Asset Pricing In a World of Imperfect Foresight

Peter Bossaerts, Felix Fattinger,

Frans van den Bogaerde, Wenhao Yang

Peter Bossaerts (peter.bossaerts@unimelb.edu.au) is with the University of Melbourne and the University of Cambridge. Felix Fattinger (felix.fattinger@wu.ac.at) is with the Vienna University of Economics and Business. Frans van den Bogaerde (fransvandenbogaerde@gmail.com) is with IFM Investors, Wenhao Yang (yangwenhao@cuhk.edu.cn) is with the Chinese University of Hong Kong (Shenzhen).

Bossaerts is the corresponding author; his University of Melbourne correspondance address is: Faculty of Business and Ecooomics, 198 Berkeley Street, Parkville VIC 3050, Australia; phone +61.3.3095.3257.

Acknowledgement: Bossaerts acknowledges funding from the Australian Research Council (DP180102284) and from his R@MAP Chair at the University of Melbourne. The paper benefited from presentations at conferences (Society for Experimental Finance 2019 Meetings in Copenhagen, 2019 Econometric Society Meetings in Perth, 2019 FIRN Annual Meeting in Byron Bay, 2019 Behavioral Finance and Capital Markets Conference in Melbourne, 2020 FIRN Virtual Seminar) and universities (University of St. Gallen, Monash University, University of Melbourne).

Asset Pricing In a World of Imperfect Foresight

Abstract

We consider a canonical asset pricing model, where agents with quadratic preferences are allowed to re-trade a limited set of securities over multiple periods, after which these securities expire, and agents consume their liquidation values. A key assumption in this model is that agents have perfect foresight: for all future contingencies, they correctly foresee the corresponding equilibrium prices. We show that, under myopia, prices generically are as if agents had perfect foresight. Yet their choices are wrong, because of neglected re-trading opportunities. In an experiment, we find both prices and choices to be consistent with myopia.

JEL classification: D51, D84, G12, G14

Keywords: Rational Expectations, Perfect Foresight, Myopia, Dynamic Asset Pricing

I. INTRODUCTION

The canonical model for dynamic asset pricing is one where agents are allowed to re-trade a limited set of securities for a number of periods, after which these securities expire and agents consume their liquidation values. Because of the possibility to re-trade, agents can attain significantly better final allocations (allocations that entail higher utility) than if they had only been allowed to trade once. Effectively, with the right number of securities, agents can generically trade to final allocations that are *as good as* if they had been able to trade far more securities.

In technical terms, re-trade in a few crucial securities makes the market *dynamically* complete, meaning that it allows agents to reach the Pareto optimal allocations of a complete market (Duffie and Huang, 1985). The principles behind re-trade are well understood: they are underlying, e.g., the Black-Scholes-Merton option pricing model (Black and Scholes, 1973; Merton, 1973b) and the Ho-Lee/Heath-Jarrow-Morton term structure models (Ho and Lee, 1986; Heath et al., 1992), and are widely used in practice.

The power of dynamic completeness comes with an important caveat, though. Nothing less but *perfect foresight* is required to reach Pareto optimality through continuous re-trading of a limited set of securities. Perfect foresight is the ability to correctly foresee the right (equilibrium) prices for every possible future contingency.

The requirement is not innocuous: it is hard to imagine how, without substantial repetition in a stationary environment, agents can acquire perfect foresight. Knowledge of economic constraints generally does not constrain prices enough to deduce ("educe") the right prices (Guesnerie, 1992). It is therefore interesting to investigate what happens before agents acquire perfect foresight (if ever) and how potential deviations therefrom impact prices *and* allocations in equilibrium, which is the aim of this paper.

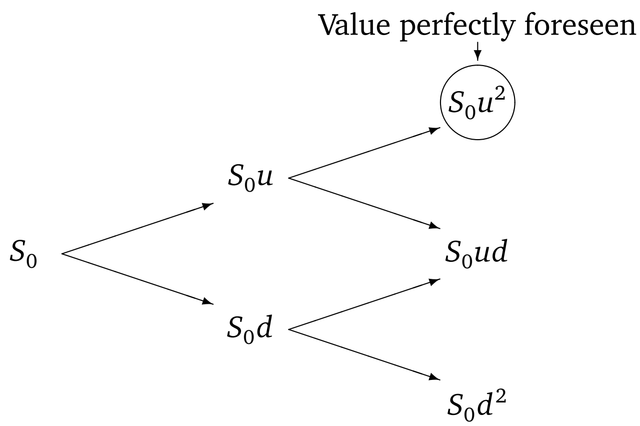

(a) Perfect Foresight in a Binomial Asset Pricing Model

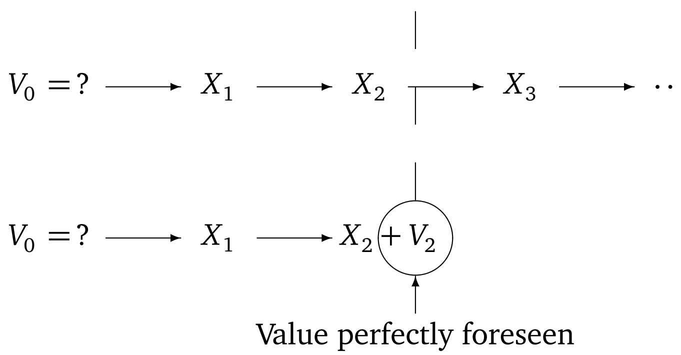

(b) Perfect Foresight in the Modigliani-Miller Model

Figure 1: Examples of Perfect Foresight in Seminal Finance Theories

Yet the requirement of perfect foresight is often not made explicit. Even advanced textbooks in asset pricing discuss the requirement only in "notes" to chapters. As an

See the end-of-chapter discussions in standard textbooks of asset pricing, such as Duffie (1988), Chapter 16; Magill and Quinzii (2002), Chapter 2; or LeRoy and Werner (2014), Chapter 21.

example, consider the discrete version of the seminal Black-Scholes-Merton model, i.e., the Cox et al. (1979) binomial option pricing framework. For a simple two-period model, Figure 1 (top) displays the usual binomial lattice on which the price of the underlying stock (S) is assumed to "live." The arbitrage-free option price obtains even if agents do not know the chances of the up or down ticks. However, they have to know the exact size of the multiplicative up (u) and down (d) movements, i.e., they have to know the price level in each node. Moreover, if one thinks about the problem in *general* rather than *partial* equilibrium, imperfect foresight about possible price levels cannot be sufficiently addressed by "simply" adding parameter uncertainty.

Figure 1 (bottom) displays a timeline of infinite firm profits as underlying the Modigliani-Miller (MM) result that the firm value (V) does not depend on its capital structure (Modigliani and Miller, 1958). Specifically, the MM irrelevance proposition relies on an absence of arbitrage argument. However, since arbitrage opportunities are delicate to define for infinite horizons,[†] the result effectively requires one to un-wind positions at some finite point in the future, instead of counting on converging returns ($^1/_T\,\Sigma_t^T X_t$) across "homogeneous" firms. Hence, the argument only works if this future value is sufficiently close to its arbitrage-free prediction, whose determination requires perfect foresight. Furthermore, under realistic assumptions, there is considerable uncertainty about future arbitrage-free prices even within finite horizons (Long et al., 1990; Shleifer and Vishny, 1997).

[†]For an infinite horizon, multiple arbitrage-free prices can exist (Delbaen and Schachermayer, 1994).

To preclude a common misconception, a crucial distinction is helpful here. The term "perfect foresight" is easily mis-interpreted. It does *not* mean that agents are fully prescient, i.e., that they can predict the future. It "only" means that agents know prevailing equilibrium prices conditional on the realized state of nature. They generally do not even have to agree on the chances of these states, but they do have to agree on prices *in* each state.[‡] The term "perfect foresight" is standard terminology in general equilibrium theory, to which dynamic asset pricing theory belongs.[§]

Acknowledging the inherent difficulty of achieving perfect foresight in the real world, the literature has proposed alternative concepts. For instance, it has been suggested that *Temporary Equilibrium* may explain prices and allocations in the interim (Radner, 1974; Grandmont, 1977). The idea is that agents posit provisional (future) prices, optimize with these prices in mind, and adjust expectations (of future prices) as experience grows. Specific ways in which agents build expectations upon repetition include online regression (Marcet and Sargent, 1988) and Bayesian updating (Jordan and Radner, 1982).

The problem with the Temporary Equilibrium is that it is silent about which prices (or even range of prices) agents could reasonably hypothesize at the beginning. At best, they happen to guess the true future equilibrium prices, in which case the Temporary Equilibrium

[‡]Some would call it rational expectations; we refrain from doing so use, in order not to confuse with the assumption of correct beliefs about the occurrence of states. Lucas (1978), for instance, requires both correct beliefs about state probabilities and perfect foresight. Under perfect foresight, the former is not needed. Agents may even disagree on the chance that a state (contingency) obtains (Anderson and Bossaerts, 2019).

[§]Perfect foresight is analogous to Subgame Perfection in game theory: there too, players are supposed to know how the game continues in every future node. Continuation is restricted to Nash (equilibrium) play (Fudenberg and Levine, 1983).

and the Perfect-Foresight Equilibrium coincide. It is unclear what the worst-case scenario looks like.

An alternative theory does provide clean predictions, namely the *Myopic Equilibrium*. In a simplified case, it can be defined as the equilibrium that obtains if agents only optimize over the next period, ignoring any future possibility of re-trading, obviating the need to form predictions about prices of currently non-traded securities. As a benchmark, we derive the implications of myopia within the discrete-time/discrete-state version of the Black-Scholes-Merton model as pioneered by Cox et al. (1979) and illustrated in Figure 1 (top).

Myopia is not really a bias, but rather an expression of bounded rationality. Myopia is a mild form of *narrow framing* (Barberis et al., 2006), a well-documented heuristic that humans adopt in the face of the complexity of decision-making under uncertainty (Tversky and Kahneman, 1974). Intuitively, agents ignore future trading opportunities because it is too difficult to guess what prices they could trade at. Under its most common interpretation, narrow framing is even more extreme than myopia: the agent ignores holdings of items that she cannot trade when making decision about today's trades. In contrast, under myopia, the agent does take into account quantities of goods she holds but cannot trade momentarily. We shall see that the latter is crucial for pricing to be *as if* agents had perfect foresight.

In our theoretical analysis, we assume quadratic utility. This is equivalent to assuming normally distributed portfolio returns, in the sense that both are consistent with mean-variance optimization. Intertemporal models such as Merton's ICAPM (Merton, 1973a) or the Black-Scholes model (Black and Scholes, 1973) effectively assume quadratic utility given that changes in asset prices over short (infinitesimal) intervals are normally distributed. Linear asset pricing models such as the CAPM or multi-factor models (Fama and French, 2004) directly assume mean-variance preferences.

Our theory explains how prices can be "right" even if allocations are "wrong." Empiricists often implicitly assume that this is possible when testing for asset pricing models despite allocations (holdings) being clearly at odds with model predictions. For instance, linear asset pricing models such as the CAPM imply that every investor should be fully diversified, while, empirically, most investors clearly are not (Odean, 1999). Evidently, an implicit acceptance of a theory does not make it true. Here, we could take one additional step by looking for confirmatory (or contradictory) evidence in archival holdings data from the field, which would be in line with the traditional approach of testing finance theories. However, corresponding evidence could only confirm that the model equations fit, and not that they fit *because* investors are indeed myopic. Moreover, field data only contains holdings for realized states, whereas testing the theory requires to observe holdings in all contingencies, even those that did not occur in the sample at hand.

To test the theory in a controlled way, we need an experiment in which we control (induce) participants' preferences across states and verify that prices can be consistent with perfect foresight, even if participants' choices reflect myopia. Therefore, we experimentally test the theory by running several trading sessions, while fully controlling for any confounding effects of risk aversion. In particular, absent any extrinsic uncertainty, we incentivize participants so that choices are as if governed by the rules of expected utility.[¶]

In the experiment, we deliberately did *not* induce quadratic utility. Instead, we induced square-root utility. Under square-root utility, prices are almost but not exactly identical in the Myopic and Perfect-Foresight Equilibria. Our design feature allowed us to test for myopia against perfect foresight, not only by investigating allocations, but also by studying prices. Testing power is enhanced because we can verify predictions across the two equilibria in both allocations *and* prices.

We find strong support in favor of the Myopic Equilibrium: when market participants have to trade assets sequentially, prices and allocations are better explained by the Myopic Equilibrium than the Perfect-Foresight Equilibrium. When participants can trade all assets simultaneously, prices and allocations are better explained by the Walrasian Equilibrium. Price and allocation predictions in the Walrasian Equilibrium are identical to those in the Perfect-Foresight Equilibrium. In the Walrasian Equilibrium, allocations are Pareto optimal.

It is customary to question the external validity of markets experiments of the type we discuss here. The question one asks is whether the results would still obtain if we were to scale them to the level of field markets. Our study does not shed light on this.[‖] Instead, it investigates to what extent it is possible *at all* that prices could be right despite bounded rationality among market participants. If this is generally believed to be true

[¶]We induce preferences that are isomorphic with expected utility, i.e., equilibrium prices and allocations are as if we had introduced extrinsic uncertainty and agents indeed maximized expectations of the assumed utility function.

[‖]We could, but that would require infrastructure and funding far beyond current levels.

in field markets (as mentioned above, dynamic asset pricing theory invariably assumes that prices are right), it should be true at a small scale and in a much better controlled environment as well. If the theory fails in our experiment, what would make it work in the field, where perfect foresight is even harder to obtain?

In Section II., we state the main result and provide the intuition behind it. A formal derivation is delegated to Section III.. In Section IV., we describe the experiment we ran to test the theory. Section V. presents the experimental evidence. In Section VI., we discuss the implications of our theory for (i) the empirical evidence on various types of asset pricing theory, (ii) the relevance of theory for the practice of finance, and (iii) the identification of possible domains for successful active investing. Finally, Section VII. concludes.

II. The Main Result

We now state the main result and the intuition behind it. The formal derivation is delegated to the next section.

Theoretical Result. *In the binomial model, the Myopic Equilibrium produces precise, generic predictions. Assuming quadratic utility, we show that prices will be exactly the same as if agents had perfect foresight. We will refer to this prediction as "prices are right." Choices (allocations), however, may be far from optimal. That is, "allocations are wrong."*

The result may be surprising, but the intuition as to why prices can be right even if agents are myopic is simple. Myopia does *not* mean that agents are ignoring contingent endowments that they cannot trade away from in the current trading period. Instead, they do take these endowments into account; only, they assume that these are (permanently) non-tradeable. For instance, an agent with a high contingent, but currently non-traded, endowment assumes that she cannot ever sell it, which she will signal indirectly through orders that reveal that she is interested in portfolios that pay much in other contingencies (she wants to smooth consumption across contingencies). If everyone has a high endowment in the same non-traded contingency, then prices of currently traded contingencies will reflect this. This is not specific to quadratic utility, but with quadratic utility, prices will be *exactly* as if all contingencies are traded at once.

Therefore, through demand for currently traded assets, market prices under myopia indirectly reflect knowledge of the scarcity or abundance of currently non-traded endowments, and as such prices behave as if the latter had been available for trade.

III. Theoretical Details

We start from the following notation and assumptions.

- As in Figure 1 (top), we assume three terminal states ($S = 3$). Beliefs about state chances are denoted α_s, $s \in \{1, 2, 3\}$. Beliefs of agents are homogeneous.
- There is a numeraire, with price equal to unity. We refer to it as "cash" C, which pays \$1 in all states. The remaining two assets are Arrow-Debreu (AD) securities that pay in state $s = 1$ ("asset 1") or $s = 3$ ("asset 3"), respectively.

- Agent i chooses (trades towards) holdings x_k^i of asset k, $k \in \{C, 1, 3\}$ and starts with an endowment of asset k equal to e_k^i.

Throughout our analysis, we contrast (i) a complete-markets case with the (potential) myopia implications of (ii) an incomplete-markets case whose structure allows for dynamic completeness (Duffie and Huang, 1985). Importantly, in the presence of perfect foresight, the two cases are equivalent.

(i) In the complete-markets case, both assets and cash are traded *simultaneously*. Let p_k denote the price of asset k, $k \in \{C, 1, 3\}$. Obviously, $p_C = 1$.

(ii) In the incomplete-markets case, trading happens *sequentially*. In round 1, assets 1 and cash are traded, while asset 3 and cash are traded in round 2. To disambiguate prices, let q_1 be the price of asset 1 in round 1 and q_3 the price of asset 3 in round 2. Cash is always priced at \$1.

Furthermore, we make the following assumption regarding the agents.

- All agents have state-independent and state-separable quadratic preferences.

As is well known (Rubinstein, 1974), the last assumption implies the existence of a representative agent. The formal proof of our Theoretical Result in the Appendix does not rely on this result, however. Instead, its analysis starts at the individual level. Here, we broadly sketch the arguments behind the proof at the level of the representative agent.

The representative agent has differentiable utility $u(w_s)$, where w_s is final wealth in state s. Utility is state-independent and quadratic: $u(w_s) = -\frac{b_0}{2}(w_s)^2 + b_1 w_s + b_2$. Across states, utility is separable, so outcomes are evaluated based on *expected* utility $\sum_s \alpha_s u(w_s)$.

Working with the representative agent, we can first write down the respective budget constraints for the two cases. Under complete markets, there is only one budget constraint:

$$p_1 x_1 + x_C + p_3 x_3 = p_1 e_1 + e_C + p_3 e_3.$$

When markets are incomplete, we have two separate budget constraints, corresponding to the two trading rounds. In round 1, the budget constraint is:

$$q_1 x_1 + x_C = q_1 e_1 + e_C.$$

In round 2, the budget constraint is:

$$y_C + q_3 x_3 = x_C + q_3 e_3,$$

where y_C denotes the updated cash holding decision for round 2.

Second, we can write down the simultaneous optimality conditions for the complete-markets case, namely,

$$\begin{aligned}
\alpha_1 u'(x_1 + x_C) &= \lambda p_1, \\
\alpha_1 u'(x_1 + x_C) + \alpha_2 u'(x_C) + \alpha_3 u'(x_3 + x_C) &= \lambda, \\
\alpha_3 u'(x_3 + x_C) &= \lambda p_3.
\end{aligned}$$

Notice the second equation: the Lagrange multiplier (λ) is to equal the expected marginal utility across all states.

We then compare these conditions to the (sequential) optimality conditions for the incomplete-markets case assuming myopia. For round 1, the optimality conditions are:

$$\begin{aligned}
\alpha_1 u'(x_1 + x_C) &= \mu q_1, \\
\alpha_1 u'(x_1 + x_C) + \alpha_2 u'(x_C) + \alpha_3 u'(e_3 + x_C) &= \mu.
\end{aligned}$$

Notice the restriction that the holdings of asset 3 equal the endowment. Also, the Lagrange multiplier (μ) is again to equal the expected marginal utility across *all* states. Despite the fact that asset 3 cannot be traded, the endowment of asset 3 still influences trading decisions in round 1. As such, relative scarcity or abundance of non-traded AD securities will already affect trade, and hence, prices in round 1. This is the key reason why myopia has no influence on prices (under quadratic utility).

For round 2, the optimality conditions are straightforward:

$$\begin{aligned}\alpha_1 u'(x_1 + y_C) + \alpha_2 u'(y_C) + \alpha_3 u'(x_3 + y_C) &= \pi, \\ \alpha_3 u'(x_3 + y_C) &= \pi q_3.\end{aligned}$$

The question is whether the *same* prices that clear the complete markets can also clear the incomplete markets under myopia. That is, letting p_k^* denote complete-markets equilibrium prices, is it possible that there are myopic incomplete-markets equilibrium prices q_1^o, q_C^o and q_3^o such that:

$$\begin{aligned}q_1^o &\stackrel{?}{=} p_1^*, \\ q_C^o &\stackrel{?}{=} p_C^*, \\ q_3^o &\stackrel{?}{=} p_3^*.\end{aligned}$$

The answer is affirmative. To illustrate the equilibrium implications of our result, we end this section with a simple numerical example.

Example. There are three types of investors in the economy, each with equal relative mass. Their choices exhibit quadratic utility with common parameters: $b_0 = 0.1$, $b_1 = 1.0$, and $b_2 = 0$. Beliefs are common, with $\alpha_s = 1/3$ for $s \in \{1, 2, 3\}$. There exist three securities, two Arrow-Debreu (state) securities that pay \$1 in states $s = 1$ and $s = 3$, respectively, and cash, which is worth \$1 in every state. Agents differ in initial endowments. Type-1 agents start with 6 units of state security 1; type-2 agents start with 2 units of cash; type-3 agents start with 2 units of state security 3 (see top panel of Table 1).

Using cash as numeraire, the complete-markets equilibrium prices are as follows:

$$\begin{bmatrix} p_1^* \\ p_C^* \\ p_3^* \end{bmatrix} = \begin{bmatrix} 0.2895 \\ 1 \\ 0.3421 \end{bmatrix}.$$

Notice that payments in state 1 are priced lower than in state 3, which is expected since state security 3 is in lowest supply. The complete-markets equilibrium allocations are reported in the middle panel of Table 1.

Next, we construct an *incomplete* set of markets by allowing investors to trade only state security 1 (against cash) in round 1, and state security 3 in round 2. No additional information is released in-between trading rounds.

Illustrating our main result, complete-markets prices also clear the incomplete markets, i.e,

$$\begin{bmatrix} q_1^o \\ q_C^o \\ q_3^o \end{bmatrix} = \begin{bmatrix} 0.2895 \\ 1 \\ 0.3421 \end{bmatrix} \stackrel{!}{=} \begin{bmatrix} p_1^* \\ p_C^* \\ p_3^* \end{bmatrix}.$$

The above equilibrium prices lead to the end-of-round-1 (for state security 1, these are final allocations) and end-of-round-2 allocations reported in the bottom panel of Table 1. Note that we do not impose short-sale constraints; otherwise the first-order conditions

Table 1: Numerical Example of Initial Endowments and Equilibrium Allocations

Table reports exemplary initial endowments and corresponding equilibrium allocations for both the complete-markets and the incomplete-markets (sequential) case. All agent types display the same quadratic utility of state-dependent wealth w_s, i.e., $u(w_s) = -\frac{b_0}{2}(w_s)^2 + b_1 w_s + b_2$, with $b_0 = 0.1$, $b_1 = 1.0$, and $b_2 = 0$, respectively. Initial endowments differ as indicated in the top panel. There are an equal number of agents of each type.

	Type 1	Type 2	Type 3
INITIAL ENDOWMENTS			
State security 1 (e_1)	6	0	0
Cash (e_C)	0	2	0
State security 3 (e_3)	0	0	2
COMPLETE MARKETS			
State security 1 (x_1^*)	1.9383	1.8765	2.1852
Cash (x_C^*)	0.9547	1.2428	-0.1975
State security 3 (x_3^*)	0.6461	0.6255	0.7284
INCOMPLETE MARKETS			
	Round 1		
State security 1 (x_1^o)	1.6169	1.5654	2.8177
Cash (x_C^o)	1.2688	1.5469	-0.8157
	Round 2		
Cash (y_C^o)	1.1029	1.3863	-0.4892
State security 3 (x_3^o)	0.4849	0.4695	1.0456

listed before would not apply. Our agent 3 sells short in both the complete-markets and incomplete-markets equilibrium. Importantly, the incomplete-markets equilibrium allocations are *very different* from the complete-markets ones. Prices, however are *identical*.

One remark is in order. The result fails to hold exactly for other utility functions. However, as we shall demonstrate below (relying on the identical market structures), price differences are generally small, whereas allocations continue to differ substantially.

IV. The Experiment

It is ultimately an empirical question which of the two theories – perfect foresight, or myopia (narrow framing) – better explains actual market outcomes. We therefore design an experiment that allows us to clearly distinguish between the two theories. The design closely follows the binomial model of our theoretical analysis. To reduce complexity as much as possible, we limit our attention to two trading rounds and three states.

To avoid confounding factors, we eliminate all extrinsic uncertainty, however. In the standard interpretation of the binomial model, we would pay participants based on the realization of a single state, with elimination of one state after each trading round. Moreover, we assume choices to be governed by the rules of expected utility, which means that participants are to choose as if maximizing a weighted average of utilities across states, where weights equal the chances of the occurrence of respective states.

In our design, participants are paid the weighted average of nonlinear transformations of the payoffs on three traded assets. Starting from heterogeneous endowments, participants first trade one state-dependent asset (called "Steel") and the numeraire asset (a cash equivalent called "Plastic"), followed by a second round when they trade a second state-dependent asset (called "Wood") against the numeraire. Crucially, this way, there is *no* uncertainty except about the price at which the second asset will trade. This is illustrated in Figure 2 (bottom).

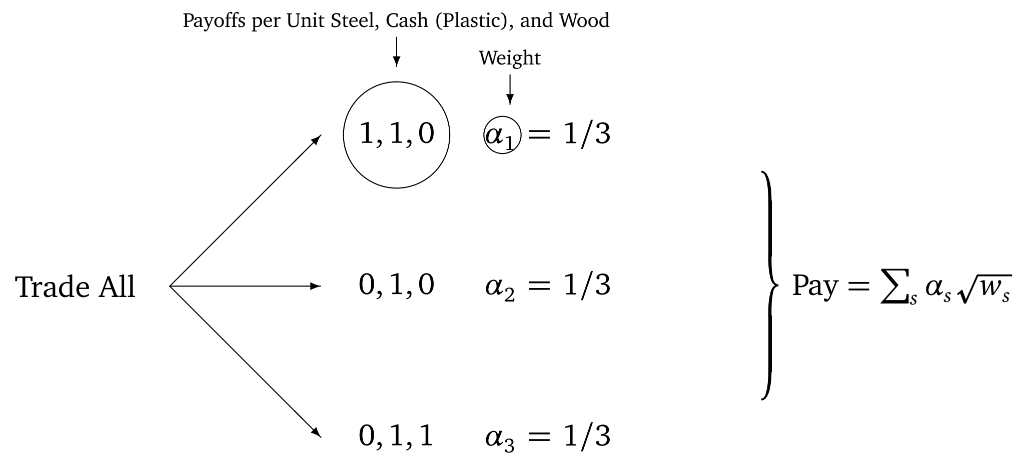

(a) Simultaneous Markets Treatment

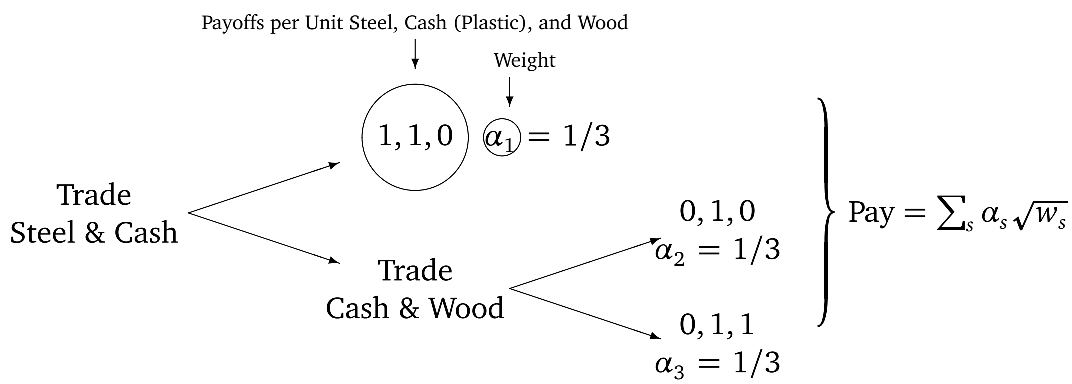

(b) Sequential Markets Treatment

Figure 2: Schematic Representation of Trade and Payoffs

Another way to see this is as follows. If we had paid based on the drawing of a state s whose chance α_s was known to participants, then for the theory to apply, we would have to maintain the auxiliary assumption that all participants choose on the basis of expected utility. A rejection of the theoretical predictions could then either be attributed to violations of expected utility, or to the absence of myopia among participants. In addition, we would have to estimate the utility function $u(\cdot)$ revealed in participants' choices. Any resulting estimation error would have reduced the power of our tests even further.

Our design can be considered a "certainty equivalent" version of the two-period, three-state binomial model. The use of certainty equivalent settings has proven useful to understand, among others, off-equilibrium evolution of prices and states in financial markets (Asparouhova et al., 2020). In the certainty-equivalent version, neither are participants required to exhibit expected utility preferences nor does the experimenter need to know their risk aversion; participant choices merely need to reflect non-satiation (for money).

We compare outcomes in our two-round trading setting, where, sequentially, only *one* asset is traded for the numeraire in any given round, against a single-round trading setting, where *all* three assets (including the numeraire) are traded simultaneously. This is illustrated in Figure 2 (top). We refer to the former as the *sequential treatment* and to the latter as the *simultaneous treatment*. Importantly, while any given market in the sequential treatment is incomplete in isolation, taken together they *dynamically* complete each other.

The appropriate notion of equilibrium that applies to the simultaneous treatment is the *Walrasian Equilibrium*, where demands equal supplies given equilibrium prices.

Extrinsic uncertainty is known to generate enormous heterogeneity in choices, both across participants, and for a given participant over time (Bossaerts et al., 2007), making it impossible to compare outcomes across treatments. Indeed, in a precursor to this experiment where payoffs were determined by the random drawing of one out of three states (Bossaerts et al., 2008), there were no significant differences in either prices or choices between (i) a treatment where three securities with independent payoffs were traded once, and (ii) a treatment where one of these securities could not be transacted at all, two trading rounds were introduced, and one of the states was excluded after one round. These results could be interpreted as providing support for perfect foresight. However, they could also have been the result of lack of power, with preferences for extrinsic uncertainty that cannot be captured using expected utility with time-invariant beliefs and risk aversion.

The standard notion of equilibrium for the sequential treatment is the *Perfect-Foresight Equilibrium*, where, in round 1, agents optimize dynamically to determine current choices given correct (perfect) foresight of the (Walrasian) Equilibrium prices to prevail in round 2. The Perfect-Foresight Equilibrium is the almost exclusive equilibrium notion of dynamic asset pricing theory.[††]

Given the dynamic completeness of our sequential treatment, its corresponding Perfect-Foresight Equilibrium coincides with the Walrasian Equilibrium of the simultaneous treatment. Crucially, if participants' myopia prohibits perfect foresight, final holdings in the sequential treatment will not equilibrate towards Perfect-Foresight Equilibrium allocations and hence will also differ from Walrasian Equilibrium allocations of the simultaneous treatment.

Finally, in our experiment, we induce square-root utility instead of quadratic utility, for two reasons. First, this avoids issues of decreasing utility beyond a certain level of securities holdings. Second, and more importantly, the Myopic Equilibrium prices of sequential markets are no longer identical to those in the Walrasian Equilibrium of simultaneous markets, and hence, to those of the Perfect-Foresight Equilibrium of sequential markets. As such, square-root utility provides a stronger test of the theory: in the sequential treatment, *both* allocations *and* prices are expected to be different from the Perfect-Foresight Equilibrium. Had we chosen quadratic utility, only predicted choices would have been different.

[††]See Radner (1972), Duffie and Huang (1985), Anderson and Sonnenschein (1985), and Schraeder (2015). It is also used in the disagreement literature, see, e.g., Scheinkman and Xiong (2003).

Instructions for the experiment can be found in Appendix 2. The Appendix also includes snapshots of the trading interface as well as the online spreadsheet with which participants could compute widget production (transformed payoffs) as a function of all traded input goods.

V. EXPERIMENTAL EVIDENCE

We ran eight experimental sessions involving 117 participants. Participants were undergraduate and postgraduate students from University of Melbourne.[‡‡] We randomly alternated between sequential and simultaneous markets, generating between five and six replications per session. In three sessions, endowments (and hence, equilibria) changed across replications, while in the remaining five sessions, they remained constant across replications.[§§]

At the end of each session, participants were paid the payoffs (in AUD) from one randomly selected replication. Payoffs ranged from \$40 to \$57, with a mean of \$50.74. No session lasted longer than two and a half hours. Given the deliberate exclusion of any exogenous uncertainty, the low variation in payoffs (standard deviation of \$3.65) reflects a high degree of competitiveness. Hence, our key assumption of non-satiation seems justified.

[‡‡]The study was approved by the University of Melbourne Human Research Ethics Committee (Ethics ID: 1749620.1) and was conducted in accordance with the World Medical Association Declaration of Helsinki. All participants provided written informed consent.

[§§]Note, even with fixed endowments, initial holdings varied across participants.

A. Descriptive Results

Here, we illustrate the results for one exemplary session with variable endowments. Figure 3 shows trading prices for both assets steel (asset 1) and wood (asset 3) for the whole session. Markets are grouped according to their treatment, with sequential (simultaneous) markets displayed in the top (bottom) row. Additionally, markets differ in terms of initial allocations, resulting in different equilibrium predictions. The plots in Figure 3 are vertically aligned in accordance with their initial allocations; e.g., in the first column, Market 1 and Market 6 have identical initial allocations.

We observe that prices in the sequential treatment are myopic, whereas, in the simultaneous treatment, they are economically indistinguishable from the Walrasian Equilibrium. This is consistent with our hypothesis of myopia for sequential markets. Moreover, Figure 3 illustrates that the Walrasian Equilibrium actually works for multiple simultaneous markets, which, considering the required cognitive efforts, is far from self-evident.

Myopic Equilibrium prices (top row) are generally very similar to their Walrasian counterparts (bottom row). Due to square-root utility (instead of quadratic utility – see above), equilibrium prices for steel do measurably deviate between Market 2 and Market 5 (third column), however. The respective trading prices thus provide support for myopic behavior under the sequential treatment.

In contrast to prices, allocations predicted by the Myopic vs. the Perfect-Foresight Equilibrium are distinctly different for the early traded asset (steel) in sequential markets.

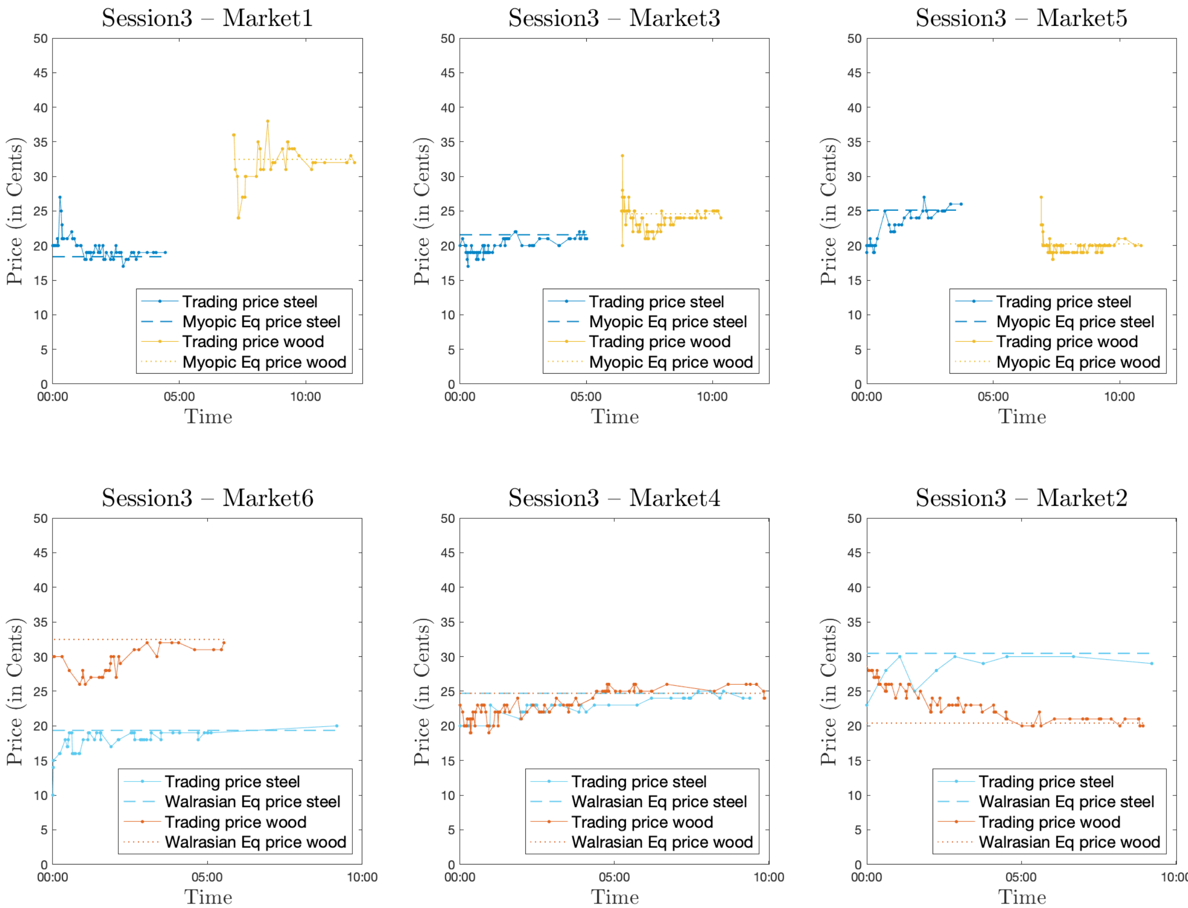

Figure 3: Trading Prices of One Session

For the sequential treatment, Figure 4 (left) shows histograms of the absolute differences in final steel holdings relative to the Myopic ("M-Eq") and the Perfect-Foresight ("PF-Eq") Equilibrium, respectively. Accordingly, for the simultaneous treatment, Figure 4 (right) shows histograms of absolute holding differences relative to the analogous Myopic ("M-Eq") and Walrasian ("W-Eq") Equilibrium, respectively.

While the simultaneous treatment generates holdings that are close to the (Pareto optimal) allocations of the Walrasian Equilibrium, Figure 4 (left) indicates that the sequential treatment generates holdings that are much closer to the (inferior) myopic allocations.

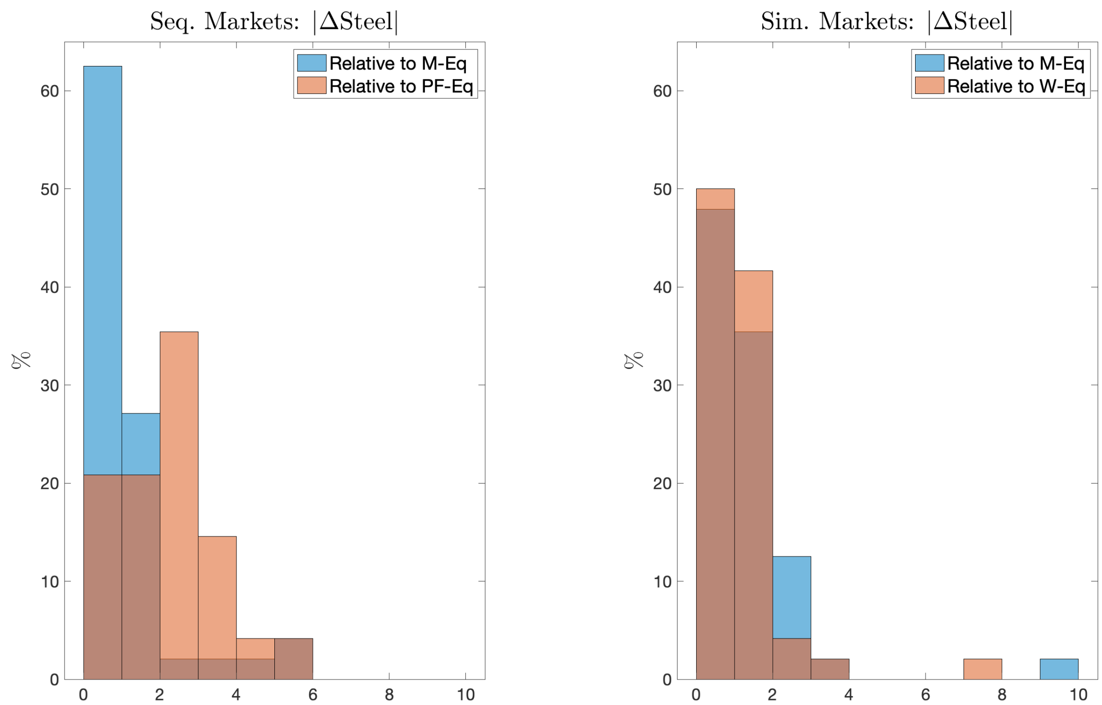

(a) Sequential Markets

(b) Simultaneous Markets

Figure 4: Distribution of Final Allocations

Focusing on sequential market replications, our descriptive results also reject the Temporary Equilibrium, which would agents deliberately have budget their intermediate holdings in anticipation of later trading opportunities (based on price forecasts that are not necessarily correct).

B. Formal Evidence

The following is based on an exhaustive analysis of price and holdings data across all sessions. In each case, we only report the model results that survived a strict elimination strategy using standard information criteria, i.e., the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC).[¶] Various potential confounding factors are considered, including session, market type (sequential vs. simultaneous), replication, and participant-specific random effects (intercepts). Since the Walrasian Equilibrium of any given simultaneous market coincides with the Perfect-Foresight Equilibrium of the corresponding sequential market, we below solely refer to the latter for both treatments.

Prices. Table 2 reports the results from pooled price level regressions. To control for the collinearity between the Perfect-Foresight and the Myopic Equilibrium, we consider two model specifications: starting from Perfect-Foresight Equilibrium prices, we add either (1) the differences between Myopic Equilibrium and corresponding Perfect-Foresight Equilibrium prices, or (2) the orthogonalized residuals from regressing those differences on the latter.

For both specifications, the best model explains steel and wood prices in terms of (i) the Perfect-Foresight Equilibrium prices, (ii) the (orthogonalized) differences between Myopic Equilibrium and corresponding Perfect-Foresight Equilibrium prices, and (iii) a correction to accommodate a better fit of Perfect-Foresight Equilibrium prices during simultaneous markets ($D_{\text{SIM}} = 1$). Included are random effects at the level of sessions, market types, and replications. Note, the two models perform equally well according to both information criteria.

The coefficient estimates reported in Table 2 imply that Perfect-Foresight Equilibrium prices do not fully predict trade prices (the corresponding coefficients are significantly

[¶]See Bossaerts and Hillion (1999) for an early application in finance.

Table 2: Trade Price Regressions

Table reports coefficient estimates (with t-stats in parentheses) of generalized linear mixed models of steel and wood trade prices (levels). Model selection is based on the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). Only the best models according to both criteria are shown. "PF-Eq price" denotes Perfect-Foresight Equilibrium prices, "Δ M-Eq price" denotes differences between Myopic Equilibrium prices and corresponding Perfect-Foresight Equilibrium prices, and "Orth. Δ M-Eq price" denotes orthogonalized residuals of "Δ M-Eq price" with respect to "PF-Eq price." "D_{SIM}" refers to a dummy variable indicating simultaneous market replications. To capture potential confounding factors, random effects per session, market type (sequential vs. simultaneous), replication, and participant are considered for model selection. [1] indicates t-stats for the null hypothesis of a slope equal to one.

	(1)	(2)
Intercept	1.299	-0.382
	(0.91)	(-0.26)
PF-Eq price	0.831	0.872
	(-22.43)[1]	(-14.08)[1]
Δ M-Eq price	0.764	
	(13.48)	
Orth. Δ M-Eq price		0.764
		(13.48)
PF-Eq price × D_{SIM}	0.073	0.073
	(14.58)	(14.58)
Session RE	YES	YES
Market type RE	YES	YES
Replication RE	YES	YES
Participant RE	NO	NO
Observations	4,119	4,119
AIC	18,766	18,766
BIC	18,817	18,817

lower than one), but they perform significantly better ($p < 0.001$) in the simultaneous treatment (when $D_{\text{SIM}} = 1$). Furthermore, the increment towards Myopic Equilibrium prices (whether orthogonalized or not) exhibits strong additional explanatory power ($p < 0.001$). Altogether, while prices are closest to the Perfect-Foresight Equilibrium in the simultaneous treatment, myopia explains a significant fraction of the deviations.

Table 3: Final Holdings Regressions

Table reports coefficient estimates (with t-stats in parentheses) of generalized linear mixed models of final steel holdings. Model selection is based on the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC). Only the best models according to both criteria are shown. "PF-Eq holdings" denotes Perfect-Foresight Equilibrium holdings, "Δ M-Eq holdings" denotes differences between Myopic Equilibrium holdings and corresponding Perfect-Foresight Equilibrium holdings, and "Orth. Δ M-Eq holdings" denotes orthogonalized residuals of "Δ M-Eq holdings" with respect to "PF-Eq holdings." "D_{SIM}" and 'D_{VAR}" refer to dummy variables indicating simultaneous market replications and sessions with non-stationary (variant) endowments, respectively. "D_{TYPE}" denotes a dummy variable indicating participants with (relatively) low cash endowments. To capture potential confounding factors, random effects per session, market type (sequential vs. simultaneous), replication, and participant are considered for model selection. [1] indicates t-stats for the null hypothesis of a slope equal to one.

	(1)	(2)
Intercept	-2.550	-0.711
	(-3.03)	(-1.09)
PF-Eq holdings	0.996	0.852
	(-0.77)[1]	(-3.21)[1]
Δ M-Eq holdings	2.634	
	(5.80)	
Δ M-Eq holdings $\times$ D_{SIM}	-0.910	
	(-2.26)	
Δ M-Eq holdings $\times$ D_{VAR}	-0.958	
	(-3.30)	
Orth. Δ M-Eq holdings		15.945
		(6.30)
Orth. Δ M-Eq holdings $\times$ D_{SIM}		-3.930
		(-2.08)
Orth. Δ M-Eq holdings $\times$ D_{VAR}		-6.087
		(-2.90)
D_{SIM}	1.948	1.439
	(3.27)	(3.61)
D_{TYPE}	6.091	4.727
	(4.62)	(5.05)
$D_{\text{SIM}} \times D_{\text{TYPE}}$	-4.027	-2.910
	(-3.37)	(-3.92)
Session RE	NO	NO
Market type RE	NO	NO
Replication RE	NO	NO
Participant RE	YES	YES
Observations	900	900
AIC	4,314	4,302
BIC	4,362	4,350

Holdings. Table 3 reports the results from pooled final holdings regressions. We focus on final steel holdings (asset 1), as the differences between predicted wood holdings (asset 3) are negligible. The best models explain final steel holdings in terms of (among others) both

Perfect-Foresight Equilibrium holdings *and* the corresponding increments towards Myopic Equilibrium holdings.

According to the regression in column (1), the Perfect-Foresight Equilibrium provides unbiased predictions of steel holdings even in the (baseline) sequential-markets treatment (slope coefficient is only insignificantly different from one). The results in column (2), however, attribute this finding to correlation with holdings predicted under the Myopic Equilibrium. Once the contributions of the two equilibria are disentangled by replacing the difference with the orthogonalized difference, the Perfect Foresight Equilibrium holdings underestimate the outcomes in the baseline (coefficient significantly below one); the (orthogonalized) difference explains the remainder (highly significant, positive coefficient). As expected, interaction with the dummy variable for simultaneous treatment significantly reduces the predictive power of the Myopic Equilibrium. The reduction is even bigger in the non-stationary treatment ($D_{\text{VAR}} = 1$), but it does not completely eliminate the pulling power of the Myopic Equilibrium in the baseline. Adding the two interaction coefficients ($-3.930 - 6.087 = -10.017$) shows that the Myopic Equilibrium still provides explanatory power even in simultaneous non-stationary markets (total coefficient $= 15.945 - 10.017 = 5.928$).

Finally, besides participant-specific random effects, the best model fits both include confounding effects from dummy covariates for market type (simultaneous markets) and endowment type (relatively low cash endowment), as well as their interaction. While the effect of the former is due to a slightly different number of replication types with non-stationary aggregate steel endowments, the latter captures differential trading incentives given initial endowments. The interaction term is simple to interpret: in the simultaneous treatment, participants with low cash endowment (and hence, high steel endowment), sell more of their holdings. In the baseline, they keep more of their endowment (of steel), all things equal.

VI. Implications

A. On The Empirical Record of Asset Pricing Theory

It has long been observed that there is a discrepancy in empirical support (in historical data from the field) for asset pricing theory depending on whether it explains prices in terms of benchmark portfolio values or whether the theory explains prices in terms of covariation with choices such as consumption (Campbell and Cochrane, 2000). Generally, benchmark portfolio models (multi-factor models) generate far more support. Such models only require correct prices. Models that tie prices to choices require choices to be right too, which implies perfect foresight. Our theory shows that prices can be right, while choices are wrong, because they reflect myopia. Therefore, our theory may explain why benchmark portfolio models fit the historical record better than consumption-based models.

There are other puzzles that our theory sheds light on. For instance, it has been observed in a number of studies that option prices were of similar (informational) quality before and after the publication of the Black-Scholes-Merton model (Moore and Juh, 2006; Chambers and Saleuddin, 2019). It seems as if traders magically had already figured out how to correctly price options long before theorists derived the right model. This is particularly bothersome if this evidence dates back to a time when the core mathematical framework, Itô calculus, had not been developed yet. Our theory predicts that this is possible: prices can be right even in the absence of correct (perfect) foresight.

B. The Relevance of Theory for the Practice of Finance

A practical implication of our findings could be phrased in the form of a question: *Why would the industry bother hiring university graduates who are trained in the theory?* Indeed, evidence abounds that market participants do not have perfect foresight and that it is hard to reconcile market prices with theory-implied fundamentals, such as consumption choices. In view of this evidence, it may seem pointless endowing future finance practitioners with theoretical knowledge, especially if the theory relies on "correct prices," as it is ubiquitously the case in both asset pricing and corporate finance. A better approach, it is claimed, would be to teach them the psychology of corporate decision-making (see, e.g., Baker and Wurgler (2013)).

Our answer is that these objections may be unjustified when the theory only relies on correct pricing. Capital structure theory, hedging techniques in options analysis, or duration and immunization analysis in term structure theory "only" require prices to be right, while individual choices may be wrong.

C. Where Prices May Not Be Right

Our main theoretical result relies on myopia. The experiment shows that myopia is a reasonable working hypothesis. We do not expect myopia to obtain in all circumstances, however. Indeed, there are situations in which agents are forced to predict future prices. That is, agents *have to speculate*. For instance, intermediate dividend payments may require agents to think about future prices, if only because "home-made dividends" (synthetic dividends constructed purely from buying and selling securities) may be taxed less heavily than actual dividends. In the face of such market frictions, agents have to consider whether price changes upon future dividend payments are such that collecting dividends is not beneficial. Speculation about ex-dividend prices may lead to mis-pricing, which could be exploited. Therefore, mis-pricing may emerge in areas such as, e.g., high-dividend-yield stocks.

Two recent experiments demonstrate how forced speculation leads to wrong prices. In the first experiment, participants were forced to think about future dividend endowments because their utility depended on how much cash flows today's investments would generate in the future (Asparouhova et al., 2016). In a second experiment, participants were paid only if they assured the right amount of cash flows from future dividends (Asparouhova, Besliu, and Lemmon, 2016). In both experiments, participants had to speculate about future prices, and significant mis-pricing emerged because of lack of perfect foresight.

Likewise, prices will be wrong when agents *want to speculate*. The desire to speculate, and the negative impact on price quality, has been documented in a number of circumstances, such as for technology stocks (Brunnermeier and Nagel, 2004), or when stocks exhibit high betas (Hong and Sraer, 2016).

When prices are right, there is no benefit from active portfolio management: strategic asset allocation in the form of passive investments in a number of indices will eventually outperform actively managed portfolios. However, in situations where investors are forced to think about future re-sale prices of their investments, mis-pricing may emerge. The same may occur when investors voluntarily engage in speculation. Active portfolio management may become beneficial. As such, our findings suggest domains where active portfolio management should concentrate.

Interestingly, prices may be "right" even if agents do not attempt to be right. However, prices may become wrong once agents start speculating about the future, i.e., when they attempt to be "right." Crowd ("swarm") intelligence (Kennedy, 2006) relies on well-adapted simplicity of the members, and in our setting this appears to be true as well. In an attempt to improve their personal welfare, members' efforts to better understand the system as a whole actually destroy the system's intelligence, and hence, the system's ability to make everyone better off. This is disconcerting, because experimental evidence therefore seems to suggest that participant simplicity is necessary for market intelligence.

VII. Conclusion

In this paper, we study traders' ability to attain perfect foresight, a key assumption of dynamic asset pricing theory that is not sufficiently made explicit in the literature. We first show theoretically that, under quadratic preferences, Myopic Equilibrium prices generically are as if agents had perfect foresight, even if equilibrium allocations are different. Second, with a markets experiment, we find strong empirical evidence for myopia in both observed prices and choices.

Our findings have important implications for finance. We provide a foundation for asset pricing that does not require the (empirically rejected) level of rationality and foresight inherent in the traditional premise of market efficiency. Moreover, our findings could explain why asset pricing models that relate prices of risky securities to each other (e.g., multi-factor models) generally fit historical data from field markets better than models that relate prices to choices (e.g., consumption-based models). Finally, our result of price efficiency obtains in the absence of a fully rational marginal investor. However, given price efficiency, trading towards Pareto optimal allocations – although hard – is still feasible at the individual level.

References

Anderson, R. and P. Bossaerts (2019). Modelling asset prices under heterogeneous beliefs. SSRN working paper 3465423.

Anderson, R. M. and H. Sonnenschein (1985). Rational expectations equilibrium with econometric models. *The Review of Economic Studies 52*(3), 359–369.

Asparouhova, E., C. Besliu, and M. Lemmon (2016). Payout policy, investor rationality, and market efficiency: Evidence from laboratory experiments. Available from author's webpage.

Asparouhova, E., P. Bossaerts, and J. Ledyard (2020). Price formation in continuous double auctions; with implications for finance. SSRN working paper 2774301.

Asparouhova, E., P. Bossaerts, N. Roy, and W. Zame (2016). "Lucas" in the Laboratory. *The Journal of Finance 71*, 2727–2780.

Baker, M. and J. Wurgler (2013). Behavioral corporate finance: An updated survey. In *Handbook of the Economics of Finance*, Volume 2, pp. 357–424. Elsevier.

Barberis, N., M. Huang, and R. H. Thaler (2006). Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *The American Economic Review 96*(4), 1069–1090.

Black, F. and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy 81*(3), pp. 637–654.

Bossaerts, P. and P. Hillion (1999). Implementing statistical criteria to select return forecasting models: what do we learn? *Review of Financial Studies 12*(2), 405–428.

Bossaerts, P., D. Meloso, and W. Zame (2008). Dynamically-complete experimental asset markets. SSRN working paper 1103309.

Bossaerts, P., C. Plott, and W. Zame (2007). Prices and portfolio choices in financial markets: Theory, econometrics, experiments. *Econometrica 75*, 993–1038.

Brunnermeier, M. and S. Nagel (2004). Hedge funds and the technology bubble. *The Journal of Finance 59*(5), 2013–2040.

Campbell, J. Y. and J. H. Cochrane (2000). Explaining the poor performance of consumption-based asset pricing models. *The Journal of Finance 55*(6), 2863–2878.

Chambers, D. and R. Saleuddin (2019). Commodity option pricing efficiency before black, scholes, and merton. *The Economic History Review 73*(2), 540–564.

Cox, J. C., S. A. Ross, and M. Rubinstein (1979). Option pricing: a simplified approach. *Journal of Financial Economics 7*(3), 229–263.

Delbaen, F. and W. Schachermayer (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen 300*(1), 463–520.

Duffie, D. (1988). *Security markets: Stochastic models*. Academic Press New York.

Duffie, D. and C.-F. Huang (1985). Implementing arrow-debreu equilibria by continuous trading of few long-lived securities. *Econometrica 53*(6), 1337–1356.

Fama, E. and K. French (2004). The capital asset pricing model: Theory and evidence. Available from author's webpage.

Fudenberg, D. and D. Levine (1983). Subgame-perfect equilibria of finite-and infinite-horizon games. *Journal of Economic Theory 31*(2), 251–268.

Grandmont, J. M. (1977). Temporary general equilibrium theory. *Econometrica 45*(3), pp. 535–572.

Guesnerie, R. (1992). An exploration of the eductive justifications of the rational expectations hypothesis. *The American Economic Review 82*(5), 1254–1278.

Heath, D., R. Jarrow, A. Morton, et al. (1992). Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica 60*(1), 77–105.

Ho, T. S. and S.-B. Lee (1986). Term structure movements and pricing interest rate contingent claims. *The Journal of Finance 41*(5), 1011–1029.

Hong, H. and D. A. Sraer (2016). Speculative betas. *The Journal of Finance 71*(5), 2095–2144.

Jordan, J. S. and R. Radner (1982). Rational expectations in microeconomic models: An overview. *Journal of Economic Theory 26*(2), 201–223.

Kennedy, J. (2006). Swarm intelligence. In *Handbook of nature-inspired and innovative computing*, pp. 187–219. Springer.

LeRoy, S. F. and J. Werner (2014). *Principles of financial economics*. Cambridge University Press.

Long, J. B. D., A. Shleifer, L. H. Summers, and R. J. Waldmann (1990). Noise trader risk in financial markets. *The Journal of Political Economy 98*(4), 703–738.

Lucas, R. E. J. (1978). Asset prices in an exchange economy. *Econometrica 46*(6), 1429–1445.

Magill, M. and M. Quinzii (2002). *Theory of incomplete markets*, Volume 1. Mit press.

Marcet, A. and T. J. Sargent (1988). The fate of systems with "adaptive" expectations. *The American Economic Review 78*(2), 168–172.

Merton, R. C. (1973a). An intertemporal capital asset pricing model. *Econometrica 41*(5), pp. 867–887.

Merton, R. C. (1973b). Theory of rational option pricing. *The Bell Journal of Economics and Management Science 4*(1), 141–183.

Modigliani, F. and M. H. Miller (1958). The cost of capital, corporation finance and the theory of investment. *The American Economic Review 48*(3), 261–297.

Moore, L. and S. Juh (2006). Derivative pricing 60 years before black–scholes: evidence from the johannesburg stock exchange. *The Journal of Finance 61*(6), 3069–3098.

Odean, T. (1999). Do investors trade too much? *American economic review 89*(5), 1279–1298.

Radner, R. (1972). Existence of equilibrium of plans, prices, and price expectations in a sequence of markets. *Econometrica 40*(2), 289–303.

Radner, R. (1974). Market equilibrium and uncertainty: concepts and problems. *Frontiers of Quantitative Economics 2*, 43–105.

Rubinstein, M. (1974). An aggregation theorem for securities markets. *Journal of Financial Economics 1*(3), 225–244.

Scheinkman, J. A. and W. Xiong (2003). Overconfidence and speculative bubbles. *Journal of Political Economy 111*(6), 1183–1220.

Schraeder, S. (2015). Information processing and non-bayesian learning in financial markets. *Review of Finance 20*(2), 823–853.

Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *The Journal of Finance 52*(1), 35–55.

Tversky, A. and D. Kahneman (1974, September). Judgment under Uncertainty: Heuristics and Biases. *Science 185*(4157), 1124–1131.

Appendix 1: Proof of the Main Result

There are two parts to the proof. The first part deals with the case where trading is sequential but no information is released between trading rounds. This case is relevant for the experiment, because there is no uncertainty; probabilities are merely weights with which agent earnings are calculated. In the second part, we consider the case where information is released between trading rounds, and hence, weights (probabilities) are updated.

Case I: No Information Release Between Trading Rounds

We first work with the optimality conditions (first-order conditions) at the individual level (without explicitly using indices that track the individual) and then sum across agents to obtain market-wide predictions. We use quadratic utility, with $b_0 > 0$ (risk aversion) and $b_1 > 0$ large enough so that an interior optimum exists and the first-order conditions are valid. Both parameters may differ across individuals.

For an individual, write the optimality conditions in matrix form, and add the budget constraint for round 1. Define $B := 1/b_0$.

$$\begin{bmatrix} \alpha_1 & \alpha_1 & Bq_1^o \\ \alpha_1 & 3(\alpha_1+\alpha_2+\alpha_3) & B \\ q_1^o & 1 & 0 \end{bmatrix} \begin{bmatrix} x_1^o \\ x_C^o \\ \mu^o \end{bmatrix} = \begin{bmatrix} \alpha_1 B b_1 \\ -\alpha_3 e_3 + 3(\alpha_1+\alpha_2+\alpha_3) B b_1 \\ q_1^o e_1 + e_C \end{bmatrix} \tag{1}$$

Since these are optimality and budget constraints for the representative agent, these will also be equilibrium conditions. That is, prices also need to satisfy these equations. Round-2 optimality conditions and budget constraint are as follows:

$$\begin{bmatrix} \alpha_3 & \alpha_3 & Bq_3^o \\ 3(\alpha_1+\alpha_2+\alpha_3) & \alpha_3 & B \\ 1 & q_3^o & 0 \end{bmatrix} \begin{bmatrix} y_C^o \\ x_3^o \\ \pi^o \end{bmatrix} = \begin{bmatrix} \alpha_3 B b_1 \\ -\alpha_1 x_1^o + 3(\alpha_1+\alpha_2+\alpha_3)B b_1 \\ x_C^o + q_3^o e_3 \end{bmatrix} \tag{2}$$

Complete-markets optimality conditions, for comparison, are the following:

$$\begin{bmatrix} \alpha_1 & \alpha_1 & 0 & Bp_1^* \\ \alpha_1 & 3(\alpha_1+\alpha_2+\alpha_3) & \alpha_3 & B \\ 0 & \alpha_3 & \alpha_3 & Bp_3^* \\ p_1^* & 1 & p_3^* & 0 \end{bmatrix} \begin{bmatrix} x_1^* \\ x_C^* \\ x_3^* \\ \lambda^* \end{bmatrix} = \begin{bmatrix} \alpha_1 B b_1 \\ 3(\alpha_1+\alpha_2+\alpha_3)B b_1 \\ \alpha_3 B b_1 \\ p_1^* e_1 + e_C + p_3^* e_3 \end{bmatrix} \tag{3}$$

We assume that prices and choices in the complete-markets case are such that equilibrium obtains. *Our question is: if we set $q_1^o = p_1^*$, $q_3^o = p_3^*$ in (1) and (2), will we be able to clear markets?* That is, for the proposed prices, would excess demands equal zero?

To answer this question, we first add the first-order conditions (including budget constraints) across agents. To simplify notation, variables will now refer to sums of demands, sum of endowments, and risk-tolerance weighted Lagrange multipliers. E.g.,

$$x_3^0 := \sum_k x_3^{k,0},$$

where k indexes agents,

$$e_C := \sum_k e_C^k,$$

and

$$\lambda^* := \sum_k \lambda^{k,*} B^k \quad (= \sum_k \lambda^{k,*}/b_0^k).$$

After summation across agents, subtract complete-markets equations from corresponding incomplete-markets equations:

$$\begin{aligned}
\alpha_1(x_1^o - x_1^*) + \alpha_1(x_C^o - x_C^*) + (\mu^o - \lambda^*)p_1^* &= 0 && (4)\\
\alpha_3(y_C^o - x_C^*) + \alpha_3(x_3^o - x_3^*) + (\pi^o - \lambda^*)p_3^* &= 0 && (5)\\
\alpha_1(x_1^o - x_1^*) + 3(\alpha_1 + \alpha_2 + \alpha_3)(x_C^o - x_C^*) + \alpha_3(e_3 - x_3^*) + (\mu^o - \lambda^*) &= 0 && (6)\\
\alpha_1(x_1^o - x_1^*) + 3(\alpha_1 + \alpha_2 + \alpha_3)(y_C^o - x_C^*) + \alpha_3(x_3^o - x_3^*) + (\pi^o - \lambda^*) &= 0 && (7)
\end{aligned}$$

Since $x_1^* = e_1$, $x_C^* = e_C$, $x_3^* = e_3$ (we assume that we are in the complete-markets equilibrium), we can substitute e_1, e_C and e_3 for x_1^*, x_C^* and x_3^*. After this, multiply (6) by p_1^*, subtract the result from (4), and eliminate terms equal to zero:

$$\alpha_1(1 - p_1^*)(x_1^o - e_1) + (\alpha_1 - 3(\alpha_1 + \alpha_2 + \alpha_3)p_1^*)(x_C^o - e_C) = 0. \tag{8}$$

Do the analogous for (5) and (7), taking into account that $x_1^o = e_1$ since agents cannot trade asset 1 in round 2:

$$(\alpha_3 - 3(\alpha_1 + \alpha_2 + \alpha_3)p_3^*)(y_C^o - e_C) + \alpha_3(1 - p_3^*)(x_3^o - e_3) = 0. \qquad (9)$$

Now add the incomplete-markets budget constraints in order to generate a system of equations in excess demands. Instead of blindly adding the round-2 budget constraint, we add it to the round-1 budget constraint in order to eliminate the round-1 choice of cash (which enters as an endowment in the round-2 budget constraint). So, add lines 3 in (1) and (2) (at complete-markets equilibrium prices):

$$p_1^*(x_1^o - e_1) + (x_C^o - e_c) + (y_C^o - x_C^o) + p_3^*(x_3^o - e_3)$$

$$= p_1^*(x_1^o - e_1) + (y_C^o - e_C) + p_3^*(x_3^o - e_3) = 0.$$

The system of equations defining the equilibrium becomes:

$$\begin{bmatrix} \alpha_1(1-p_1^*) & \alpha_1 - 3(\alpha_1+\alpha_2+\alpha_3)p_1^* & 0 & 0 \\ 0 & 0 & \alpha_3 - 3(\alpha_1+\alpha_2+\alpha_3)p_3^* & \alpha_3(1-p_3^*) \\ p_1^* & 1 & 0 & 0 \\ p_1^* & 0 & 1 & p_3^* \end{bmatrix} \begin{bmatrix} x_1^o - e_1 \\ x_C^o - e_C \\ y_C^o - e_C \\ x_3^o - e_3 \end{bmatrix}$$

$$= \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

Since the coefficient matrix is of full rank, there is only one solution to this homogeneous system of equations, namely, the vector of zeros. That is, markets clear, while demands are optimal.

Case II: Information Release Between Trading Rounds

We now consider a situation where information is released between trading rounds 1 and 2. This leads to updates of the weights, from α_s to α'_s (all s), since these are probabilities. The new weights constitute random variables: they depend on the particulars of the information release. In the Perfect-Foresight Equilibrium, agents are assumed to be Bayesian. This implies the following restrictions on weights:

$$\alpha_s = E[\alpha'_s].$$

That is, round-1 weights equal expected round-2 weights. We will not need this restriction in our proof, however.

Before, we compared allocations and prices in the Myopic Equilibrium against those of the complete-markets (simultaneous-trading) equilibrium. Now, we compare the Myopic

Equilibrium to the Perfect-Foresight Equilibrium. This is possible because Duffie and Huang (1985) proved that the complete-markets equilibrium allocations and prices can be implemented in a Perfect-Foresight Equilibrium. We thus assume that the Perfect-Foresight Equilibrium prices p_1^+ and $p_3'^+$ and equilibrium choices x_1^+, x_C^+, $y_C'^+$ and $x_3'^+$ implement the complete-markets equilibrium prices p_1^* and p_3^* and equilibrium allocations x_1^*, x_C^* and x_3^*. Notice that $p_3'^+$, $y_C'^+$ and $x_3'^+$ depend on the information released between trading rounds; there are as many outcomes for these variables as there are possible information releases. See Duffie and Huang (1985) for details.

Let us start with remembering the first-order conditions for optimal allocations in the Perfect-Foresight Equilibrium. They look like those in a Myopic Equilibrium, for one crucial exception. For general utility functions (strictly concave, so an interior solution is possible), round-1 conditions are as follows:

$$\begin{aligned}
\alpha_1 u'(x_1^+ + x_C^+) &= \lambda^+ p_1^+, \\
\alpha_1 u'(x_1^+ + x_C^+) + \alpha_2 u'(x_C^+) + \alpha_3 u'(e_3 + x_C^+) &= \lambda^+, \\
p_1^+ x_1^+ + x_C^+ &= p_1^+ e_1 + e_C.
\end{aligned}$$

Generic round-2 conditions are:

$$
\begin{aligned}
\alpha_1' u'(x_1^+ + y_C'^+) + \alpha_2' u'(y_C'^+) + \alpha_3' u'(x_3'^+ + y_C'^+) &= \theta'^+, \\
\alpha_3' u'(x_3'^+ + y_C'^+) &= \theta'^+ p_3'^+, \\
y_C'^+ + p_3'^+ x_3'^+ &= x_C^+ + p_3'^+ e_3.
\end{aligned}
$$

The are as many round-2 condition triplets as there are possible information releases. Importantly, x_1^+ and x_C^+ are determined by *all* conditions. In a Myopic Equilibrium, they are determined *only* by round-1 conditions. In that sense, the conditions are very different for myopia.

Translated into quadratic utility, and using matrix notation, these equations become:

$$
\begin{bmatrix} \alpha_1 & \alpha_1 & Bp_1^+ \\ \alpha_1 & 3(\alpha_1+\alpha_2+\alpha_3) & B \\ p_1^+ & 1 & 0 \end{bmatrix}
\begin{bmatrix} x_1^+ \\ x_C^+ \\ \lambda^+ \end{bmatrix}
=
\begin{bmatrix} \alpha_1 B b_1 \\ -\alpha_3 e_3 + 3(\alpha_1+\alpha_2+\alpha_3) B b_1 \\ p_1^+ e_1 + e_C \end{bmatrix}
\tag{10}
$$

$$
\begin{bmatrix} \alpha_3' & \alpha_3' & Bp_3^+ \\ 3(\alpha_1'+\alpha_2'+\alpha_3') & \alpha_3' & B \\ 1 & p_3^+ & 0 \end{bmatrix}
\begin{bmatrix} y_C'^+ \\ x_3'^+ \\ \theta'^+ \end{bmatrix}
=
\begin{bmatrix} \alpha_3' B b_1 \\ -\alpha_1' x_1^+ + 3(\alpha_1'+\alpha_2'+\alpha_3') B b_1 \\ x_C^+ + p_3'^+ e_3 \end{bmatrix}
\tag{11}
$$

The Myopic-Equilibrium conditions are the same as before, except that round-2 restrictions apply separately to each possible information release. So we do not repeat them here. We will add “primes” to round-2 decision variables, in order to make clear that these variables depend on information release.

As before, we first add the first-order conditions (including budget constraints) across agents. To simplify notation, variables will now refer to sums of demands, sum of endowments, and risk-tolerance weighted Lagrange multipliers. After summation across agents, subtract perfect-foresight equations from corresponding incomplete-markets equations:

$$\alpha_1(x_1^o - x_1^+) + \alpha_1(x_C^o - x_C^+) + (\mu^o - \lambda^+)p_1^+ = 0 \quad (12)$$

$$\alpha_3'(y_C'^o - y_C'^+) + \alpha_3'(x_3'^o - x_3'^+) + (\pi'^o - \theta'^+)p_3'^+ = 0 \quad (13)$$

$$\alpha_1(x_1^o - x_1^+) + 3(\alpha_1 + \alpha_2 + \alpha_3)(x_C^o - x_C^+) + (\mu^o - \lambda^+) = 0 \quad (14)$$

$$\alpha_1'(x_1^o - x_1^+) + 3(\alpha_1' + \alpha_2' + \alpha_3')(y_C'^o - y_C'^+) + \alpha_3'(x_3'^o - x_3'^+) + (\pi'^o - \theta'^+) = 0 \quad (15)$$

Since $x_1^+ = e_1$, $x_C^+ = e_C$, $y_C'^+ = e_C$, $x_3'^+ = e_3$ (we assume that we are in the Perfect-Foresight Equilibrium), we can substitute e_1, e_C, e_C and e_3 for x_1^+, x_C^+, $y_C'^+$ and $x_3'^+$, respectively. After this, multiply (14) by p_1^+, subtract the result from (12), and eliminate terms equal to zero:

$$\alpha_1(1 - p_1^+)(x_1^o - e_1) + (\alpha_1 - 3(\alpha_1 + \alpha_2 + \alpha_3)p_1^+)(x_C^o - e_C) = 0.$$

Do the analogous for (13) and (15), taking into account that $x_1^o = e_1$ since agents cannot trade asset 1 in round 2:

$$(\alpha_3' - 3(\alpha_1' + \alpha_2' + \alpha_3')p_3'^{+})(y_C'^{o} - e_C) + \alpha_3'(1 - p_3'^{+})(x_3'^{o} - e_3) = 0.$$

Formally, we have obtained the same equations as in the proof for the case without information release; see Equations (8) and (9). The remainder of the proof therefore proceeds as before.

APPENDIX 2: INSTRUCTIONS AND SCREEN SNAPSHOTS

THE UNIVERSITY OF MELBOURNE

BRAIN MIND AND MARKETS LABORATORY

INSTRUCTIONS

Summary

You will trade 'inputs' (wood, steel, and plastic) in an online market with other participants. The goal is to collect inputs which will let you produce Widgets. You should try to produce as many Widgets as possible because your earnings will depend on it. You will have access to a spreadsheet which tells you how many Widgets you can get for a given amount of wood, steel, and plastic.

Plastic is a special input because it can be used to replace wood or steel in the production of Widgets.

Steel can be traded for plastic in the 'steel market' and wood can be traded for plastic in the 'wood market'. You cannot trade steel and wood for one-another directly but must trade through plastic. So, plastic acts as "cash", and we will often refer to it as cash.

This game will be replicated several times, switching between situations where you can simultaneously trade in the steel and wood markets, and situations where you must first trade in the steel market and then in the wood market.

Online Platform

In this experiment, you are asked to trade with other players online through a platform called flex-e-markets. You can access the platform by typing http://www.flexemarkets.com into your browser. Once on the site, click 'sign in' on the top right corner of the page. On the sign-in page, you will need to type the account name **'smiley-chum'** in the 'Account' field. Your unique email and password will be provided to you on a separate piece of paper. Type your unique email in the 'E-Mail' field and your unique password in the 'Password' field. Click sign in once this has been done.

Trading Game

In the trading game, the aim is to trade with other participants to collect inputs called steel, plastic, and wood. These inputs are used to produce 'Widgets'. The number of Widgets you produce determines your performance (see performance, below). You will have access to two markets: the steel market where you can trade plastic for steel, and the wood market where you can trade plastic for wood. Note that steel and wood cannot be traded directly for one another, but only through plastic. Plastic does not have its own market. So, **plastic acts as cash**, and will be called **cash** in flex-e-markets.

Additionally, plastic is a versatile input and can replace either wood or steel to produce Widgets. Plastic therefore affects production of Widgets in two ways: directly (because part of the widget is made of plastic) and indirectly (because using plastic obviates the need to use steel or wood). Therefore, steel and wood should in general be worth less than half as much as plastic because plastic is a more flexible input in producing Widgets. That is, the price of steel and wood in terms of plastic (cash) should generally be less than 0.5.

Sessions

Once you have signed in, you will see two marketplaces called 'InputsForWidgets--Practice' and 'InputsForWidgets-Real'. InputsForWidgets-Practice is a practice market which will give you time to get used to trading. InputsForWidgets-Real is the market where the real trading will occur. Within this marketplace we will run as many sessions as we can in the allotted time. The sessions will be of two distinct types:

1. The first type of session will be referred to as **SIMULTANEOUS**. In the simultaneous sessions, you will be allowed to trade in both markets (steel and wood) at the same time. Each session will last at least 10-minutes.
2. The second type of session will be referred to as **SEQUENTIAL**. Sequential sessions will be broken into two periods which will last at least 5 minutes each. In the first period, you trade ONLY in steel. In the second period, you trade ONLY in wood. Note that **you will be penalised if you trade in a restricted market** (e.g. if we are in the first period, where the wood market is closed, any trade you make in the wood market will reduce your final payment by $2.00).

We will start with a number of sequential sessions and then switch to simultaneous sessions. At the end of each session, we will record your performance for that session and reset your holdings. Please note that before we begin a new session, you should refresh the browser so that your holdings update accordingly. (You may want to do the same during trading if you notice that your holdings do not update correctly after a trade.)

Performance

Your performance is based on the number of Widgets you produce – the more Widgets you make, the more money you earn. As you increase one type of input, however, its effect on the number of Widgets decreases. Plastic has more of an effect because it is at once an input in its own right and it can be used to replace wood or steel.

You will be given access to an **online spreadsheet** (see separate piece of paper) that calculates your performance based on the number of inputs you currently hold. The performance calculator also allows you to see how your performance will change if you execute a specific trade. You will be given sufficient time to get familiar with the performance calculator during the training period. The performance calculator can be accessed by typing the link under 'Performance Calculator', on the piece of paper with your unique email and password, into your preferred web browser.

Your payment will be $20 plus the number of Widgets you produce in 1 session that we will choose at random at the end of the experiment. That is, each Widget you produce translates into one (1) dollar.

We recommend that you open the performance calculator and the marketplace next to one another, so that you can save time switching back and forth during the experiment.

Trading

Trading takes place as follows. You submit limit orders: orders to buy a chosen number of inputs at a chosen price (or lower), or to sell a chosen number of inputs at a chosen price (or higher). Transactions take place from the moment a buy order with a higher price crosses a sell order with a lower price or the other way around. All trade occurs at the price specified by the best standing order. In other words, if a trade occurs, the price of the earlier best order determines the price. Orders at a better price execute first. Given a price, orders arriving earlier execute first. Orders remain valid until you cancel them or the marketplace closes. You will be given sufficient time to practice submitting and cancelling orders.

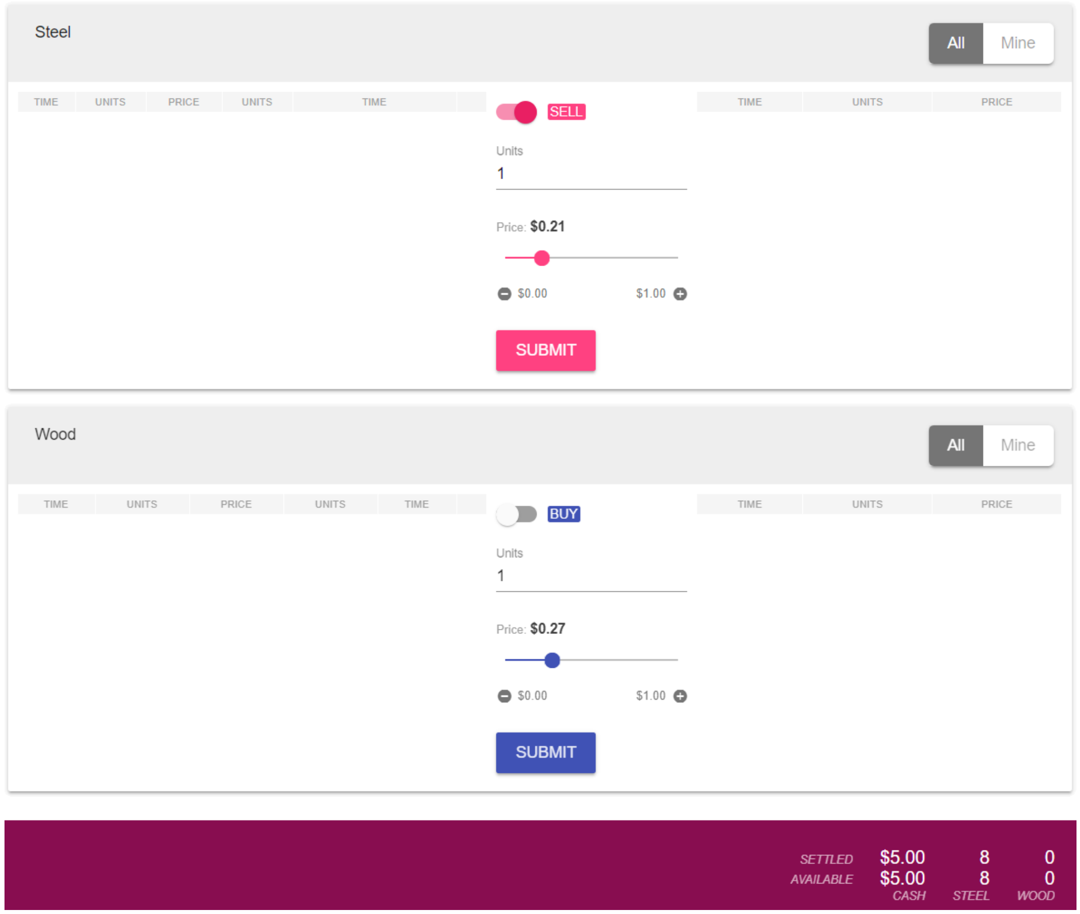

Figure 5: Trading Interface Snapshot

There is a market for steel (top) and wood (bottom). The bottom right displays participant settled holdings (assets they currently hold) and available holdings (which takes into account standing orders). There is no market for cash (plastic) since it is the numeraire. Participant holdings of cash are displayed to the bottom right. There is a slider which allows participants to select a price at which they wish to trade either wood or steel, ranging from $0.00 to $1.00, in $0.01 increments. There is a switch which participants click to change between the sell and buy side. Orders are limit orders, that is, participants select the number of units and the price at which they are willing to trade. Orders are executed on a price-time priority, that is, standing orders at the best price are executed first. If there are several orders at the best price then the longest-standing order is executed first. To the left of the "submit" button is the (still empty) book of orders, to the right is the (empty) list of past transactions. More information on the trading interface can be obtained from http://www.adhocmarkets.com.

Initial Holdings	
Plastic (Cash)	5
Steel	8
Wood	0
Performance	**$4.35**

Current Holdings	
Plastic (Cash)	5
Steel	8
Wood	0
Performance	**$4.35**

Possible Future Trades				
	Buy		**Sell**	
	Quantity	**Price (Plastic)**	**Quantity**	**Price (Plastic)**
Steel				
Wood				

Post-Trade Holdings	
Plastic (Cash)	5
Steel	8
Wood	0
Performance	**$4.35**

Perf. Change	**$0.00**

Figure 6: Online Spreadsheet Aid to Determine Production of Widgets as a Function of Inputs

Example of the performance calculator which helped participants trade. The "Initial Holdings" section was fixed and could not be altered by participants. The "Current Holdings" section allowed participants to determine the performance of their current holdings. The "Possible Future Trades" section allowed participants to input a buy/sell order for a chosen number of steel or wood units at a chosen price. The "Post-Trade Holdings" section showed participants how their holdings and performance would change if they executed the trades specified in "Possible Future Trades." The "Perf. Change" showed participants whether the trades specified in "Possible Future Trades" would increase or decrease their overall performance.