DETECTING ASSET VALUE DISLOCATIONS IN MULTI-AGENT MODELS OF MARKET MICROSTRUCTURE

Vikram Krishnamurthy and Anup Aryan

The University of British Columbia
Department of Electrical and Computer Engineering
2332 Main Mall, Vancouver, BC, V6T 1Z4, Canada

ABSTRACT

Consider a financial market participant observing the trade flow of an asset traded through a limit order book. Trades are driven by an agent-based model where individual agents observe the trading decisions of previous agents, as well as their private signal on the value of the asset and then execute a trading decision. Given trading decisions of agents, how can a market observer detect a shock to the underlying value of the traded asset? The distribution of shock times is assumed to be phase-type distributed to allow for a general set of change time probabilities beyond geometric change times. We show that this problem is equivalent to change detection with social learning. We provide structural results that allow the optimal detection policy to be characterized by a single threshold policy.

***Index Terms*—** Computational Finance, Agent-based Models, Social Learning, Quickest Change Detection

1. INTRODUCTION

Models of agent based computational finance view the market through its microstructure as the net effect of interacting and learning (boundedly-rational) agents. Agent based models (ABMs) have been used to capture empirical stylized facts observed in markets such as "fat tails", correlation of returns and volatility clustering which implies long term memory [1, 2, 3, 4, 5]. A particular area of interest are models of agents that are performing social learning. In models of social learning, agents adapt their behavior based on the trading decisions of previous agents. The study of social learning in markets have lead to many interesting results such as herding and trend-following behavior [6, 7, 8, 9, 10].

We consider agents trading in a dealer market with a monopolist market maker [11, 12, 13]. In this framework we are interested in studying a market event where a sudden change in the underlying value of the asset occurs. The aim is to detect the change time with minimal cost. Given local decisions of trading agents, how can a market observer achieve quickest time change detection to determine a shock in the asset value? Quickest detection is important for market participants to enable timely risk management. We will show that the optimal decision policy has multiple thresholds and the stopping regions are in general non-convex thus making global decisions (stop or continue) based on local decisions (from social learning) non-trivial. Problems that deal with the interaction of local and global decision makers are of independent interest in signal processing and multi-agent systems.

Corresponding author. Email: vikramk@ece.ubc.ca

2. MULTI-AGENT MODEL AND QUICKEST DETECTION OF MARKET SHOCKS

We model the market microstructure as a discrete time dealer market motivated by algorithmic and high-frequency tick-by-tick trading [14]. There is a single traded asset and a countable number of trading agents. Each agent acts once in a predetermined sequential order indexed by k. We assume that the asset has a true underlying value Z which is known to all traders and the market observers. Once trading progresses, the market observer does not receive direct information about Z. It is only able to observe the public buy/sell actions of agents. Agents receive noisy private observations of the underlying value. At a random time τ^0, the asset experiences a jump change in its value to a new value λZ where $\lambda > 0$. The aim of the market observer is to detect the change time (global decision) with minimal cost.

Let $y_k \in \mathbb{Y} = \{1, 2, \dots, Y\}$ denote the local (private) observation of agent k and $a_k \in \mathbb{A} = \{1 \text{ (buy order)}, 2 \text{ (no trade)}, 3 \text{ (sell order)}\}$ denote the local decision agent k takes. Define the σ-algebras:

$$\begin{aligned} \mathcal{H}_k &:= \sigma\text{-algebra generated by } (a_1, \dots, a_{k-1}, y_k), \\ \mathcal{G}_k &:= \sigma\text{-algebra generated by } (a_1, \dots, a_{k-1}, a_k). \end{aligned} \quad (1)$$

The social learning model [15, 8] comprises of the following ingredients:

1. *Shock in Asset Value*: The state variable x_k represents the underlying asset value. At time τ^0 the asset experiences a jump change (shock) in its value due to factors which are exogenous to the system. We model the change point τ^0 by a

phase type (PH) distribution. PH distributed jump times have been used as a model for price shocks and credit events such as corporate default [16, 17]. To construct the PH change time, assume the underlying state x_k evolves as a Markov chain on the finite state space $\mathbb{X} = \{1, \ldots, X\}$. Here state '1' is an absorbing state and denotes the state after the jump change. The states $2, \ldots, X$ can be viewed as a single composite state that x resides in before the jump. These states can also be labelled with the unit indicator vectors $e_1, e_2, \ldots, e_X$. The initial distribution is $\pi_0 = (\pi_0(i), i \in \mathbb{X})$, $\pi_0(i) = P(x_0 = i)$. We are only interested in the case where the change occurs after a least one time step, so assume $\pi_0(1) = 0$. The transition probability matrix P is of the form

$$P = \begin{bmatrix} 1 & 0 \\ \underline{P}_{(X-1)\times 1} & \bar{P}_{(X-1)\times(X-1)} \end{bmatrix} \tag{2}$$

Let the "change time" τ^0 denote the time at which x_k enters the absorbing state 1, i.e., $\tau^0 = \inf\{k : x_k = 1\}$. It is at this random time τ^0 that the asset experiences a jump change in its value. To ensure that τ^0 is finite, we assume states $2, 3, \ldots, X$ are transient. In the special case when x is a 2-state Markov chain, the change time τ^0 is geometrically distributed.

The vector g denotes the asset value pre and post jump. The value associated with state 1 and states $2, \ldots, X$ is $g = (\lambda Z,\ Z)'$, where $\lambda > 0$ Recall Z is the true underlying value of the asset. The constant λ models the relative size of the jump change in the asset value at time τ^0.

2. *Agent's Private Observation*: Agent k's local observation $y_k \in \mathbb{Y} = \{1, \ldots, Y\}$ is a noisy estimate of the true value of the asset [13, 18]. The noisy observations are a function of the underlying asset state and is modeled through the observation likelihood distribution

$$B_{xy} = P(y_k = y | x_k = x). \tag{3}$$

The states $2, 3, \ldots., X$ are fictitious and are defined to generate the PH-distributed change time τ^0. States $2, 3, \ldots., X$ are indistinguishable in terms of the observation y, that is, $P(y|2) = P(y|3) = \cdots = P(y|X)$ for all $y \in \mathbb{Y}$.

3. *Private belief*: Using local observation y_k, trading agent k updates its private belief (posterior) π_k^P defined as

$$\pi_k^P = (\pi_k^P(i),\ i \in \mathbb{X}), \quad \pi_k^P(i) = P(x_k = i | a_1, \ldots, a_{k-1}, y_k) \tag{4}$$

initialized by π_0. The private belief is the posterior distribution of the underlying state given the past trading decisions and current observation [10]. It is computed by agent k according to the following Bayesian Hidden Markov Model filter:

$$\pi_k^P = T(\pi_{k-1}, y_k),\ T(\pi, y) = \frac{B_y P' \pi}{\sigma(\pi, y)},\ \sigma(\pi, y) = \mathbf{1}' B_y P' \pi. \tag{5}$$

$$B_y = \text{diag}(B_{1y}, \ldots, B_{Xy}) \quad (X \times X \text{ diagonal matrix})$$

Also π_{k-1} denotes the public belief available at time $k - 1$ (defined in Step 6 below). For a PH distributed change time, the states $2, \ldots, X$ are fictitious, that is $B_{2y} = \ldots = B_{Xy}$.

4. *Agent's Trading Decision*: Agent k then makes a trading decision $a_k \in \mathbb{A} = \{$ 1 (buy order), 2 (no trade), 3 (sell order) $\}$ to minimize its expected trading loss. To formulate this, let $c(i, a)$ denote the non-negative cost incurred if the agent picks local decision a when the underlying state is $x = i$. Denote the local decision X-dimensional cost vector

$$c_a = \begin{bmatrix} c(1, a) & c(2, a) & \cdots & c(X, a) \end{bmatrix}. \tag{6}$$

Then agent k chooses local decision a_k greedily to minimize its expected cost of trading: $a_k = a(\pi_{k-1}, y_k) = \arg\min_{a \in \mathbb{A}} \{c_a' \pi_k^P\}$. In quickest change detection, since states $2, 3, \ldots, X$ are indistinguishable in terms of observation y, we assume that $c(2, a) = c(3, a) = \cdots = c(X, a)$ for each $a \in \mathbb{A}$. Based on the posted bid/ask prices, the local cost vectors c_a for action $a \in \{1, 2, 3\}$ are $c_1 = g - \bar{p}$, $c_2 = 0$ and $c_3 = \underline{p} - g$. Here $\underline{p}$ and $\bar{p}$ are the bid and ask price vectors set by the market maker as described below.

At the beginning, we assume that the bid/ask prices bracket the intrinsic value, $Z \in [\underline{p}, \bar{p}]$. Furthermore, we assume that the jump size is strictly greater than the bid/ask spread, $|1 - \lambda| > (\bar{p} - \underline{p})/Z$. If this condition is violated then action 2 dominates in all information states. Under these two assumptions, the costs of one action is dominated by the other two and can therefore be neglected. In the case where $\lambda < 1$, costs for action 1, c_1, are dominated by the cost vector of actions 2 and 3. In the case where $\lambda < 1$, it is action 3 which is dominated. The associated cost matrix restricted to the two action reduction is sub-modular.

5. *Market Making Mechanism*: At each time period, the market maker sets bid and ask prices, $\theta_k = (\underline{p}_k, \bar{p}_k)$, at which he will respectively buy and sell one unit of the asset. The bid/ask prices may be thought of as activation levels for the agent's actions. The dealer sets prices based on its belief of the true value Z. The knowledge of Z is driven only by the public information available, that is the past actions of agents. We assume that the market maker earns zero expected profits and sets prices according to $\underline{p}_k = E[Z | a_{k-1} = \text{sell}]$ and $\bar{p}_k = E[Z | a_{k-1} = \text{buy}]$, see [11]. In a special case of change detection, we assume that the time scale of the price update by the market maker is much slower than the time scale of arriving agents. This allows us to consider θ as a constant over the time horizon of interest to change detection.

6. *Social Learning and Public Belief*: Agent k's local decision a_k is recorded in the order book and hence broadcast publicly. Subsequent agents $\bar{k} > k$ use decision a_k to update their public belief of the underlying state x_k as follows: Define the public belief π_k as the posterior distribution of the state x given all local decisions taken up to time k.

$$\pi_k = \mathbb{E}\{x_k | \mathcal{G}_k\} = (\pi_k(i),\ i \in \mathbb{X}), \\ \pi_k(i) = P(x = i | a_1, \ldots a_k) \tag{7}$$

initialized by π_0. Then agents $\bar{k} > k$ update their public belief according to the following "social learning Bayesian filter":

$$\pi_k = T^{\pi_{k-1}}(\pi_{k-1}, a_k), \text{ where } T^\pi(\pi, a) = \frac{R_a^\pi P'\pi}{\sigma(\pi, a)},$$
$$\sigma(\pi, a) = \mathbf{1}'_X R_a^\pi P'\pi \quad (8)$$

We use the notation $T^\pi(\cdot)$ to point out that the above Bayesian update map depends explicitly on the belief state π. This is a key difference compared to the HMM filter (5) where the Bayesian update map $T(\cdot)$ does not depend explicitly on belief state π. In (8), R_a^π denotes the diagonal matrix $R_a^\pi = \text{diag}(R_{i,a}^\pi, i \in \mathbb{X})$ where $R_{i,a}^\pi = P(a_k = a | x_k = i, \pi_{k-1} = \pi)$ denotes the conditional probability that agent k chose local decision a given state i. We call $R_{i,a}^\pi$ as the *local decision likelihood probabilities* in analogy to observation likelihood probabilities B_{iy} (3) in classical filtering.

Clearly, observing the local decision a_k taken by agent k yields information about its local observation y_k. The local decision likelihood probability matrix R^π in the social learning Bayesian filter (8) is computed as $R^\pi = BM^\pi$ where

$$M_{y,a}^\pi \triangleq P(a|y, \pi) = \prod_{\tilde{a} \in \mathbb{A} - \{a\}} I(c_a' B_y P'\pi < c_{\tilde{a}}' B_y P'\pi). \quad (9)$$

The main implication of is that the social learning Bayesian filter (8) is discontinuous in the belief state π, due to the presence of indicator functions.The likelihood probabilities R^π are an explicit function of the belief state π – this is stark contrast to the standard quickest detection problems where the observation distribution is not an explicit function of the posterior distribution. Note that the public belief belongs to the unit $X - 1$ dimensional simplex.

2.1. Market Observer's Detection Strategy

With the above social learning based local decision framework, we turn our attention to the market observer (global decision maker). The market observer seeks to achieve quickest detection balancing cost of delayed detection while minimizing false alarms. The detection of change point may signal an arbitrage opportunity or other trading action to the market observer. The formulation presented considers a general parameterization of the costs associated with detection delay and false alarms. The specific costs based on a particular the course of action to be taken by the global decision maker are not considered.

At each time k, given the public belief π_k, let u_k denote the decision made by the market observer: $u_k = \mu(\pi_k)$ which may be either 1) announce change and stop or 2) continue. Thus the global decision u_k is $\mathcal{G}_k$ measurable, where $\mathcal{G}_k$ is defined in (1). The policy μ belongs to the class of stationary decision policies denoted $\boldsymbol{\mu}$. Below we formulate the costs incurred when taking these global decisions u_k.

(i) *Cost of announcing change and stopping*: If global decision $u_k = 1$ is chosen, then the social learning protocol terminates. The false alarm event $\cup_{i \geq 2}\{x_k = i\} \cap \{u_k = 1\} = \{x_k \neq 1\} \cap \{u_k = 1\}$ represents the event that a change is announced before the change happens at time τ^0. To evaluate the *false alarm penalty*, let $f_i I(x_k = i, u_k = 1)$ denote the cost of a false alarm in state i, $i \in \mathbb{X}$, where $f_i \geq 0$. Of course, $f_1 = 0$ since a false alarm is only incurred if the stop action is picked in states $2, \ldots, X$. The expected false alarm penalty is

$$\bar{C}(\pi_k, u_k = 1) = \sum_{i \in \mathbb{X}} f_i \mathbb{E}\{I(x_k = i, u_k = 1) | \mathcal{G}_k\} = \mathbf{f}'\pi_k, \quad (10)$$

where $\mathbf{f} = (f_1, \ldots, f_X)'$, $f_1 = 0$. The false alarm vector $\mathbf{f}$ is chosen with increasing elements so that states further from state 1 incur larger penalties.

(ii) *Delay cost of continuing*: If global decision $u_k = 2$ is taken then the social learning protocol continues to time $k+1$ and a delay cost is incurred. The expected delay cost is

$$\bar{C}(\pi_k, u_k = 2) = d\,\mathbb{E}\{I(x_k = 1, u_k = 2) | \mathcal{G}_k\} = d e_1' \pi_k \quad (11)$$

where $d > 0$ denotes the delay cost.

2.2. Quickest Time Detection Objective

For any $\pi_0 \in \Pi(X)$, and policy $\mu \in \boldsymbol{\mu}$, there exists a (unique) probability measure $\mathbb{P}_{\pi_0}^\mu$ on $(\Omega, \mathcal{F})$ [19]. Let $\mathbb{E}_{\pi_0}^\mu$ denote the expectation with respect to the measure $\mathbb{P}_{\pi_0}^\mu$.

Let τ denote a stopping time adapted to the sequence of σ-algebras $\mathcal{G}_k, k \geq 1$. With decision policy u_k $\tau = \inf\{k : u_k = 1\}$. For each initial distribution $\pi_0 \in \Pi(X)$, and policy μ, the following cost is associated:

$$J_\mu(\pi_0) = \mathbb{E}_{\pi_0}^\mu \{\sum_{k=1}^{\tau-1} \rho^{k-1} \bar{C}(\pi_k, u_k = 2)$$
$$+ \rho^{\tau-1} \bar{C}(\pi_\tau, u_\tau = 1)\}. \quad (12)$$

Here $\rho \in [0, 1]$ denotes a discount factor. Since $\bar{C}(\pi, 1)$, $\bar{C}(\pi, 2)$ are non-negative and bounded for all $\pi \in \Pi(X)$, stopping is guaranteed in finite time.

The goal of the global decision maker is to determine the change time τ^0 with minimal cost, that is, compute the optimal global decision policy $\mu^* \in \boldsymbol{\mu}$ to minimize (12), where $J_{\mu^*}(\pi_0) = \inf_{\mu \in \boldsymbol{\mu}} J_\mu(\pi_0)$.

2.3. Market Observer's Optimal Strategy

The optimal stationary policy $\mu^* : \Pi(X) \to \{1, 2\}$, and associated value function $\bar{V}(\pi)$, are the solution of "Bellman's dynamic programming equation": $J_{\mu^*}(\pi_0) = \bar{V}(\pi_0)$

$$\mu^*(\pi) = \arg\min\{\bar{C}(\pi, 1),\ \bar{C}(\pi, 2) \quad (13)$$
$$+ \rho \sum_{a \in \mathbb{A}} \bar{V}\left(T^\pi(\pi, a)\right) \sigma(\pi, a)\}$$

with the value function $\bar{V}(\pi)$ defined by the associated minimization.

Define, $C = de_1 - (I - \rho P)\mathbf{f}$, with elements denoted as C_j, $j = 1, \ldots, X$. Let $\mathcal{Q}_l$, $l = 1, \ldots, 2^Y$, denote the elements of the power set of $\mathbb{Y}$ (excluding, of course, the empty set). Define the following 2^Y convex polytopes $\bar{\mathcal{P}}_l$, $l = 1, 2, \ldots, 2^Y$:

$$\bar{\mathcal{P}}_l = \left\{ \pi \in \Pi(X) : \begin{cases} (c_1 - c_2)' B_y P' \pi < 0 & y \in \mathcal{Q}_l \\ (c_1 - c_2)' B_y P' \pi \geq 0 & y \in \mathbb{Y} - \mathcal{Q}_l \end{cases} \right\} \tag{14}$$

Although in general there are 2^Y possible R^π matrices, we now show that by introducing assumptions (A1), (A2) and (S) below, there are only $Y + 1$ distinct local decision likelihood matrices R^π. This forms an important preliminary step for characterizing the optimal global decision policy. We list the following assumptions.

- (A1) The observation distribution $B_{xy} = p(y|x)$ is TP2 i.e., all second order minors of matrix B are non-negative [20].
- (A2) The transition probability matrix P is TP2.
- (A3) The elements of vector C are decreasing.
- (S) The local decision cost vector c_a in (6) is submodular.

Theorem 1 *Under (A1), (A2), (S), the belief state space $\Pi(X)$ can be partitioned into at most $Y + 1$ non-empty polytopes denoted $\mathcal{P}_1, \ldots, \mathcal{P}_{Y+1}$ where*

$$\begin{aligned} \mathcal{P}_1 &= \{\pi \in \Pi(X) : (c_1 - c_2)' B_1 P' \pi \geq 0\} && (15) \\ \mathcal{P}_l &= \{\pi \in \Pi(X) : (c_1 - c_2)' B_{l-1} P' \pi < 0 && (16) \\ &\qquad \cap \ (c_1 - c_2)' B_l P' \pi \geq 0\}, \ l = 2, \ldots, Y \\ \mathcal{P}_{Y+1} &= \{\pi \in \Pi(X) : (c_1 - c_2)' B_Y P' \pi < 0\} \end{aligned}$$

On each such polytope, the local decision likelihood matrix R^π is a constant with respect to belief state π. ■

The proof of the above theorem is in [21]. As a consequence of Theorem 1, there are only $Y + 1$ possible decision likelihood matrices R^π, one per polytope $\mathcal{P}_l$, $l = 1, \ldots, Y + 1$.

3. NUMERICAL RESULTS

In this section we provide numerical results to demonstrate the results presented. We consider geometrically and phase type distributed change times. In this example, the underlying price starts at $Z = 2$ and at a random time the value drops to $Z = 1$. The following parameters were used in the simulation; geometrically distributed change time, $p = .99$, $\underline{p} = 1.5$, $\bar{p} = 2.1$, quantized Gaussian observation noise with standard-deviation $\Sigma = 1$, $f = 0.25$, $d = 0.1$ and $\rho = 0.75$. Figure 1(a) shows optimal policy. Note the multiple switching behavior of the optimal policy. Figure 1(b) shows the associated value function $\bar{V}$.

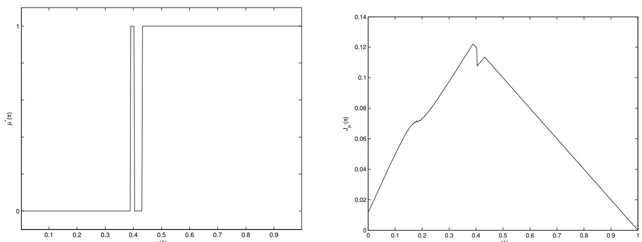

(a) (b)

Fig. 1. (a) Optimal decision policy $\mu^*(\pi)$ for social learning based quickest time change detection for geometric distributed change time. (b)Value function $\bar{V}(\pi)$ is non-concave and discontinuous.

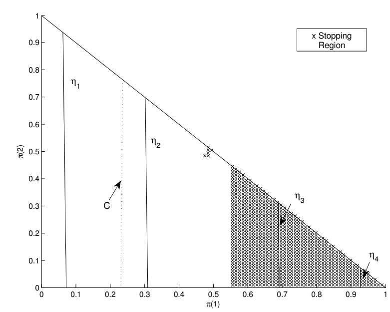

Fig. 2. Optimal Decision Policy under fixed bid/ask prices under phase-type distributed change time.

In Figure 2 we illustrate the optimal detection policy for a phase distributed change time. All parameters remain unchanged except for the transition matrix $\bar{P}_{(2\times 2)} = [p, 0; 1 - p, p]$. The plot shows the Euclidean projection of the polytope. The axes represent belief state components $\pi(1)$ and $\pi(2)$. The boundaries of the convex polytopes $\mathcal{P}_l$ (1) are shown, denoted by the hyperplanes η, as well as the cost hyperplane C. Once again, note the multiple threshold nature of the optimal policy.

4. RELATION TO PRIOR WORK

In this work, we seek to characterize optimal trading signal detection in multi-agent models with social learning through stochastic control techniques. Previous work has concentrated on empirical results of multi-agent simulations [1, 3, 4] and market making in ABMs [11, 13]. This work is a novel application of multi-agent signal processing and control methods to the problem of change detection in dynamical models affected by social learning. The proof of the structure of the market maker's optimal strategy involves submodularity on the lattice of posterior Bayesian distributions, see [22, 21].

5. REFERENCES

[1] V Alfi, M Cristelli, L Pietronero, and A Zaccaria, "Minimal agent based model for financial markets I," *The European Physical Journal B*, vol. 67, no. 3, pp. 385–397, Jan. 2009.

[2] D. Challet, M. Marsili, and Y.C. Zhang, *Minority Game: Interacting Agents in Financial Markets*, Oxford University Press, 2005.

[3] Giorgio Fagiolo, Alessio Moneta, and Paul Windrum, "A Critical Guide to Empirical Validation of Agent-Based Models in Economics: Methodologies, Procedures, and Open Problems," *Computational Economics*, vol. 30, no. 3, pp. 195–226, Sept. 2007.

[4] M Cristelli, L Pietronero, and A Zaccaria, "Critical Overview of Agent-Based Models for Economics ," *ArXiv*, Jan. 2011.

[5] R Cont, "Volatility Clustering in Financial Markets: Empirical Facts and Agent–Based Models," *Long memory in economics; edited by A Kirman and G Teyssiere (Springer)*, 2007.

[6] R Cont and J-P Bouchaud, "Herd Behavior And Aggregate Fluctuations In Financial Markets," *Macroeconomic Dynamics*, vol. 4, pp. 170–196, May 2000.

[7] T Lux and M Marchesi, "Volatility Clustering In Financial Markets - A Microsimulation Of Interacting Agents," *Int. J. Theor. Appl. Finance*, vol. 3, no. 4, pp. 675–702, 2000.

[8] C.P. Chamley, *Rational Herds: Economic Models of Social Learning*, Cambridge University Press, 2004.

[9] A Park and H Sabourian, "Herding and Contrarian Behaviour in Financial Markets," *Econometrica*, vol. 79, no. 4, pp. 973–1026, Nov. 2011.

[10] M. Cipriani and A. Guarino, "Estimating a Structural Model of Herd Behavior in Financial Markets," *IMF Working Paper*, Oct. 2010.

[11] L Glosten, "Insider Trading, Liquidity, and the Role of the Monopolist Specialist," *Journal of Business*, vol. 62, no. 2, pp. 211–235, 1989.

[12] Sanmay Das, "A learning market-maker in the Glosten–Milgrom model," *Quantitative Finance*, vol. 5, no. 2, pp. 169–180, Apr. 2005.

[13] S Das and M Magdon-Ismail, "Adapting to a Market Shock: Optimal Sequential Market-Making," *In NIPS*, pp. 361–368, Sept. 2008.

[14] A Cartea and S Jaimungal, "Modeling asset prices for algorithmic and high frequency trading," *Preprint University of Toronto*, pp. 1–29, Sept. 2011.

[15] S Bikhchandani, D Hirshleifer, and I Welch, "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades," *J. Political Economy*, vol. 100, no. 5, pp. 992–1026, 1992.

[16] S Asmussen, F Avram, and M Pistorius, "Russian and American put options under exponential phase–type Levy models," *Stochastic Processes and their Applications*, vol. 109, no. 1, pp. 79–111, Oct. 2004.

[17] S Asmussen, D Madan, and M Pistorius, "Pricing Equity Default Swaps under an approximation to the CGMY L evy Model," *Journal of Computational Finance*, vol. 11, no. 2, pp. 79–93, Dec. 2007.

[18] J Doyne Farmer, P Patelli, and I Zovko, "The Predictive Power of Zero Intelligence in Financial Markets," *PNAS*, vol. 102, no. 6, pp. 2254–2259, Dec. 2005.

[19] O. Hernandez-Lerna and J.B. Laserre, *Discrete-Time Markov Control Processes: Basic Optimality Criteria*, Springer-Verlag, N.Y., 1996.

[20] S Karlin and Y Rinott, "Classes of Orderings of Measures and Related Correlation Inequalities. I. Multivariate Totally Positive Distributions," *Journal of Multivariate Analysis*, vol. 10, pp. 467–498, 1980.

[21] V. Krishnamirthy, "Quickest Detection POMDPs With Social Learning: Interaction of Local and Global Decision Makers," *IEEE Transactions on Information Theory*, vol. 58, no. 8, pp. 5563–5587, Aug. 2012.

[22] S. Athey, "Monotone comparative statics under uncertainty," *The Quarterly Journal of Economics*, vol. 117, no. 1, pp. 187–223, 2002.