R}$ is also unaffected by the policy. We then have the result that:

Corollary 1.1 *When investors are restricted from holding portfolios with weights that fluctuate at some set of frequencies $j \in \mathcal{R}$, then prices at those frequencies, p_j, become completely uninformative about dividends. The informativeness of prices for $j \notin \mathcal{R}$ about dividends is unchanged. More formally, $Var\left[d_j \mid p_j\right]$ for $j \notin \mathcal{R}$ is unaffected by the restriction. For $j \in \mathcal{R}$, $Var\left[d_j \mid p_j\right] = Var\left[d_j\right]$.*

3.2.2 Price informativeness across dates

The fact that prices remain equally informative at some frequencies does not mean that they remain equally informative for any particular *date*. They are linked through

$$Var\left(D_t \mid P\right) = \frac{1}{T}\sum_{j=1}^{T} Var\left[d_j \mid p_j\right]. \tag{34}$$

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.[18] So when uncertainty rises at some set of frequencies, the informativeness of prices for fundamentals on every date falls by an equal amount.

Corollary 1.2 *Investment restrictions reduce price informativeness for fundamentals on all dates by equal amounts, and by an amount that weakly increases with the number of frequencies that are restricted.*

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 due to a restriction on high-frequency investment, a natural response would be to examine an average of fundamentals and prices over time (across maturities of futures contracts).

Corollary 1.3 *Under the asymptotic result for variance matrices, the informativeness of prices for the sum of fundamentals depends only on informativeness at the lowest frequency:*

$$Var\left(T^{-1}\sum_{t=1}^{T} D_t \mid P\right) = Var\left[T^{-1/2} d_0 \mid p_0\right], \tag{35}$$

where d_0 is the lowest frequency portfolio – with equal weight each date – and p_0 is its price.

Result 1.3 follows immediately from the definition of d_0 and the independence across frequencies in the solution. It shows that the informativeness of prices for moving averages of fundamentals depends only on the very lowest frequency. So even if prices have little or no information at high

[18] This result is proven in appendix 4.2.

frequencies – $Var[d_j \mid p_j]$ is high 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).

More concretely, going back to our example of oil futures, when investors are not allowed to use high-frequency investment strategies, 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. But 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 under restrictions on high-frequency strategies. It is possible to derive a similar result for shorter moving averages; in that case the weights on the frequencies are given by the Fejér kernel.

In the case where fundamentals are stationary in terms of growth rates instead of levels, the results in this section also hold, but replacing D_t by its first difference. In particular, result 1.3 then states that $Var(D_T \mid P)$ becomes arbitrarily close to the variance of the lowest frequency portfolio. This is unsurprising since, as we had previously noted, in the difference-stationary case, the lowest frequency portfolio is the one that places weight only on D_T. In that case, the prediction of the model is that $Var(D_T \mid P)$ is unaffected by restrictions on high-frequency investment.

When low-frequency investment strategies are restricted, on the other hand, as in the case of a trading desk that cannot have exposure to cycles lasting longer than a day, then it is natural to examine the informativeness of *differences* in prices across dates. As an example, we can consider the variance of the first difference of fundamentals.

Corollary 1.4 *The variance of an estimate of the change in fundamentals across dates conditional on observing the vector of prices is*

$$Var[D_t - D_{t-1} \mid P] = \sum_{j=1}^{T} 2\left(1 - \cos\left(\omega_{\lfloor j/2 \rfloor}\right)\right) Var[d_j \mid p_j]. \tag{36}$$

The function $2 - 2\cos(\omega)$ is equal to 0 at $\omega = 0$ and rises smoothly to 4 at the highest frequency, $\omega = \pi$. So period-by-period changes in fundamentals are driven primarily by high-frequency variation. Reductions in price informativeness at low frequencies then have relatively large effects on moving averages and small effects on changes, while the reverse is true for reductions in informativeness at high frequencies.

To summarize, any restriction on investment reduces price informativeness for any particular date. But when high-frequency investment is restricted, there is little change in the behavior of moving averages of prices. 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 week or a month, then the model predicts that there need not be adverse consequences.

3.2.3 Return volatility

Corollary 1.5 *Given an information policy $f_{avg,j}^{-1}$, the variance of returns at frequency j, $r_j \equiv d_j - p_j$ is*

$$Var\left(r_j\right) = \begin{cases} f_{D,j} + \frac{f_{Z,j}}{k^2} \text{ for } j \in \mathcal{R} \\ \min\left(\psi, \lambda_j\left(0\right)\right) \text{ otherwise} \end{cases}. \tag{37}$$

Moreover, the variance of returns at restricted frequencies satisfies $Var(r_j) > f_{D,j} + \frac{f_{Z,j}}{(k+\rho f_{D,j}^{-1})^2}$, which is the variance that returns would have at the same frequency if investment were unrestricted but agents were uninformed.

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 uninformed active investors have risk-bearing capacity ($\rho > 0$), they absorb some of the exogenous demand by simply trading against prices, buying when prices are below their means and selling when they are above. The greater is the risk-bearing capacity, the smaller is the effect of sentiment volatility on return volatility. Thus, the restriction affects return volatilities through its effects on both liquidity provision and information acquisition.

Restricting sophisticated investors from following high-frequency strategies 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 (though those fluctuations in prices are, literally, variations in prices across maturities for different futures contracts on date 0). Sophisticated traders typically play a role of smoothing prices across maturities, essentially intermediating between excess demand in one minute and excess supply in the next. When they are restricted from holding positions in futures that fluctuate from minute to minute, they can no longer provide that intermediation service, and volatility at high frequencies increases.

Finally, we note that the results in this section could be extended fairly easily to account for more general types of restrictions, such as placing restrictions only on the trade of a subset of agents, or perhaps bounding the size of the positions of some agents at certain frequencies.

4 Investor outcomes

This section studies how restricting high-frequency trade affects low-frequency investors. We obtain two main results, which initially appear to be in conflict:

1. The entrance of high-frequency investors reduces the profits of low-frequency investors.

2. Restricting high-frequency investment reduces the profits and utility of low-frequency investors.

So while low-frequency investors are worse off when high-frequency investors enter the market, cutting off high-frequency trade neither restores the old equilibrium, nor does it make the low-frequency investors better off.

4.1 Who are high- and low-frequency investors?

We define a high-frequency investor as one whose portfolio is driven relatively more by high-frequency fluctuations, while a low-frequency investor holds a portfolio that is driven relatively more by low-frequency fluctuations. That definition can be formalized by a variance decomposition, using the fact that

$$Var\left(\tilde{Q}_{i,t}\right) = \sum_{j=1}^{T} Var\left(\tilde{q}_{i,j}\right). \tag{38}$$

Furthermore, the component of the variance of $\tilde{Q}_{i,t}$ that is driven by fluctuations at frequency j is increasing in the precision of the signals agent i acquires at frequency j:

$$\frac{d}{df_{i,j}^{-1}}\left[Var\left(\tilde{q}_{i,j}\right)\right] > 0. \tag{39}$$

So if two investors have the same total variance of their positions, $Var\left(\tilde{Q}_{1,t}\right) = Var\left(\tilde{Q}_{2,t}\right)$, but one of them has higher-precision signals at high frequencies, i.e. $f_{1,j}^{-1} > f_{2,j}^{-1}$ for j above some cutoff, then variation in that investor's position is driven relatively more by high-frequency components.

(39) shows that $Var\left(\tilde{q}_{i,j}\right)$ is increasing in the precision of the signals that agent i receives. When an investor has more precise signals at a given frequency, they trade more aggressively for two reasons. First, since their signals are more precise, their demand is more sensitive to their own signals. Second, the quality of their signals also means that they can worry less about adverse selection, so they trade more strongly to accommodate demand shocks from noise traders.

For two investors with positions that have the same unconditional variance, the high-frequency investor – whose fluctuations happen relatively faster – is the one with relatively more precise signals about the high-frequency features of fundamentals. That is, *high-frequency investors have high-frequency information, and low-frequency investors have low-frequency information.* As an extreme case, we will take high-frequency investors as people whose signals have positive precision only for j above some cutoff j_{HF}, and low-frequency investors have signals with positive precision only for j below some j_{LF} with $j_{HF} > j_{LF}$.

4.2 Investor profits and utility

Result 2 *Let $R = D - P$ be the vector of returns in the time domain; investor i's average discounted profits are*

$$E_{-1}\left[\tilde{Q}_i' R\right] = \sum_{j=1}^{T} (1 - ka_2)\left(-E_{-1}\left[z_j r_j\right]\right) + ka_1 E_{-1}\left[r_j d_j\right] + \rho\left(f_{i,j}^{-1} - f_{avg,j}^{-1}\right) Var_{-1}\left[r_j\right] \tag{40}$$

and expected profits at each frequency are nonnegative,

$$E_{-1}\left[\tilde{q}_{i,j} r_j\right] \geq 0 \text{ for all } i, j \tag{41}$$

with equality only if $f_{i,j} = 0$ *and* $f_{D,j}^{-1} = \rho f_{avg,j}^{-1} f_{Z,j}^{-1} k$ *(i.e. in a knife-edge case).*

Each investor's expected discounted profits depend on three terms. The first represents the profits earned from noise traders. $E[z_j r_j] = -a_2 f_Z^{-1} < 0$ since the sophisticated investors imperfectly accommodate their demand. When the noise traders have high demand (that is, when z is high), they drive prices up and expected returns down. The sophisticated investors earn profits from trading with that demand.

The second term represents the profits that the informed investors earn by buying from the noise traders when they have positive signals on average. The coefficient $ka_{1,j}$ represents the slope of the supply curve that the informed investors face.

Finally, the third term in (40) represents a reallocation of profits from the less to the more informed sophisticated investors. An investor who has highly precise signals about fundamentals at frequency j can accurately distinguish periods when prices are high due to strong fundamentals to those when prices are high due to high sentiment. That allows them to earn relatively more profits on average than an uninformed investor.

That said, an uninformed sophisticated investor does not earn negative expected profits at any frequency, even with $f_{i,j}^{-1} = 0$. There are always, except in a knife-edge case, profits to be earned by trading with noise traders.

Result 2 therefore gives us two key insights. First, all investors, no matter their information, have the ability to earn profits at all frequencies through liquidity provision. Second, all else equal, investors who are informed about a particular frequency earn the most money from investing at these frequencies. High-frequency investors – those with relatively more information about high-frequency fundamentals – earn relatively higher returns at high frequencies, while low-frequency investors earn relatively higher returns at low frequencies.

4.2.1 The entrance of high-frequency investors

The two main results of this section follow from result 2. First, consider a scenario in which $f_{avg,j}^{-1} = 0$ for high frequencies (i.e. for all j greater than j_{HF}). That is, there are initially no high-frequency investors, perhaps because an unmodeled cost of acquiring high-frequency information is prohibitively large. Existing investors may trade at the frequencies $j > j_{HF}$, but in an uninformed manner. What is the effect of the initial entry of high-frequency investors, i.e. a marginal increase in $f_{avg,j}^{-1}$, holding all other parameters fixed?

Result 3 *Starting from* $f_{avg,j}^{-1} = 0$ *for* $j > j_{HF}$, *an increase in* $f_{avg,j}^{-1}$ *at one of those frequencies* *reduces profits and utility of an investor for whom $f_{i,j}^{-1}$ remains unchanged. Specifically,*

$$\left.\frac{d}{df_{avg,j}^{-1}}E_{-1}\left[\tilde{q}_{LF,j}r_j\right]\right|_{f_{avg,j}=0} < 0 \quad (42)$$

$$\left.\frac{d}{df_{avg,j}^{-1}}E_{-1}\left[\sum_t \tilde{Q}_{LF,t}\left(D_t - P_t\right)\right]\right|_{f_{avg,j}=0} < 0 \quad (43)$$

$$\left.\frac{d}{df_{avg,j}^{-1}}E_{-1}\left[U_{LF,0}\right]\right|_{f_{avg,j}=0} < 0 \quad (44)$$

where the LF subscripts denote positions and utility of a low-frequency investor. Concretely, in an economy populated only by low-frequency investors, the entrance of high-frequency investors increases $f_{avg,j}^{-1}$ for $j > j_{HF}$ and therefore reduces the expected profits at all frequencies, total expected profits, and the utility of low-frequency investors.

The source of that result is the fact that investors with low-frequency information may still trade at high frequencies. Suppose, for example, that not only does $f_{LF,j}^{-1} = 0$ for high j, but also that $f_{avg,j}^{-1}$ does also – nobody has high-frequency information. In that setting obviously any sophisticated investor will be happy to accommodate transitory fluctuations in noise trader demand. More concretely, an investor who has information that the long-term value of a stock is \$50 will be willing to provide liquidity in the short-run, buying when the price is below \$50 and selling when the price is higher. That liquidity provision will have high-frequency components when liquidity demand (here noise trader demand) has high-frequency components (i.e. $f_{z,j} > 0$ for $j > j_{HF}$). That is, if there are short-run variations in sentiment, then there will be short-run variation in the low-frequency investor's position.

The entry of investors with high-frequency information hurts those with low-frequency information because the new investors are better at providing high-frequency liquidity. Result 2 and corollary 3 formalize that idea and shows that how high-frequency investors hurt low-frequency investors – by crowding out their ability to provide liquidity. It is critical to note, though, that result 2 still shows that the entry of high-frequency investors never reduces the profits earned by low-frequency investors to zero, even at high frequencies.

It should also be noted that these results do not change the incentives of low-frequency investors to acquire information at low frequencies. While they lose money from a decrease in liquidity provision at high frequencies, their choices at low frequencies are unaffected, so if one's primary concern is price informativeness at low frequencies, the entry of high-frequency traders will have no effect.

4.2.2 Restoring the previous equilibrium

If the entrance of high-frequency investors hurts the incumbent low-frequency investors, a natural question (at least to the incumbents) might be how to restore the old equilibrium. We consider three possible policies: restricting or eliminating high-frequency investment, taxing trade (or variation in positions), and restricting high-frequency information acquisition.

First, consider a restriction on high-frequency investment that says that no sophisticated investor may set $q_{i,j} \neq 0$ for j above some cutoff, as in the previous section. A concrete example of such a policy would be an infrequent batch auction mechanism (e.g. Budish, Cramton, and Shim (2015)). Restricting trade above the daily frequency would correspond to having an auction once per day. Result 2 shows that such a restriction would, rather than restoring the profits and utility of the low-frequency investors, actually reduce them *further*. The result follows from the fact that restricting trade eliminates the terms in the summation for j above the cutoff, which are all nonnegative. While high-frequency investors make liquidity provision at low frequencies more difficult, outlawing high-frequency trades simply makes it impossible. So eliminating high-frequency investment does not restore the old equilibrium – it actually compounds the effect of the entrance of high-frequency investors.

Imposing a tax on changes in positions, specifically, a tax on $(Q_{i,t} - Q_{i,t-1})^2$, will have similar effects to a restriction on high-frequency investment in that the tax is most costly for high-frequency strategies. The next section provides a more complete derivation of that result. But such a tax will be broadly similar to a blanket restriction on high-frequency investment.

The final policy response would be to somehow limit the acquisition of high-frequency information. In the context of the model, this would represent a restriction on the ability of investors to learn about period-to-period variation in fundamentals. Very loosely, one might think of this as representing a restriction on the ability of investors to receive economic news continuously, and rather force them to receive announcements that are clustered at lower frequencies. For example, news might be released only once per day instead of on a continuous basis (which is already partially implemented through the practice of major announcements coming outside regular equity trading hours).

In the context of the model, a restriction on information acquisition would exactly restore the equilibrium that exists in the absence of the high-frequency investors. Since the low-frequency investors do not acquire high-frequency information, the restriction has no effect on them, while the high-frequency investors can be made essentially irrelevant to the equilibrium by the removal of their information. In terms of result 2, such a restriction would reduce $f_{avg,j}^{-1}$ to zero at high frequencies, thus restoring the profits earned by the low-frequency investors from liquidity provision.

5 Quadratic trading costs

The restriction that investors have *exactly* zero exposure at certain frequencies is a natural one to study in the model. But there are other ways of imposing limits on investors' exposures across frequencies. We now examine the equilibrium when there are quadratic costs of trading. We argue that, relative to the frictionless benchmark, introducing these costs has analogous effects to the more abstract restriction $q_{i,j} = 0$ for $j \in \mathcal{R}$. Changes in trading costs could be caused either by the imposition of a quadratic tax on shares traded, or by changes in the trading technology.

The model obviously does not have trade over time. However, the exposures that investors choose in the futures market can be replicated through a commitment to trade (at a fixed price) the fundamental on future dates. That is, define a date-t equity claim to be an asset that pays dividends equal to the fundamental on each date from $t+1$ to T. Since the futures contracts involve exchanging money only at maturity, the date-t cost of an equity claim is $P_t^{equity} = \sum_{j=1}^{T-t} \beta^{-j} P_{t+j}$. An investor's exposure to fundamentals on date t, $Q_{i,t}$ can be acquired either by buying $Q_{i,t}$ units of forwards on date 0 or by holding $Q_{i,t}^{EQ}$ units of equity entering date t. In the latter case, the volume of trade by investor i would be equal to the change in $Q_{i,t}$ over time. That is, $\Delta Q_{i,t}^{EQ} = \Delta Q_{i,t}$.

We assume that investors now maximize the following objective:

$$U_{0,i} = \max_{\{Q_{i,t}\}} E_{0,i}\left[T^{-1}\sum_{t=1}^{T} Q_{i,t}\left(D_t - P_t\right)\right] - \frac{1}{2}cT^{-1}E_{0,i}\left[QV\left\{Q_i\right\}\right] - \frac{1}{2}bT^{-1}E_{0,i}\left[\sum_{t=1}^{T} Q_{i,t}^2\right], \quad (45)$$

where $b > 0$ is a cost of holding large positions in the assets, $c \geq 0$ is a cost incurred from quadratic variation in positions, with quadratic variation defined as:

$$QV\left\{Q_i\right\} \equiv \left[\sum_{t=2}^{T}\left(Q_{i,t} - Q_{i,t-1}\right)^2 + \left(Q_{i,1} - Q_{i,T}\right)^2\right]. \quad (46)$$

The term involving b in (45) replaces the aversion to variance in the benchmark setting. That change is made for the sake of tractability, but its economic consequences are minimal (see, e.g., Kasa, Walker, and Whiteman (2013)). We also set discount rates to zero here to maintain tractability.

Appendix F shows that:

$$T^{-1}QV\left\{Q_i\right\} = 2\sum_{j=1}^{T}\sin^2\left(\omega_{\lfloor j/2\rfloor}/2\right) q_{i,j}^2. \quad (47)$$

Note that we have defined quadratic variation as the sum of the squared changes in $Q_{i,t}$ between $t = 2$ and T plus $\left(Q_{i,1} - Q_{i,T}\right)^2$. Without the final term, there would be no cost to investors of entering and exiting very large positions at the beginning and end of the investment period. This term helps account for that, and has the added benefit of yielding the simple closed-form expression in the frequency domain reported above. The right-hand side shows that the quadratic variation in the volume induced by an investor depends on their squared exposures at each frequency scaled by $\sin^2\left(\omega_{\lfloor j/2\rfloor}/2\right)$, which rises from 0 to 1 as j rises. Intuitively, when $c > 0$, holding exposure to higher frequency fluctuations in fundamentals is more costly because it requires more frequent portfolio rebalancing.

The equilibrium of the model is described in detail in Appendix F. Here, we highlight key results and explain how they relate to the previous results on restricting trade frequencies.

Result 4 *When $c > 0$, all else equal, investors' equilibrium signal precision is higher at lower frequencies.*

With the assumption of fixed quadratic trading costs, the marginal benefit of increasing precision at frequency j is given by:

$$\frac{1}{2}(c\sin^2\left(\omega_{\lfloor j/2\rfloor}/2\right)+b)^{-1}Var\left[d_j \mid p_j, y_{i,j}\right]^2. \tag{48}$$

In particular, it is declining with both the signal precision and the frequency of exposure. Given that the marginal cost of information is the same across frequencies, investors choose higher signal precisions at lower frequencies, all else equal.

The main result regarding the effect of the quadratic trading cost is the following.

Result 5 *A small increase in trading costs, when starting from zero, reduces information acquisition at all frequencies except frequency* 0*. The effect is larger at higher frequencies. As a corollary, the effect of an increase in trading costs on price informativeness is weaker at longer horizons.*

The first part of this result suggests that if the goal is to reduce high-frequency trade, then a quadratic tax is a more blunt instrument than placing an explicit restriction on trade at the targeted frequencies. A tax on volume affects all investors, regardless of the strategy that they follow. However, the second part of the result suggests that trading costs affect the highest frequencies most strongly. The quadratic cost thus leads, endogenously, to the same changes in information acquisition studied in the main model; namely, the variance of dividends conditional on prices, $Var(d_j|p_j)$, falls more at higher frequencies. The corollary regarding price informativeness refers to the fact that the variance of moving averages of the form:

$$Var\left(\frac{1}{n}\sum_{m=0}^{n-1} D_{t+m} \mid P\right) \tag{49}$$

increases less as a result of the increase in trading costs for longer horizons n. In the extreme case of $n = T$, which corresponds to the frequency 0 component of the signals, the increase in trading costs has in fact *no* effect on equilibrium signal precision and thus price informativeness. This can be seen from the expression for the marginal benefit of signal precision above, which is independent of c when $j = 0$.

Thus, overall, the message of the model with quadratic costs is consistent with the previous analysis. Increasing trading costs leads to less informed trading and the effect is tilted toward high frequencies; at lower frequencies, information acquisition decisions are less impacted. As a result, the effect of the increase on the informativeness of prices for fundamentals at long horizons is limited.

6 Conclusion

The aim of this paper is to understand how regulations that restrict the types of strategies that investors may pursue affect price informativeness and investor profits and utility. We are specifically interested in regulations that affect the speed with which investors may turn over their positions. In order to study that question, we need to have a setting in which investors can make meaningful decisions about investment strategies and in which they have an endogenous information choice. We develop a simple rotation of the standard noisy rational expectations equilibrium that incorporates trade and information acquisition in a futures market.

Our first key result is that such a policy has precisely zero effect on the informativeness of prices or the profitability of trading at the untargeted frequencies. This result is a natural consequence of the independence of the problem across frequencies. Another important byproduct of this independence is that restrictions on high-frequency investment have a diminishing impact on price informativeness as the forecast horizon increases.

Second, we show that while the entry of high-frequency *investors* reduces the utility and profits of low-frequency investors, restricting high-frequency *investment* in response to that entry does not make low-frequency investors better off. A buy-and-hold investor is able to provide the market short-term liquidity – a person with a price target of \$50 should be willing to accommodate transitory demand shocks that drive the price above their target. High-frequency investors are better at such liquidity provision; this is why their entry makes buy-and-hold investors worse off. But eliminating all high-frequency *investment* does not solve the problem. In fact, it makes it worse, by eliminating entirely rents from liquidity provision for all investors. These findings make the result that low-frequency price informativeness is unaffected by high-frequency restrictions all the more surprising – those restrictions hurt low-frequency investors, but do not change their incentives for information acquisition.

Our results on restricting high-frequency investment are thus mixed. Implementing these restrictions does not change how informative equilibrium prices are about the slow-moving components of fundamentals. However, these restrictions affect the profits that any investor (high-frequency or buy-and-hold) can earn from liquidity provision. So, while these restrictions might get rid of high-frequency investors, they do not restore the status-quo for buy-and-hold types. In fact, they make them worse off by erasing any previously earned rents from short-term liquidity provision.

References

Auxier, Richard, Len Burman, Jim Nunns, Ben Page, and Jeff Rohaly, "An Updated Analysis of Hillary Clinton's Tax Proposals," October 2016.

Baldauf, Markus and Joshua Mollner, "High-Frequency Trading and Market Performance," 2017. Working paper.

Bandi, Federico and Andrea Tamoni, "Business-cycle consumption risk and asset prices," *Working Paper*, 2014.

Banerjee, Snehal and Brett Green, "Signal or noise? Uncertainty and learning about whether other traders are informed," *Journal of Financial Economics*, 2015, *117* (2), 398–423.

Bernhardt, Dan, Peter Seiler, and Bart Taub, "Speculative dynamics," *Economic Theory*, 2010, *44* (1), 1–52.

Binsbergen, Jules H Van and Ralph SJ Koijen, "The term structure of returns: Facts and theory," *Journal of Financial Economics*, 2017, *124* (1), 1–21.

Brillinger, David R., *Time Series: Data Analysis and Theory*, McGraw Hill, 1981.

Brock, William A and Allan W Kleidon, "Periodic market closure and trading volume: A model of intraday bids and asks," *Journal of Economic Dynamics and Control*, 1992, *16* (3), 451–489.

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.

Budish, Eric, Peter Cramton, and John Shim, "The high-frequency trading arms race: Frequent batch auctions as a market design response," *The Quarterly Journal of Economics*, 2015, *130* (4), 1547–1621.

Cartea, Álvaro and José Penalva, "Where is the value in high frequency trading?," *The Quarterly Journal of Finance*, 2012, *2* (03), 1250014.

Chacko, George and Luis M Viceira, "Dynamic consumption and portfolio choice with stochastic volatility in incomplete markets," *The Review of Financial Studies*, 2005, *18* (4), 1369–1402.

Chaudhuri, Shomesh E. and Andrew W. Lo, "Spectral Portfolio Theory," 2016. Manuscript, Massachusetts Institute of Technology.

Chinco, Alex and Mao Ye, "Investment-Horizon Spillovers," 2017. Manuscript, University of Illinois.

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.

Dew-Becker, Ian and Stefano Giglio, "Asset Pricing in the Frequency Domain: Theory and Empirics," *Review of Financial Studies*, 2016. Forthcoming.

Du, Songzi and Haoxiang Zhu, "What is the Optimal Trading Frequency in Financial Markets?," *The Review of Economic Studies*, 2017, p. rdx006.

Dugast, Jerome and Thierry Foucault, "Data Abundance and Asset Price Informativeness," 2017. Working paper.

Farboodi, Maryam and Laura Veldkamp, "Long Run Growth of Financial Technology," 2017. Manuscript, New York University and Princeton University.

Foucault, Thierry, Johan Hombert, and Ioanid Roşu, "News trading and speed," *The Journal of Finance*, 2016, *71* (1), 335–382.

Goldstein, Itay and Liyan Yang, "Information diversity and complementarities in trading and information acquisition," *The Journal of Finance*, 2015, *70* (4), 1723–1765.

Hendershott, Terrence and Ryan Riordan, "Algorithmic trading and the market for liquidity," *Journal of Financial and Quantitative Analysis*, 2013, *48* (4), 1001–1024.

—— **, Charles M Jones, and Albert J Menkveld**, "Does algorithmic trading improve liquidity?," *The Journal of Finance*, 2011, *66* (1), 1–33.

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, "A rational theory of mutual funds' attention allocation," *Econometrica*, 2016, *84* (2), 571–626.

Kasa, Kenneth, Todd B Walker, and Charles H Whiteman, "Heterogeneous beliefs and tests of present value models," *Review of Economic Studies*, 2013, *81* (3), 1137–1163.

Makarov, Igor and Oleg Rytchkov, "Forecasting the forecasts of others: Implications for asset pricing," *Journal of Economic Theory*, 2012, *147* (3), 941–966.

Menkveld, Albert J, "High frequency trading and the new market makers," *Journal of Financial Markets*, 2013, *16* (4), 712–740.

Osipovich, Alexander and Dennis K. Berman, "Silicon Valley Vs. Wall Street: Can the New Long-Term Stock Exchange Disrupt Capitalism?," *Wall Street Journal*, 2017, *October 17.*

Ostrovsky, Michael, "Information aggregation in dynamic markets with strategic traders," *Econometrica*, 2012, *80* (6), 2595–2647.

Rondina, Giacomo and Todd Walker, "Confounding Dynamics," 2017. Working paper.

Shao, Xiaofeng and Wei Biao Wu, "Asymptotic Spectral Theory for Nonlinear Time Series," *The Annals of Statistics*, 2007, *35* (4), 1773–1801.

Shleifer, Andrei and Robert Vishny, "Equilibrium Short Horizons of Investor and Firms," *American Economic Review, Papers and Proceedings*, 1990, *80* (2), 148–153.

Shumway, Robert H. and David T. Stoffer, *Time Series Analysis and Its Applications*, New York: Springer, 2011.

Tobin, James, "A proposal for international monetary reform," *Eastern economic journal*, 1978, *4* (3/4), 153–159.

Townsend, Robert M, "Forecasting the forecasts of others," *The Journal of Political Economy*, 1983, pp. 546–588.

van Binsbergen, Jules H., Michael W. Brandt, and Ralph Koijen, "On the Timing and Pricing of Dividends," *American Economic Review*, 2012, *102*, 1596–1618.

Vayanos, Dimitri, "Strategic trading and welfare in a dynamic market," *The Review of Economic Studies*, 1999, *66* (2), 219–254.

——, "Strategic trading in a dynamic noisy market," *The Journal of Finance*, 2001, *56* (1), 131–171.

Weller, Brian, "Does Algorithmic Trading Deter Information Acquisition?," 2017. Forthcoming, Review of Financial Studies.

Zhang, X. Frank, "The effect of high-frequency trading on stock volatility and price discovery," 2010. Working paper.

A Noise trader demand

We assume that noise traders have preferences similar to those of sophisticates, but they have different information. They receive signals about fundamentals, and believe that the signals are informative, although the signals are actually random. The signals are also perfectly correlated across the noise traders, so that they do not wash out in the aggregate. They can be therefore thought of as common sentiment shocks among noise traders. Furthermore, the noise traders assume that prices contain no information about fundamentals.

The noise traders optimize

$$\max_{\{N_t\}_{t=1}^T} T^{-1} \sum_{t=1}^{T} \beta^t N_t E_{0,N} \left[D_t - P_t \right] - \frac{1}{2} \left(\rho T \right)^{-1} Var_{0,N} \left[\sum_{t=1}^{T} \beta^t N_t \left(D_t - P_t \right) \right] \tag{50}$$

where N_t is the demand of the noise traders and $E_{0,N}$ and $Var_{0,N}$ are their expectation and variance operators conditional on their signals.

We model the noise traders as being Bayesians who simply misunderstand the informativeness of their signals, and ignore prices. Their prior belief, before receiving signals, is that

$$D \sim N\left(0, \Sigma_N^{prior}\right). \tag{51}$$

They then receive signals that they believe (incorrectly) are of the form

$$S \sim N\left(D, \Sigma_N^{signal}\right). \tag{52}$$

The usual Bayesian update then yields the distribution of D conditional on S,

$$\begin{aligned} D \quad | \quad S \sim N\left(\Sigma_N \left(\Sigma_N^{signal}\right)^{-1} S, \Sigma_N\right) \quad (53) \\ \text{where } \Sigma_N \quad \equiv \quad \left(\left(\Sigma_N^{signal}\right)^{-1} + \left(\Sigma_N^{prior}\right)^{-1}\right)^{-1}. \quad (54) \end{aligned}$$

So we have

$$\begin{aligned} E_{0,N}[D] &= \Sigma_N \left(\Sigma_N^{signal}\right)^{-1} S \quad (55) \\ Var_{0,N}[D] &= \Sigma_N \quad (56) \end{aligned}$$

Define $\tilde{N}_t \equiv \beta^t N_t$ and $\tilde{N} = [N_1, ..., N_T]'$. The optimization problem then becomes

$$\max_{\tilde{N}} T^{-1} \tilde{N}' \left(\Sigma_N \left(\Sigma_N^{signal}\right)^{-1} S - P\right) - \frac{1}{2} (\rho T)^{-1} \tilde{N}' \Sigma_N \tilde{N}. \tag{57}$$

This has the solution:

$$\begin{aligned} \tilde{N} &= \rho^{-1} \Sigma_N^{-1} \left(\Sigma_N \left(\Sigma_N^{signal}\right)^{-1} S - P\right) \quad (58) \\ &= \rho^{-1} \left(\left(\Sigma_N^{signal}\right)^{-1} S - \Sigma_N^{-1} P\right). \quad (59) \end{aligned}$$

For the sake of simplicity, we assume that $\Sigma_N = k^{-1} I$, where I is the identity matrix and k is a parameter. (This can be obtained, for instance, by assuming that $\Sigma_N^{signal} = \Sigma_N^{prior} = 2kI$). We then have

$$\tilde{N} = \rho^{-1} \left(\Sigma_N^{signal}\right)^{-1} S - kP, \tag{60}$$

so that the vector $Z = (Z_1, ..., Z_T)'$ from the main text is:

$$Z \equiv \rho^{-1} \left(\Sigma_N^{signal}\right)^{-1} S, \tag{61}$$

and the true variance of S, Σ_S, can always be chosen to yield any particular $\Sigma_Z \equiv Var(Z)$ by setting

$$\Sigma_S = \rho^2 \Sigma_N^{signal} \Sigma_Z \Sigma_N^{signal}. \tag{62}$$

B Time horizon and investment

At first glance, the assumption of mean-variance utility over cumulative returns over a long period of time ($T \to \infty$) 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 $T \to \infty$ 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

$$\left\{2\pi j/2^k\right\}_{j=0}^{2^{k-1}} \tag{63}$$

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

$$\left\{2\pi j/2^{k-1}\right\}_{j=0}^{2^{k-2}} = \left\{2\pi\left(2j\right)/2^k\right\}_{j=0}^{2^{k-2}} \tag{64}$$

That is, when T falls from 2^k to 2^{k-1}, the effect is to simply eliminate alternate frequencies. Reducing 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 $T \to \infty$, 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.

C Results on the frequency solution

C.1 Proof of lemma 1

The broad idea of the proof is as follows. Let Σ be any matrix of the form:

$$\Sigma = \begin{pmatrix} \sigma_0 & \sigma_1 & \cdots & \cdots & \sigma_{T-1} \\ \sigma_1 & \sigma_0 & \sigma_1 & \cdots & \sigma_{T-2} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \sigma_{T-1} & \cdots & \cdots & \cdots & \sigma_0 \end{pmatrix} \tag{65}$$

where $x_0 > 0$. Matrices of this type contain all the variance-covariance matrices of order T of arbitrary weakly stationary processes. The lemma follows from "approximating" Σ by the circulant matrix:

$$\Sigma_{circ} = circ(\sigma_{circ}) \quad , \quad \sigma \equiv (\sigma_0, \sigma_1 + \sigma_{T-1}, \sigma_2 + \sigma_{T-2}, ..., \sigma_{T-2} + \sigma_2, \sigma_{T-1} + \sigma_1)' , \tag{66}$$

where, for any real vector $\{x_i\}_{i=0}^{T-1}$,

$$circ(x) \equiv \begin{pmatrix} x_0 & \cdot & \cdot & x_{T-1} \\ x_{T-1} & x_0 & \cdot & x_{T-2} \\ \cdot & & & \\ x_1 & \cdot & \cdot & x_0 \end{pmatrix} . \tag{67}$$

In order to obtain this approximation, we first need the following result.

Appendix lemma 4 *For any matrix Σ of the form given above, and associated circulant matrix Σ_{circ}, the family of vectors Λ defined in the main text exactly diagonalizes Σ_{circ}:*

$$\Sigma_{circ}\Lambda = \Lambda diag\left(\left\{f_\Sigma\left(\omega_{\lfloor j/2\rfloor}\right)\right\}_{j=1}^{T}\right), \tag{68}$$

where each distinct eigenvalue in $\left\{f_\Sigma\left(\omega_{\lfloor j/2\rfloor}\right)\right\}_{j=1}^{T}$ is given by:

$$f_\Sigma(\omega_h) = \sigma_0 + 2\sum_{t=1}^{T-1} \sigma_t cos(\omega_h t), \quad \omega_h \equiv 2\pi h/T, \tag{69}$$

for some $h = 0, ..., \frac{T}{2}$.

Given that Λ is orthonormal,

$$\Lambda'\Sigma_{circ}\Lambda = diag\left(f_\Sigma\right). \tag{70}$$

The approximate diagonalization of the matrix Σ consists in writing:

$$\Lambda'\Sigma\Lambda = diag\left(f_\Sigma\right) + R_\Sigma, \tag{71}$$

where the $T \times T$ matrix R_Σ is given by:

$$R_\Sigma \equiv \Lambda'\left(\Sigma - \Sigma_{circ}\right)\Lambda. \tag{72}$$

This is an approximation in the sense that R_Σ is generically small. Specifically, it is of order T^{-1} element-wise. The following lemma proves the first result stated in lemma 1 of the main text.

Appendix lemma 5 *For any $T \geq 2$, we have:*

$$|R_\Sigma| \leq \frac{4}{\sqrt{T}} \sum_{j=1}^{T-1} |j\sigma_j|, \tag{73}$$

where $|M|$ denotes the weak matrix norm, as in the main text.

Proof. Define $\Delta\Sigma = \Sigma_{circ} - \Sigma$. First note that since:

$$\Sigma^{(i,j)} = \begin{cases} \sigma_0 & \text{if } i = j \\ \sigma_{|i-j|} & \text{otherwise} \end{cases}, \tag{74}$$

$$\Sigma_{circ}^{(i,j)} = \begin{cases} \sigma_0 & \text{if } i = j \\ \sigma_{|i-j|} + \sigma_{T-|i-j|} & \text{otherwise} \end{cases}, \tag{75}$$

we have:

$$\Delta\Sigma^{(i,j)} = \begin{cases} 0 & \text{if } i = j \\ \sigma_{T-|i-j|} & \text{otherwise} \end{cases} \tag{76}$$

where $\Sigma^{(i,j)}$ is the (i,j) element of Σ. This means that the matrix $\Delta\Sigma$ has constant and symmetric diagonals. Moreover, the first subdiagonals both contain σ_{T-1}, the second contain σ_{T-2}, and so on. That is,

$$\Delta\Sigma = \begin{pmatrix} 0 & \sigma_{T-1} & \sigma_{T-2} & & \sigma_2 & \sigma_1 \\ \sigma_{T-1} & \ddots & \ddots & \ddots & & \sigma_2 \\ \sigma_{T-2} & \ddots & \ddots & \ddots & \ddots & \\ & \ddots & \ddots & \ddots & \ddots & \sigma_{T-2} \\ \sigma_2 & & \ddots & \ddots & \ddots & \sigma_{T-1} \\ \sigma_1 & \sigma_2 & & \sigma_{T-2} & \sigma_{T-1} & 0 \end{pmatrix} \tag{77}$$

Therefore,

$$\sum_{i=1}^{T} \sum_{j=1}^{T} |\Delta\sigma_{i,j}| = 2 \sum_{j=1}^{T-1} |j\sigma_j|. \tag{78}$$

Let λ_k denote the k-th column of the matrix Λ. For any $(l,m) \in [1,T]^2$, we have:

$$\begin{aligned}
\left|R_\Sigma^{(l,m)}\right| &= |\lambda_l' \Delta\Sigma \lambda_m| \\
&= \left| \sum_{i=1}^{T} \sum_{j=1}^{T} \lambda_{i,l} \lambda_{j,m} \Delta\sigma_{i,j} \right| \\
&\leq \sum_{i=1}^{T} \sum_{j=1}^{T} |\lambda_{i,l}| \, |\lambda_{j,m}| \, |\Delta\sigma_{i,j}| \\
&\leq \sum_{i=1}^{T} \sum_{j=1}^{T} \frac{\sqrt{2}}{\sqrt{T}} \frac{\sqrt{2}}{\sqrt{T}} |\Delta\sigma_{i,j}| \\
&= \tfrac{4}{T} \sum_{j=1}^{T-1} |j\sigma_j|.
\end{aligned} \tag{79}$$

This implies that:

$$||R_\Sigma||_\infty \leq \frac{4}{T} \sum_{j=1}^{T-1} |j\sigma_j|, \tag{80}$$

where $||.||_\infty$ is the element-wise max norm. The inequality for the weak norm follows from the fact that the weak norm and the element-wise max norm satisfy $|.| \leq \sqrt{T}\, ||.||_\infty$. ■

C.2 Derivation of solution 1

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 (rather than vectors, which is what the unsubscripted variables represent in the main text).

C.2.1 Statistical inference

We guess that prices take the form

$$p = a_1 d + a_2 z \tag{81}$$

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

$$\begin{bmatrix} d \\ y_i \\ p \end{bmatrix} \sim N\left(0, \begin{bmatrix} f_D & f_D & a_1 f_D \\ f_D & f_D + f_i & a_1 f_D \\ a_1 f_D & a_1 f_D & a_1^2 f_D + a_2^2 f_Z \end{bmatrix}\right) \tag{82}$$

The expectation of fundamentals conditional on the signal and price is

$$E[d \mid y_i, p] = \begin{bmatrix} f_D & a_1 f_D \end{bmatrix} \begin{bmatrix} f_D + f_i & a_1 f_D \\ a_1 f_D & a_1^2 f_D + a_2^2 f_Z \end{bmatrix}^{-1} \begin{bmatrix} y_i \\ p \end{bmatrix} \tag{83}$$

$$= [1, a_1] \begin{bmatrix} 1 + f_i f_D^{-1} & a_1 \\ a_1 & a_1^2 + a_2^2 f_Z f_D^{-1} \end{bmatrix}^{-1} \begin{bmatrix} y_i \\ p \end{bmatrix} \tag{84}$$

and the variance satisfies

$$\tau_i \equiv Var[d \mid y_i, p]^{-1} = f_D^{-1} \left(1 - \begin{bmatrix} 1 & a_1 \end{bmatrix} \begin{bmatrix} 1 + f_i f_D^{-1} & a_1 \\ a_1 & a_1^2 + a_2^2 f_Z f_D^{-1} \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ a_1 \end{bmatrix} \right)^{-1} \tag{85}$$

$$= \frac{a_1^2}{a_2^2} f_Z^{-1} + f_i^{-1} + f_D^{-1} \tag{86}$$

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

$$E[d \mid y_i, p] = \tau_i^{-1} \left(f_i^{-1} y_i + \frac{a_1}{a_2^2} f_Z^{-1} p \right) \tag{87}$$

C.2.2 Demand and equilibrium

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

$$U_i = \max_{\{Q_{i,t}\}} \rho^{-1} E_{0,i} \left[T^{-1} \tilde{Q}_i' (D - P) \right] - \frac{1}{2} \rho^{-2} Var_{0,i} \left[T^{-1/2} \tilde{Q}_i' (D - P) \right] \tag{88}$$

$$= \max_{\{Q_{i,t}\}} \rho^{-1} E_{0,i} \left[T^{-1} \tilde{q}_i' (d - p) \right] - \frac{1}{2} \rho^{-2} Var_{0,i} \left[T^{-1/2} \tilde{q}_i' (d - p) \right] \tag{89}$$

$$= \max_{\{Q_{i,t}\}} \rho^{-1} T^{-1} \sum_j \tilde{q}_{i,j} E_{0,i} [(d_j - p_j)] - \frac{1}{2} \rho^{-2} T^{-1} \sum_j \tilde{q}_{i,j}^2 Var_{0,i} [d_j - p_j], \tag{90}$$

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 $\tilde{q}_{i,j}$ independent from all others.

Taking the first-order condition associated with the last line above for a single frequency (with $\tilde{q}_i$, d, etc. again representing scalars, for any j), we obtain

$$\tilde{q}_i = \rho \tau_i E[d - p \mid y_i, p] \tag{91}$$

$$= \rho \left(f_i^{-1} y_i + \left(\frac{a_1}{a_2^2} f_Z^{-1} - \tau_i \right) p \right) \tag{92}$$

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

$$-z+k\left(a_1 d+a_2 z\right) = \int_i \rho\left(f_i^{-1} y_i+\left(\frac{a_1}{a_2^2} f_Z^{-1}-\tau_i\right)\left(a_1 d+a_2 z\right)\right) di \tag{93}$$

$$= \int_i \rho\left(f_i^{-1} d+\left(\frac{a_1}{a_2^2} f_Z^{-1}-\tau_i\right)\left(a_1 d+a_2 z\right)\right) di, \tag{94}$$

where the second line uses the law of large numbers. Matching coefficients on d and z then yields

$$\int_i \rho\left(\frac{a_1}{a_2^2} f_Z^{-1}-\tau_i\right) di = -a_2^{-1}\left(1-k a_2\right) \tag{95}$$

$$\int_i \rho f_i^{-1} a_1^{-1}+\rho\left(\frac{a_1}{a_2^2} f_Z^{-1}-\tau_i\right) di = k \tag{96}$$

and therefore

$$k-\int_i \rho f_i^{-1} a_1^{-1} = a_2^{-1}\left(k a_2-1\right) \tag{97}$$

$$\int_i \rho f_i^{-1} = \frac{a_1}{a_2} \tag{98}$$

Now define aggregate precision to be

$$f_{avg}^{-1} \equiv \int_i f_i^{-1} di \tag{99}$$

We then have

$$\tau_i = \frac{a_1^2}{a_2^2} f_Z^{-1}+f_i^{-1}+f_D^{-1} \tag{100}$$

$$\tau_{avg} \equiv \int \tau_i di=\left(\rho f_{avg}^{-1}\right)^2 f_Z^{-1}+f_{avg}^{-1}+f_D^{-1} \tag{101}$$

Inserting the expression for τ_i into (95) yields

$$a_1 = \frac{\tau_{avg}-f_D^{-1}}{\tau_{avg}+\rho^{-1} k} \tag{102}$$

$$a_2 = \frac{a_1}{\rho f_{avg}^{-1}} \tag{103}$$

The expression for a_1 can be written more explicitly as:

$$a_1=\frac{\tau_{avg}-f_D^{-1}}{\tau_{avg}+\rho^{-1} k} = \frac{\frac{a_1^2}{a_2^2} f_Z^{-1}+f_{avg}^{-1}+f_D^{-1}+\rho^{-1} k-\rho^{-1} k-f_D^{-1}}{\frac{a_1^2}{a_2^2} f_Z^{-1}+f_{avg}^{-1}+f_D^{-1}+\rho^{-1} k} \tag{104}$$

$$= 1-\frac{\rho^{-1} k+f_D^{-1}}{\left(\rho f_{avg}^{-1}\right)^2 f_Z^{-1}+f_{avg}^{-1}+\rho^{-1} k+f_D^{-1}}. \tag{105}$$

The expression for a_2 is invalid in the case when $f_{avg}^{-1}=0$. In that case, we have

$$a_2 = \frac{1}{\rho f_D^{-1}+k}. \tag{106}$$

C.3 Proof of Proposition 1

We use the notation $\bar{O}$ to mean that, for any matrices A and B,

$$|A-B| = \bar{O}\left(T^{-1/2}\right) \Longleftrightarrow |A-B| \leq bT^{-1/2} \tag{107}$$

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

$$\begin{aligned}
|cA-cB| &= \bar{O}\left(T^{-1/2}\right) && (108)\\
\left|A^{-1}-B^{-1}\right| &= \bar{O}\left(T^{-1/2}\right) && (109)\\
|(A+C)-(B+D)| &= \bar{O}\left(T^{-1/2}\right) && (110)\\
|AC-BD| &= \bar{O}\left(T^{-1/2}\right) && (111)
\end{aligned}$$

In other words, convergence in weak norm carries through under addition, multiplication, and inversion. Following the time domain solution (252), A_1 and A_2 can be expressed as a function of the Toeplitz matrices Σ_D, Σ_Z and Σ_{avg} using those operations. it follows that $|A_1 - \Lambda diag\left(a_1\right)\Lambda'| \leq c_1 T^{-\frac{1}{2}}$ for some constant c_1, and the same holds for A_2 for some constant c_2.

For the variance of prices, we define

$$\begin{aligned}
R_1 &\equiv A_1 - \Lambda diag\left(a_1\right)\Lambda', && (112)\\
R_2 &\equiv A_2 - \Lambda diag\left(a_2\right)\Lambda'. && (113)
\end{aligned}$$

In what follows, we use the strong norm $||.||$, defined as:

$$||A|| = \max_{x'x=a}\left(x'A'Ax\right)^{\frac{1}{2}}. \tag{114}$$

Finally, we use the following property of the weak norm: for any two square matrices A, B of size $T\times T$,

$$|AB| \leq \sqrt{T}\,|A|\,|B|. \tag{115}$$

The proof for this inequality is standard and reported at the end of this section. We then have the following:

$$\begin{aligned}
|Var\,[P-\Lambda p]| &= \left|Var\left[(A_1-\Lambda a_1\Lambda')D+(A_2-\Lambda a_2\Lambda')Z\right]\right| &(116)\\
&\leq \left|R_1\Sigma_D R_1'\right|+\left|R_2\Sigma_Z R_2'\right| &(117)\\
&\leq \sqrt{T}\left(|R_1\Sigma_D|\,|R_1|+|R_2\Sigma_Z|\,|R_2|\right) &(118)\\
&\leq \sqrt{T}\left(\|\Sigma_D\|\,|R_1|^2+\|\Sigma_Z\|\,|R_2|^2\right) &(119)\\
&\leq \sqrt{T}K\left(|R_1|^2+|R_2|^2\right). &(120)
\end{aligned}$$

The second line follows from the triangle inequality. The third line comes from property (115). The fourth line uses the fact that for any two square matrices G, H, $||GH|| \leq \|G\|\,|H|$; for a proof, see Gray (2006), lemma 2.3. The last line follows from the assumption that the eigenvalues of Σ_D and Σ_Z are bounded. Indeed, since Σ_D and Σ_Z are symmetric and real, they are Hermitian; following Gray (2006), eq. (2.16), we then have $\|\Sigma_Z\| = \max_t |\alpha_{Z,t}|$ and $\|\Sigma_D\| = \max_t |\alpha_{D,t}|$, where $\alpha_{X,t}$ denotes the eigenvalues of the matrix X.

Given that $|R_1| \leq c_1 T^{-\frac{1}{2}}$ and $|R_2| \leq c_2 T^{-\frac{1}{2}}$, this implies:

$$\begin{aligned}
|Var\,[P-\Lambda p]| &\leq K\sqrt{T}\left(c_1^2+c_2^2\right)T^{-1} &(121)\\
&= c_P T^{-\frac{1}{2}}. &(122)
\end{aligned}$$

A similar proof establishes the result for $\left|Var\left[\tilde{Q}-\Lambda\tilde{q}\right]\right|$.

To prove inequality (115), note that:

$$\begin{aligned}
|AB|^2 &= \tfrac{1}{T}\sum_{m,n}\left(\sum_{t=1}^{T} a_{mt}b_{tn}\right)^2\\
&\leq \tfrac{1}{T}\sum_{m,n}\left(\sum_{t=1}^{T} a_{mt}^2\right)\left(\sum_{t=1}^{T} b_{tn}^2\right)\\
&= \tfrac{1}{T}\left(\sum_{m,t} a_{mt}^2\right)\left(\sum_{n,t} b_{nt}^2\right)\\
&= T\left(\tfrac{1}{T}\left(\sum_{m,t} a_{mt}^2\right)\right)\left(\tfrac{1}{T}\left(\sum_{n,t} b_{nt}^2\right)\right)\\
&= T\,|A|^2\,|B|^2,
\end{aligned} \qquad (123)$$

so that $|AB| \leq \sqrt{T}\,|A|\,|B|$. In this sequence of inequalities, going from the second to the third line uses the Cauchy-Schwarz inequality.

C.4 Proof of lemma 2

First, since the trace operator is invariant under rotations,

$$tr\left(\Sigma_i^{-1}\right) = \sum_j f_{i,j}^{-1}. \tag{124}$$

The information constraint is linear in the frequency-specific precisions. Investors also face a technical constraint that the elements of $f_{i,j}$ corresponding to paired sines and cosines must have the same value. That is, if $\lfloor j/2 \rfloor = \lfloor k/2 \rfloor$, then $f_{i,j} = f_{i,k}$; this condition is necessary for $\varepsilon_{i,t}$ to be stationary.

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

$$E_{-1}\left[U_{i,0}\right] \equiv \frac{1}{2} E\left[T^{-1} \sum_j \tau_{i,j} E\left[d_j - p_j \mid y_{i,j}, p_j\right]^2\right] \tag{125}$$

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

$Var\left[E\left[d_j - p_j \mid y_{i,j}, p_j\right]\right]$ 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

$$Var\left[d_j - p_j\right] = Var\left[E\left[d_j - p_j \mid y_{i,j}, p_j\right]\right] + E\left[Var\left[d_j - p_j \mid y_{i,j}, p_j\right]\right] \tag{126}$$

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

$$Var(r_j) = Var\left[d_j - p_j\right] = \left(1 - a_{1,j}\right)^2 f_{D,j} + \frac{a_{1,j}^2}{\left(\rho f_{avg,j}^{-1}\right)^2} f_{Z,j}. \tag{127}$$

So then

$$E_{-1}\left[U_{i,0}\right] = \frac{1}{2} T^{-1} \sum_j \left[\left(\left(1 - a_{1,j}\right)^2 f_{D,j} + \frac{a_{1,j}^2}{\left(\rho f_{avg}^{-1}\right)^2} f_{Z,j}\right) \tau_{i,j} - 1\right]. \tag{128}$$

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 $f_{i,j}^{-1}$; see appendix equation 100). Now define

$$\lambda_j\left(f_{avg,j}^{-1}\right) \equiv (1 - a_{1,j})^2 f_{D,j} + \left(\frac{a_{1,j}}{\rho f_{avg,j}^{-1}}\right)^2 f_{Z,j} = Var(r_j). \tag{129}$$

From equations (101)-(102), λ_j can be re-written as:

$$\lambda_j\left(f_{avg,j}^{-1}\right)=\frac{f_{D,j}\left(f_{D,j}^{-1}+\rho^{-1}k\right)^2+(\rho f_{avg,j}^{-1})^2 f_{Z,j}^{-1}+f_{Z,j}\rho^{-2}}{\left((\rho f_{avg,j}^{-1})^2 f_{Z,j}^{-1}+f_{D,j}^{-1}+\rho^{-1}k+f_{avg,j}^{-1}\right)^2}, \tag{130}$$

which can be further decomposed as:

$$\begin{aligned}\lambda_j\left(f_{avg,j}^{-1}\right) &= \frac{1}{\left((\rho f_{avg,j}^{-1})^2 f_{Z,j}^{-1}+f_{D,j}^{-1}+\rho^{-1}k+f_{avg,j}^{-1}\right)^2}\\ &+ \frac{f_{Z,j}-\frac{f_{avg,j}^{-1}}{\rho}}{\left((\rho f_{avg,j}^{-1})^2 f_{Z,j}^{-1}+f_{D,j}^{-1}+\rho^{-1}k+f_{avg,j}^{-1}\right)^2}\\ &+ \frac{\rho^{-1}k(1+f_{D,j}^{-1}\rho^{-1}k)}{\left((\rho f_{avg,j}^{-1})^2 f_{Z,j}^{-1}+f_{D,j}^{-1}+\rho^{-1}k+f_{avg,j}^{-1}\right)^2}\end{aligned} \tag{131}$$

Each of these three terms is decreasing in $f_{avg,j}^{-1}$, so that the function $\lambda_j(\cdot)$ is decreasing.

D Results on price informativeness with restricted frequencies

D.1 Result 1 and corollaries 1.1 and 1.5

When there are no active investors and just exogenous supply, we have that $0=z_j+kp_j$ and so:

$$p_j = k^{-1}z_j, \tag{132}$$

$$r_j = d_j-k^{-1}z_j. \tag{133}$$

Because of the separability of information choices across frequencies, the coefficients $a_{1,j}$ and $a_{2,j}$ are unchanged at all other frequencies. Moreover, it is clear that $Var(d_j|p_j)=Var(d_j)$ at the restricted frequencies, since prices now only carry information about supply, which is uncorrelated with dividends.

Note that for any $j\in\mathcal{R}$,

$$Var(r_j)=f_{D,j}+\frac{f_{Z,j}}{k^2}. \tag{134}$$

Additionally, if trading at that frequency were not restricted, but the investors endogenously chose not to allocate any attention to the frequency, the return volatility would be:

$$Var_{unrestr.}(r_j)=\lambda_j(0)=f_{D,j}+\frac{f_{Z,j}}{\left(k+\rho f_{D,j}^{-1}\right)^2}<Var(r_j). \tag{135}$$

D.2 Corollary 1.2 and result 1.4

Under the diagonal approximation, we have:

$$D\mid P\sim N\left(\bar{D},\Lambda diag\left(\tau_0^{-1}\right)\Lambda'\right) \tag{136}$$

where τ_0 is a vector of frequency-specific precisions conditional on prices, as of time 0. Given the independence of prices across frequencies, the j-th element of τ_0 is:

$$\tau_{0,j}^{-1} = Var(d_j \mid p_j). \tag{137}$$

Using this expression, we can compute:

$$\begin{aligned}
Var\left(D_t \mid P\right) &= 1_t' \Lambda diag\left(\tau_0^{-1}\right) \Lambda' 1_t & (138)\\
&= \left(\Lambda' 1_t\right)' diag\left(\tau_0^{-1}\right)\left(\Lambda' 1_t\right) & (139)\\
&= \sum_j \lambda_{t,j}^2 Var(d_j \mid p_j) & (140)\\
&= \lambda_{t,0}^2 Var(d_0|p_0) + \lambda_{t,\frac{T}{2}}^2 Var(d_{\frac{T}{2}} \mid p_{\frac{T}{2}}) + \sum_{k=1}^{T/2-1}\left(\lambda_{t,2k}^2 + \lambda_{t,2k+1}^2\right) Var(d_k \mid p_k) & (141)
\end{aligned}$$

where 1_t is a vector equal to 1 in its t-th element and zero elsewhere, and $\lambda_{t,j}$ is the t,j element of Λ. The last line follows from the fact that all the spectra have $f_{X,2k} = f_{X,2k+1}$ for $0 < k < T/2 - 1$. Furthermore, note that for $0 < k < T/2 - 1$,

$$\begin{aligned}
\lambda_{t,2k}^2 + \lambda_{t,2k+1}^2 &= \frac{2}{T}\cos\left(\omega_k\left(t-1\right)\right)^2 + \frac{2}{T}\sin\left(\omega_k\left(t-1\right)\right)^2 & (142)\\
&= \frac{2}{T} & (143)
\end{aligned}$$

which yields equation (34). Result 3 immediately follows from this expression.

Result 1.4 uses the fact that

$$\begin{aligned}
Var\left(D_t - D_{t-1} \mid P\right) &= \left(\lambda_{t,1} - \lambda_{t-1,1}\right)^2 \tau_{0,1}^{-1} + \left(\lambda_{t,T} - \lambda_{t-1,T}\right)^2 \tau_{0,\frac{T}{2}}^{-1} & (144)\\
&+ \sum_{k=1}^{T/2-1}\left[\left(\lambda_{t,2k} - \lambda_{t-1,2k}\right)^2 + \left(\lambda_{t,2k+1} - \lambda_{t-1,2k+1}\right)^2\right]\tau_{0,k}^{-1}, & (145)
\end{aligned}$$

and the fact that $(\cos(x) - \cos(y))^2 + (\sin(x) - \sin(y))^2 = 4\sin\left(\frac{1}{2}(x-y)\right)^2 = 2\left(1 - \cos(x-y)\right)$.

E Results on investor outcomes

E.1 Result 2

Expression (38) in the main text follows from the steps used in appendix D.2. Recall from (92) that, omitting the j notation,

$$\begin{aligned}
\tilde{q}_i &= \rho\left(f_i^{-1}y_i + \left(\frac{a_1}{a_2^2}f_Z^{-1} - \tau_i\right)p\right) &(146)\\
&= \rho f_i^{-1}\varepsilon_i + \rho\left(f_i^{-1} + \left(\frac{a_1}{a_2^2}f_Z^{-1} - \tau_i\right)a_1\right)d + \rho\left(\frac{a_1}{a_2^2}f_Z^{-1} - \tau_i\right)a_2 z &(147)
\end{aligned}$$

Recall also that:

$$\tau_i = \left(\frac{a_1}{a_2}\right)^2 f_Z^{-1} + f_D^{-1} + f_i^{-1}, \tag{148}$$

so that:

$$\tilde{q}_i = \rho\left(\tau_i - \left(\frac{a_1}{a_2}\right)^2 f_Z^{-1} - f_D^{-1}\right)\varepsilon_i + \rho\left(f_i^{-1} + \left(\frac{a_1}{a_2^2}f_Z^{-1} - \tau_i\right)a_1\right)d + \rho\left(\frac{a_1}{a_2^2}f_Z^{-1} - \tau_i\right)a_2 z \tag{149}$$

Moreover,

$$\begin{aligned}
f_i^{-1} - a_1\tau_i + \left(\frac{a_1}{a_2}\right)^2 f_Z^{-1} &= \tau_i - \left(\frac{a_1}{a_2}\right)^2 f_Z^{-1} - f_D^{-1} - a_1\tau_i + \left(\frac{a_1}{a_2}\right)^2 f_Z^{-1} &(150)\\
&= (1-a_1)\tau_i - f_D^{-1}. &(151)
\end{aligned}$$

Therefore

$$\rho^{-1}\tilde{q}_i = \left(\tau_i - \left(\frac{a_1}{a_2}\right)^2 f_Z^{-1} - f_D^{-1}\right)\varepsilon_i + \left((1-a_1)\tau_i - f_D^{-1}\right)d + \left(\frac{a_1}{a_2}f_Z^{-1} - a_2\tau_i\right)z, \tag{152}$$

so that

$$\rho^{-2}Var\left(\tilde{q}_i\right) = \left(\tau_i - \left(\frac{a_1}{a_2}\right)^2 f_Z^{-1} - f_D^{-1}\right) + \left((1-a_1)\tau_i - f_D^{-1}\right)^2 f_D + \left(\frac{a_1}{a_2}f_Z^{-1} - a_2\tau_i\right)^2 f_Z. \tag{153}$$

(where the first term uses the fact that $Var\left(f_i^{-1}\varepsilon_i\right) = f_i^{-1}$). The derivative of this expression with respect to τ_i is:

$$\begin{aligned}
\rho^{-2}\frac{\partial Var\left(\tilde{q}_i\right)}{\partial \tau_i} &= 2\tau_i\left((1-a_1)^2 f_D + a_2^2 f_Z\right) - 1 \\
&\geq 2\left(f_D^{-1} + \left(\tfrac{a_1}{a_2}\right)^2 f_Z^{-1}\right)\left((1-a_1)^2 f_D + a_2^2 f_Z\right) - 1 \\
&= 2\left((1-a_1)^2 + a_2^2 f_Z f_D^{-1} + (1-a_1)^2\left(\tfrac{a_1}{a_2}\right)^2 f_Z^{-1} f_D + a_1^2\right) - 1 \\
&= 2\left(1 - 2a_1(1-a_1) + a_2^2 f_Z f_D^{-1} + (1-a_1)^2\left(\tfrac{a_1}{a_2}\right)^2 f_Z^{-1} f_D\right) - 1 \\
&= 2\left(-2a_1(1-a_1) + a_2^2 f_Z f_D^{-1} + (1-a_1)^2\left(\tfrac{a_1}{a_2}\right)^2 f_Z^{-1} f_D\right) + 1 \\
&= 2\left((1-a_1)\left(\tfrac{a_1}{a_2}\right)(f_Z^{-1} f_D)^{\frac{1}{2}} - a_2(f_Z^{-1} f_D)^{-\frac{1}{2}}\right)^2 + 1 \\
&> 0,
\end{aligned} \tag{154}$$

where to go from the first to the second line, we used the fact that $\tau_i \geq \left(\frac{a_1}{a_2}\right)^2 f_Z^{-1} + f_D^{-1}$, and where we also used the fact that $a_1 \leq 1$. Since τ_i is a monotonic transformation of f_i^{-1}, this establishes equation (39) from the main text.

For result 2, first note that $E_{-1}\left[\tilde{Q}_i' R\right] = E_{-1}\left[\tilde{q}_i' \Lambda' \Lambda r\right] = E_{-1}\left[\tilde{q}_i r\right] = \sum_j E_{-1}\left[\tilde{q}_{i,j} r_j\right]$, where the last equality follows from the diagonal approximation. Moreover, straightforward but tedious algebra shows that:

$$f_i^{-1} + \left(\frac{a_1}{a_2^2} f_Z^{-1} - \tau_i\right) a_1 = \rho(f_i^{-1} - f_{avg}^{-1})(1-a_1) + k a_1, \tag{155}$$

$$\left(\frac{a_1}{a_2^2} f_Z^{-1} - \tau_i\right) a_2 = -\rho(f_i^{-1} - f_{avg}^{-1}) a_2 + (k a_2 - 1). \tag{156}$$

We can use these expressions, and the fact that $r = (1-a_1)d - a_2 z$ to re-write $\tilde{q}_i$ as:

$$\tilde{q}_i = \rho f_i^{-1}\varepsilon_i + \rho\left(f_i^{-1} - f_{avg}^{-1}\right) r + k a_1 d + (k a_2 - 1)\, z. \tag{157}$$

Therefore,

$$E_{-1}\left[\tilde{q}_i r\right] = \rho\left(f_i^{-1} - f_{avg}^{-1}\right) Var\left(r\right) + k a_1 E_{-1}\left[rd\right] + (k a_2 - 1)\, E_{-1}\left[rz\right], \tag{158}$$

which is the decomposition from result 2.

The result that expected profits are nonnegative is a simple consequence of the investors' objective:

$$\max_{\{\tilde{q}_{i,j}\}} \rho^{-1} T^{-1} \sum_j E_{0,i}\left[\tilde{q}_{i,j}\left(d_j - p_j\right)\right] - \frac{1}{2}\rho^{-2} T^{-1} \sum_j Var_{0,i}\left[\tilde{q}_{i,j}\left(d_j - p_j\right)\right] \tag{159}$$

Since the variance is linear in $\tilde{q}_{i,j}^2$, if $E_{0,i}[\tilde{q}_{i,j} r_j] < 0$, utility can always be increased by setting $\tilde{q}_{i,j} = 0$ (or, even more, by reversing the sign of $\tilde{q}_{i,j}$). In order for $E_{-1}[\tilde{q}_{i,j} r_j] = 0$, it must be the case that $Var_{-1,i}[E_{0,i}[d_j - p_j]] = 0$, since any deviation of $E_{0,i}[d_j - p_j]$ will cause the investor to optimally take a position. We have, from above,

$$a_1 = \frac{\tau_{avg} - f_D^{-1}}{\tau_{avg} + \rho^{-1} k} = \frac{\left(\rho f_{avg}^{-1}\right)^2 f_Z^{-1} + f_{avg}^{-1}}{\left(\rho f_{avg}^{-1}\right)^2 f_Z^{-1} + f_{avg}^{-1} + f_D^{-1} + \rho^{-1} k} \tag{160}$$

$$a_2 = \frac{a_1}{\rho f_{avg}^{-1}} \tag{161}$$

$$\tau_{avg} \equiv \left(\rho f_{avg}^{-1}\right)^2 f_Z^{-1} + f_{avg}^{-1} + f_D^{-1} \tag{162}$$

The expression for a_2 is invalid in the case when $f_{avg}^{-1} = 0$. In that case, we have

$$E[d \mid y_i, p] = \tau_i^{-1} \left(f_i^{-1} y_i + \frac{a_1}{a_2^2} f_Z^{-1} p \right) \tag{163}$$

$$E[d - p \mid y_i, p] = \tau_i^{-1} f_i^{-1} y_i + \left(\tau_i^{-1} \frac{a_1}{a_2^2} f_Z^{-1} - 1 \right) (a_1 d + a_2 z) \tag{164}$$

$$Var[E[d - p \mid y_i, p]] = \left(\tau_i^{-1} f_i^{-1} + \left(\tau_i^{-1} \frac{a_1}{a_2^2} f_Z^{-1} - 1 \right) a_1 \right)^2 f_D + \left(\tau_i^{-1} \frac{a_1}{a_2^2} f_Z^{-1} - 1 \right)^2 a_2^2 f_Z \tag{165}$$

Now first we must have $\tau_i^{-1} \frac{a_1}{a_2^2} f_Z^{-1} - 1 = 0$ in order for the third term to be zero. But if that is true, then for the first term to be zero we must have $f_i^{-1} = 0$ (since τ_i^{-1} is always positive). Combining $f_i^{-1} = 0$ with $\tau_i^{-1} \frac{a_1}{a_2^2} f_Z^{-1} - 1 = 0$, we obtain

$$f_D^{-1} = \rho f_{avg}^{-1} f_Z^{-1} k. \tag{166}$$

E.2 Corollary 3

We drop the notation j for clarity. Assume that low-frequency agents are initially uninformed about the frequency; then $f_i^{-1} = 0$, for all i so:

$$\tau_i = \left(\frac{a_1}{a_2} \right)^2 f_Z^{-1} + f_D^{-1}. \tag{167}$$

Using expression (152), we then have

$$\rho^{-1} \tilde{q}_{LF,i} = \left((1 - a_1) \left(\frac{a_1}{a_2} \right)^2 f_Z^{-1} - a_1 f_D^{-1} \right) d + \left(\frac{a_1 (1 - a_1)}{a_2} f_Z^{-1} - a_2 f_D^{-1} \right) z. \tag{168}$$

Given that $r = (1 - a_1)d - a_2 z$ and that z and d are independent,

$$\begin{aligned}
\rho^{-1} E_{-1}\left[\tilde{q}_{LF,i} r\right] &= \left((1-a_1)\left(\tfrac{a_1}{a_2}\right)^2 f_Z^{-1} - a_1 f_D^{-1} \right)(1-a_1) f_D - \left(\tfrac{a_1(1-a_1)}{a_2} f_Z^{-1} - a_2 f_D^{-1} \right) a_2 f_Z \\
&= (1-a_1)^2 \left(\tfrac{a_1}{a_2}\right)^2 f_Z^{-1} f_D - 2a_1(1-a_1) + a_2^2 f_Z f_D^{-1} \\
&= \left((1-a_1)\left(\tfrac{a_1}{a_2}\right)(f_Z^{-1} f_D)^{\frac{1}{2}} - a_2 (f_Z f_D^{-1})^{\frac{1}{2}} \right)^2
\end{aligned} \tag{169}$$

For any $f_{avg}^{-1} > 0$, where $a_1/a_2 = \rho f_{avg}^{-1}$, the derivative of this expression with respect to f_{avg}^{-1} is

$$\begin{aligned}
\rho^{-1} \tfrac{dE_{-1}\left[\tilde{q}_{LF,i} r\right]}{df_{avg}^{-1}} &= 2\left((1-a_1)\left(\tfrac{a_1}{a_2}\right)(f_Z^{-1} f_D)^{\frac{1}{2}} - a_2 (f_Z f_D^{-1})^{\frac{1}{2}} \right) \\
&\times \left\{ \rho \left[(1-a_1)(f_Z^{-1} f_D)^{\frac{1}{2}} - a_1 (f_Z f_D^{-1})^{\frac{1}{2}} \right] - \left[(f_Z^{-1} f_D)^{\frac{1}{2}} + (f_Z f_D^{-1})^{\frac{1}{2}} \right] \rho \tfrac{\partial a_1}{\partial f_{avg}^{-1}} f_{avg}^{-1} \right\}
\end{aligned} \tag{170}$$

Moreover, when $f_{avg}^{-1} > 0$,

$$\frac{\partial a_1}{\partial f_{avg}^{-1}} f_{avg}^{-1} = a_1(1-a_1) + (1-a_1) \frac{(\rho f_{avg}^{-1})^2 f_Z^{-1}}{(\rho f_{avg}^{-1})^2 f_Z^{-1} + f_{avg}^{-1} + f_D^{-1} + \rho^{-1} k}. \tag{171}$$

The following limits follow from the discussion in Appendix C.2.2:

$$\lim_{f_{avg}^{-1} \to 0^+} a_1 = 0, \qquad \lim_{f_{avg}^{-1} \to 0} a_2 = \frac{1}{\rho f_D^{-1} + k}. \tag{172}$$

Using these limits and the expressions just derived, we arrive at

$$\lim_{f_{avg}^{-1} \to 0^+} \frac{\partial E_{-1}\left[\tilde{q}_{LF,i} r\right]}{\partial f_{avg}^{-1}} = -2\rho \frac{(f_Z f_D^{-1})^{\frac{1}{2}} (f_Z^{-1} f_D)^{\frac{1}{2}}}{f_D^{-1} + \rho^{-1} k} < 0. \tag{173}$$

Re-introducing the notation j, for the frequency at which entry takes place, we then have

$$\frac{d}{df_{avg,j}^{-1}} E_{-1}\left[\sum_t \tilde{Q}_{LF,t}\left(D_t - P_t\right) \right] = \frac{d}{df_{avg,j}^{-1}} \sum_k E_{-1}\left[\tilde{q}_{LF,k} r_k\right] = \frac{d}{df_{avg,j}^{-1}} E_{-1}\left[\tilde{q}_{LF,j} r_j\right] < 0; \tag{174}$$

that is, all the effect of entry on total profits is concentrated on frequency j, where entry reduces profits, as just established.

For the last result, we again use the frequency separability,

$$\frac{d}{df_{avg,j}^{-1}} E_{-1}\left[U_{LF,0}\right] = \frac{d}{df_{avg,j}^{-1}} E_{-1}\left[u_{LF,0,j}\right], \tag{175}$$

where

$$E_{-1}[u_{LF,0,j}] \equiv \frac{1}{2} T^{-1} \left[\left((1-a_{1,j})^2 f_{D,j} + a_{2,j}^2 f_{Z,j} \right) \tau_{i,j} - 1 \right] \tag{176}$$

is the component of utility which fluctuates at frequency j. This latter definition uses expression

(128), derived in Appendix C.4. Omitting the j notation for clarity, the derivative of this expression with respect to f_{avg} assuming that $f_i^{-1} = 0$ is:

$$\begin{aligned} 2T\tfrac{dE_{-1}[u_{LF,0}]}{df_{avg}^{-1}} &= \left((1-a_1)^2 f_D + a_1^2(\rho f_{avg}^{-1})^2 f_Z\right) 2\rho^2 f_Z^{-1} f_{avg}^{-1} \\ &+ \left(-2(1-a_1)\tfrac{\partial a_1}{\partial f_{avg}^{-1}} f_D + 2a_1 \tfrac{\partial a_1}{\partial f_{avg}^{-1}} \left(\rho f_{avg}^{-1}\right)^2 f_Z + 2a_1^2 \rho^2 f_Z f_{avg}^{-1}\right) \left(\left(\rho f_{avg}^{-1}\right)^2 f_Z^{-1} + f_D^{-1}\right) \end{aligned} \tag{177}$$

Given that:

$$\lim_{f_{avg}^{-1} \to 0^+} a_1 = 0, \tag{178}$$

the only term in this expression for which the limit may not be 0 as $f_{avg}^{-1} \to 0^+$ is:

$$-2(1-a_1)\frac{\partial a_1}{\partial f_{avg}^{-1}} f_D + 2a_1 \frac{\partial a_1}{\partial f_{avg}^{-1}} \rho f_{avg}^{-1} f_Z. \tag{179}$$

However, given equation (171), we have that:

$$\lim_{f_{avg}^{-1} \to 0^+} \frac{\partial a_1}{\partial f_{avg}^{-1}} f_{avg}^{-1} = 0, \tag{180}$$

and so the second term in (179) goes to 0 as $f_{avg}^{-1} \to 0^+$. For the second term, note that, using (171) we have that:

$$\frac{\partial a_1}{\partial f_{avg}^{-1}} = \frac{a_1}{f_{avg}^{-1}} + o(1) = \frac{1 + (\rho f_{avg}^{-1}) f_Z^{-1}}{(\rho f_{avg}^{-1})^2 f_Z^{-1} + f_D^{-1} + f_{avg}^{-1} + \rho^{-1} k} + o(1). \tag{181}$$

Therefore,

$$\lim_{f_{avg}^{-1} \to 0^+} 2T \frac{dE_{-1}[u_{LF,0}]}{df_{avg}^{-1}} = -2\frac{f_D}{f_D^{-1} + \rho^{-1} k} = -2 f_D a_2 < 0, \tag{182}$$

which proves the last statement of corollary 3.

F Quadratic costs

F.1 Frequency domain expressions for trading costs

Using $Q_i = \Lambda q_i$, each agent's position at time t can be written as

$$Q_{i,t} = \sum_j \left[\begin{array}{c} q_j \cos(2\pi j t/T) \\ +q_{j'} \sin(2\pi j t/T) \end{array} \right]. \tag{183}$$

Trading costs are then written in terms of $(Q_{i,t} - Q_{i,t-1})^2$ as:

$$QV\{Q_i\} \equiv \sum_{t=2}^{T} (Q_{i,t} - Q_{i,t-1})^2 + (Q_{i,1} - Q_{i,T})^2. \tag{184}$$

We can write that as

$$QV\left\{Q_i\right\} = (DQ)'(DQ) \tag{185}$$

where D is a matrix that generates first differences,

$$D \equiv \begin{bmatrix} -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & -1 & 1 \\ 1 & 0 & \cdots & 0 & -1 \end{bmatrix}. \tag{186}$$

Using again the fact that $Q_i = \Lambda q_i$,

$$QV\left\{Q_i\right\} = q'\Lambda' D' D\Lambda q \tag{187}$$

In what follows, we will need to evaluate the matrix $\Lambda' D' D\Lambda$. The m, n element of that matrix is the inner product of the m and n columns of $D\Lambda$. Each column of $D\Lambda$ contains the first difference of the corresponding column of Λ, with the exception of the last element, $(D\Lambda)_{m,T}$, which is equal to $\Lambda_{m,t} - \Lambda_{n,T}$. We have the following standard trigonometric results: for $m \neq n$:

$$\begin{aligned} \sum_{t=1}^{T} \left(\cos\left(\omega_m t\right) - \cos\left(\omega_m\left(t-1\right)\right)\right)\left(\cos\left(\omega_n t\right) - \cos\left(\omega_n\left(t-1\right)\right)\right) &= 0, & (188) \\ \sum_{t=1}^{T} \left(\cos\left(\omega_m t\right) - \cos\left(\omega_m\left(t-1\right)\right)\right)\left(\sin\left(\omega_n t\right) - \sin\left(\omega_n\left(t-1\right)\right)\right) &= 0, & (189) \\ \sum_{t=1}^{T} \left(\sin\left(\omega_m t\right) - \sin\left(\omega_m\left(t-1\right)\right)\right)\left(\sin\left(\omega_n t\right) - \sin\left(\omega_n\left(t-1\right)\right)\right) &= 0, & (190) \end{aligned}$$

where recall that $\omega_m = \frac{2\pi m}{T}$, and:

$$\begin{aligned} \sum_{t=1}^{T} \left(\cos\left(\omega_m t\right) - \cos\left(\omega_m\left(t-1\right)\right)\right)^2 &= 2T\sin^2\left(\omega_m/2\right), & (191) \\ \sum_{t=1}^{T} \left(\sin\left(\omega_m t\right) - \sin\left(\omega_m\left(t-1\right)\right)\right)^2 &= 2T\sin^2\left(\omega_m/2\right), & (192) \\ \sum_{t=1}^{T} \left(\cos\left(\omega_m t\right) - \cos\left(\omega_m\left(t-1\right)\right)\right)\left(\sin\left(\omega_m t\right) - \sin\left(\omega_m\left(t-1\right)\right)\right) &= 0. & (193) \end{aligned}$$

These results immediately imply that the off-diagonal elements of $\Lambda' D' D\Lambda$ are equal to zero and the jth element of the main diagonal is $2T\sin^2\left(\omega_{\lfloor j/2 \rfloor}/2\right)$.

We then have

$$
\begin{aligned}
QV\{Q_i\} &= q\Lambda' D' D\Lambda q && (194)\\
&= \sum_{j=1}^{T} 2T\sin^2\left(\omega_{\lfloor j/2\rfloor}/2\right) q_{i,j}^2 && (195)
\end{aligned}
$$

Total holding costs can be written as:

$$
\sum_{t=1}^{T} Q_t^2 = \sum_{j=1}^{T} q_j^2, \tag{196}
$$

which is just Parseval's theorem.

F.2 Equilibrium of the trading cost model

Throughout the analysis, unless it is necessary, we omit the index j of the particular frequency in order to simplify notation.

F.2.1 Investment and equilibrium

The first-order condition for frequency j is

$$
\begin{aligned}
0 &= E\left[d_j - p_j \mid y_{i,j}, p_j\right] - 2c\sin^2\left(\omega_{\lfloor j/2\rfloor}/2\right) q_j - bq_j && (197)\\
q &= \frac{E\left[d_j - p_j \mid y_{i,j}, p_j\right]}{\gamma_j} && (198)\\
&= \gamma_j^{-1}\tau_i^{-1}\left(f_i^{-1} y_i + \left(\frac{a_1}{a_2^2} f_Z^{-1} - \tau_i\right) p\right) && (199)
\end{aligned}
$$

where

$$
\gamma_j \equiv 2c\sin^2\left(\omega_{\lfloor j/2\rfloor}/2\right) + b \tag{200}
$$

is the marginal cost of q_j. We can then solve for the coefficients a_1 and a_2 as before.

Inserting the formula for the conditional expectation and integrating across investors yields

$$
\begin{aligned}
\int_i \gamma_j^{-1}\tau_i^{-1}\left(f_i^{-1} y_i + \left(\frac{a_1}{a_2^2} f_Z^{-1} - \tau_i\right)(a_1 d - a_2 z)\right) di &= z_j && (201)\\
\int_i \gamma_j^{-1}\tau_i^{-1}\left(f_i^{-1} d + \left(\frac{a_1}{a_2^2} f_Z^{-1} - \tau_i\right)(a_1 d - a_2 z)\right) di &= z_j && (202)
\end{aligned}
$$

Matching coefficients then yields

$$\int_i \gamma_j^{-1} \tau_i^{-1} \left(\frac{a_1}{a_2^2} f_Z^{-1} - \tau_i \right) di = -a_2^{-1} \tag{203}$$

$$\int_i \gamma_j^{-1} \tau_i^{-1} \left(f_i^{-1} + \left(\frac{a_1}{a_2^2} f_Z^{-1} - \tau_i \right) a_1 \right) di = 0 \tag{204}$$

Combining those two equations, we obtain

$$\int_i \gamma_j^{-1} \tau_i^{-1} f_i^{-1} di = \frac{a_1}{a_2} \tag{205}$$

Now put the definition of τ_i into that equation for f_i^{-1}

$$\int_i \gamma_j^{-1} \tau_i^{-1} \left(\tau_i - \frac{a_1^2}{a_2^2} f_Z^{-1} - f_D^{-1} \right) di = \frac{a_1}{a_2} \tag{206}$$

$$\gamma_j^{-1} \int_i 1 - \left(\frac{a_1^2}{a_2^2} f_Z^{-1} - f_D^{-1} \right) \tau_i^{-1} di = \frac{a_1}{a_2} \tag{207}$$

F.2.2 Expected utility

At any particular frequency,

$$U_{i,j} = q_{i,j} E_{0,i} \left[d_j - p_j \right] - \frac{1}{2} q_{i,j}^2 2c \sin^2 \left(\omega_{\lfloor j/2 \rfloor}/2 \right) - \frac{1}{2} b q_{i,j}^2 \tag{208}$$

$$= \frac{1}{2} \frac{E \left[d_j - p_j \mid y_{i,j}, p_j \right]^2}{\gamma_j} \tag{209}$$

Expected utility prior to observing signals is then

$$EU_{i,j} \equiv \frac{1}{2} E \left[\frac{E \left[d_j - p_j \mid y_{i,j}, p_j \right]^2}{\gamma_j} \right] \tag{210}$$

$E \left[E \left[d_j - p_j \mid y_{i,j}, p_j \right]^2 \right]$ is the variance of the part of the return on portfolio j explained by $y_{i,j}$ and p_j, while $\tau_{i,j}$ is the residual variance. We know from the law of total variance that

$$Var \left[d_j - p_j \right] = Var \left[E \left[d_j - p_j \mid y_{i,j}, p_j \right] \right] + E \left[Var \left[d_j - p_j \mid y_{i,j}, p_j \right] \right] \tag{211}$$

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

$$Var \left[d_j - p_j \right] = Var \left[\left(1 - a_1 \right) d_j + a_2 z_j \right] \tag{212}$$

$$= \left(1 - a_{1,j} \right)^2 f_{D,j} + a_2^2 f_{Z,j} \tag{213}$$

So then

$$EU_{i,j} = \frac{1}{2}\frac{Var\left[d_j - p_j\right] - \tau_{i,j}^{-1}}{\gamma_j} \tag{214}$$

What we end up with is that utility is decreasing in $\tau_{i,j}^{-1}$. That is,

$$EU_{i,j} = -\frac{1}{2}\frac{\tau_{i,j}^{-1}}{\gamma_j} + \text{constants}. \tag{215}$$

F.2.3 Information choice

With the linear cost on precision, agents maximize

$$\begin{aligned}
& -\frac{1}{2}\frac{\tau_{i,j}^{-1}}{\gamma_j} - \psi f_{i,j}^{-1} && (216)\\
= \quad & -\frac{1}{2}\left(\frac{a_1^2}{a_2^2} f_{Z,j}^{-1} + f_{i,j}^{-1} + f_{D,j}^{-1}\right)^{-1} \gamma_j^{-1} - \psi f_{i,j}^{-1} && (217)
\end{aligned}$$

The FOC for $f_{i,j}^{-1}$ is

$$\begin{aligned}
\psi &= \frac{1}{2}\tau_{i,j}^{-2}\gamma_j^{-1} && (218)\\
\tau_{i,j} &= \frac{1}{\sqrt{2}}\psi^{-1/2}\gamma_j^{-1/2} && (219)
\end{aligned}$$

But τ has a lower bound of $\frac{a_1^2}{a_2^2} f_Z^{-1} + f_D^{-1}$, so it's possible that this has no solution. That would be a state where agents do no learning. Formally,

$$\tau_{i,j} = \max\left(\frac{a_1^2}{a_2^2} f_Z^{-1} + f_D^{-1}, \frac{1}{\sqrt{2}}\psi^{-1/2}\gamma_j^{-1/2}\right) \tag{220}$$

Note that, unlike in the other model, the equilibrium is unique here – all agents individually face a concave problem with an interior solution.

Frequencies with no learning Now using the result for a_1/a_2 from above, at the frequencies where nobody learns, $f_i^{-1} = 0$, we have

$$\begin{aligned}
\frac{a_1}{a_2} &= \int_i \gamma_j^{-1}\tau_i^{-1} f_i^{-1} di && (221)\\
&= 0 && (222)
\end{aligned}$$

which then implies

$$\tau_{i,j} = \max\left(f_D^{-1}, \frac{1}{\sqrt{2}}\psi^{-1/2}\gamma_j^{-1/2}\right) \tag{223}$$

To get a_2, we have

$$\int_i \left(cj^2 + b\right) \tau_i^{-1} \left(\frac{a_1}{a_2^2} f_Z^{-1} - \tau_i\right) di = -a_2^{-1} \tag{224}$$

$$\gamma_j = a_2 \tag{225}$$

So the sensitivity of the price to supply shocks is increasing in the cost of holding inventory, b, and the trading costs, c. It is also higher at higher frequencies – it is harder to temporarily push through supply than to do it persistently.

Frequencies with learning At the frequencies at which there is learning, where

$$f_D^{-1} < \frac{1}{\sqrt{2}} \psi^{-1/2} \gamma_j^{-1/2} \tag{226}$$

we have, just by rewriting the τ equation,

$$f_i^{-1} = \tau_i - \frac{a_1^2}{a_2^2} f_Z^{-1} - f_D^{-1} \tag{227}$$

Using the second equation from above,

$$\int_i \gamma_j^{-1} \tau_i^{-1} \left(\frac{a_1}{a_2^2} f_Z^{-1} - \tau_i\right) di = -a_2^{-1} \tag{228}$$

$$\int_i \gamma_j^{-1} \tau_i^{-1} \left(\frac{a_1}{a_2} f_Z^{-1} - a_2 \tau_i\right) di = -1 \tag{229}$$

$$\int_i \gamma_j^{-1} \left(\tau_i^{-1} \frac{a_1}{a_2} f_Z^{-1} - a_2\right) di = -1 \tag{230}$$

Under the assumption of a symmetric strategy, this is

$$\tau^{-1} \frac{a_1}{a_2} f_Z^{-1} - a_2 = -\gamma_j \tag{231}$$

$$\frac{a_1}{a_2} = \tau f_Z \left(-\gamma_j + a_2\right) \tag{232}$$

Using the other equilibrium condition, we have

$$\int_i \gamma_j^{-1} \tau_i^{-1} \left(\tau_i - \frac{a_1^2}{a_2^2} f_Z^{-1} - f_D^{-1}\right) di = \frac{a_1}{a_2} \tag{233}$$

$$\int_i \gamma_j^{-1} \left(1 - \tau_i^{-1} \frac{a_1}{a_2} f_Z^{-1} \frac{a_1}{a_2} - \tau_i^{-1} f_D^{-1}\right) di = \frac{a_1}{a_2} \tag{234}$$

$$1 - \left(-\gamma_j + a_2\right) \frac{a_1}{a_2} - \tau_i^{-1} f_D^{-1} = \left(cj^2 + b\right) \frac{a_1}{a_2} \tag{235}$$

$$1 - \tau_i^{-1} f_D^{-1} = a_1 \tag{236}$$

Plugging in the formula for τ_i when there is learning,

$$1-\sqrt{2}\psi^{1/2}\gamma_j^{1/2}f_D^{-1}=a_1. \tag{237}$$

The expression for a_2 can be obtained from:

$$\frac{a_1}{\tau f_Z}=\left(-\gamma_j+a_2\right)a_2. \tag{238}$$

Since $a_1/\tau f_Z>0$, we know that there is only one solution to this equation for $a_2>0$. The positive root is

$$a_2=\frac{\gamma_j+\sqrt{\gamma_j^2+4\frac{a_1}{\tau f_Z}}}{2} \tag{239}$$

G Results when fundamentals are difference-stationary

In the main text, we assume that the level of fundamentals is stationary. Here we examine an extension in which fundamentals are stationary in terms of first differences and show that the results go through nearly identically, with the primary difference being in how the low-frequency portfolio is defined.

G.1 Informed investors under difference stationarity

We assume that D_0 is known to investors when making decisions, and without loss of generality normalize $D_0=0$. Define Δ to be the first difference operator so that

$$\Delta D_t=D_t-D_{t-1} \tag{240}$$

and define the vector $\Delta D\equiv[\Delta D_1,\Delta D_2,...\Delta D_T]'$. We assume that

$$\Delta D\sim N\left(0,\Sigma_D\right). \tag{241}$$

For any given allocation to the futures contracts, there is an allocation to claims on ΔD that gives an identical payoff. Specifically, an allocation $Q_i'D$ can be exactly replicated by

$$\begin{aligned} Q_i'D &= Q_i'L_1\Delta D && (242)\\ &= \left(L_1'Q_i\right)'\Delta D && (243)\end{aligned}$$

where L_1 is a matrix that creates partial sums,

$$L_1 \equiv \begin{bmatrix} 1 & 0 & 0 & \cdots \\ 1 & 1 & 0 & \\ 1 & 1 & 1 & \\ \vdots & & & \ddots \end{bmatrix} \tag{244}$$

So an allocation of Q_i to the futures is equivalent to an allocation of $L_1' Q_i$ to claims on the first differences of fundamentals, which we will call the growth rate futures. Define the notation

$$Q_{\Delta,i} \equiv L_1' Q_i \tag{245}$$

Furthermore, the prices of the growth rate futures are simply the vector ΔP (by the law of one price). We can therefore rewrite the optimization problem equivalently as

$$\max T^{-1} \sum_{t=1}^{T} \beta^t Q_{\Delta,i,t} E_{0,i} \left[\Delta D_t - \Delta P_t\right] - \frac{1}{2} \left(\rho T^{-1}\right) Var_{0,i} \left[\sum_{t=1}^{T} \beta^t Q_{\Delta,i,t} \left(\Delta D_t - \Delta P_t\right)\right] \tag{246}$$

Now suppose for the moment that we are able to solve the entire model in terms of first differences (that is not obvious as we will need to ensure that noise trader demand is also difference stationary). So we have an allocation $Q_{\Delta D,i}$. An allocation to the first differences is then equivalent to an allocation of $\left(L_1'\right)^{-1} Q_{\Delta,i}$ to the levels (which follows trivially from the definition of $Q_{\Delta,i}$ in (245)).

Since our maintained assumption is that we will solve the model in first differences in the same way we did in the main text for levels, that means that we will continue to use the rotation Λ, but now in first differences. So the frequency domain allocations in terms of first differences will be

$$\tilde{Q}_{\Delta D,i} = \Lambda \tilde{q}_{\Delta,i} \tag{247}$$

where $\tilde{Q}_{\Delta D,i,t} \equiv Q_{\Delta D,i,t} \beta^t$. $\tilde{q}_{\Delta,i}$ now represents the allocations to different frequencies of growth in fundamentals. The key question, then, is what that implies for the behavior of portfolios in terms of levels. We have

$$\tilde{Q}_i = \left(L_1'\right)^{-1} \tilde{Q}_{\Delta,i} \tag{248}$$
$$= \left(L_1'\right)^{-1} \Lambda \tilde{q}_{\Delta,i} \tag{249}$$

So in terms of levels, the basis vectors, instead of being Λ, are $\left(L_1'\right)^{-1} \Lambda$.

For $(L_1')^{-1}$ we have

$$\left(L_1'\right)^{-1} \equiv \begin{bmatrix} 1 & -1 & 0 & \cdots & 0 \\ 0 & 1 & -1 & & \vdots \\ 0 & 0 & 1 & \ddots & 0 \\ \vdots & & & \ddots & -1 \\ 0 & \cdots & 0 & 0 & 1 \end{bmatrix} \tag{250}$$

So the way that $(L_1')^{-1}$ transforms a matrix is to take a forward difference of each column, and then retaining the value of the final row. A way to see the implications of that transformation is to approximate the finite differences of the sines and cosines as derivatives. The columns of $(L_1')^{-1}\Lambda$ are equal to $(L_1')^{-1} c_j$ and $(L_1')^{-1} s_j$, which can be written using standard trigonometric formulas as:

$$\left(L_1'\right)^{-1} c_j \quad \approx \quad \begin{bmatrix} 2\sin\left(\frac{1}{2}\omega_j\right)\sqrt{\frac{2}{T}}\left\{\sin\left(\omega_j\left(t-\frac{1}{2}\right)\right)\right\}_{t=2}^{T} \\ \sqrt{\frac{2}{T}}\cos\left(\omega_j(T-1)\right) \end{bmatrix} \tag{251}$$

$$\left(L_1'\right)^{-1} s_j \quad \approx \quad \begin{bmatrix} -2\sin\left(\frac{1}{2}\omega_j\right)\sqrt{\frac{2}{T}}\left\{\cos\left(\omega_j\left(t-\frac{1}{2}\right)\right)\right\}_{t=2}^{T} \\ \sqrt{\frac{2}{T}}\sin\left(\omega_j(T-1)\right) \end{bmatrix} \tag{252}$$

The column c_j represents a portfolio in terms of the first differences of fundamentals with weights equal to a cosine fluctuating at frequency ω_j. $(L_1')^{-1} c_j$ measures the loadings of that portfolio on claims to the level of fundamentals. These loadings also fluctuate at frequency ω_j, with the only difference being the replacement of the cosine with a sine function. (Intuitive, the loadings are approximately equal to the derivative of the columns of Λ with respect to time; taking derivatives does not affect the characteristic frequency of fluctuations.)

So consider a relatively high-frequency investor, whose portfolio weights are all close to zero except for a large value in the vector $q_{\Delta,i}$ at some large value of j. By assumption, that investor holds a portfolio whose loadings on the first differences of fundamentals fluctuate at frequency ω_j. What the approximations in (251–252) show, though, is that that investor's positions measured in terms of the *level* of fundamentals (i.e. $\tilde{Q}_i$) has loadings that also fluctuate at frequency ω_j.

One subtlety is in the lowest-frequency portfolio, $(L_1'^{-1}\left(\frac{1}{\sqrt{2}}c_0\right)$. That portfolio puts equal weight on growth in fundamentals on all dates – it is a bet on the sample mean growth rate. In terms of levels, note that $(L_1')^{-1}\left(\frac{1}{\sqrt{2}}c_0\right) = \left[0,0,0,...,\sqrt{2/T}\right]$. A person who wants to bet on the mean growth rate between dates 1 and T can do that by buying a claim to fundamentals only on date T.[19]

[19]The highest frequency portfolio, $(L_1'^{-1}\left(\frac{1}{\sqrt{2}}c_{\frac{T}{2}}\right)$, is given by $1/\sqrt{T}\,(2,-2,...,2,1)'$, and therefore fluctuates at the highest sample frequency.

G.2 Noise traders under difference stationarity

Last, we need to show that noise trader demand will also take a form such that the entire model can be solved in terms of first differences (and then shifted back into levels for interpretation). First, as above, since the model expressed in first differences is just a linear transformation of the levels, the noise traders' optimization problem can be written in terms of first differences,

$$\max T^{-1} \sum_{t=1}^{T} \beta^t N_{\Delta,t} E_{0,N} \left[\Delta D_t - \Delta P_t\right] - \frac{1}{2} \left(\rho T^{-1}\right) Var_{0,N} \left[\sum_{t=1}^{T} \beta^t N_{\Delta,t} \left(\Delta D_t - \Delta P_t\right)\right] \tag{253}$$

where $N_{\Delta,t}$ is the demand of the noise traders for the claims on first differences.

We assume that the noise traders understand that fundamentals have a unit root and that they therefore have priors and signals that refer to the change in fundamentals. The analogs to (51) and (52) are then

$$\Delta D \sim N\left(0, \Sigma_{N\Delta}^{prior}\right) \tag{254}$$

$$S \sim N\left(\Delta D, \Sigma_{N\Delta}^{signal}\right) \tag{255}$$

and the Bayesian update is

$$\Delta D \quad | \quad S \sim N\left(\Sigma_{N\Delta}\left(\Sigma_{N\Delta}^{signal}\right)^{-1} S, \Sigma_{N\Delta}\right) \tag{256}$$

$$\text{where } \Sigma_{N\Delta} \equiv \left(\left(\Sigma_{N\Delta}^{signal}\right)^{-1} + \left(\Sigma_{N\Delta}^{prior}\right)^{-1}\right)^{-1} \tag{257}$$

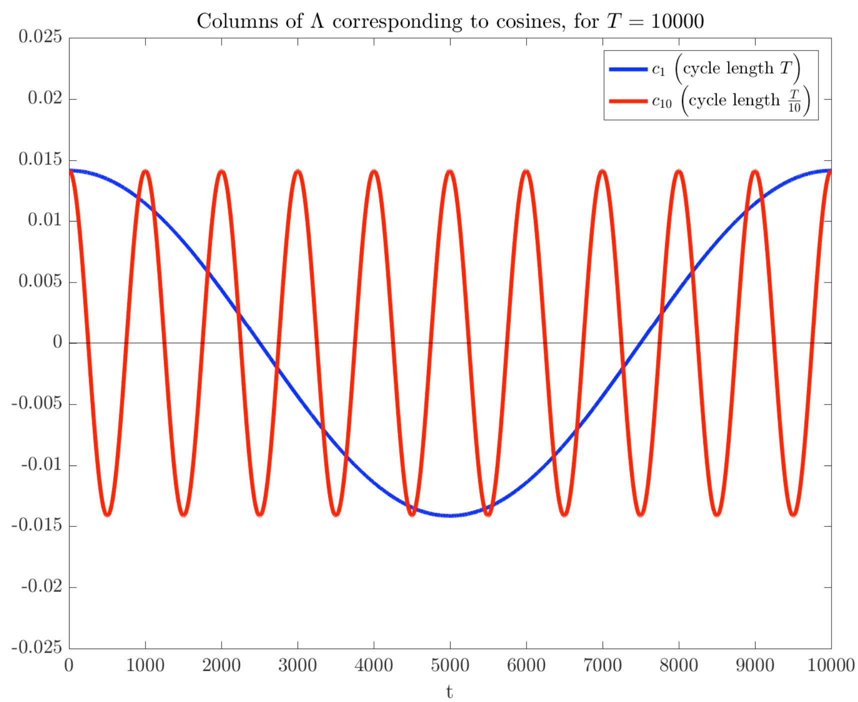

Figure 1: Portfolio weights for the cosing frequency portfolios c_1 and c_{10}, as defined in the main text. The horizontal axis is time, or the maturity of the corresponding futures contract. The vertical axis is the weight which each portfolio puts on that futures contract.