Learning Agents in Financial Markets: Consensus Dynamics on Volatility

Extended Abstract

Tushar Vaidya
SUTD
Singapore
tushar_vaidya@mymail.sutd.edu.sg

Carlos Murguia
University of Melbourne
Melbourbe, Victoria
carlos.murguia@unimelb.edu.au

Georgios Piliouras
SUTD
Singapore
georgios@sutd.edu.sg

ABSTRACT

Black-Scholes (BS) is the standard mathematical model for European option pricing in financial markets. Option prices are calculated using an analytical formula whose main inputs are strike (at which price to exercise) and volatility. The BS framework assumes that volatility remains constant across all strikes, however, in practice it varies. How do traders come to learn these parameters? We introduce and analyze the convergence properties of natural models of learning agents, in which they update their beliefs about the true implied volatility based on the opinions of other traders.

KEYWORDS

opinion dynamics; volatility smiles; trading agents

ACM Reference Format:
Tushar Vaidya, Carlos Murguia, and Georgios Piliouras. 2018. Learning Agents in Financial Markets: Consensus Dynamics on Volatility. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

A European option is the right to buy or sell an underlying asset at a fixed point in the future at a fixed price, also known as the strike. A call option gives the right to buy an asset and a put option gives the right to sell an asset at the agreed price. On the opposite side of the buyer is the seller who has relinquished his control of exercise. Buyers of puts and calls can exercise the right to buy or sell. The payoff of a buyer of a call option with stock price S_T at expiry time T and exercise price K is $\max\{S_T - K, 0\}$, whereas for a put option is $\max\{K - S_T, 0\}$.

To get a price we input the current stock price S_0 (e.g. $101), the exercise price K (e.g. $90), the expiry T (e.g. three months from today) and the volatility σ in the Black-Scholes (BS) formula and out comes the answer, the quoted price of the instrument [8].

$$\text{Price} = BS(S_0, K, T, \sigma).$$

Volatility, which captures the beliefs about how turbulent the stock price will be, is left up to the market. This parameter is so important that in practice the market trades European calls and puts by quoting volatilities.[1] How does the market decide about what the quoted volatility should be (e.g. for a stock index in 3 months from now) is a critical, but not well understood, question.

If the underlying asset and the time to exercise T (e.g. 3 months) are the same, one would expect the volatility to be the same at different strikes. In practice, however, markets after the 1987 crash have evolved to exhibit different volatilities. This rather strange phenomenon is referred to as the smile, or smirk (see figure 1). Depending on the market, these smirks can be more or less pronounced [14].

Related work. Multiple learning models have been studied in economics [16] and machine learning/AI [18, 25, 31]. We focus on simultaneous observational learning and not sequential learning common in many theoretical economics models [32]. Specific models on order placement by traders are developed in [22, 34].

Discrete dynamical systems are now standard tools in opinion formation [20]. As in [3, 23] we use a linear dynamical systems framework. The work closest to us in spirit is that of [4] though our focus is on building novel tractable models for volatility smiles. Our models are quite distinct from previous works.

2 CONSENSUS (AGENT DYNAMICS)

We assume that the agents are able to learn how far off they are from the true volatility by informational channels in the marketplace. There are many avenues, platforms and private online chat rooms that provide quotes for option prices; some of these are stale and some are fresh. The agents' learning ability determines the quality of the feedback from all these sources. We aggregate all of this information in the form of a feedback controller. If the agents are fast learners, they adjust their volatility estimates quickly.

2.1 Consensus with Feedback

We model this feedback by introducing an extra driving term into the opinion dynamics. In particular, we feedback the difference between the agents' opinion and the true volatility $\sigma(K, T)$ scaled by a *learning coefficient* $\epsilon_i \in (0, 1)$. We assume that $\sigma(K, T)$ is invariant, i.e., for some fixed $\bar{\sigma} \in (0, 1)$, $\sigma(K, T) = \bar{\sigma}$ for some fixed strike K and maturity M. The opinion $x_t^i \in \mathbb{R}$ of the i-th agent is given by

$$x_t^i = \sum_{j=1}^{n} a_{ij} x_{t-1}^j + \epsilon_i(\bar{\sigma} - x_{t-1}^i), \tag{1}$$

or in matrix form

$$X_t = AX_{t-1} + \mathcal{E}(\bar{\sigma}\mathbf{1}_n - X_{t-1}), \tag{2}$$

[1]Using the Black-Scholes formula with particular implied volatility, traders obtain a dollar value price.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

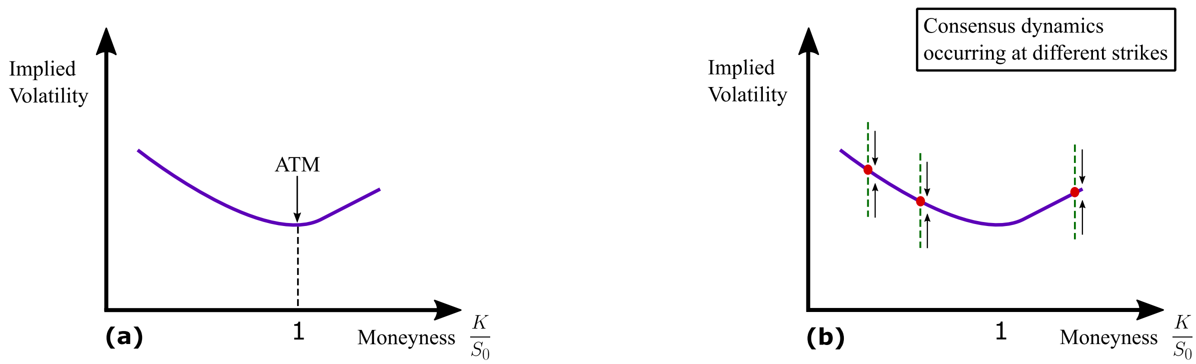

Figure 1: (a) A typical implied volatility smile for varying strikes K divided by fixed spot price. Moneyness is K/S_0. ATM denotes at-the-money where K equals S_0, (b) Consensus occurs as all investors' opinions of the implied volatility converge, round by round, to a distinct value for varying strikes.

where $\mathcal{E} := \mathrm{diag}(\epsilon_1, \ldots, \epsilon_n)$ and $A := a_{ij} \in \mathbb{R}^{n\times n}$ is a *row-stochastic matrix*. Then, we have the following result.

Theorem 2.1. *Consider the opinion dynamics* (2) *and assume that* $\epsilon_i \in (0, a_{ii})$, $i = \{1, \ldots, n\}$*; then, consensus to* $\bar{\sigma}$ *is reached, i.e.,* $\lim_{t\to\infty} X_t = \bar{\sigma}\mathbf{1}_n$.

2.2 Consensus with an unknown leader

One criticism of model (2) is that feedback, even if it is not perfect, has to be learned. In practice, there might not be a helpful mechanism that provides feedback. An alternative is to have an unknown leader embedded in the set of traders. The agents are unsure who the leader is but by taking averages of other traders, they all arrive at the opinion of the leader. Such behaviour is called an absorbing state. The leader guides the system to the true value.

Without loss of generality, we assume that the first agent (with corresponding opinion x_t^1) is the leader; it follows that $x_1^1 = \bar{\sigma}$, $a_{1i} = 0$, $i \in \{2, \cdots, n\}$, and $a_{11} = 1$. Then, in this configuration, the opinion dynamics are

$$X_t = AX_{t-1},\ A = \begin{pmatrix} 1 & 0 & \ldots & 0 \\ a_{21} & a_{22} & \ldots & a_{2n} \\ \vdots & \vdots & \ldots & \vdots \\ a_{n1} & a_{n2} & \ldots & a_{nn} \end{pmatrix} =: \begin{pmatrix} 1 & \mathbf{0} \\ * & \tilde{A} \end{pmatrix}, \tag{3}$$

with $a_{ij} \geq 0$, $\sum_{j=1}^{n} a_{ij} = 1$, $a_{ii} > 0$ for all $1 \leq i \leq n$, and for at least one i, $\sum_{j=2}^{n} a_{ij} < 1$.

Theorem 2.2. *Consider the opinion dynamics* (3) *and assume that the matrix* $\tilde{A}$ *is substochastic and irreducible. It holds that* $\lim_{t\to\infty} X_t = \bar{\sigma}\mathbf{1}_n$*, i.e., consensus to* $\bar{\sigma}$ *is reached.*

These results extend to a multidimensional setting, where each agent holds opinions across a range of strikes.

2.3 Arbitrage Bounds

True volatility is exogenous, however, we require no static arbitrage, by which we mean that all the quotes in volatility which translate to option prices are such that one cannot trade in the different strikes to create a profit. Checking whether a volatility surface is indeed arbitrage free is nontrivial, nevertheless some sufficient conditions are well known [7]. As long as the volatility surface satisfies them our analysis implies global stability towards an arbitrage free smile. How these arbitrage-free curve volatility conditions are developed is not an easy task: see an account by [27].

3 CONNECTIONS AND CONCLUSION

In actuality, for each smile there are only a few strikes that are actively traded. This means whatever model a trader chooses, there is flexibility in choosing different parameters to produce a perfect fit. In this reality, opinion dynamics seems clear. Model choice is addressed in [6, 9, 21] but the issue of opinions or interaction is not considered. Even with perfect calibration there will be periods when one model outperforms another [15]. In such situations, it makes sense to be close to the crowd! Market-makers optimize to make profits [10] but even here one has to be confident in one's estimates. Our models convey how such confidence may develop between transactions.

The smile itself is a conundrum and there have even been articles questioning whether it can be solved [5]. The traditional way from the ground up is to develop a stochastic process for the volatility and asset price [19].

While such models have been successfully developed the time is ripe to incorporate multi-agent models with arbitrage free curves.

Thus far, we proved convergence to equilibrium. A natural step forward would be to look at the beliefs as probability measures, where each measure corresponds to a different option pricing model. Our learning models focus on interaction between agents. Actually, traders can be interpreted as algorithms. Each algorithm corresponding to a particular belief of a pricing model.

ACKNOWLEDGMENTS

We acknowledge SUTD Presidential fellowship, National Research Foundation (NRF) grant (Award No. NRF2014NCR-NCR001-40), SUTD grant SRG ESD 2015 097, MOE AcRF Tier 2 grant 2016-T2-1-170 and NRF 2018 Fellowship NRF-NRFF2018-07.

REFERENCES

[1] Jacob Abernethy, Peter L Bartlett, Rafael Frongillo, and Andre Wibisono. 2013. How to hedge an option against an adversary: Black-scholes pricing is minimax optimal. In *Advances in Neural Information Processing Systems*. 2346–2354.

[2] Jacob Abernethy, Rafael M Frongillo, and Andre Wibisono. 2012. Minimax option pricing meets Black-Scholes in the limit. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM, 1029–1040.

[3] Daron Acemoglu and Asuman Ozdaglar. 2011. Opinion dynamics and learning in social networks. *Dynamic Games and Applications* 1, 1 (2011), 3–49.

[4] Itai Arieli and Manuel Mueller-Frank. 2015. Multi-Dimensional Social Learning. (2015).

[5] Elie Ayache, Philippe Henrotte, Sonia Nassar, and Xuewen Wang. 2004. Can anyone solve the smile problem? *The Best of Wilmott* (2004), 229.

[6] Hans Buehler. 2006. Expensive martingales. *Quantitative Finance* 6, 3 (2006), 207–218.

[7] Peter Carr and Dilip B Madan. 2005. A note on sufficient conditions for no arbitrage. *Finance Research Letters* 2, 3 (2005), 125–130.

[8] Neil Chriss. 1996. *Black Scholes and beyond: option pricing models*. McGraw-Hill.

[9] Laurent Cousot. 2007. Conditions on option prices for absence of arbitrage and exact calibration. *Journal of Banking & Finance* 31, 11 (2007), 3377–3397.

[10] Sanmay Das. 2008. The effects of market-making on price dynamics. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 887–894.

[11] Guillaume Deffuant, David Neau, Frederic Amblard, and Gérard Weisbuch. 2000. Mixing beliefs among interacting agents. *Advances in Complex Systems* 3, 01n04 (2000), 87–98.

[12] Morris H DeGroot. 1974. Reaching a consensus. *J. Amer. Statist. Assoc.* 69, 345 (1974), 118–121.

[13] Peter DeMarzo, Ilan Kremer, and Yishay Mansour. 2006. Online trading algorithms and robust option pricing. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*. ACM, 477–486.

[14] Emanuel Derman and Michael B Miller. 2016. *The Volatility Smile*. John Wiley & Sons.

[15] Moritz Duembgen and LCG Rogers. 2014. Estimate nothing. *Quantitative Finance* 14, 12 (2014), 2065–2072.

[16] Benjamin Golub and Evan D Sadler. 2017. Learning in social networks. (2017).

[17] Roger A. Horn and Charles R. Johnson. 2012. *Matrix Analysis* (2nd ed.). Cambridge University Press, New York, NY, USA.

[18] Brian Hrolenok, Byron Boots, and Tucker Hybinette Balch. 2017. Sampling Beats Fixed Estimate Predictors for Cloning Stochastic Behavior in Multiagent Systems.. In *AAAI*. 2022–2028.

[19] Michael Kamal and Jim Gatheral. 2010. Implied volatility surface. *Encyclopedia of Quantitative Finance* (2010).

[20] Ulrich Krause. 2015. Positive dynamical systems in discrete time: theory, models, and applications. (2015).

[21] Jean Paul Laurent, D Leisen, et al. [n. d.]. Building a consistent pricing model from observed option prices. In *Quantitative Analysis in Financial Markets: Collected Papers of the New York University Mathematical Finance Seminar*, Vol. 2. 216–238.

[22] Zhuoshu Li and Sanmay Das. 2016. An Agent-Based Model of Competition Between Financial Exchanges: Can Frequent Call Mechanisms Drive Trade Away from CDAs?. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 50–58.

[23] Jan Lorenz. 2005. A stabilization theorem for dynamics of continuous opinions. *Physica A: Statistical Mechanics and its Applications* 355, 1 (2005), 217–223.

[24] Tung Mai, Ioannis Panageas, and Vijay V. Vazirani. 2017. Opinion Dynamics in Networks: Convergence, Stability and Lack of Explosion. *44th International Colloquium on Automata, Languages, and Programming (ICALP)* (2017).

[25] Luc Moreau. 2005. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on automatic control* 50, 2 (2005), 169–182.

[26] Elchanan Mossel and Omer Tamuz. 2017. Opinion Exchange Dynamics. *Probability Surveys* (2017).

[27] Michael Roper. 2010. Arbitrage free implied volatility surfaces. *preprint* (2010).

[28] Ernesto Salinelli and Franco Tomarelli. 2014. *Discrete dynamical systems: one-step scalar equations*. Springer International Publishing, Cham, 85–124.

[29] Martin Schweizer and Johannes Wissel. 2008. Arbitrage-free market models for option prices: The multi-strike case. *Finance and Stochastics* 12, 4 (2008), 469–505.

[30] U.S. Securities and Exchange Commission. 2017. U.S. ETP Cancel to Trade Ratio. (2017). https://www.sec.gov/marketstructure/datavis.html#.WgvqJ2iCyUk [Online; accessed 14-November-2017].

[31] Shahin Shahrampour, Sasha Rakhlin, and Ali Jadbabaie. 2013. Online Learning of Dynamic Parameters in Social Networks. In *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 2013–2021. http://papers.nips.cc/paper/4976-online-learning-of-dynamic-parameters-in-social-networks.pdf

[32] Lones Smith and Peter Sørensen. 2000. Pathological outcomes of observational learning. *Econometrica* 68, 2 (2000), 371–398.

[33] Alan Tsang and Kate Larson. 2014. Opinion dynamics of skeptical agents. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 277–284.

[34] Xintong Wang and Michael P Wellman. 2017. Spoofing the limit order book: An agent-based model. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 651–659.