# Information, Liquidity, and Market Dynamics

Aryan Ayyar

September 22, 2025

# Contents

# Chapter 1

# Introduction

This book develops a rigorous yet intuitive framework for understanding sequential Kyle games, a class of models in financial market microstructure that explain how asymmetric information affects price formation over time. We begin by reviewing the foundational Kyle (1985) model, and then progressively introduce extensions involving multi-period settings, partially informed traders, information leakage, stochastic signals, and Bayesian learning by market makers.

## 1.1 Overview

In his book ”Elements of Pure Economics”, Leon Walsrus established the conceptual framework of a general equilibrium. In so far as the investor is concerned, market prices play two important roles, namely allocation of scarce resources and being vehicles of information. It is today well known that economics equilibrium is a system-wide phenomenon and is not isolated to individual markets.Arrow and Debreu (1954) provided the first conceptual proof of the existence of a general equilibrium. Later, as it were to be, Debreu(1959) werre to present the classic Arrow-Debrew framework with not just unparalleled mathematical rigour but with clarity and generality.

What was revolutionary was the 1980 paper by Sanford Grossman and Joseph Stiglitz ”On the Impossibility of Informationally Efficient Markets” . This challenged the fundamental assumption of costless, symmetric information in the Arrow-Debreu general equilibrium and showed how incorporating the cost of information may lead to profound paradoxes. Formally, they proved a fundamental impossibility theorem which states that ”perfectly informationally efficient markets are impossible if information is costly to acquire”. Hence, they introduced the concept of a Rational Expectations Equilibrium. A Rational Expectations Equilibrium is a state in which, once all market participants have observed the equilibrium price $p^*$, no one has an incentive to revise their portfolio choice. In this equilibrium, all agents agree that $p^*$ is optimal given their information set and that further adjustments would not improve expected payoff.This directly contrasts with a Walrasian equilibrium: a price decline not only clears markets (the Walrasian effect) but also reduces perceived fundamental value (the REE effect).[1]

## 1.2 Financial Market Equilibrium

To illustrate this paradox, let's consider a representative agent endowed with $I$ shares of a risky asset and $I_f$ units of a risk-free asset. The risk-free asset yields a gross return $1 + r_f$, while the risky asset pays a random payoff $F$ at time $T$. If the agent demands $X$ units of the risky asset at price $p$, then initial wealth at $t = 0$ is

$$W_0 \;=\; I\,p \;+\; I_f. \tag{1.1}$$

At time $T$, terminal wealth is

$$W_t \;=\; (X + I)\,F \;+\; \bigl(I_f - Xp\bigr)\,(1 + r_f). \tag{1}$$

The agent maximizes expected utility $U(W_t)$ of terminal wealth, with $U'(W_t) = \frac{\mathrm{d}U}{\mathrm{d}W_t}$. The first-order condition for optimal demand $X$ is

$$\mathbb{E}\bigl[U'(w)\,(F - p\,(1 + r_f))\bigr] \;=\; 0.$$

Using $\mathbb{E}[AB] = \mathbb{E}[A]\,\mathbb{E}[B] + \mathrm{Cov}(A, B)$ and Stein's lemma yields

$$\mathbb{E}\bigl[U'(w)\bigr]\,\mathbb{E}\bigl[F - p\,(1 + r_f)\bigr] + \mathbb{E}\bigl[U''(w)\bigr]\,(I + X)\,\mathrm{Var}(F) \;=\; 0. \tag{2}$$

Rearranging (2) gives the equilibrium price:

$$p \;=\; \frac{1}{1 + r_f}\left(\frac{\mathbb{E}\bigl[U''(w)\bigr]\,(I + X)\,\mathrm{Var}(F)}{\mathbb{E}\bigl[U'(w)\bigr]} + \mathbb{E}[F]\right). \tag{3}$$

Under CARA utility $U(W_t) \;=\; -e^{-A\,W_t}$, one has $U'(W_t) \;=\; -A\,U(W_t)$ and $U''(W_t) \;=$

[1] Kyle (1989) introduced imperfect competition among informed traders, demonstrating that prices reveal at most half of their private information, so even risk-neutral informed agents trade less aggressively and an REE exists.

$A^2U(W_t)$, so (3) simplifies to

$$p \;=\; \frac{1}{1+r_f}\Big(A\,(I+X)\,\mathrm{Var}(F)+\mathbb{E}[F]\Big). \tag{4}$$

Thus the risky asset's price equals the discounted expected payoff plus a risk premium proportional to position size and risk aversion. Consequently, the expected gross return satisfies

$$\mathbb{E}[r] = r_f + \frac{A\,(I+X)\,\mathrm{Var}(F)}{p}. \tag{5}$$

### 1.2.1 Capital Asset Pricing Model

In the case of multiple risky assets, we can now derive the CAPM smoothly (see Sharpe, 1964). Starting from (3), a slightly modified and generalized version reads

$$p_i \;=\; \frac{1}{1+r_f}\left[\mathbb{E}[F_i]+\frac{\mathbb{E}\big[U''(w)\big]}{\mathbb{E}\big[U'(w)\big]}\;\mathrm{Cov}\big(w,F_i\big)\right]. \tag{6}$$

Hence the expected gross return on asset $i$ satisfies

$$\mathbb{E}[R_i] \;=\; r_f \;-\; \frac{\mathbb{E}\big[U''(w)\big]}{\mathbb{E}\big[U'(w)\big]}\;\mathrm{Cov}\big(w,r_i\big). \tag{7}$$

Consider the market portfolio $M$, whose price is $p_M = \sum_i p_i X_i$ and whose return is $R_M$. Define the value-weights

$$w_i = \frac{p_i X_i}{p_M}.$$

Weight-averaging (7) gives

$$\sum_i w_i\,\mathbb{E}[R_i] = \sum_i w_i\Big(r_f - \tfrac{\mathbb{E}[U''(w)]}{\mathbb{E}[U'(w)]}\;\mathrm{Cov}(w,r_i)\Big) \tag{8}$$

$$\mathbb{E}[R_M] - r_f = -\frac{\mathbb{E}[U''(w)]}{\mathbb{E}[U'(w)]}\sum_i w_i\;\mathrm{Cov}(w,r_i) = -\frac{\mathbb{E}[U''(w)]}{\mathbb{E}[U'(w)]}\;\mathrm{Cov}(w,R_M). \tag{9}$$

Under the CARA specification $U(w) = -e^{-Aw}$, one shows that $\frac{\mathbb{E}[U''(w)]}{\mathbb{E}[U'(w)]} = A$, and noting that $\mathrm{Cov}(w,R_M) = p_M^{-1}\,\mathrm{Var}(R_M)\,p_M$ yields, after substitution into (7),

$$\mathbb{E}[R_i] - r_f \;=\; \frac{\mathrm{Cov}(R_i,R_M)}{\mathrm{Var}(R_M)}\;\big[\mathbb{E}[R_M]-r_f\big]. \tag{10}$$

**This is precicely the Sharpe–Lintner Capital Asset Pricing Model.**

[2]

**Proposition 1.1** (Representative-agent pricing identity)**.** *Let there be a risk-free asset with gross return* $R_f = 1 + r_f > 0$ *and a single risky asset with payoff* $F$ *at date* $T$*. A representative agent with strictly increasing, twice continuously differentiable utility* $U$ *over terminal wealth* $W_T$ *holds* $I$ *initial units of the risky asset, chooses demand* $X$*, and faces price* $p$ *at* $t = 0$*. If* $F$ *is integrable and* $\mathrm{Var}(F) < \infty$*, then any competitive equilibrium price* $p$ *satisfies*

$$p \;=\; \frac{1}{R_f}\left(\mathbb{E}[F] \;+\; \frac{\mathbb{E}[U''(W_T)]}{\mathbb{E}[U'(W_T)]}\,(I+X)\;\mathrm{Var}(F)\right),$$

*where expectations are taken under the objective probability measure and* $W_T = (X+I)F + (W_0 - pX)R_f$*.*

*Proof.* The first-order condition is $\mathbb{E}\big[U'(W_T)\,(F - pR_f)\big] = 0$. Using $\mathbb{E}[AB] = \mathbb{E}[A]\mathbb{E}[B] + \mathrm{Cov}(A, B)$ and $\mathrm{Cov}(U'(W_T), F) = \mathbb{E}[U''(W_T)](I + X)\,\mathrm{Var}(F)$ by the law of iterated expectations and linearity of $W_T$ in $F$, one gets

$$\mathbb{E}[U'(W_T)]\,(\mathbb{E}[F] - pR_f) \;+\; \mathbb{E}[U''(W_T)](I+X)\,\mathrm{Var}(F) \;=\; 0,$$

which rearranges to the stated identity. □

**Proposition 1.2** (CARA–Normal specialization and risk premium)**.** *Under the conditions of the previous proposition, suppose* $U(w) = -\exp(-Aw)$ *with* $A > 0$ *and* $F$ *is independent of* $W_0$ *with variance* $\mathrm{Var}(F)$*. Then* $\mathbb{E}[U''(W_T)]/\mathbb{E}[U'(W_T)] = -A$ *and the equilibrium price satisfies*

$$p \;=\; \frac{1}{R_f}\Big(\mathbb{E}[F] \;-\; A\,(I+X)\;\mathrm{Var}(F)\Big),$$

*so the expected gross return on the risky asset obeys*

$$\mathbb{E}[R] \;=\; \mathbb{E}\left[\frac{F}{p}\right] \;=\; R_f \;+\; \frac{A\,(I+X)\;\mathrm{Var}(F)}{p},$$

*which identifies a positive risk premium proportional to risk aversion, position size, and payoff variance.*

[2]We must remember that this insight of Arrow and Debreu was under clearly specified and relatively general conditions. In my opinion, this is arguably the most foundational paper of the concept. This was supported by Lckinzie (1954)'s independent and nearly coherent work with Arrow and Debreu who reinforced the possibility of a coherent competitive equilibrium system by specifically focusing on the Gale-Nikaido-Debreu Lemma. Yes, one may agree that there were significant limitations to the Arrow-Debrew framework, but it can be shown that the aggregate excess demand function can behave almost arbitrarily (see, Sonnenschein (1972, 1973), Mantel (1974), Debreu (1974) ). One must also note that the Gale–Nikaido–Debreu Lemma is a mathematical tool used to prove the existence of a competitve equilibrium. Specifically, it provides the conditions under which a system of inequalities has a deterministic solution.

*Proof.* For CARA, $U'(w) = Ae^{-Aw}$ and $U''(w) = -A^2e^{-Aw}$, hence $\mathbb{E}[U''(W_T)]/\mathbb{E}[U'(W_T)] = -A$. Substitute into the pricing identity and divide by $p$ to obtain the return expression. □

## 1.3 Asymmetric Information

A classical platform to startoff would be Akerlof (1970). In markets with unobservable product quality (e.g., used cars), he finds asymmetric information between buyers and sellers causes adverse selection. Sellers of low-quality goods (”lemons”) drive out high-quality goods because buyers cannot distinguish quality and only offer average prices.Hence, under these conditions, markets may unravel entirely or operate at suboptimal equilibria. Similarly, Stiglitz and Rothschild (1976) discuss formallized screening as a solution to asymmetric information. Hence, one thing is clear - information asymmetry can cause pareto inefficiency even in competitive markets. Hence, we move on to consider that not all agents possess the same information. There seem to be two distinct groups of traders: risk-averse agents and noise (liquidity) traders. Each agent's demand is $X_i \in \{X_I, X_U, X_N\}$, and the population sizes are $N_i \in \{N_I, N_U, N_N\}$. Noise traders submit $\hat{x} \sim \mathcal{N}(0, \sigma_x^2)$. There are no initial endowments and the risk-free rate is normalized to zero. The prior for the asset's fundamental value is $\hat{F} \sim \mathcal{N}(\bar{F}, \sigma_F^2)$. Priors represent beliefs before observing new information; posteriors incorporate private signals via Bayes' rule. At time $t = 0$, each agent receives a noisy signal $S \sim \mathcal{N}(F, \sigma_s^2)$.

**Proposition 1.3** (Gaussian conjugate update for a scalar signal)**.** *Let the prior for a scalar fundamental $F$ be $F \sim \mathcal{N}(\bar{F}, \sigma_F^2)$, and let a private signal satisfy $S \mid F \sim \mathcal{N}(F, \sigma_S^2)$, independent of other randomness. Then the posterior is Gaussian with*

$$\mathbb{E}[F \mid S] \;=\; \bar{F} \;+\; \frac{\sigma_F^2}{\sigma_F^2 + \sigma_S^2}\,(S - \bar{F}), \qquad \mathrm{Var}(F \mid S) \;=\; \frac{\sigma_F^2\,\sigma_S^2}{\sigma_F^2 + \sigma_S^2}.$$

*Proof.* Complete the square in the joint normal density of $(F, S)$ or apply the linear regression formula $\mathbb{E}[F \mid S] = \bar{F} + \frac{\mathrm{Cov}(F,S)}{\mathrm{Var}(S)}(S - \mathbb{E}[S])$ with $\mathrm{Cov}(F, S) = \sigma_F^2$ and $\mathrm{Var}(S) = \sigma_F^2 + \sigma_S^2$; the conditional variance follows from the Schur complement. □

# Chapter 2

# The One-Period Kyle Model

Classical competitive equilibrium assumes that prices fully reflect all available information, ensuring informational efficiency. Yet, as Grossman and Stiglitz famously argued, this ideal cannot be sustained: if markets were perfectly revealing, no investor would have an incentive to incur the costs of acquiring information, and trade would vanish altogether **?**. This paradox highlighted the inherent tension between incentives for information acquisition and the possibility of fully efficient markets.

## 2.1 Background and Motivation

In response, the market microstructure literature made the trading process itself explicit. Models such as Glosten and Milgrom demonstrated how order flow can act as a conduit for private information, and how adverse selection endogenously generates trading costs and bid–ask spreads even when dealers are risk-neutral and competitive **??**. These quote-driven frameworks explain spreads trade-by-trade, attributing them directly to the presence of better-informed traders. This then, should naturally raise the question: *why Kyle?* Kyle's (1985) auction-style formulation provides a complementary perspective. Instead of spreads, it emphasizes linear price impact and endogenous market depth as order-flow-based measures of illiquidity **??**. The framework became a workhorse for analyzing price discovery under asymmetric information, not only because of its tractability but also because later extensions preserved the linear structure while introducing stochastic noise-trading volatility, thereby capturing state-dependent liquidity and the empirically observed links between volume, volatility, and impact **?**. At its core, Kyle's contribution was to embed a strategic, risk-neutral insider into a rational expectations setting with competitive market makers who observe only aggregate order flow, while noise trading sustains volume and camouflages informed trades. The resulting equilibrium is linear: prices equal the conditional expectation of fundamentals given total flow, and the constant price impact parameter succinctly captures adverse selection.

## 2.2 Model

We consider a single risky asset whose terminal value $v$ is uncertain. Before trading begins, this fundamental value $v$ is drawn from a normal distribution with parameters $v \sim \mathcal{N}(\mu_0, \sigma_v^2)$, where $\mu_0 \in \mathbb{R}$ represents the common prior expectation about the asset's value and $\sigma_v^2 > 0$ captures the degree of fundamental uncertainty. These parameters are publicly known, i.e. they reflect the collective assessment of market participants about the asset before any private information acquisition takes place. The key innovation of Kyle's framework is the presence of an **informed trader** who, unlike other market participants, observes the true realization of $v$ before trading. This trader essentially possesses perfect information about the fundamental value, creating a stark information asymmetry. However, this informational advantage comes with a strategic challenge: how to exploit private knowledge without fully revealing it through trading behavior. To make informed trading viable, Kyle introduces noise (liquidity) traders whose order $u$ is distributed as $u \sim \mathcal{N}(0, \sigma_u^2)$, independent of the fundamental $v$. These traders represent participants who trade for reasons unrelated to the asset's fundamental value, they might be selling to meet liquidity needs, rebalancing portfolios, or responding to other non-informational motives. Crucially, the parameter $\sigma_u^2 > 0$ is known to all participants. Noise trading serves three essential functions in the model. First, it provides camouflage for informed orders: when market makers observe total order flow, they cannot perfectly distinguish between informed and uninformed components. Second, it ensures market viability: without noise, any order would immediately reveal the informed trader's signal, making information valueless. Third, it creates equilibrium depth: the presence of noise trading allows for a linear price impact that doesn't completely eliminate informed trading profits.

### 2.2.1 The Linear Equilibrium Ansatz

Competitive, risk-neutral market makers observe only the aggregate order flow $y = x + u$, where $x$ is the informed trader's order. They cannot observe $x$ and $u$ separately—please note that this observational limitation is crucial for maintaining the information asymmetry that drives the model. Being competitive, market makers earn zero expected profits in equilibrium. Being risk-neutral, they set prices to equal their conditional expectation of the asset's value given the information available to them. This leads to the semi-strong efficient pricing condition

$$P(y) = \mathbb{E}[v \mid y]$$

This pricing rule reflects rational expectations: market makers use all available information (the order flow $y$) to form the best possible estimate of the fundamental value, and they set the price equal to this estimate. The model seeks a linear equilibrium where strategies take simple, tractable forms. The parameter $\beta$ represents the trading intensity or how aggressively the informed trader responds to deviations of the fundamental from its prior mean. The parameter $\lambda$ is the price impact coefficient, or how much prices move in response to each unit of order flow. The reciprocal $1/\lambda$ measures market depth i.e. the order size needed to move prices by one unit.

### 2.2.2 The Informed Trader's Problem

Now we turn to the informed trader's optimization problem, which embodies the central tension in the model: the desire to profit from private information versus the concern about moving prices adversely. The informed trader knows the true value $v$ and conjectures that market makers will set prices according to $P(y) = \mu_0 + \lambda y$ with some positive $\lambda$. Given this pricing rule, the trader's profit from submitting order $x$ is:

$$\begin{aligned} \pi &= x(v - P) \\ &= x\big(v - \mu_0 - \lambda(x + u)\big) \end{aligned}$$

The trader profits $x(v - \mu_0)$ from the difference between the true value and the prior expectation, but suffers a cost $\lambda x^2$ from the price impact of their own trade, plus a random component $\lambda x u$ from the interaction with noise trading. Taking the conditional expectation given $v$ (so that $\mathbb{E}[u \mid v] = 0$), we obtain

$$\mathbb{E}[\pi \mid v] = x(v - \mu_0) - \lambda x^2$$

This is a quadratic objective in $x$. The first term represents the expected gain from trading on the information advantage, while the second term represents the expected cost of price impact. The optimal trade balances these forces: trade more when the fundamental deviates further from the prior mean, but moderate the trade size to avoid excessive price impact.nThe first-order condition $\frac{\partial \mathbb{E}[\pi|v]}{\partial x} = 0$ yields:

$$v - \mu_0 - 2\lambda x = 0$$

Solving for $x$ gives the insider's best response:

$$x(v) = \frac{v - \mu_0}{2\lambda} \tag{2.1}$$

**Proposition 2.1** (Insider best response under linear pricing)**.** *Fix a conjectured linear pricing rule $P(y) = \mu_0 + \lambda y$ with $\lambda > 0$. Then the insider's optimal order given $v$ is*

$$x(v) = \frac{v - \mu_0}{2\lambda},$$

*so in any linear equilibrium one must have $\beta = \frac{1}{2\lambda}$.*

This reveals the key insight: the informed trader's optimal strategy is indeed linear in the fundamental, with trading intensity $\beta = \frac{1}{2\lambda}$. The trader trades more aggressively (higher $\beta$) when price impact is low (low $\lambda$), and more conservatively when price impact is high. The factor of $\frac{1}{2}$ emerges from the quadratic nature of the price impact cost—this is the familiar result from monopolistic pricing where the markup is half the demand slope.

### 2.2.3 Market Makers and Linear Bayesian Updating

Having established the informed trader's optimal strategy, we now turn to the market makers' problem. Market makers must infer the fundamental value from the order flow they observe, knowing that this flow contains both informed and noise components.

When the informed trader uses the strategy $x(v) = \beta(v - \mu_0)$, the total order flow becomes $y = \beta(v - \mu_0) + u$. This creates a linear relationship between the unobservable fundamental $v$ and the observable order flow $y$, contaminated by the noise term $u$.

Since both $v$ and $u$ are normally distributed and independent, the joint distribution of $(v, y)$ is bivariate normal. This Gaussian structure allows us to apply the linear projection formula for conditional expectations. The market makers' optimal pricing rule is to set the price equal to the conditional expectation of the fundamental given the observed order flow:

$$P(y) = \mathbb{E}[v \mid y] = \mu_0 + \lambda y$$

To determine the slope coefficient $\lambda$, we use the fact that for jointly normal random variables, the conditional expectation is linear with slope equal to the ratio of covariance to variance:

$$\lambda = \frac{\mathrm{Cov}(v, y)}{\mathrm{Var}(y)}$$

Under the informed trading strategy $y = \beta(v - \mu_0) + u$, we can compute these moments. The covariance between $v$ and $y$ is $\mathrm{Cov}(v, y) = \beta\sigma_v^2$, since $u$ is independent of $v$. The variance of the order flow is $\mathrm{Var}(y) = \beta^2\sigma_v^2 + \sigma_u^2$, reflecting both the variability induced by informed trading and the exogenous noise.

Therefore, the price impact coefficient is:

$$\lambda = \frac{\beta\sigma_v^2}{\beta^2\sigma_v^2 + \sigma_u^2}$$

The equilibrium requires that both the informed trader's best response and the market makers' pricing rule be mutually consistent. We have two equations: $\beta = \frac{1}{2\lambda}$ from the informed trader's optimization, and $\lambda = \frac{\beta\sigma_v^2}{\beta^2\sigma_v^2+\sigma_u^2}$ from the market makers' inference problem. These two conditions must be satisfied simultaneously.
Substituting the first into the second yields a quadratic equation that can be solved to obtain the unique positive solution:

$$\beta = \frac{\sigma_u}{\sigma_v}, \qquad \lambda = \frac{\sigma_v}{2\sigma_u}$$

These equilibrium values reveal important economic intuitions. The trading intensity $\beta$ increases with noise variance $\sigma_u^2$ and decreases with fundamental variance $\sigma_v^2$. More noise provides better camouflage, encouraging more aggressive informed trading. Conversely, higher fundamental uncertainty makes each unit of information less precise, leading to more cautious trading.
The price impact $\lambda$ decreases with noise variance and increases with fundamental variance. Market depth, measured by $1/\lambda = 2\sigma_u/\sigma_v$, is higher when there is more noise trading relative to fundamental uncertainty. This captures the intuitive idea that markets with more noise trading can absorb informed orders with less price movement.

### 2.2.4 Price Informativeness and Learning

A crucial question in any model of asymmetric information is how much private information gets revealed through the trading process. In Kyle's model, this can be measured by comparing the prior uncertainty about the fundamental with the posterior uncertainty after observing the order flow. The posterior variance of the fundamental given the order flow is calculated using the standard formula for conditional variance in the bivariate normal case

$$\text{Var}(v \mid y) = \sigma_v^2 - \frac{\text{Cov}(v, y)^2}{\text{Var}(y)}$$

Substituting the equilibrium values, we find:

$$\text{Cov}(v, y) = \beta\sigma_v^2 = \frac{\sigma_u}{\sigma_v} \cdot \sigma_v^2 = \sigma_u\sigma_v$$

$$\text{Var}(y) = \beta^2\sigma_v^2 + \sigma_u^2 = \frac{\sigma_u^2}{\sigma_v^2} \cdot \sigma_v^2 + \sigma_u^2 = 2\sigma_u^2$$

Therefore:

$$\text{Var}(v \mid y) = \sigma_v^2 - \frac{(\sigma_u \sigma_v)^2}{2\sigma_u^2} = \sigma_v^2 - \frac{\sigma_v^2}{2} = \frac{\sigma_v^2}{2}$$

This remarkable result shows that a single trading round in the Kyle model reveals exactly half of the prior variance, regardless of the noise level $\sigma_u^2$. This invariance property is a distinctive feature of the Kyle equilibrium and reflects the endogenous adjustment of trading intensity to noise levels.

### 2.2.5 Profitability and Market Impact

The informed trader's expected profit provides another lens through which to understand the equilibrium. From the quadratic optimization problem, the conditional expected profit given the fundamental realization is:

$$\mathbb{E}[\pi \mid v] = \frac{(v - \mu_0)^2}{4\lambda}$$

Taking expectations over the fundamental gives the ex ante expected profit

$$\mathbb{E}[\pi] = \frac{\mathbb{E}[(v - \mu_0)^2]}{4\lambda} \\ = \frac{\sigma_v^2}{4\lambda}$$

The expression reveals that expected profits increase with both fundamental uncertainty (more valuable information) and noise trading (better camouflage). The profit is proportional to the geometric mean of the two variance parameters, highlighting the complementary nature of information value and camouflage. The comparative statics of market depth deserve special attention. Since $\lambda = \sigma_v/(2\sigma_u)$, price impact decreases with noise variance, while market depth $1/\lambda = 2\sigma_u/\sigma_v$ increases linearly with noise variance. This endogenous relationship between noise trading and market liquidity is central to understanding how markets self-organize around information asymmetries.

**Matcha with Ayyar**

*Over a warm cup of matcha, let's pause and think about something that often confuses students...*

**Question:** "I keep hearing about different sources of trading costs in market microstructure models. In the Kyle model, where exactly do these costs come from? Are they the same as the inventory costs I read about in other papers?"

Great question! This is actually a subtle but important distinction that gets to the heart of what drives spreads and price impact in different market structures. In Kyle's model, the trading costs arise purely from **adverse selection**. Here's what's happening: the market makers know that some of the orders they see come from informed traders who know more about the asset's true value. This creates a classic "winner's curse" problem; when market makers get hit by a large order, it's more likely to be coming from someone who knows bad news (if it's a sell order) or good news (if it's a buy order). To protect themselves from this adverse selection, market makers build the expected cost into their pricing. This shows up as the price impact parameter $\lambda = \frac{\sigma_v}{2\sigma_u}$, which measures how much the price moves per unit of order flow. The key insight is that this impact exists even though market makers are risk-neutral and competitive, they're not worried about holding inventory per se, they're worried about being picked off by better-informed traders.

Now, **inventory costs** are a different animal entirely. They arise when market makers are risk-averse and worry about the risk of holding positions. Consider a dealer with CARA utility $U(w) = -\exp(-\gamma w)$ who holds inventory $q$. Their certainty equivalent from this position is

$$CE = q(\mu_0 - P) - \frac{\gamma}{2} q^2 \sigma_v^2$$

That second term $\frac{\gamma}{2} q^2 \sigma_v^2$ is pure inventory cost. It increases quadratically with position size and reflects the dealer's aversion to bearing risk. Crucially, this cost exists even if there's no asymmetric information at all! The beauty of Kyle's framework is that it isolates the adverse selection channel cleanly. The risk-neutral assumption strips away inventory concerns, leaving us with a pure laboratory to study how private information gets impounded into prices through strategic trading. In real markets, of course, both effects are likely present, but understanding them separately is crucial for empirical work that tries to decompose bid-ask spreads into their components.

# Chapter 3

# Dynamic Information Revelation

The single-period Kyle model elegantly captures how informed trading, market making, and noise provision interact to incorporate private information into prices. Yet real markets are dynamic: information arrives over time, traders adapt their strategies, and strategic interactions evolve. Extending Kyle's framework to multiple periods transforms a static snapshot into a dynamic theory of information-based price discovery. Kyle himself introduced the multi-period extension in his 1985 Econometrica paper, showing that the insider's problem becomes one of dynamic programming: current profits must be balanced against the information revealed to market makers, which alters future opportunities. This temporal trade-off introduces genuine intertemporal strategy—far more than a repetition of the single-period game. Kyle's discrete-time formulation demonstrated linear equilibria through recursive difference equations.

## 3.1 Multiperiod Kyle

The breakthrough came with Back (1992), who proved that as trading intervals shrink, the discrete model converges to a tractable continuous-time limit. This insight provided the mathematical foundation for a generation of advances in dynamic microstructure theory. In continuous time, the Kyle framework has inspired extensive research: multiple insiders, dynamic information acquisition, stochastic noise volatility, funding constraints, disclosure requirements, and correlated signals across assets. The unifying theme is that informed traders manage information intertemporally—trading less aggressively early on to preserve private information, then accelerating as horizons shorten. The model predicts rich dynamics: market depth typically increases as the terminal date approaches; price informativeness rises as uncertainty resolves; and trading intensity follows time-varying patterns shaped by noise, horizon, and signal precision. Unlike the static case where price impact depends only on the signal-to-noise ratio, in multi-period settings impact itself becomes a forward-looking process, influenced by expectations of future order flow and information release. Crucially, the multi-period Kyle model retains linear equilibrium structure, allowing for closed-form characterizations despite the dynamic complexity. This combination of tractability and depth explains why it remains a cornerstone of market microstructure, with implications for optimal execution, high-frequency trading, and the design of modern electronic markets.

## 3.2 Model Setup

The multiperiod Kyle model preserves the three-player structure of the single-period version-informed trader, noise traders, and competitive market makers while introducing intertemporal dynamics that fundamentally alter strategy. Its power lies in combining simple Gaussian-linear assumptions with dynamic optimization, allowing tractable analysis of how information gets revealed and prices adjust over time. A finite horizon is assumed, both for analytical convenience (backward induction via dynamic programming) and for economic realism: private information often expires (e.g., earnings announcements, merger outcomes, or patent approvals), creating urgency and shaping trading incentives. The model runs for $T$ discrete periods, with a risky asset of terminal value $v \sim \mathcal{N}(\bar{v}, \sigma_v^2)$. At time zero, the informed trader learns the true $v$, while the market only knows the prior. Each period, the informed trader chooses an order $x_t$, noise traders submit independent demands $u_t \sim \mathcal{N}(0, \sigma_u^2)$, and market makers observe the total flow $y_t = x_t + u_t$. Prices update via conditional expectation, $p_t = \mathbb{E}[v \mid y_1, \dots, y_t]$. The informed trader's challenge is dynamic: trading too aggressively reveals information and reduces future profits, while trading too cautiously underutilizes the informational advantage. The problem is thus an optimal control problem balancing immediate gains against preserving information rents across time.

### 3.2.1 Strategic Interaction

The informed trader solves a dynamic programming problem, choosing the sequence $(x_1, \dots, x_T)$ to maximize expected cumulative profit. Current trades affect both immediate returns and the informativeness of future prices, creating intertemporal externalities. Optimal strategies typically imply declining trading intensity as the horizon shortens. Noise traders supply the camouflage that sustains informed trading. Their period-by-period independent orders represent liquidity needs unrelated to fundamentals, providing the randomness that prevents perfect inference by market makers. Market makers, observing only aggregate flows, update beliefs using Bayesian inference. Thanks to the Gaussian-linear structure, this process admits closed-form characterization through Kalman filter recursions. Prices form a martingale that gradually converges to the true value as information is revealed. The outcome is a dynamic process of price discovery that predicts how market depth, informativeness, and liquidity evolve over time.

**Matcha with Ayyar**

*Let me pour some matcha and think about what changes when we go dynamic...*

**Question:** "I understand the one-period Kyle model, but I'm confused about the multiperiod version. If the informed trader knows $v$ from the beginning, why doesn't he just trade his entire position immediately in period 1 to maximize profits?"

Think of it this way: if the informed trader dumps his entire desired position in period 1, the massive order flow would cause a huge price movement. The market makers, seeing this large order, would infer that someone has very strong information about the asset's value. This would cause prices to move most of the way to the fundamental value immediately! So while the trader gets high profits per unit traded in period 1 (since the price hasn't moved much yet), he's "killed the golden goose"—there's no information advantage left for periods 2 through $T$. The optimal strategy involves a delicate balance: trade enough today to capture some profits, but not so much that you give away all your informational advantage. It's like being a poker player who knows everyone's cards—you want to win money, but if you bet too aggressively on every hand, everyone will figure out that you're cheating! This creates a beautiful dynamic optimization problem where the informed trader is essentially deciding how fast to reveal his private information to the market.

### 3.2.2 Price Process and Information Revelation

A central insight of the multiperiod Kyle model is that prices are martingales under the public filtration generated by order flow. Let $\mathcal{F}_t = \sigma(y_1, \ldots, y_t)$. With competitive, risk–neutral market makers,

$$p_t \;=\; \mathbb{E}x[v \mid \mathcal{F}_t] \quad \Rightarrow \quad \mathbb{E}[p_t \mid \mathcal{F}_{t-1}] = p_{t-1},$$

and under linear–Gaussian structure the pricing rule takes the form

$$p_t \;=\; p_{t-1} + \lambda_t\, y_t, \qquad y_t \;=\; x_t + u_t, \qquad u_t \sim \mathcal{N}(0, \sigma_u^2),$$

where $\lambda_t > 0$ is the period-$t$ price–impact (inverse depth). The cumulative decomposition

$$v - \bar{v} \;=\; \sum_{t=1}^{T} \lambda_t\, y_t \;+\; \varepsilon_T$$

holds with a terminal residual $\varepsilon_T$ that reflects remaining (posterior) uncertainty about $v$ after $T$ periods. The linear–Gaussian setting implies Kalman–filter updates for the posterior variance:

$$\sigma_t^2 \;\equiv\; \mathrm{Var}(v \mid \mathcal{F}_t) \;=\; \sigma_{t-1}^2 \;-\; \frac{\mathrm{Cov}(v, y_t \mid \mathcal{F}_{t-1})^2}{\mathrm{Var}(y_t \mid \mathcal{F}_{t-1})} \;=\; \sigma_{t-1}^2\,\frac{\sigma_u^2}{\beta_t^2\sigma_{t-1}^2 + \sigma_u^2},$$

once we specify the insider's linear strategy $x_t = \beta_t\,(v - p_{t-1})$. Equivalently, the period-$t$ *information revelation rate* is

$$\rho_t \;\equiv\; \frac{\sigma_{t-1}^2 - \sigma_t^2}{\sigma_{t-1}^2} \;=\; \frac{\beta_t^2\sigma_{t-1}^2}{\beta_t^2\sigma_{t-1}^2 + \sigma_u^2} \in (0,1),$$

so $\sigma_t^2 = \sigma_{t-1}^2(1-\rho_t)$ and hence $\sigma_T^2 = \sigma_v^2 \prod_{t=1}^{T}(1-\rho_t)$.

### 3.2.3 The Informed Trader's Dynamic Problem

Let $\Delta_t \equiv v - p_t$ and $W_t(p_t, v)$ be the insider's continuation value from period $t$. A single trade in period $t$ at order $x_t$ yields *expected* one–period profit

$$\mathbb{E}_t\Big[(v - p_t)x_t - \lambda_t x_t^2 - \lambda_t x_t u_t\Big] \;=\; \Delta_t x_t \;-\; \lambda_t x_t^2,$$

and pushes the next price via $p_{t+1} = p_t + \lambda_t(x_t + u_t)$, so $\Delta_{t+1} = \Delta_t - \lambda_t(x_t + u_t)$. With discount factor $\delta \in (0,1]$, the Bellman equation is

$$W_t(p_t, v) \;=\; \max_{x_t}\; \mathbb{E}_t\Big[\,\Delta_t x_t - \lambda_t x_t^2 + \delta\, W_{t+1}(p_{t+1}, v)\,\Big].$$

**Proposition 3.1** (Quadratic value function and optimality condition)**.** *There exist coefficients $\{A_t, B_t\}_{t=1}^{T+1}$ with terminal condition $A_{T+1} = B_{T+1} = 0$ such that*

$$W_t(p_t, v) \;=\; A_t\,\Delta_t^2 + B_t.$$

*Given $A_{t+1}$ and $\lambda_t$, the insider's period-$t$ optimal order is linear,*

$$x_t^\star \;=\; \beta_t\,\Delta_t, \qquad \beta_t \;=\; \frac{2\delta A_{t+1}\lambda_t - 1}{2\lambda_t\,(1 - \delta A_{t+1}\lambda_t)}\,,$$

*and the value–function coefficient satisfies the backward recursion*

$$A_t \;=\; \frac{1}{4\,\lambda_t\,(1 - \delta A_{t+1}\lambda_t)}.$$

*Moreover,* $B_t = \delta B_{t+1} + \frac{\delta A_{t+1}\lambda_t^2(\delta A_{t+1}\lambda_t - 1)}{1 - \delta A_{t+1}\lambda_t}\,\sigma_u^2.$

*Proof.* Plug the quadratic ansatz into the Bellman equation; take expectations using $\mathbb{E}[u_t] = 0$ and $\mathbb{E}[u_t^2] = \sigma_u^2$; maximize the resulting quadratic in $x_t$. The first–order condition yields the stated $\beta_t$. Substituting $x_t^\star$ back gives the recursions for $A_t$ and $B_t$. □

### 3.2.4 Market Maker Pricing and Dynamic Consistency

With $x_t = \beta_t \Delta_{t-1}$, the covariance and variance terms are

$$\mathrm{Cov}(v, y_t \mid \mathcal{F}_{t-1}) = \beta_t \sigma_{t-1}^2, \qquad \mathrm{Var}(y_t \mid \mathcal{F}_{t-1}) = \beta_t^2 \sigma_{t-1}^2 + \sigma_u^2,$$

hence competitive pricing implies

$$\lambda_t \;=\; \frac{\beta_t \sigma_{t-1}^2}{\beta_t^2 \sigma_{t-1}^2 + \sigma_u^2} \quad \text{and} \quad p_t \;=\; p_{t-1} + \lambda_t y_t.$$

Equations in Proposition 3.1 together with the pricing and variance updates

$$\sigma_t^2 \;=\; \sigma_{t-1}^2 \frac{\sigma_u^2}{\beta_t^2 \sigma_{t-1}^2 + \sigma_u^2}$$

jointly characterize equilibrium via backward–forward recursion.

## 3.3 Equilibrium Characterization

**Theorem 3.2** (Linear equilibrium: existence, uniqueness, and dynamics)**.** *Fix $\delta \in (0, 1]$, $\sigma_v^2 > 0$, and $\sigma_u^2 > 0$. There exists a unique linear equilibrium with strategies $x_t = \beta_t(v - p_{t-1})$ and prices $p_t = p_{t-1} + \lambda_t y_t$ such that for $t = 1, \dots, T$:*

$$\beta_t \;=\; \frac{2\delta A_{t+1}\lambda_t - 1}{2\lambda_t\,(1 - \delta A_{t+1}\lambda_t)}, \quad A_t \;=\; \frac{1}{4\,\lambda_t\,(1 - \delta A_{t+1}\lambda_t)}, \quad \lambda_t \;=\; \frac{\beta_t \sigma_{t-1}^2}{\beta_t^2 \sigma_{t-1}^2 + \sigma_u^2}.$$

*In the canonical case $\delta = 1$ with homoscedastic noise $\sigma_u^2$ and a single insider:*

1. *Trading intensity is* increasing *over time: $\beta_1 < \beta_2 < \cdots < \beta_T$.*

2. *Price impact is* decreasing *over time: $\lambda_1 > \lambda_2 > \cdots > \lambda_T$.*

3. *Information revelation accelerates: $\rho_1 < \rho_2 < \cdots < \rho_T$ and $\rho_T = \frac{1}{2}$.*

4. *Period-$t$ expected profit is*

$$\mathbb{E}[\pi_t] \;=\; \mathbb{E}[\Delta_{t-1} x_t - \lambda_t x_t^2] \;=\; \frac{\beta_t \sigma_{t-1}^2\, \sigma_u^2}{\beta_t^2 \sigma_{t-1}^2 + \sigma_u^2},$$

*so total expected profit is* $\sum_{t=1}^{T} \mathbb{E}[\pi_t]$.

**Intuition.** Early on, the insider protects future rents (trades cautiously), but the large residual uncertainty $\sigma_{t-1}^2$ makes each unit of order flow less informative, leading to higher depth (lower $\lambda_t$) later only after enough information is revealed. By the terminal period, the insider behaves as in a one–shot Kyle game with remaining variance $\sigma_{T-1}^2$, revealing exactly half of that variance.

**Matcha with Ayyar**

"Why does $\beta_t$ rise over time while $\lambda_t$ falls?"

Urgency grows as the horizon shrinks, pushing the insider to trade more aggressively (rising $\beta_t$). At the same time, previous trading has already reduced posterior variance, so each unit of new order flow is *less* masked by noise relative to the shrinking uncertainty set. Market makers therefore need *less* slope to extract the same information (falling $\lambda_t$), and the terminal step always reveals half of what remains. The two forces—urgency vs. remaining uncertainty—jointly generate rising intensity but falling impact.

### Two-Period Model ($T = 2$)

Backward induction yields:

$$\lambda_2 = \frac{\sigma_1}{2\sigma_u}, \qquad \beta_2 = \frac{\sigma_u}{\sigma_1}, \qquad \sigma_1^2 = \sigma_v^2 \frac{\sigma_u^2}{\beta_1^2\sigma_v^2 + \sigma_u^2}, \qquad \lambda_1 = \frac{\beta_1\sigma_v^2}{\beta_1^2\sigma_v^2 + \sigma_u^2}.$$

The optimal $\beta_1$ maximizes $\mathbb{E}[\pi_1] + \mathbb{E}[\pi_2]$, delivering $\beta_2 > \beta_1$ and $\lambda_2 < \lambda_1$. Closed forms follow from the first–order condition but are omitted for brevity.

### Continuous-Time Limit

Let $\Delta t = 1/T \to 0$. The discrete model converges to a continuous–time Kyle economy (Back, 1992) in which

$$dP_t = \Lambda(t)\, dY_t, \qquad dY_t = \beta(t)\,(v - P_t)\, dt + dU_t,$$

with $\{U_t\}$ a Brownian motion with variance rate $\sigma_u^2$. The residual variance $\sigma(t)^2 = \mathrm{Var}(v \mid \mathcal{F}_t)$ solves a Riccati–type ODE, and $\Lambda(t) = \frac{1}{2\sigma_u}\,\sigma(t)$ while $\beta(t)$ increases as time to maturity shrinks.