Journal of Economic Behavior & Organization 68 (2008) 153–166

JOURNAL OF Economic Behavior & Organization

www.elsevier.com/locate/econbase

Don't let your robots grow up to be traders: Artificial intelligence, human intelligence, and asset-market bubbles

Ross M. Miller *

State University of New York (SUNY) at Albany and Miller Risk Advisors, 2255 Algonquin Road, Niskayuna, New York 12309, USA

Received 18 June 2003; received in revised form 31 July 2007; accepted 10 August 2007
Available online 20 May 2008

Abstract

Researchers who have examined markets populated by "robot traders" have claimed that the high level of allocative efficiency observed in experimental markets is driven largely by the "intelligence" implicit in the rules of the market. Furthermore, they view the ability of agents (artificial or human) to process information and make rational decisions as unnecessary for the efficient operation of markets. This paper presents a new series of market experiments that show that markets populated with standard robot traders are no longer efficient if time is a meaningful element, as it is in all asset markets. While simple two-season markets with human subjects reliably converge to an efficient equilibrium, markets with minimally intelligent robot traders fail to attain this equilibrium. Instead, these markets overshoot the equilibrium and then crash below it. In addition to firmly establishing the role of trader intelligence in asset-market equilibrium, these experiments also provide insights into why bubbles and crashes are consistently observed in many asset-market laboratory experiments using human subjects.

Keywords: Market bubbles; Intertemporal competitive equilibrium; Experimental asset markets; Autonomous market agents

1. Introduction

A particularly striking result from the early days of experimental economics was Smith's (1962) discovery that a simple auction mechanism patterned after the one used on the floor of the New York Stock Exchange (NYSE) leads to the efficient competitive equilibrium market allocation under a wide range of circumstances. The principal difference between the mechanism that Smith used, which he referred to as a double-oral auction, and the one used on the floor of the exchange was that the "specialist" who managed the auction (a role subsumed by the experimenter) was not required to step in and provide liquidity to the market by trading for his own account.

In experiment after experiment spanning several years and using subjects ranging from high school students to seasoned commodities traders, results similar to Smith's have been obtained, with the market quickly converging to the competitive equilibrium price and quantity. Furthermore, the resulting allocation tended toward 100 percent efficiency relative to the total surplus available in the market. Hence, subjects were taking away from these experiments almost all of the money that was available to them. Moreover, all this occurred without providing

* Tel.: +1 518 346 0781.
E-mail address: millerrm@alumni.caltech.edu.

0167-2681/$ – see front matter
doi:10.1016/j.jebo.2007.08.010

subjects the "perfect knowledge" that economists, with the notable exception of Hayek, had once assumed was necessary for markets to operate efficiently. (Subjects in these experiments knew only their own supply and demand parameters.)

As Smith and his colleagues gained more experience with experimental markets, it became clear that the ability of the markets to converge to the competitive equilibrium did not require that subjects involved in the experiment have an understanding of what they were doing. Indeed, it was common for one or more subjects in an experiment to behave in an erratic manner, either intentionally or from misunderstanding the rules, yet only minimally affect the market. Such behavior by subjects, which was beyond the experimenter's control, was part of the "human element" that the experiment was designed to investigate.

The computerization of experimental markets (Williams, 1980) made it easy to insert preprogrammed "robots" into experiments. With computer-mediated interactions, it is simple to conceal the presence of robots from human subjects when necessary. Such robots were also suited to game-theoretic experimentation in nonmarket environments because many strategies for common games could easily be implemented in machine form. Pitting man against machine enabled experimenters to elicit the human response to a variety of strategies. Artificial intelligence in robot form, however, became an end in itself, with the challenge being not to mimic humans, but rather to make better (and more profitable) decisions than them.

Following the success of Axelrod (1984) prisoners' dilemma tournament, Rust et al. (1993) conducted a series of double-oral auctions robot tournaments at the Santa Fe Institute beginning in March 1990. These tournaments set robots programmed in a variety of procedural computer languages against one another. Each robot was capable of observing every bid, offer, and transaction in the auction market and could condition its actions on the past (and projected future) states of the market.

As one might expect, the robots that made the most money for their programmers in the Santa Fe Institute tournament were "aware" of their environment. The overall winner of the tournament employed a variant of a simple strategy known as "steal the deal." This strategy involved leaving the price discovery process (and its associated costs) to the other robots, waiting until the last possible moment to trade on favorable terms.[1] Of course, such a strategy, which is individually "rational," only works if there are other, not-so-rational, robots to facilitate price discovery in the first place.[2]

The Santa Fe Institute tournament also served to stimulate interest in how well human experiments involving double-oral auctions might be replicated by markets populated only by robots. Such markets turned out to share the tendency of human markets to converge on the competitive equilibrium price and quantity. Human and robots, however, generally take much different paths to equilibrium. Moreover, neither of these paths could be fully reproduced under controlled simulations involving archetypical robots (Rust et al., 1993; Cason and Friedman, 1993, 1996).

The baseline for all robot archetypes is zero-intelligence (ZI) agent. This simple robot is programmed to generate bids and offers selected randomly from a uniform distribution subject only to the constraint it cannot "deliberately" lose money. ZI agents are designed to be oblivious to their environment; all their bids and offers were absolutely determined without regard to the past or expected future actions of other market participants. Moreover, ZI agents do not control the timing of their actions; when they placed or accepted bids and offers is also randomly determined for them. ZI agents thereby lack the intelligence to respond to their environment or even to take actions that compensate for their inability to respond.

Gode and Sunder (1993) explored the properties of markets composed solely of ZI agents prior to the Santa Fe Institute tournament. They found that these markets generated average prices and quantities that approached the competitive equilibrium. Market with ZI agents were also allocatively efficient, on average extracting more than 95 percent of the available surplus from the market. This is in line with what was found both in experimental markets with human traders and in the most efficient runs of the Santa Fe Institute tournament.

The intentional randomness of the ZI agents had the direct effect of injecting substantial volatility into the market. Moreover, this randomness induced negative serial correlation in the first-order difference in prices from transaction to transaction within a trading period. Markets with ZI agents failed to exhibit most of the subtleties associated withmarkets

[1] This practice has become well-established on Internet auction sites, and specialized "bots" have been developed to facilitate last-second bidding.
[2] See Rust et al. (1993, pp. 187–193) for a discussion of stability issues that arise with the "steal the deal" strategy.