# Private Information and Trade Timing

*By* LONES SMITH*

This paper investigates the Bayesian decision-theoretic foundations of the popular Wall Street refrain that "timing is everything." I explore the purely informational motives for trade timing: If public information is being released over time, when is a financially and informationally small and partially informed trader's "informational edge" over the market at its zenith? In other words, if a trader will not affect the price by his transaction, and he has a fixed investment budget, when should he invest?

One might think that with no strategic or other reasons to time his trade, and so long as all information is conditionally independent, he wishes to act at once on any private information—because delay sees only its depreciation. This paper questions this intuition by exploring an informational difference between the calculus of prices and their reciprocals, returns. The paper builds on a timing irrelevance result for Arrow securities, to deduce economically interpretable sufficient conditions for timing impatience of general securities.

Investors trade among themselves because they have either different information or preferences, usually diverse risk-sharing needs. The most important and interesting problems in financial timing concern how this drama unfolds in an equilibrium setting, where markets recognize the informational content in the order flow. One thread of research, as exemplified by Jiang Wang (1994), takes inspiration from the asymmetric information rational expectation pricing literature. More recently, strategic effects have received the lion's share of attention. A sequence of papers starting with Albert S. Kyle (1985) have investigated trade timing given a few large agents; in these settings, trade volume conveys private information about the asset's value, and thereby affects current and future prices. Prime examples of this work insofar as it affects timing include: the analyses with many informed insiders in Douglas Foster and S. Viswanathan (1996), or timing for purely risk-sharing reasons in Dimitri Vayanos (1999), or by uninformed investors in Anat Admati and Paul Pfleiderer (1988).

It is in the context of such advanced and increasingly applied equilibrium analyses that I ask a very simple and idealized timing question. Absent any price effects, when is it best to spend a dollar? Properly understanding this purely decision theory question may possibly allow deeper insight into equilibrium timing questions still unresolved: for instance, the speed that private information actually gets impounded into prices, or the incentives to engage in technical (uninformed) trading.

Optimal financial timing is often associated with the added complexity of American over European style put or call options. There, the problem is when to sell or buy a share, in advance of a given date. This paper in essence attacks the dual problem of when to spend a dollar. It thus concerns the calculus of returns, rather than share prices. Moreover, the interaction between private and public information has been a nonissue in the option pricing literature, but is the central focus of this paper.

To motivate the question I examine, consider the following purely illustrative scenario. An employee of a biotech firm has gotten wind of an announcement next Friday of the unexpected FDA approval of his firm's soon-to-be major drug. The insider wishes to profit from this

* Department of Economics, University of Michigan, Ann Arbor, MI 48109 (e-mail and web page: econtheorist@earthling.net and www.umich.edu/~lones/). This paper is a renaming and radical revision of a March 1995 mimeo, "On the Irrelevance of Trade Timing." I owe special thanks to Max Amarante for reasoning that led me to the current Jensen proof, to Susan Athey for kindly showing me how to use her theorem, and to Terry Tao (UCLA) for help with some correlation notions. I am also grateful for comments on the usenet group sci.math.research—including Peter Matthews (University of Maryland, Baltimore County), Dan Cass (St. John Fisher College in Rochester), and Dražen Pantić. John Geanakoplos, Thomas Piketty, Stephen Ross, Peter Sørensen, Angel Serrat Tubert, Jiang Wiang, and Richard Zeckhauser offered helpful feedback. Three referees greatly impacted how I wrote this paper. I acknowledge the National Science Foundation for financial support. Any errors and opinions expressed are mine alone.

information, but for whatever reason he cannot buy on margin—a restriction that will ring truer when I consider only imperfect information; similarly, with unfavorable information, assume he likewise faces a short-sale constraint. Let him have a given amount of money with which to invest by buying stocks, and assume he wishes to maximize his *return* (or his "bang per buck"). Further assume that his purchases do not affect the price. When should he purchase?

It so turns out that under classical assumptions—Arrow securities and conditionally independent private and public information—provided he trades in advance of the public announcement, his timing does not matter. To illustrate and illuminate this unexpected result, let us consider the following simple two-period, two-state model. An asset pays off, respectively, 1 or 0 in the "high" and "low" states, $H$ and $L$. Discounting is ignored and traders are assumed risk neutral, so that the initial asset price $p_0$ is the prior public chance of $H$. After the insider receives private information that the payoff is 1, he wishes to buy as many shares as his wealth permits, taking the price as given. If he purchases now, his return per dollar is $1/p_0$. If he waits until the next period's partially informative public signal is revealed, then the expected price is intuitively higher. But his expected return is unchanged. Why is this? The more strongly that a given public signal indicates state $H$, the higher the resulting price. For instance, if a signal doubles the price, then it clearly must be twice as likely in state $H$ as it is *ex ante*; this means that the insider is twice as confident that his return will halve as would be an uninformed market observer. More generally, the insider's future return is exactly inversely related to his relative confidence versus the market in that event; therefore, his expected return (i.e., knowing state $H$) is the same as if the public signal was totally uninformative. Thus timing does not matter, as the return per dollar invested is expected to remain constant.

The paper generally explores trade timing by building on this insight. Obviously, trade timing is not always moot. A purely graphical argument establishes it for any partially informed (informationally small) trader—so long as the public and his private information are conditionally independent, that is, given the asset value.[1] The second assumption is crucial, and restrictive—the asset must be an *Arrow security*, paying a dividend of 1 in one state of the world and 0 elsewhere. For compound securities, the inverse relationship between an insider's future return and his confidence relative to the market only separately holds for each state. To understand the general timing problem, therefore, one must decompose the public-private informational interaction for general securities across states of nature.

The graphical apparatus also shows that with assets bundling two Arrow securities, generically there is a strict incentive to trade at once on any private information. But for general securities, timing is ambiguous, because it depends on the possibly rich interaction of private and public information across states of nature. To make sense of this relationship, public signals and private information must be stochastically ordered.

My main proposition gives sufficient conditions for timing impatience: If (i) private information is either good or bad news about the asset payoff and (ii) public information has a conditional density family given any positive asset value ordered by the *monotone likelihood ratio property* (MLRP), then one should trade immediately on any private information. Consistent with the preceding results, these two conditions are automatically satisfied by Arrow securities with equality, and condition (ii) by two-state securities. Together, (i) and (ii) force a comovement of the private and public estimates of the asset value, damping the incentive to wait and learn from the public information. This structure emerges rather naturally in many settings, since many common parametric density families used by economists happen to satisfy the MLRP (see Susan Athey, 1996). Timing impatience thus has not only intuitive appeal, but also a compelling economic basis for it.

The next section outlines an elementary model of informational trade timing. The third section offers a diagrammatic joint proof of the one- and two-state results that highlight the role

[1] Conditional independence is automatically satisfied in any perfect information context like the previous example, because conditioning on the state reveals nothing more than conditioning on one's private signal.

of conditional independence. I then give the timing impatience proposition and an example violating its assumptions with delay of trading.

## I. The Formal Model

### A. *Information Structure*

I build on the now-standard model of Paul Milgrom and Nancy Stokey (1982) and partition the state space $\Omega$ into two statistically related components $\Omega = \Theta \times \Sigma$. Think of each $\theta \in \Theta$ as the *state*, with $\Theta$ a finite set; $\theta$ alone determines the ultimate payoff. The set $\Sigma$ consists of the range of the payoff irrelevant signals and is partitioned into finitely many events $\mathcal{S}$. All signals are then random variables that map the payoff-relevant state space $\Theta$ into the signal outcomes $\mathcal{S}$. In this way, Bayesian traders observing the signal can make inferences about $\theta$. Let $\mathcal{P}$ and $\mathcal{P}(\cdot \mid \cdot)$ be the common prior and resulting conditional probability on $\Omega$.

Although I restrict to a single round of public signals, this is for ease of presentation only, and is not at all critical. The small trader with the starring role in this paper is endowed in period zero with a *private signal* $\tau$ with possible range $\{T, \Omega \backslash T\}$, and $T$ realized. He then is afforded an opportunity to trade. In period one, a *public signal* about the state of the world is realized; it has range $\mathcal{S}_1 \subseteq \mathcal{S}$, where $\mathcal{S}_1 = \{S_1, S_2, \dots, S_M\}$ partitions $\Omega$. (See Figure 1.) Finally, there is another trading session, and then the asset value is realized.

The private-public informational interaction is decisive, which requires some stochastic monotonicity relations: Signal $\sigma$ is *good/bad news* for the random variable $x$ if its conditional probability $\mathcal{P}(\sigma|x)$ rises/falls in $x$. Also, the density family $\{f(x|\sigma_m)\}$ has the monotone likelihood ratio property (MLRP) if $f(x_H|\sigma_m)/f(x_L|\sigma_m)$ rises in $m$ if $x_H > x_L$.

Some assumptions are maintained throughout this paper. First, for simplicity and to rule out trivialities, (i) *all states $\theta$ are possible, or $\mathcal{P}(\theta) > 0$*, and (ii) *no private or public signal $\sigma$ is perfectly revealing*, $0 < \mathcal{P}(\sigma|\theta) < 1$ for all $\theta$. Second and critically, *private and public signals are independent, conditional on each state $\theta$.* The alternative is that private and public information are conditionally correlated. Although technically intractable, some plausible information structures produce just this result. For instance, if traders learn independently of different divisions of a firm, and if profits are the sum of its separate parts, then the separate signals are correlated given the value of the firm (i.e., the state).

### B. *Payoffs, Prices, and Returns*

Let the asset pay $\pi(\theta)$ in state $\theta$, realized and observed after the final trading session. Recall that an Arrow security is an asset with payoff $\pi(\theta) = 1$ in some state $\theta = \theta_0$, and 0 otherwise. *Compound securities* are linear combinations of these building blocks. Discounting adds nothing to the story, and so I shall ignore it. The price of the asset is thus its publicly expected payoff; in notation, we have $p_0 \equiv \mathscr{E}[\pi]$ after period 0 and $p_1 \equiv \mathscr{E}[\pi|\mathcal{S}_1]$ after period one. Embedded here is the standard assumption that the asset is efficiently priced given risk-neutral traders. In other words, $p_0 = \mathscr{E}[p_1]$ if there are no arbitrage opportunities using public information alone.

Given private information $\tau$, if the trader is initially more optimistic than the market, then we must be sure that he remains so after any new information arrivals $\sigma$; otherwise, he might switch sides of the market. A standard assumption precludes this.

LEMMA 1: *If the private signal $\tau$ is good or bad news, then the trader never will switch sides of the market after any possible public signal $\sigma$: If $\mathscr{E}[\pi|\tau] \gtreqless \mathscr{E}[\pi]$, then $\mathscr{E}[\pi|\sigma, \tau] \gtreqless \mathscr{E}[\pi|\sigma]$.*

Observe that the no "crossover" conclusion of Lemma 1 (see Appendix for proof) is always satisfied for an Arrow security, simply because any private signal prompting one to buy (sell) must be good (bad) news. For more general securities, the good or bad news assumption is not without loss of generality.

This paper assumes that the privately informed trader invests a given amount, and so cares about the return per dollar. This holds even for short-selling—buying negative units of the share worth \$1, say, and later netting out by buying the shares back after the asset value is realized. To be definite, I focus on the purchase scenario.

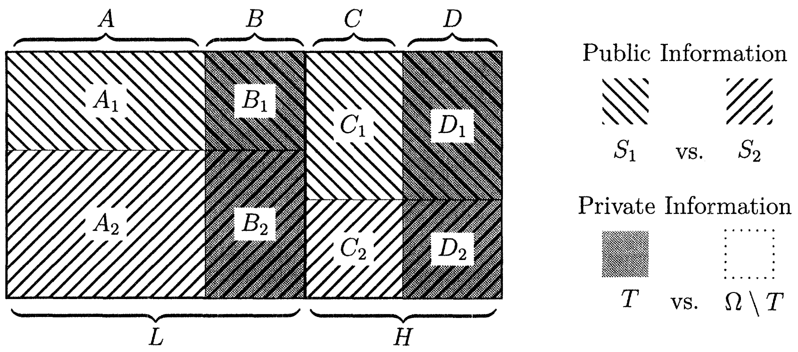


FIGURE 1. ILLUSTRATION OF STATE AND SIGNAL SPACE.

*Notes:* Consider the rectangle above with area 1, corresponding to the total probability mass. Let $A = A_1 \cup A_2$, etc. The state is either $L = A \cup B$ or $H = C \cup D$. Private information reveals the partition elements $T = B \cup D$ or $\Omega \backslash T = A \cup C$, and public information reveals $S_1 = A_1 \cup B_1 \cup C_1 \cup D_1$ or $S_2 = A_2 \cup B_2 \cup C_2 \cup D_2$. That $\{S_1, S_2\}$, $\{T, \Omega \backslash T\}$ are geometrically orthogonal partitions in $L$, $H$ implies conditional independence.

In this stylized setting, we wish to know whether the return of an immediate investment of \$1 exceeds, equals, or is below the expected payoff from deferring this investment:

(1) Current return

$$= \mathscr{E}[\pi|T]/p_0 \gtreqless \mathscr{E}[\mathscr{E}[\pi|\mathscr{S}_1, T]/p_1]|T]$$

$$= \text{expected future return.}$$

Observe how the trader and market alike learn from the public information next period. Since the trader is privately informed, he must average over all his possible future expected returns based on all available information, weighted by his initial privately assessed conditional probability. His private information thus twice feeds into this expectation.

## II. Motivation and Results

### A. *Graphical Motivation for Arrow and Two-State Securities*

To begin with, suppose there are two states and two public signals, as depicted in Figure 1. The associated algebra is in fact valid for many signals.

For definiteness, assume an asset pays off 1 in state $H$, and $\ell < 1$ in state $L$. The initial price is thus $(\ell(A + B) + C + D)/(A + B + C + D)$, whereas the trader's more refined private estimate is $(\ell B + D)/(B + D)$. So the initial expected *net* private return on a purchase is

(2)
$$\frac{(\ell B + D)/(B + D)}{(\ell(A + B) + C + D)/(A + B + C + D)} - 1$$

$$= \frac{(AD - BC)(1 - \ell)}{(B + D)(\ell(A + B) + C + D)}.$$

As is visually apparent, $AD - BC > 0$, and thus private information favors state $H$; buying is then privately optimal given favorable information $\tau = T$. Weighting net returns like (2) by the perceived public signal chances, the expected next private return likewise reduces to:

(3)
$$\sum_i \frac{B_i + D_i}{B + D} \frac{(A_i D_i - B_i C_i)(1 - \ell)}{(B_i + D_i)[(A_i + B_i)\ell + (C_i + D_i)]}$$

$$= \sum_i \frac{(AD - BC)(1 - \ell)}{(B + D)[(A + B)\ell/\rho_i + (C + D)/\psi_i]},$$

where the last equality holds by a cancellation of $(B_i + D_i)$, and given the graphically clear implications of conditional independence, that $A_i/A = B_i/B \equiv \psi_i$ and $C_i/C = D_i/D \equiv \rho_i$.

Namely, the conditional chances of signal $S_i$ in states $L$ and $H$ do not depend on $T$ versus $\Omega \backslash T$.

For an Arrow security, the payoff in state $L$ is $\ell = 0$, and the initial net return (2) reduces to the final one (3), given $\psi_1 + \psi_2 = 1$: Public signal $S_i$ only impacts the return in a state by its chance $\psi_i$, independent of private information. But for two-state securities, I introduce the net return function $f(x) = a/(c + b/x)$ for buying, with $a = (AD - BC)(1 - \ell)$, $b = (A + B)\ell(B + D)$, and $c = (C + D)(B + D)$. As $f$ is concave, Jensen's inequality applies:

$$\begin{aligned} \text{RHS (3)} &= \psi_1 f(\rho_1/\psi_1) + \psi_2 f(\rho_2/\psi_2) \\ &\le f(\psi_1(\rho_1/\psi_1) + \psi_2(\rho_2/\psi_2)) \\ &= f(1) = \frac{a}{b+c} = \text{RHS (2)}. \end{aligned}$$

The expected final return is then at most the initial one, and immediate trade is best. In fact, so long as public information nuances between states $L$ and $H$ (i.e., $\psi_1 \neq \rho_1$, so that the asset is not informationally equivalent to an Arrow security), this preference is strict. But if the buyer has a small chance of crossing over to sell, then (3) strictly understates the true expected return, and there may well be a strict incentive to delay.

### B. *The Sufficient Conditions for Timing Impatience with General Securities*

Timing does not matter with Arrow securities, as the informational premium remains constant. With two-state securities impatience is the rule. What can we learn from this? The main lesson is that the cross-interaction of private and public information across states is important—conditionally independent private and public signals need not be identically distributed. One trades at once when one expects to lose more by having one's private information indirectly revealed than one profits by learning of the public information. Since this occurs when the private estimate rises *proportionately* more than the public estimate, this is a standard intuition that suggests the importance of informational *log-complementarity*. That $f$ is *log-supermodular* jointly in $(x, \sigma)$ is equivalent to $\langle f(x|\sigma)\rangle$ having the MLRP. The proof of Proposition 1 is presented in the Appendix.

PROPOSITION 1: *Timing impatience is optimal if the private signal $\tau$ is good news for the asset payoff $\pi$ (if buying; bad news if selling), and the public information and asset payoff distribution $\langle \mathcal{P}(\pi|S_m)\rangle$ jointly have the MLRP, across all states $\theta$ with positive payoffs.*

The assumptions ensure that private and public information work in concert to reveal the asset payoff and stifle any incentive to delay. The MLRP inequality is only weakly satisfied for an Arrow security, since all but one state has zero payoff. With two states, one can always order the public signals to give the MLRP; this proves our earlier results.

COROLLARY 1: *Any partially informed trader with good/bad news, and a fixed amount to invest (a) is indifferent about how he times his profitable trades of an Arrow security, and (b) has a (strict) incentive to trade at once for (generic) assets paying on two states.*

Observe that the likelihood ratio $(1 - p)/p$ of $L$ to $H$ is the *net return* on the Arrow security paying on state $H$. Thus, in the perfect information example of the introduction, timing irrelevance for Arrow securities is simply a financial interpretation of the classical result that likelihood ratios are a martingale if one conditions on the true state (see Chow et al., 1971). This martingale motivation has independently led Peter Bossaerts (1996) through similar analysis as in part (*a*) in an unrelated study of survivorship bias in asset returns.

### C. *An Example with Delay*

I now provide a minimal three-state example of delay when the assumptions are violated. Let payoffs in states $\theta_1$, $\theta_2$, and $\theta_3$ be 1, 2, 3, respectively. *Ex ante*, all three states are equilikely. Altogether, this yields an initial asset price of 2. The trader is assumed to have information $\tau = T$, resulting in posterior conditional beliefs (1/3, 1/6, 1/2) skewed toward $\theta_3$. So the private signal realization $T$ is neither good nor bad news for $X$. The expected private return per dollar is then $[(1/3)1 + (1/6)2 + (1/2)3]/2 = (13/6)/2 = 13/12$.

Let the realized public signal be $S_1$ or $S_2$ with respective state-dependent chances (1/6, 2/3, 1/6) or (5/6, 1/3, 5/6); here, $S_2$ tips beliefs toward $\theta_1$ and $\theta_3$, and $S_1$ toward $\theta_2$—and the MLRP fails. As shown in the Appendix, the expected final return is

$$\sum_m \frac{\sum_\ell \mathcal{P}(S_m|\theta_\ell)\mathcal{P}(\theta_\ell|T)\pi(\theta_\ell)\ \sum_\ell \mathcal{P}(S_m|\theta_\ell)\mathcal{P}(\theta_\ell)}{\sum_\ell \mathcal{P}(S_m|\theta_\ell)\mathcal{P}(\theta_\ell)\pi(\theta_\ell)}. \tag{4}$$

Of course, this formula is only valid provided that the trader will never switch sides of the market, which is no longer guaranteed, since $\tau = T$ is not good or bad news. Still, one can check that the returns following both signals strictly exceed 1. Plugging the preceding data into (4) yields an unchanged expected return of 13/12. Because it is a knife-edge, it is strictly optimal to delay until after the public information is revealed, for mildly perturbed data. For instance, the expected return rises to 854/754 if $\pi(\theta_3) = 4$.

## III. Conclusion

This paper has studied the behavior of a trader who plausibly has finite wealth and limited borrowing capacity. The substance of my insight is that the trader then cares about the expected reciprocal prices, and that these behave in a surprising fashion. A referee has suggested a different type of constraint, that the trader with a bad-news scoop is endowed with a finite number of stock shares; this is perhaps a more natural budget constraint when one is selling stock. More commonly, a fixed-share, rather than monetary, budget arises if one possesses put or call options on the stock that one must exercise by some deadline. Here, one can see that such a trader should simply sell at once, as the price is expected to fall: For, since the market price is a martingale to an uninformed observer, it is a supermartingale (downward drifting) to someone possessing bad news; the proof of Lemma 1 (see Appendix) also implies this. This difference shows how trade-timing decisions depend on the precise nature of the budget constraint.

## Appendix

PROOF OF LEMMA 1:

Assume the ordering $X_1 < X_2 < \cdots$. We want to show that if $\tau = T$ is good news, then $\sum_k X_k \mathcal{P}(X_k|T, \sigma) > \sum_k X_k \mathcal{P}(X_k|\sigma)$ for all public signals $\sigma$. This will be true certainly if $\mathcal{P}(X_k|T, \sigma)/\mathcal{P}(X_k|\sigma)$ is increasing in $k$, which is true, given conditional independence:

$$\frac{\mathcal{P}(X_k|T, \sigma)}{\mathcal{P}(X_k|\sigma)} = \frac{\mathcal{P}(T|X_k, \sigma)}{\mathcal{P}(T|\sigma)} = \frac{\mathcal{P}(T|X_k)}{\mathcal{P}(T|\sigma)}.$$

PROOF OF PROPOSITION 1:

Assume a purchase. Let the random variable $\mathbb{I}_T$ indicate the event $T$, and define the random variable $X$ by $X_k = \pi(\theta_k)$. Define random variables $Y$ and $Z$ on the public information partition by $Y_m = \mathcal{E}(X|S_m)$ and $Z_m = \mathcal{E}(X \cdot \mathbb{I}_T|S_m)$. By modifying (4), the timing impatience inequality in (1) holds if and only if $\mathcal{E}(Z)/\mathcal{E}(Y) > \mathcal{E}(Z/Y)$, or when $Z/Y$ and $Y$ are positively correlated, given $Z \equiv (Z/Y)Y$.

A simple sufficient condition for a positive correlation of sequences $\langle Y_m \rangle$ and $\langle Z_m/Y_m \rangle$ is that both be increasing. Because the density family $\langle \mathcal{P}(X|S_m) \rangle$ has the MLRP by assumption, $\langle Y_m \rangle$ is increasing. Next, with the dummy variable $t \in \{0, 1\}$, define a density $h(X|S_m, t) = \mathcal{P}(X|S_m)$ for $t = 0$ and $\mathcal{P}(X, T|S_m)$ for $t = 1$. Now, $\langle Z_m/Y_m \rangle$ is increasing if and only if $H(S_m, t) \equiv \sum xh(x|S_m, t)$ is log-supermodular. By Corollary 5.2.1 of Athey (1996) or, more simply, Theorem 3-5.1 in Karlin's (1968) classic, this follows if $Xh(X|S_m, t)$ is jointly log-supermodular in $(X, S_m, t)$, since summation preserves this property. For this, the pairwise log-supermodularity of the density $h(X|S_m, t)$ suffices. In other words, $\mathcal{P}(X, T|S_m)/\mathcal{P}(X|S_m)$ is increasing in $m$, and the families $\langle \mathcal{P}(X|S_m) \rangle$ and $\langle \mathcal{P}(X, T|S_m) \rangle$ have the MLRP. The first proviso holds if $\tau = T$ is good news for $X$. Indeed,

$$\mathcal{P}(X, T|S_m)/\mathcal{P}(X|S_m) = \mathcal{P}(T|X, S_m) = \mathcal{P}(T|X),$$

by conditional independence, while $\mathcal{P}(T|X)$ is increasing in $X$ when $\tau = T$ is good news. The second condition is given, while the last one follows, because

$$\frac{\mathcal{P}(X, T|S_{m+1})}{\mathcal{P}(X, T|S_m)}$$

$$= \frac{\mathcal{P}(S_{m+1}|X, T)\mathcal{P}(X, T)/\mathcal{P}(S_{m+1})}{\mathcal{P}(S_m|X, T)\mathcal{P}(X, T)/\mathcal{P}(S_m)}$$

$$= \frac{\mathcal{P}(S_{m+1}|X)\mathcal{P}(S_m)}{\mathcal{P}(S_m|X)\mathcal{P}(S_{m+1})}$$

by conditional independence, and this is increasing in $X$ if $\langle \mathcal{P}(X|S_m) \rangle$ has the MLRP.

PROOF OF EQUATION (4):

$$\mathscr{E}[\mathscr{E}[\pi|T, \mathscr{S}_1]/p_1|T]$$

$$= \sum_m \mathcal{P}(S_m|T) \frac{\sum_\ell \mathcal{P}(\theta_\ell|T, S_m)\pi(\theta_\ell)}{\sum_\ell \mathcal{P}(\theta_\ell|S_m)\pi(\theta_\ell)}$$

$$= \sum_m \frac{\sum_\ell \mathcal{P}(S_m, \theta_\ell|T)\pi(\theta_\ell)}{\sum_\ell \mathcal{P}(S_m|\theta_\ell)\mathcal{P}(\theta_\ell)\pi(\theta_\ell)/\mathcal{P}(S_m)}.$$

Next, substitute $\mathcal{P}(S_m, \theta_\ell|T) = \mathcal{P}(S_m|\theta_\ell, T)\mathcal{P}(\theta_\ell|T) = \mathcal{P}(S_m|\theta_\ell)\mathcal{P}(\theta_\ell|T)$ by the conditional independence of $\sigma$ and $\tau$, while $\mathcal{P}(S_m) = \sum_\ell \mathcal{P}(S_m|\theta_\ell)\mathcal{P}(\theta_\ell)$ enters the numerator.

## REFERENCES


**Admati, Anat and Pfleiderer, Paul.** "A Theory of Intraday Patterns: Volume and Price Variability." *Review of Financial Studies*, Spring 1988, *1*(1), pp. 3–40.

**Athey, Susan.** "Comparative Statics Under Uncertainty: Single Crossing Properties and Log-Supermodularity." Working Paper No. 96-22R, Massachusetts Institute of Technology, 1996.

**Bossaerts, Peter.** "Martingale Restrictions of Equilibrium Prices of Arrow-Debreu Securities Under Rational Expectations and Consistent Beliefs." Working paper, California Institute of Technology, 1996.

**Chow, Y. S.; Robbins, Herbert and Siegmund, David.** *Great Expectations: The Theory of Optimal Stopping*. Boston: Houghton Mifflin, 1971.

**Foster, Douglas and Viswanathan, S.** "Strategic Trading When Agents Forecast the Forecasts of Others." *Journal of Finance*, September 1996, *51*(4), pp. 1437–78.

**Karlin, Samuel.** *Total Positivity*. Stanford, CA: Stanford University Press, 1968.

**Kyle, Albert S.** "Continuous Auctions and Insider Trading." *Econometrica*, November 1985, *53*(6), pp. 1315–35.

**Milgrom, Paul and Stokey, Nancy.** "Information, Trade, and Common Knowledge." *Journal of Economic Theory*, February 1982, *26*(1), pp. 17–27.

**Vayanos, Dimitri.** "Strategic Trading in a Dynamic Noisy Market." *Review of Economic Studies*, April 1999, *66*(2), pp. 219–54.

**Wang, Jiang.** "A Model of Competitive Stock Trading Volume." *Journal of Political Economy*, February 1994, *102*(1), pp. 127–68.