Evaluating Multi-Agent Commodity Market Simulation using Global Coherence

[redacted]

Formatted: Highlight

Multi-Agent systems (MAS) have been used in the study of economic models for many decades. MAS can provide a framework in which to model economic systems by introducing autonomous actors (agents) that are able to individually achieve some goal. This models real-world economic systems in that rational actors – individuals participating in commerce, each working towards their own objective, collectively control the behavior of the system. Events like global economic panic can be triggered by the actions of a small number of actors influencing the individual actions of other actors. As such, designing economic systems models for the study of complex interactions can be made easier through the use of MAS to achieve an emergent behavior. This paper evaluates global coherence as influenced by local decisions while modeling complex economic interactions. We consider the question: what role do self-interested agents have on the emergence of price stability in a commodity market? We evaluate the effectiveness of using REPAST to model such an environment to study this behavior and we address the issues of using an economic model to study emergent behavior.

Commented [LS1]: OK.

General Terms: Multi-agent Systems

Additional Key Words and Phrases: Multi-agent Systems, Agents, Game Theory, Economic Models, Auction, First-price Sealed Auction, Bellman Equation, Dynamic Programming

1. INTRODUCTION

Using computational analysis to study the behavior of economic systems is a complicated endeavor because of the sheer complexity and number of interactions that take place in such systems. Multi-Agent Systems (MAS) can be useful in reducing the perceived computational complexity of such systems by providing an environment in which to simulate the actions of rational actors in isolation. This approach allows systems designers to focus on the behavior of individual actors in the system, resulting in the emergence of complex interactions within the model that can be studied. One could imagine that this approach could be beneficial to studying the complex behaviors of economic systems. In an MAS – much like a real economic system – the individual actions of the system actors will collectively lead to the emergence of some behavior of the overall system; for example the global emergence of credit limitation as a result of individual bank evaluation of asset values[1]. The emergent behavior of a multi-agent system could be exploited to simulate a complicated system property by modeling the simpler actions of system actors that will lead to the behavior to be studied.

Commented [LS2]: For your footnote, do you have a reference to back it up?

This paper seeks to evaluate the usefulness of MAS in studying complex economic behavior using the emergent global coherence of agents making individual decisions by evaluating price stability in a simulated commodity market. Our simulation will attempt to demonstrate whether or not the actions of agent types can actually produce the expected emergent system behavior in our simple economic model. In order to identify expected emergent behavior, we will consider the trading price of oil in a system made up of oil producing agents who sell oil, oil consuming agents who need to purchase some amount of oil, and speculator agents who will purchase oil intending to sell it for a profit. The interactions of the various agents, we hypothesize, will lead to the emergence of either price fluctuation or stability as the availability of oil in the system changes. In evaluating our model, we hypothesize that the following intuitive properties will hold:

- **Hypothesis 1:** The trading price of oil is inversely proportional to the total number of produced barrels in the system $\sum P_i$. If oil is a scarce commodity, producers will command a high price for

[1] Among many other factors in 2008, individual banks put limits on the amount of credit they would extend due to the questionable values placed on assets – mortgage backed securities – which had the unintended consequence of placing further constraints on the investment banks holding the assets, forcing devaluation of those institutions stock, which devalued assets further and prevented banks from loaning. The individual decisions of banks resulted in a lack of available capital within the system. While a gross over simplification of what was actually going on in the markets at the time, this example serves to demonstrate the emergence of some behavior (system not providing loans) based on individual agent actions (devaluing of assets).

offered barrels as a result of the consumer need to purchase some number of barrels before the simulation ends. For small values of $\sum P_i$ trading price will be large and for large values of $\sum P_i$, the trading price will be low.

- **Hypothesis 2:** The trading price of oil is proportional to total capital available to all consumer agents and all speculator agents $\sum (S_i + C_i)$. Because agents are limited in their ability to value oil greater than the capital with which they have to purchase it, the increased availability of capital will allow prices to rise and increase the trading price of oil. For large values of $\sum (S_i + C_i)$, the trading price will quickly converge to $\sum (S_i + C_i)/P$.

Commented [LS3]: What is P?

- **Hypothesis 3:** Convergence on a maximum trading price of $\sum (S_i + C_i)/P$, will have the same effect as an abundance of oil. Once the trading price of oil hits a maximum value, agents will no longer be able to highly value it, due to their inability to buy at that price (i.e. the limiting factor of the value function is *S* or *C*) and trading prices will fall. As a result of the lowering price, agents will decreasingly value oil, causing prices to "bottom out" at $\sum P_i / \sum N_i$.

Commented [LS4]: What is N_i?

- **Hypothesis 4:** The number of speculator agents participating in the exchange will decrease the stability of the trading price where stability is measured as a consistent trend. Large numbers of speculator agents buying oil demanded by consumer agents will result in rapid trading price increases until convergence is met and prices fall. As prices fall, speculator agents will continue to compete with consumer agents to purchase oil to drive the price back up in an effort to "sell high". This will result in rapid price increase and decrease; thereby decreasing stability in trading price.

Commented [LS5]: OK.

In designing agents that, through individual behavior, will have the hypothesized effect on trading price, we will demonstrate the efficacy of using a MAS to study complex emergent behavior. The remainder of this paper is structured as follows: section two presents our system design, section three discusses how we constructed experiments to test whether our system was able to produce the desired emergent behavior, section four presents our results, and section five closes with the conclusions we draw from our data and future work.

2. SYSTEM DESIGN

A basic commodities market allows actors with some commodity to sell the commodity to other actors who demand it. While there are very complex rules governing real commodity markets, we have limited the scope of our market for the purposes of this paper; our model will simulate very basic agent interactions that impact trading price. Our model is implemented as a turn-based system in which producer agents offer some number of barrels in an auction to consumers who can purchase them for a price. At the end of each turn, all auctions are settled, and sold barrels and paid amounts are transferred between producer and consumer agents. The simulation is governed by the following basic control parameters:

1. the total number of turns that will take place in a simulation T;
2. the number of agents of each type that will be participating in trades are α, β, and γ for producer, consumer and speculator agents, respectively;
3. the number of produced barrels P, where P_i represents the number of produced barrels randomly assigned a given producer agent i satisfying the condition $\sum P_i <= P$;
4. the total production cost PC associated with producing P barrels and assigned randomly to producer agents i satisfying the condition $\sum PC_i <= P \mid \forall\ PC_i, Prod_{min} <= PC_i <= Prod_{max}$, where $prod_{max}$ and $prod_{min}$ are the maximum and minimum amounts respectively it can cost to produce a single barrel of oil;

5. the number of demanded barrels D where D_i represents the number of demanded barrels randomly assigned a given consumer agent i satisfying the condition $\sum D_i <= D$;
6. the capital available for consumer and speculator agents, C and S respectively, with C_i and S_i representing the amount of capital randomly assigned consumer and speculator agents i and satisfying the condition $\sum (C_i + S_i) <= C+S \mid \forall\ C_i, Pay_{min} <= C_i <= Pay_{max}, S_i, Pay_{min} <= S_i <= Pay_{max}$, where Pay_{max} and Pay_{min} are the maximum and minimum amounts respectively an agent can pay per barrel;
7. the contract length CL which is the maximum number of time units a given consumer has to fulfill their demand for oil where CL_i is the random number of time units that a consumer agent i is assigned for a contract length satisfying the condition $\sum CL_i <= CL \mid \forall\ CL_i, CL_{min} <= CL_i <= CL_{max}$, where CL_{max} and CL_{min} are the maximum and minimum lengths for a contract respectively, $CL_{min} = CL_{max} = T$ effectively removes a contract length from the simulation;
8. the cost per turn associated with storing a barrel of oil C_{store};
9. the maximum percentage of inventory that a producer can offer in a given auction, $offer_{max}$;
10. the discount factor to be applied to the current trading price when computing the new trading price under a trading volume of 0, *NS*;
11. the number of previous turns to consider when computing the price trend, *wind;*
12. the global discount factor, δ, that will be applied to an agent's estimation of future reward.

Commented [LS6]: Hmm … why cap it?

Commented [LS7]: Is this realistic? Do we have this in real-world?

Commented [LS8]: Pretty good set.

The logic that governs the turns of the simulation is separated into four primary units of functionality: the trading floor, producer agents, consumer agents, and speculator agents. These functional units are created by the simulation and initialized according to provided parameters. The simulation will then extract information at each tick count from the elements in order to populate data output components.

2.1 Trading Floor

The primary functional unit for a simulation is the Trading Floor, which is used to allow agents in the system to report price, facilitate transfer of capital for oil, and provide an environment in which agents could discover sales of oil. In addition to agent coordination[2], the Trading Floor keeps track of the current price for oil, the trading volume, price trends, and reports variable data back to the simulation for data output. The Trading Floor is also responsible for reporting the basic trading rules of the system. In this way, the Trading Floor does have an impact on the local decisions that can be made by the agents as these decisions must be made within the context of the rules of the system.

The Trading Floor is initialized at the outset of each simulation run with values for T, NS, P, $P\times offer_{max}$, *wind*, and collection objects for α, β, and γ agents. Beyond the argument parameters, the Trading Floor provides the following calculated parameters throughout the simulation:

Commented [LS9]: This is the global information sharing … that means the agents are not very localized in their decision making …

- $price_{current}$, the current trading price pre barrel of oil;
- $price_{min}$, the lowest price paid per barrel of oil during the simulation;
- $price_{max}$, the highest price paid per barrel of oil during the simulation;
- $price_{trend}$, the trend, positive or negative, that $price_{current}$ is taking from the previous *wind* terms;
- $volume_{total}$, the total number of barrels traded on the Trading Floor during the simulation;
- $volume_i$, the number of barrels traded in the last turn;
- bid_{total}, the total amount paid for $volume_i$ barrels;
- $bid_{Average}$, the average amount offered for $volume_i$ barrels;
- T', the number of turns taken in the simulation;
- $capital_{remaining}$, the amount of capital remaining in the system;

Commented [LS10]: This is a very SIGNIFICANT global information – this will artificially control the environment, and will more likely lead to coherence. Key weakness.

[2] Coordination here is not meant to imply a loss of autonomy on the part of the agents. The Trading Floor coordinates information flow between the agents which use the information to make local decisions.

These properties are computed by the Trade Floor at the beginning and end of each round of auctions. The primary role of the Trade Floor is to hold auctions and oversee trading turns.

2.1.1 Auction

Commented [LS11]: Ok.

The principal action of the Trade Floor is facilitating sales of produced barrels by producer agents to consumer or speculator agents. This is accomplished in our market model as an auction. We have selected a first-price sealed auction as the method for offering barrels in the market[3]. The auctions are created in the manner described in section 2.1.1. Each auction is created with a random number of barrels, $offer_i$ that is randomly generated number of barrels calculated as $ceil(offer_{Max} \times P_i') \times rand(0..1))$, where P_i' is P_i less the number of barrels already sold by the agent. Consumer or speculator agents may place bids bid_i, which are composed of a number of barrels demanded and the price that the agent will pay. Bids are evaluated by the auction holder and assigned a score using the value function described below; the highest scored bid is the winner of the auction. The winning agent pays the auction holder $bid_i.price$ for $bid_i.barrels \mid bid_i.barrels <= offer_i$; if $bid_i.barrels > offer_i$ then $offer_i$ barrels are transferred at $bid_i.price$. The decision to award a winner even in the event of the requested number of barrels not being satisfied by the offer was motivated by the random assignment of agents to auctions. If the auction did not award a winner if the offered amount could not need demand, the result could be a decrease in price if no auctions were won despite the fact that producer agents had accepted offered prices. While this decision did make sense in the context of ensuring that declined auctions were motivated by only price, it did have an impact on the emergence of price stability as expected, which we discuss in section 5.

2.1.2 Simulation Turns

The Trading Floor provides three operations that are invoked by the simulation during each tick, representing one turn of trading. At the start of a tick the Trading Floor *startTrade()* method is invoked. This method sets the state for the current trade, resetting any internal values that are single turn based, and initializes the auctions for producer and speculator agents. Once trades are ready to start, the simulation will call the *trade()* method of the Trading Floor. This method will iterate over all of the consumer and speculator agents and randomly assign each participating agent to one of the auctions created in the first step. Consumer agents will participate in auctions for each turn $T_i \mid T_i <= CL_i$ and $C_i' > 0$. Speculator agents utilize more complex logic, as discussed below, to determine whether or not to participate in an auction. Participating agents are randomly assigned a single auction in order to prevent the need to complicate the agent design by introducing a function to place a value on an auction; the impact of this decision is discussed in section 5. Once all participating agents have submitted a bid in the assigned auction the *trade()* method returns. At the end of a turn the simulation invokes the *settleTrades()* method to process the auctions and settles any transactions.

To complete a given trading term, the Trading Floor must resolve each auction and transfer capital and barrels of oil between agents. It is in this action that the Tradinge Floor is able to compute the information used by agents in making decisions concerning auction participation and valuation. When the simulation calls *settleTrades()* the Trading Floor will iterate over each auction and process the bids. Accepted bids are sorted by the score assigned them by the producer holding the auction with the highest score being awarded by the Trading Floor. The Trading Floor tracks the total amount of capital that is transferred from winning agents, bid_{total} to producers, the number of barrels transferred to winning agents, $volume_i$, and the number of auctions won, win_{total}. These values are then used to determine how to adjust the current trading price. After all auctions are settled, the Trading Floor computes the new price per barrel and the price trend. The

[3] Because the focus of this paper is the evaluation of location decisions leading to global coherence, we do not address the use of other auction methods for trading barrels. For our work, the auctions were a mechanism to achieve the goal of moving barrels of oil as opposed to the driver for emergent behavior. As such, only a single auction type was used for our system.

price per barrel is computed in three possible ways for the three outcomes of the auction: 1) auctions were held and bids were accepted; 2) auctions were held and no bids were accepted; 3) auctions were held and no offers were made. Each of these different outcomes will have a different impact on price. Clearly, for outcomes 2 and 3, the price of oil will fall. However, the behavior that causes the outcome will have a different effect. In the case of 2, no accepted bids means that the current trading price is higher than what the market can sustain but there is market demand for oil and the price will fall relative to the demand. Outcome 3 indicates that no demand for oil exists during the current turn, and the price will fall by a predefined amount. The following calculations are performed to determine current price, label 1 – 3 for each outcome respectively.

Commented [LS12]: Not exactly true … it depends on the timing … brinksmanship.

$$price_{current} = price_{current} + (((bid_{total} / win_{total}) - price_{current}) * (volume_i / (P - volume_{total}))) \quad (1)$$

$$price_{current} = price_{current} + (((bid_{total} / win_{total}) - price_{current}) * NS) \quad (2)$$

$$price_{current} = price_{current} * (1 - NS) \quad (3)$$ [4]

Commented [LS13]: Pretty good thinking. But, this is actually negotiation. And in MAS, negotiation of this kind usually comes with a time factor – sometimes called "conceding factor" that is a function of time to "must sell" or "must buy" deadline.

Once the new price per barrel is computed, the Trading Floor will then update the price trend, $price_{trend}$. The price trend is computed as a weighted average of all prices in a predefined window as follows where i is the current turn:

$$price_{trend} = price_{current} + ((\sum price_n * n) / wind) \ \ \forall \ n \mid wind\text{-}n <= n <= i \quad (4)$$

$$price_{trend} = price_{current} - ((\sum price_n * n) / wind) \ \ \forall \ n \mid wind\text{-}n <= n <= i \quad (5)$$

Commented [LS14]: Hmm .. why this design, usually, one would simply just add it …. Why this difference? This could cause the price trend to fluctuate quite a bit … a bit too sensitive.

If the price has increased, the trend is computed using 4, if it is less than the current price, it is computed using 5. After computing the price trend, the Trade Floor will log the results of the trading turn and is then ready for the next round. The simulation will update its data collectors with the results of the trading turn before calling *startTrade()*.

2.2 Producer Agents

Commented [LS15]: Quite well thoughtout … but should time should have been considered in the decision making.

Producer agents are responsible for selling produced barrels of oil. They start the simulation with P_i barrels of oil and $prod_i$ production cost that are randomly assigned by the simulation as defined above. The goal of the producer agent is to sell all of its barrels of oil for a price greater than $P_i * prod_i$; in other words, to sell of its barrels for a profit. The price per barrel that the agent must earn in order to not take a loss, PM, is $P_i * 1 + margin$, where $margin$ is a predefined desired profit margin. The amount of profit that an agent will be able to earn is discounted by a storage cost, C_{store}, that is assessed per turn for the number of barrels of oil P_i' and accumulated as P_{store}. This motivates the agent over time to possibly accept a lower price in auctions in order to prevent a loss for the produced oil. The producer agent performs two tasks: it creates auctions offering some number of barrels as discussed in 2.1.1, and it can compute the reward associated with a bid price. The latter function provides that agent with the ability to score a bid for an auction.

Rewards for a given bid are calculated using a modified Bellman equation [Bellman 1957]. Simply described, the Bellman equation computes the value of an action by taking the reward associated with

[4] At the time of writing, we determined that the source code for our model contained an incorrect computation for this value. Rather than $price_{current} = price_{current} * (1 - NS)$, we were computing it as $price_{current} = (1 - NS)$. The result of this error was that our experimental data contained periods where the price would fall to near zero and then jump back up. This corresponded to the price falling to NS under no offers being made, and then returning to the previous price when offers started being made. We initially filtered these periods out of our results, which has the same effect as setting NS to 0. Based on the discovery, we re-ran our experiments with corrected code to generate output inline with our intended calculations. We note this here as our reported results may not reflect those obtained by the TA as a result of our correcting this fault.

taking the action and adding it to the value of the state transition that will occur as a result of the selection. The transition value is modified by a discount factor to provide a way to control the value an agent will put on a future occurrence. We use this basic method to compute the value of a bid by comparing the value of accepting the offer compared to the value of accepting an estimated offer for the same number of barrels on the next turn. If the value of accepting the offer exceeds the value of waiting, the calculated value is assigned as the bid score. The value of the offer is easily calculated as:

$$F(bid_i.price) = bid_i.price * bid_i.barrels \tag{6}$$

Less straightforward is determining the value of not accepting the bid. Doing so requires that the agent estimate with the trading price will be in the next round and then assume that will get at least that much as an offer. This estimated offer is discounted by the cost of storing the barrels for another turn and then is further discounted by the global discount factor δ.

$$V(i) = ((price_{current} + ((price_{current} + price_{trend}) - price_{current})) - \mathrm{C}_{store}) * bid_i.barrels \tag{7}$$

Commented [LS16]: Good.

More specifically, the value of waiting is equal to the current price plus the difference between the estimated next price and the current price less the cost of storing the barrels. Multiplied by the number of requested barrels, the value assumes that at least that many barrels will be purchased at this price to derive the value of waiting. One complicating factor in this computation is estimating the future price based on the trend. This method provides a very rudimentary[5] way to estimate the next price, however the actions of the agent will suffer when the trend cannot be accurately determined. We address this by discounting the estimated price based on the amount of data used in the trend computation. That is to say, we recognize that early on in the simulation, the moving weighted average of the prices will be less accurate due to the lack of available data to compute it. Therefore, for smaller windows of data, the computed trend is discounted by an amount relative to the number of completed turns in a trend window. The agent will, for early values until the number of turns reaches the trend window size, discount their estimate. The following modified formula shows the value of waiting including a discount of the future price based on the number of turns relative to *wind*.

Commented [LS17]: I can help you guys find some better pricing strategies, especially from papers in negotiation, if you are set to publish this.

Commented [LS18]: Okay!

$$V(i) = ((price_{current} + (((price_{current} + price_{trend}) - price_{current})) * 1+(T'/wind)) - \mathrm{C}_{store}) * bid_i.barrels \tag{8}$$

Commented [LS19]: What is this?

The value of waiting is further discounted by the overall cost to the producer for holding barrels of oil over the entire simulation. Because the producer will have a decreased profit margin for each turn in which they do not sell oil, due to the storage cost, the producer is less motivated to wait as the total storage cost increases. We discount the value of waiting accordingly.

$$V(i) = V(i) * 1 - (P_{store} / PM * P_i) \tag{9}$$

Once a producer agent has sold P_i barrels, it will no longer participate in the simulation.

2.3 Consumer Agents

Commented [LS20]: Okay. But lack of symmetry. Consumer agents should also think about future pricing too. This is because otherwise, you will have no agents that consume AND still think about future pricing when bidding. Speculator agents do not really consume.

Consumer agents are responsible for purchasing barrels of oil. They are created with some demand for oil measured in number of barrels and defined as D_i, some amount of capital with which to purchase barrels, C_i, and a contract length measured in number of turns by which all demand must be satisfied and defined as CL_i. The goal of the consumer is to purchase enough barrels to meet demand within CL_i turns at a total amount <= C_i. The consumer participates randomly in producer held auctions bidding on D_i' barrels of oil[6]. The singular operation of a consumer, then, is to place a value on purchasing D_i' barrels of oil and

[5] Because the scope of this work was to study the use of global coherence in a MAS to study economic models, we did not explore a more complex pricing methodology. In looking to published research, we would very likely find a more accurate pricing model.
[6] A consumer always will bid on oil to satisfy remaining demand. Using logic similar to that used to justify the transfer of as many barrels as is possible to meet a bid for the auction, we elected to utilize this simple approach to prevent the need to introduce additional decision logic beyond a

offering as a bid price that value. Unlike the producer, the consumer does not consider the future price of oil when determining an amount to bid. The future potential loss for a consumer is based on the number of remaining turns before $T' = CLi$. Once an agent has reached the contract length, it is no longer able to participate in auctions and can no longer fulfill demand for oil. The consumer is motivated only on fulfilling demand, there is no storage cost associated with holding purchased barrels, so the amount of capital an agent has to make purchases is fixed for the entire simulation; less the amount paid of barrels, of course. Like the producer, however, the consumer has the concept of a maximum price that can be paid without taking a loss. For the consumer, loss would be paying more than C_i for all barrels of oil; this is not possible in our model. For the consumer then, $PM = D_i / C_i$. The bid price for turn T' is computed by the consumer based on how close to the contract length it is as well as the current trading price. If $price_{current} > PM$, then the agent must offer PM as the price as that is the highest amount it can pay for a barrel of oil and still meet its demand[7]. In this case, the bid price is:

Commented [LS21]: This puts them at a disadvantage. Not quite logical either. Why not? Why can't the consumer consider the future price of oil?

$$bid_i.price = PM \qquad (10)$$

If the agent is not constrained by available capital relative to the current trading price, it will place a bid based on how close to it is to the end of its contract term, approaching the maximum value that can be offered of PM. The following time weighted formula is used to determine the bid price:

$$bid_i.price = (rand(0.5 \ldots 1) + T' / CL_i) * price_{current} \qquad (11)$$

In this case, the bid price is no less than ½ the current trading price and no more than the percent of time in the current contract times the current trading price. The idea being that the agent will be motivated to offer a higher price to fulfill demand the closer it gets to the end of its contract. Once the agent gets within $price_{current} / PM$ turns of the end of its contract, it will offer a percent of its no loss price trending to the absolute maximum offer as it approaches the end of its term.

$$bid_i.price = (T' / CL_i) * PM \qquad (12)$$

In the last turn of the contract term, the agent will offer PM, which if equation 10 is not used, will be higher than the current trading price. Once a consumer purchases enough barrels to satisfy demand D_i the agent will no longer participate in auctions.

2.4 Speculator Agents

Speculator agents represent a combination of producer and consumer agents, having the ability to both offer barrels for sale at auction as well as purchase barrels. Like consumers, producers start out the simulation with which to purchase oil, S_i. Unlike consumers, however, the speculator agents have no demand that must be fulfilled. They are motivated solely by the desire to purchase oil at the lowest price possible and then sell it for the highest price possible. When acting as a consumer, speculator agents will bid on randomly assigned auctions using a modified value function similar to the consumer. If the speculator is acting as a seller of oil, it will create an auction in a manner similar to the producer. The primary difference in the value functions of speculators vs. producers or consumers is that the value takes into account what the speculator originally purchased the barrels for. While similar to the concept of a production cost for producers, the speculator is able to control this value based on when the decided to purchase as opposed to sell. Like producers, the speculator does incur a cost, C_{store}, for each turn it holds barrels. When acting as a purchaser of oil, the speculator calculates its bid based on an offset of the current price relative to trading, or the current price relative to the number of remaining turns. The bid price is

Commented [LS22]: Speculators?

basic value function for oil. To partition bids would require that the agent be able to decide how many barrels to bid for based on parameters. The greedy approach of bidding for all barrels needed allowed us to study the system in terms of pricing alone.

[7] Because the producer does not contain decision logic that allows it to partition its oil demand, it must greedily seek of fulfill all demand, as a result it cannot offer more than PM per barrel.

influenced by the maximum that an agent can pay. This maximum, *PM*, is computed based on the amount of remaining capital, S_i', and the number of barrels currently held, *PB*. Capital can increase throughout the simulation as the result of selling oil for a gain.

$$offset = ((PM - price_{current}) * (volume_{total}/P) * (1- \delta) \quad (13)$$

The offset in 13 is the maximum price that the agent can offer without taking a loss, less the current trading price multiplied by the percentage of barrels traded and then discounted by the global discount factor[8]. The computation will create some discount window based on how many turns are left in the simulation, with the result – as seen below – being the agent's unwillingness to pay a large amount for barrels when most of the barrels have already been traded and the likelihood that they won't be sold is greater as the price is rising with the diminishing supply. The other factor considered his how far along the simulation is.

$$offset = ((PM - price_{current}) * (T'/T) * (1- \delta) \quad (14)$$

In this case, the agent considers how many turns remain in which oil can be sold. The discount applied to the price will be greater as the simulation progresses as it becomes more likely that the speculator will be left with unsold barrels as the end of the simulation, thereby negatively impacting its earned profit. Depending on which offset is greater, the bid price is calculated as:

$$bid_i.price = price_{current} + offset \quad (15)$$

The speculator will essentially offer a price that is slightly higher than the current trading price and discounted based on the risk of winning an auction at *T'*. This may seem counter intuitive given that the speculator is motivated by maximizing profit. The reasoning for making offers slightly above the current trading price is that, if the agent decides to participate in the auctions, it has determined that there is value in purchasing at that turn. It will, therefore, attempt to beat out other agents to buy the oil so that it can sell it to the same agents later for a higher price.

Commented [LS23]: Hmmm …

This function of the speculator leads to a significant difference from either the consumers or the producers. The speculator needs to be able to make the decision on how to act at the outset of each turn. Producers and consumers will always act to buy or sell oil so long as there is either supply or demand. Speculator agents on the other hand must decide at the beginning of the turn whether or not they will act as buyers or sellers. The decision as to how the speculator will act is based on the time remaining in the simulation and the current trading price. Generally speaking, speculators will buy when the price is close to the simulation low price, and sell when the price is close to the simulation high price. The closer to the end of the simulation, the more likely the speculator is to sell regardless of current price. At any point, if the cost of storing the currently held barrels exceeds the agent's capital, the agent will always sell. If the speculator has zero barrels to sell, it will always buy.

At the outset of the buy/sell decision, the speculator agent computes a price at which it would be willing to pay based on the weighted average of a randomized value no greater than 18% of the high and low prices. The so-called stable factor is calculated as:

$$BP = (((rand(0 \ldots 0.18) * price_{max}) - (rand(0 \ldots 0.18) * price_{min})) / (2+ price_{min})) \quad (16)$$

$$SF = price_{current} / BP \quad (17)$$

The stable factor is then used to compute a value for each action buy or sell as follows:

[8] It should be noted here that for the purposes of computing volume, the trading floor will ignore transfers between speculators. This is done because barrels transferred between speculators do not decrease the supply as they can eventually be sold to a consumer to meet some demand.

$TF = ((T'/T) * rand(0 \dots 0.32))$ (18)

$buy = (1 – TF) *((S_i' - (C_{store} * PB)) / SF) * BP$ (19)

The value associating with purchasing barrels weighs the amount of time that the speculator will have to hold the barrels, the trading factor or *TF*, against the likelihood of realizing a profit based on the proximity to the low trading price. The value associated with selling is similarly computed.

$sell = (TF * PB * BP)$ (20)

Commented [LS24]: I like the idea of stale factor, but I don't think your environment is setup such that the stale factor will make an impact … actually, if anything, I think this would cause the speculators to not act like what you wanted them to act.

Selling is valued based on the number barrels purchased multiplied by the relative proximity to the simulation high price discounted by the number of turns that remain for the barrels to possibly be sold. The greater of these values, buy or sell, is the decision that the agent will make for its action for a given turn. Values are randomized when making a decision to act in order to prevent each speculator from acting in the same way at each turn. Speculators value and score bids in the same way that producer agents do with production cost replaced by purchase price for all held barrels.

2.5 REPAST

The Trade Floor class is the controlling entity governing the turns in a simulation, but it does not provide any method for user input of parameters, display of data results during the simulation, data logging for analysis post simulation, and batch processing for multiple runs. Rather than implement a controller class for actual simulation ticks, we implemented a simple simulation, using REPAST [Collier 2011] that invokes the appropriate methods on the Trade Floor to execute the model. REPAST provides a mechanism to abstract the use cases for a generic simulation from our model, allowing us to focus on the design of the model without introducing logic to generate experimental data; which, for us, was as a separate activity. The result is a more generalized market simulation model that can be tailored for different kinds of execution and analysis. The results presented here, however, were generated exclusively using REPAST graphical display and data logging functionality. During development, we made use of the GUI for executing aspects of the model being developed. However, as discussed in section 3, our experimental runs did not utilize GUI functionality and relied on REPAST data loggers and batch simulation execution.

3. EXPERIMENTAL SETUP

Commented [LS25]: Clear and well-supported. Good to read. Purposeful.

In order to evaluate the ability of our model to achieve a level of global coherence that represents some sufficiently complex enough economic factor, we designed a set of three experiments designed to test hypotheses H1 – H4 as outlined in section 1. We initially hypothesized that if our model would be able to produce the desired emergent behavior that certain properties would hold relative to expected price stability. These properties can generally be tested by: 1) increasing available capital; 2) decreasing production; and 3) increasing the number of speculators in the system. Experiment 1 is designed to test H2 and H3, experiment 2 is designed to test H1, and experiment 3 is designed to test H4. For each of the experiments we used the input parameters detailed in Table 1.

Commented [LS26]: Good.

Experiment Number / Parameter Number												
	1	2	3	4	5	6	7	8	9	10	11	12
1 & 2	150	150,150,0	-	150,50	5,000,000	75, -	30,70	5.5	0.2	0.03	100	0
3	500	150,150,-	5,000,000	150,50	5,000,000	75,300	30,500	5.5	0.2	0.10	100	0

Table 1. Model parameter values for experiment sets 1, 2, and 3.

Input parameter values of '-' indicate variability across simulations in the experiment. No formal tests were run to learn optimal parameter values; the parameter-space is too large to adequately provide reasonable evaluation and justification of selection. Rather, we used 'reasonable' values based on our intuition of the system. The model was run several times manually with the parameters in order to refine assumptions and optimize the fixed parameter space. Between experiments 2 and 3 several of the fixed parameters were adjusted to allow sufficient time for the expected behavior of the speculators to emerge.

3.1 Increased Capital

To control the amount of capital in the system, we designed experiment 1 to increment the maximum allowable purchase price, Pay_{max} from \$100.00 by \$25.00 per simulation to \$300.00. Each simulation was run 100 times to generalize the impact across the random behavior of agents. Due to the way in which capital is divided among agents at the beginning of the simulation, only Pay_{max} needed to be adjusted as that increased the ceiling of available capital translating into an increase of capital available to consumers. We hypothesize in H2 that trading price is proportional to the amount of capital available to the consumers to purchase oil. As the amount of capital increases, the consumers' ability to bid high in auctions will also increase which will result in higher average bids and higher trading prices. In H3 we hypothesize that as the price converges on the maximum price that agents can pay, the price will suddenly fall as a result of consumers offering less than the trading price and producers being motivated to sell remaining supply. By observing the price as it increases, we can observe this behavior once the maximum price has been reached.

3.2 Decreased Production

To model varying amounts of oil production parameter 3 was incremented by 2 million, starting at 1 million and ending at 9 million barrels. The simulation was run at each increment value of *P* 100 times. *Each* increment was simulated 100 times. The goal of this experiment was to determine what impact a change in production had on the price of oil. We hypothesized in H1 that the trading price of oil would be inversely proportional to the produced barrels; insufficient supply would lead to higher demand, thereby driving the price up. By running the simulation with a range of values greater and less than demand, we evaluate the role that demand has on the price and can describe its relationship to production values. Each increment was run 100 times so as to generalize the impact across the random behavior of agents.

3.3 Increased Speculators

In the third experiment we test the impact that speculators have on this system. For this experiment we fixed the production capacity as well as the available system capital and varied the number of speculator agents. We also adjusted three additional parameters based on manual runs to optimize fixed parameters. Parameter 1 was increased for experiment 3 from 150 to 500. We did this to allow the speculators enough time to observe high and low prices and thereby calibrate their stable factor for bid determination. Because speculators tend to buy on downward price trends, we adjusted the value of parameter 10 to increase the speed at which the price would drop when no auctions were fulfilled. While this could have the unintended consequence of preventing speculators from buying as a result of consumers being able to purchase oil at a discount, the premium that a speculator is willing to pay above the current trading price offset the concern. Due to the increased number of turns for this experiment, we also increased the contract length to allow the consumers to remain active in the simulation in the face of speculator trades. If too many consumers dropped out of the simulation as a result of failing to meet demand in their needed window, we would not be able to adequately evaluate the impact that speculators have on the price as the model would digress to speculator trades only; we increased the maximum value for parameter 7 to equal to total number of allowed trades to open the window as wide as we could for consumers to remain in the simulation trading against speculators.

In addition to the three adjusted fixed parameters, we considered adjusting parameter 9 in order to allow the speculators access to a greater amount of oil. We decided against changing this value, however, out of concern that if producers were able to move a large amount of supply early in the simulation it would be purchased by consumers while speculators waited for a low price. This would counteract the increase in speculators by preventing the larger number of speculators from being able to purchase oil once a low value had been reached due to supply having been exhausted by consumers prior to the occurrence of the low price. Fewer purchasing speculators would prevent a clear evaluation of their role on price and so the value used for 9 was left unchanged.

For this experiment we incremented the speculator count in parameter 3 by 50 starting at zero and ending at 300. As was the case in the previous experiments, we ran this experiment 100 times in order to reduce the noise associated with some of the randomness that was built into the agents. In H4 we hypnotized that as the number of speculators increased so would the fluctuations in price, leading to price instability. As speculators agents buy up oil they drive up price, but once the price goes too high, agents will stop buying oil causing the price to fall. The speculators will gradually sell off their purchases as the price rises to consumers and other agents. As the price falls, speculators will again begin to buy oil causing the cycle to repeat. By running the simulation and increasing the number of speculators, we are able to observe price shifts relative to the number of speculators. Our ~~institution~~ intuition is that with an increasing number of speculators, the rate and degree to which the price fluctuates will increase.

4. RESULTS

The observed results from analysis of the data generated in our experiments were mixed. While, generally speaking, a price behavior emerged, it did not follow our expectations. Our experiments were focused on the observation of price, however in addressing the reasons for our failure to observe the hypothesized price trend, we observed several trends in variables other than price having had an impact on the emergence of our hypothesized behavior. The following is a discussions of the observations made from our experimental data.

4.1 Price vs. Available Capital

Commented [LS27]: Yes. You nailed this one! 8)

In our first experiment we attempted to demonstrate that price would emerge as proportional to the availability of capital with which agents could purchase oil. Agents with more money, we hypothesized, would be willing to spend more and drive up the price. Figure 1 displays the price averages over time for all of the runs in experiment 1. As is clearly indicated, the price of oil does steadily rise over the first half of the simulation which we attribute to the consumer agent's ability to freely purchase barrels at a rate unconstrained by production cost; that is to say, the producer agents gladly accept the bids of agents who are bidding beyond what it cost to produce the barrels, allowing the producers to realize a profit. Providing a measure of validation of H2, we observe that the rise in price tends follow the same general pattern regardless of how much capital is available. For the first quarter the price rises exponentially, reaching a maximum and then slowly falling in the second quarter of the simulation, before falling drastically in the second half. The more capital available to the agents in our experiments tended to push the price a bit higher and the occurrence of a maximum value out a bit farther. We do not see significant variance in the maximum and minimum prices, and we note that the increments in available initial capital were both small and subject to random distribution across agents with random distribution of demand; which is not here considered.

We also observe a single outlier where our hypothesis does not hold. The initial maximum value of $100.00 shows a steady decline throughout the entire simulation, dropping sharply at the same point when prices begin to fall for the other values. We attribute this to the relationship between the maximum purchase price and the maximum production cost per barrel. Recall from section 3 that the maximum production cost, as defined by parameter 4, is fixed at $150. The results seen below in Figure 1 suggest that there is a specific threshold of variation between the most an agent can pay for a barrel and the most that it cost an agent to produce the barrel that must be met in order for the price to rise. We observe that if the most a consumer can pay for a barrel falls far enough below the most a producer paid to produce it, that the price will steadily fall. This seems intuitive to us given that, for a large gap between these values, a significant number of producers will have paid more than the most a consumer can offer to purchase the barrel and will, as a result, not accept bids that would lead to a loss. The lack of trading volume would then drive prices down and consumer agents would respond by demonstrating less willingness to pay. This 'death spiral' of price, where the price can reach a state where it continually falls, is the result of the influence that trading price has on consumer bid offers combined with the fact that, later in the simulation, producers are willing to take a potential loss.

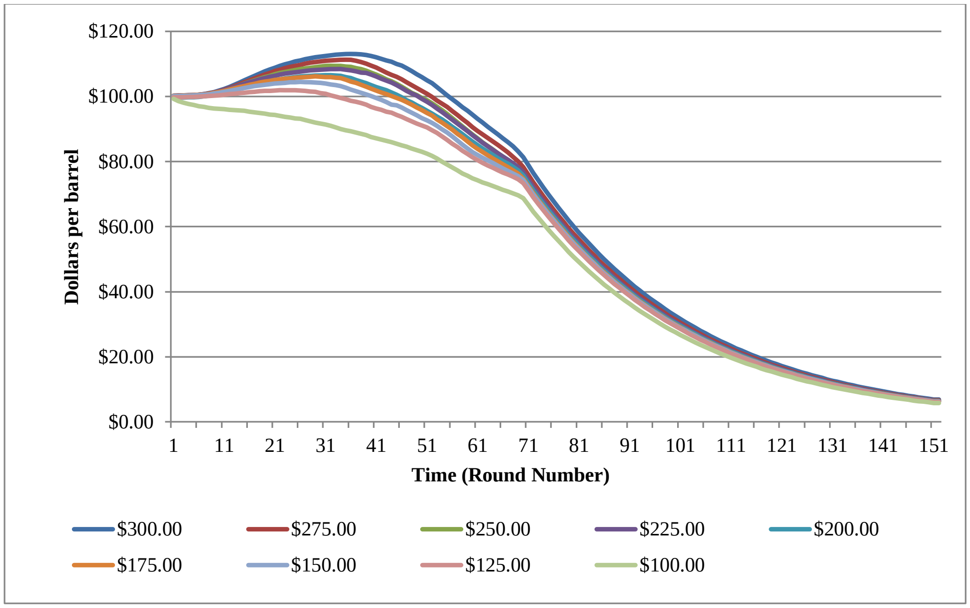

Figure 1. Price per barrel over time by total system capital

The relationship between price and available capital can be seen more clearly in Figure 2 below. This figure displays the average bids and average sale price per barrel over all of the runs of each experiment. As the amount of capital available to the consumer rises, so does price. An expected property based on H4 that we observe in this chart is the generally flat nature of the actual price paid. A stable price quickly emerges in the simulation and remains consistent throughout, growing slightly more slowly for maximum purchase prices less than the production price, and more rapidly as the maximum purchase price exceeds the production cost. Implicit in H4 is the fact that a general price range will emerge based on agents acting in a self-interested fashion. From the data in experiment 1, that seems to hold despite not having had designed the experiment to test this hypothesis. We note this and will return to it in our analysis of experiment 3. The average bid offer also fits within our intuition in H2. When consumers have less capital available than the production cost – which heavily influences initial trading price – the average bids fall significantly below the trading price; as a result of consumers offering the most they can which is still below trade. The average bids steadily increase as the available capital meets and then exceeds the production costs. We would have expected convergence between these two lines at a lower value for the maximum purchase price, however we attribute this to agents generally offering less than the trading price early in the simulation regardless of how much capital is available to them.

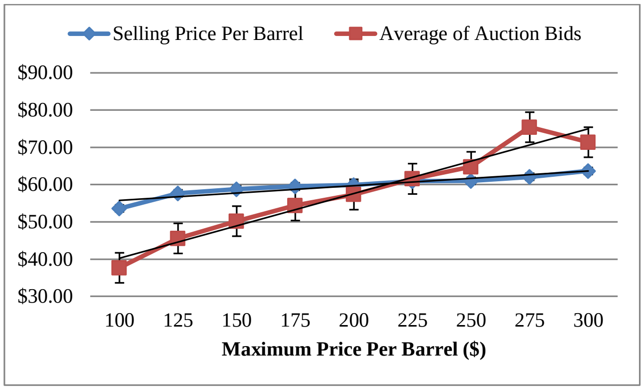

Figure 2. Comparison of selling prices and bid averages against max agent capital

Experiment 1 was designed to test both H2 and H3, and we observe that H3 appears to hold in this data as well. H3 suggests that there is a maximum price that, once reached, will not go higher and will begin to bottom out. Essentially, once the price reaches a high, consumers will no longer be able to purchase at that price and the producers will be forced to sell at lower prices in a motivation to sell off remaining barrels before the end of the simulation. Observing the data in Figure 1, this appears to be the case, with prices rising steadily and then falling steadily. We leave the observations of the second part of H3 to 4.2. Not anticipated, however, was the increase in the drop off rate observed in the second half of the simulation. In considering this, we believe it is consistent with the expected behavior of the system given random auction assignment. We validate this observation in 4.4 below as the result of auction offerings having very few barrels which leads to low trading volume, where a consumer is 'lucky' enough to be assigned to one of the few held auctions, and a price decrease when no assignment is made.

4.2 Price vs. Supply

Commented [LS28]: This is not as strong and well constructed as the 4.1. I suspect some bug in the code or flaw in the design. Your Figure 5 is puzzling and your POJI is not that convincing.

In our first hypothesis, we posited that the trading price of oil would be inversely proportional to the supply (number of produced barrels). Data gathered in experiment 2 allowed us to observe this behavior in our model.

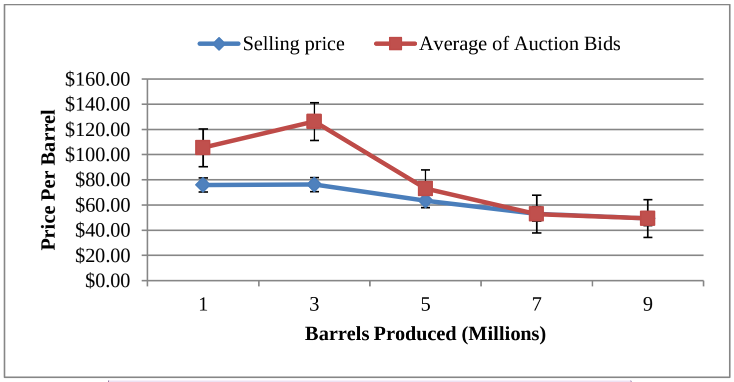

Figure 3. Comparison of selling prices and bid averages against produced barrels

Commented [LS29]: Another parameter to plot with this graph is the number of bids and the average number of agents participating in each auction. Could reveal some insights as to why your curves behaved this way.

The data in Figure 3 displays the average price at which a barrel is sold and the average bid amount per number of barrels produced in millions. Recall from section 3 that for this experiment demand was fixed at 5,000,000 barrels. When supply is less than demand, the price per barrel is higher than when supply begins to equal and then exceed demand. We suggested that price would be completely proportional in H1, however that does not hold in our observation of data. The rate at which the price falls relative to the increase in supply after convergence does not match the rate before. The price fall rate levels off after supply meets demand. In giving this further consideration, that stands to reason given that the availability of excess oil in the simulation will not be demanded as the consumers fall off and the producers will have no one to sell it to. We did expect that when looking at the trading price - Figure 3 looks at average bid awards and therefore does not consider price drops due to no offers being made or accepted – we would see a price drop rate mirrored.

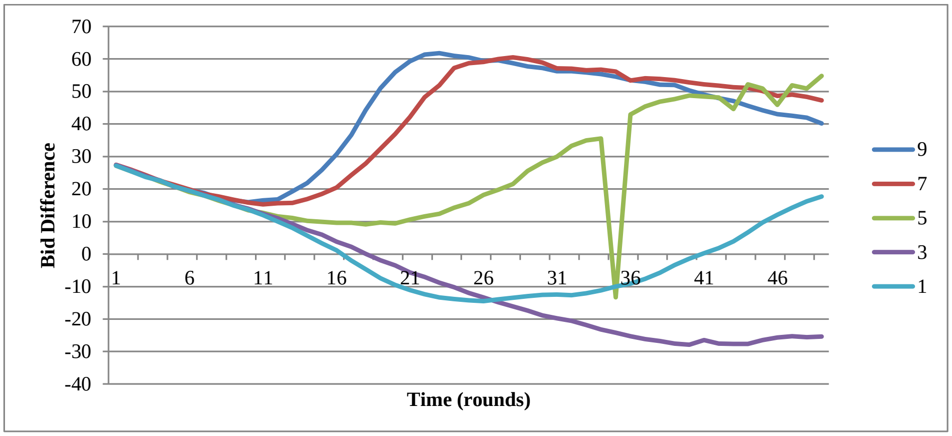

Figure 4. Average winning bid less average bid per round for each iteration of millions of barrels produced

Commented [LS30]: With the negative bid difference, that means within an auction, it is possible for the winning bid to be lower than the average bid value that the auction receives? Hmm …. Something is not quite right ?

Figure 4 depicts the shift in price over the 100 simulation runs as the difference between high and low bids for each turn based on number of barrels produced. This data again generally confirms H1. For low supply relative to demand we see the price falling steadily, and for high supply relative to demand we see it increase steadily. For larger supply, we see the peak occurring sooner in the simulation. We note that when the data are viewed more generally in terms of price we do observe the approximate mirroring of the rate change as discussed previously. We also again observe a convergence just after supply meets demand. Figure 4 clearly displays the behavior that we suggested above with the price falling for equality of supply and demand but then rising steadily. The steep decline and increase observed around tick 36 suggests a data anomaly that we believe is the result of a very small number of consumers participating in a few number of auctions. Based on the random assignment of consumers to auctions, it is not impossible that a number of turns could take place where consumers with little remaining capital could be assigned to auctions that have a high value, thus preventing a winner from being assigned. This is the only way we could account for the observed price jumps that do not appear to consistently fit with the other trends. Furthermore, we again observe that the price (both raw price and high/differential) will peak and then steadily decline, confirming the second component of H3 discussed in 4.1. Despite generally confirming our hypothesis, Figure 4 does depict a curious behavior. For 1,000,000 barrels of produced oil the price will decrease steadily, as expected, but suddenly begins to rise. We attribute this to the phenomenon of consumers successfully starting to by the limited number of produced barrels once the producing agents begin to release supply as a result of ongoing storage costs. It would be expected that there is a point at which supply is low enough that such behavior would manifest itself. Intuitively, we would expect the

Commented [LS31]: Good investigation.

Commented [LS32]: Yes. I am still suspicious …

curve for 1,000,000 to resemble the curve for 3,000,000. However, if we consider that the trade volume influences the willingness of an agent to buy or sell, we could expect to see similar behavior in the 3,000,000 curve at longer time intervals. Indeed, at the end of the simulation we do note a very slight price rate increase.

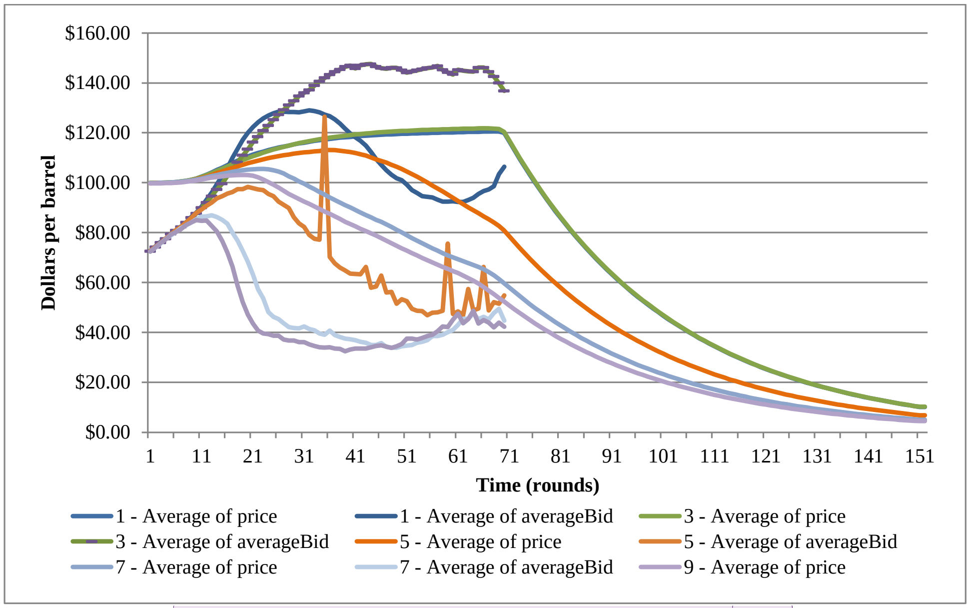

Figure 5. Bid price and average bids per round for each iteration of millions of barrels produced

Commented [LS33]: Why no 9-Average of averageBid?

Commented [LS34]: Each pair of curves …. The average of price vs. average of averageBid …Quite random except for the 7s and 9s. Why? Not sure.

Finally, Figure 5 depicts the average price and average of the bid average across all our simulations. We again observe~~r~~ a general pattern of falling and rising prices with this depiction of the data showing more noise than the previous figures. In this context, we can see the variance in price that takes place near the convergence point. We attribute this randomness to the equal distribution of supply and demand, leading to consumers experiencing greater numbers of inconsistent producer responses to bids as they are randomly assigned an auction. At convergence we do not observe the 'race' for remaining supply leading to price increases that is observed when demand exceeds supply. The oscillations between prices at convergence is the result of a sort of equilibrium offered barrels vs. demanded barrels, preventing the volume checks in the value functions from being able to 'properly' motivate the agents. Price fluctuations, in this case, would tend to become more random with the distribution of auctions.

4.3 Price vs. Speculator Count

Commented [LS35]: This is better than 4.2. More directed, more convincing.

Thus far, we have generally observed the hypothesized emergence of price trends based on input values with the minor anomalies we have discussed. Despite these anomalies, which we attribute to the randomness introduced in our model and our general intuition, our model has behaved as expected. The 3rd experiment considers our final hypothesis (H4), that the introduction of more speculators will lead to greater price fluctuations. It is for this hypothesis that our data fails to conform to our expectations. Figure 6 depicts the price as a function of the bids; we have used bid price as an indicator of price in an attempt to reduce the noise by evaluating the actions of all agents participating in the auctions as opposed to only those agents who win and therefore impact the trading price. Rather than introducing significant price fluctuations, the introduction of more speculator agents seems to have the inverse effect of maintain a fairly flat price.

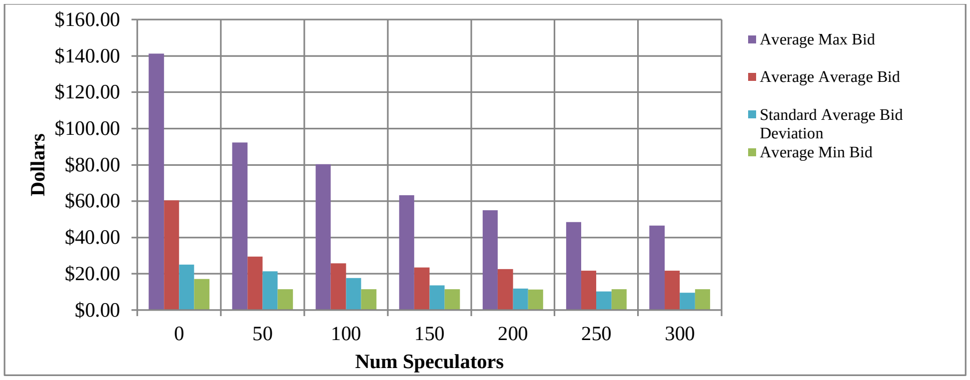

Figure 6. Simulation average bid behavior for number of speculators

The simulations that were run with no speculators have the highest standard deviation between the minimum and maximum bids, suggesting that for bids, at least, the activity tended to swing at a greater rate on a turn-by-turn basis. Initially we could not account for this as we would expect to see a higher standard deviation in bid averages when more speculators were included in the system as we hypothesized that they would drive price to the max, let it fall, and then drive it up again. This convergence of minimum and maximum bids when more speculators were present did not make sense. To provide some insight into this unexpected behavior we next considered the average selling price against the average bids. This consideration was made with the hopes of removing some of the noise that could stem from extreme agents (desperate, poor, etc) in order to see if we could observe greater fluctuations. Rather than shedding light on the issue, Figure 7 supports the observations in Figure 6. Regardless of the number of speculators that are included, average selling price remains very flat, and average bid price, once speculators are introduced, also flattens. It would seem that adding speculators actually add stability to the model with the slight effect of driving price down as more speculators are added!

Commented [LS36]: I think it is the stale factor at work.

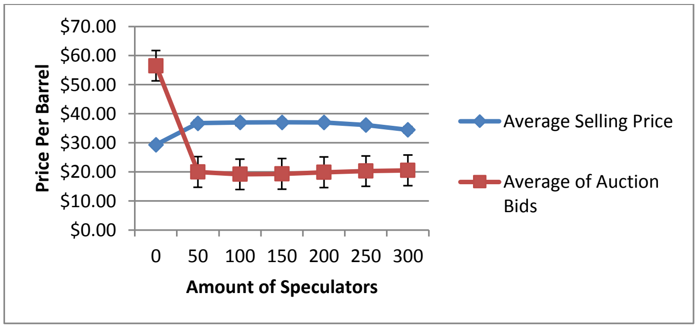

Figure 7. Average selling price against average bid amount for number of speculators

Still at a loss to explain the observed behavior, we looked at the impact that speculators had on actual trading price over time and shown in Figure 8. Price over time based on number of speculators provided us with the data to explain why our expectations of H4 were not observed in our experiments.

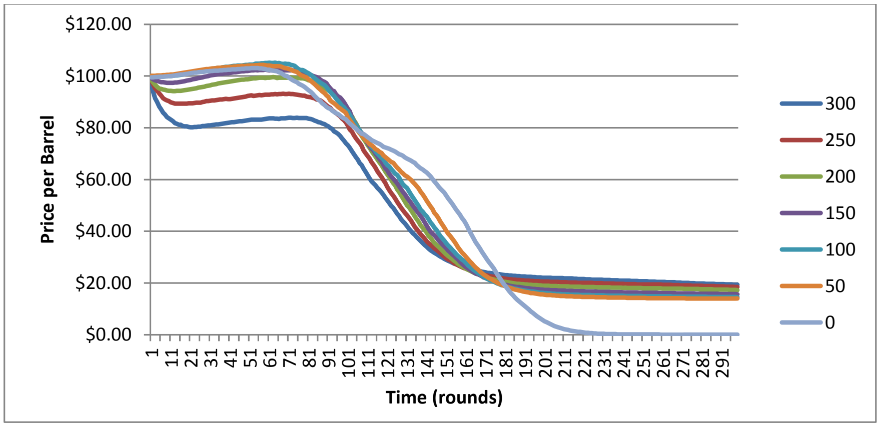

Figure 8. Price per barrel over time for number of speculators

Regardless of how many speculators are added to the system, the price follows the same general trend. It remains steady for the first quarter of the simulation, drops precipitously in the second quarter, and then remains flat in the third quarter. While the number of speculators does seem to hasten the occurrence of the price drop, this matches our observations in Figure 1. The absence of any speculators seems to draw out the eventual price decline and also prevents it from going to zero; recall that for these experiments we set the number of turns to a high value with consumers being able to stay in the trading up to the end which allows the price to fall farther than in experiments 1 and 2. This behavior gives us some insight into why speculators do not have the hypothesized effect on price. It appears from the data that our speculators are very poor speculators. At about the quarter mark it appears that they begin to buy up whatever oil has not been purchased by consumers and then proceed to hold onto it, causing the price to drop. Once the price drops to a specific point, they all appear to begin trading other speculators oil in order to dump it before the end of the simulation, but after dumping it, are willing to buy it back at the low price. Recall from section 2.4 that, unlike producers, speculators do not consider their overall expenditures on storage cost. This lack of agent 'memory' prevents a speculator who has sold barrels in the previous round from buying the back in the next. Because of this, speculator agents would seem to have a stabilizing effect on price as a result of continual trading with each other at the market low price. Since the simulation does not consider trades between speculators in its calculation of averages, the back and fourth trading has virtually no impact on price, allowing auctions to be satisfied until the end of the simulation.

Commented [LS37]: I think my hunch is probably correct – about the stale factor.

As a final validation of our assumptions made from our observations, we consider the trading volume over time as the number of speculators increases in Figure 9. From this data we can clearly see that most trades take place early in the simulation, falling rapidly as time goes on. We account for this by considering that consumers tend to greedily offer bids very near the current trading price, producers are greedily willing to accept those bids, and the initial trading price is based on the production cost. This is supported by the fact that the greatest amount of trade activity occurs earliest when no speculators are added to the simulation. As speculators are added, the early volume of trading begins to decrease. This decrease seems to be the result of early purchases that would have otherwise been made by consumers are being made by speculators and the trading floor does not consider this volume as it does not remove those barrels form supply. The decrease in early activity is directly proportional to the number of speculators added, with trading activity still leveling off in the same way. This leveling off is to be expected if we correctly assumed that at later turns in the simulation, the speculators begin trading back and fourth with each other. Figure 9 clearly depicts this expected trend. The observed initial trading activity with zero speculators also supports our observations from experiments 1 and 2 where price steadily rises early, begins to fall gradually, and then drops quickly. The significant activity early on has a dramatic impact on price, leveling off as trading slows, and then stopping virtually altogether; ceasing to trade thereby having a much more significant negative impact on price.

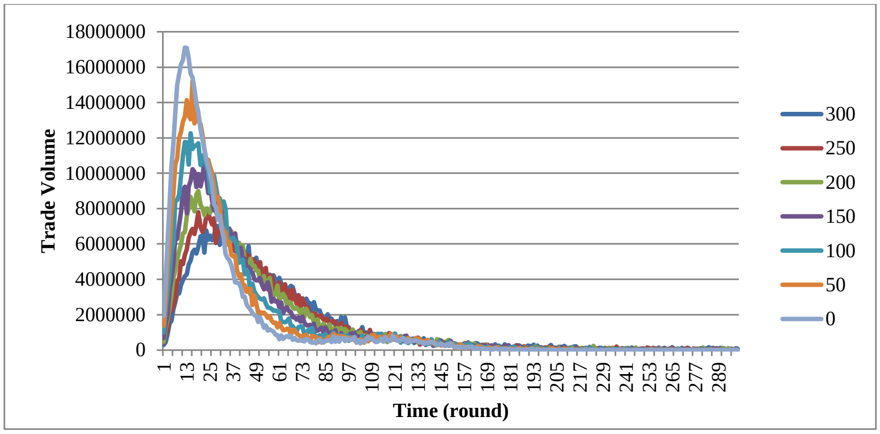

Figure 9. Barrels traded per round by number of speculators

Despite failing to validate H4, experiment 3 tends to support the observations made from experiments 1 and 2 in that we see the emergence of a very stable price in the model that conforms to the types of decisions that our agents were, for better or worse, designed to make. As such, we observe that hypotheses H1, H2, and H3 were generally confirmed in our results, while hypothesis H4 was completed rejected.

5. CONCLUSIONS AND FUTURE WORK

Our work has argued that the global coherence of a MAS as a function of local decision making can be exploited to study complex economic behaviors. We have attempted to show this by considering the local decisions that agents can make while acting in a commodities exchange and then studying the resulting emergence of price trends in the system. Despite observing mixed results in the study of our model, we can conclude that local decisions in an economic model do in fact lead to global coherence in the emergence of a complex system behavior; namely the global trending of price. Our experimental design and resulting data validated our hypotheses that price can be related in a predictable way to the behavior of agents in the system. The expected globally coherent behavior (price) is controlled by local decisions of agents (commodity valuation) that do not consider the desired global objective (price stability).

Despite holding for our implementation of a commodity market model, our observed results do pose some threat to the validity of our conclusion. It could be easily argued that our model tracks too closely to the current trading price of oil and, as a result, price stability is accidently achieved. Addressing this threat requires that additional studies be preformed. We can conclude from our observations that such a thread does indeed exist and that it is the result of flaws in our agent design. There are three primary design decisions that were made when developing our agents that create this problem. The first, there is too much randomness in the system. The random assignment of agents to auctions, valuation differentials, and consideration to future actions resulted in agents whose decisions were not consistent. Agents were more heavily influenced by the price trend than they were the factors that were designed to help them make decisions. As a result, they behaved in very similar ways, despite the introduction of random assignment, and when they did differ in their behavior, it had the potential for negative consequences on the system.

Commented [LS38]: But using the overall capital of the community – that's quite a bit of global information … making your local decisions not that "local".

Commented [LS39]: Yes.

Commented [LS40]: I think first thing would be to find out more about Figure 4 and Figure 5.

Commented [LS41]: I think this is actually good. To publish a paper like this, you would need randomness. One just would have to run a lot of runs to get statistical significance.

When all agents make the same basic decisions a "group think" mentality can occur, which has the natural consequence of the emergence of common behavior. While still globally coherent, this emergence is not necessarily desirable as may not occur if the agents were not making random decisions.

The second factor, price estimation for decision-making is poor. Even if agents were making less random decisions, the data used to make those decisions is currently based on the perception of what will happen to the price of oil on the next turn. In our implementation we take a moving weighted average of price history. This approach is somewhat flawed in that, for large moving windows, early estimations will be poor based on lack of available data. For small moving windows, decisions will always be poor because there will never be sufficient data to make a good decision. This flaw is clearly seen in the number of trades that take place early in the simulation. Producers, who cannot look far into the past when estimating the next price, will tend to react more sharply to the price changes and, as a result, are more willing to sell off supply at early prices. If the producers had some way to more accurately estimate next turn price, they may hold barrels longer into the simulation, allowing for greater and more evenly distributed trades later on. This, however, could have the consequence of significantly varying price over the simulation, which would negate our hypotheses H1 and H2 and show that global coherence does not in fact occur as the result of local economic decisions.

Commented [LS42]: YES! Use time to factor it. Conceding factor that I mentioned above.

The third factor, agents do not consider the value of future decisions nor learn from previous actions. Combined with the second factor, our agent design results in agents that make very imperfect decisions. This is more clearly demonstrated by the behavior seen when speculators are introduced into the market. Because speculators make decisions based only on the current trading price, they are not able to plan ahead and consider how their actions could drive price. A speculator should model their actions based on the impact they will have on price, as that is directly proportional to the impact on their bottom line. Speculators, for example, should be able to intentionally sit on large quantities of oil in an attempt to drive up price before selling. The speculator agents in our model, however, do not have the ability to make a decision to hold, and as a result will achieve a hold only accidently if the offer price is not sufficient to compel them to sell. As we have already discussed, the fact that decisions are so closely tied to current trading price results in the accidently hold decision to rarely, if ever, occur as offers are accepted too readily. Furthermore, as a result of agents in our model not being able to learn from their actions, a situation quickly develops whereby speculator agents will simply trade back and fourth with each other ad infinitum; their reliance on current trading price to make decisions just exacerbating the problem and leading to flat prices.

Commented [LS43]: Yep. Your consumer agents do not look ahead. Also, your speculator agents are static in their strategies.

Speculator agents are not unique in their unsophisticated implementation. Producer agents would also benefit greatly from the ability to determine how their own actions could influence the price. Just as speculators sitting on oil can drive up a price, producers could easily be given the ability to hold off until they achieve some desired profit margin. Another contributing factor to the lack of a decision to old was the formula for (8) not being accurately coded in our program which resulted in the agents significantly under valuing holding onto oil and is one explanation for the early sell offs. Consumer agents could also benefit from forethought in decisions. In our current model, they will greedily seek to fill demand with no consideration of the storage cost paid by other agents. Were they able to wait out turns in an effort to make sellers more desperate to dump supply, they would be able to achieve a lower price; this would apply to speculators as well. Essentially, all of the agents as designed are not complex enough to actually model the kind of decisions they need to make. They don't look far enough ahead; they don't strategize based on current conditions; and they are not able to consider the strategies that would be employed by their fellow agents. Speculators are the most complicated in terms of the decisions they need to make, and so the naiveté of the speculators is particularly apparent in our system as well as our observed results.

While clearly problematic, our model does, in a limited way, achieve our end goal. We were able to successfully create an economic model that demonstrates a globally coherent result which models a aspect of the system that is more complex than the local decisions that agents can make. Despite the threats to that conclusion, we believe that we have provided sufficient data to justify further study; while it isn't clear that our global coherence is able model the complex behavior we are looking for, it is also not clear that it would fail to model the behavior. Patterns in the results did emerge inline with our hypotheses, if not all of them. Correcting the issues identified with agent design would intuitively balance out the potential price fluctuations that would otherwise emerge so long as all agents were equally intelligent in their consideration of future effects of their actions. As such, it is highly likely that our emergent behavior would in fact match that seen in this study, albeit at different times in the simulation. Addressing this concern would require future work to modify the agents to introduce the ability to make more intelligent and rational decisions. It would also require that the initial input parameters undergo a more formal analysis to find near optimal values that are inline with the specific economic property that is being studied. We leave this work for a future analysis.

REFERENCES

Bellman, R. E., 1957. *Dynamic Programming,* Princeton University Press, Princeton NJ, Republished 2003: Dover.

Campbell, A., Wu, A., 2009. *On the significance of synchronicity in emergent systems*, Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems – Volume 1.

Collier, N.T., North, M.J., 2011. *Repast SC++: A Platform for Large-scale Agent-based Modeling*, in W. Dubitzky, K. Kurowski, and B. Schott, eds., Large-Scale Computing Techniques for Complex System Simulations, Wiley.

Kash, I., Friedman, E., Halpern, J., 2009. *Multiagent learning in large anonymous games,* Proceedings of the 8th International Conference on Autonomous Agents and Multiagent Systems – Volume 2.

Praca, I., Viamonte, M., Vale, Z., Ramos, C., 2008. *Agent-based simulation of electronic marketplaces with decision support,* Proceedings of the 2008 ACM symposium on Applied computing

Van Dyke Parunak, H., Brueckner, S., Savit, R., 2004. *Universality in Multi-Agent Systems*, Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems – Volume 2.

Walter, I., Gomide, F., 2008. *Electricity market simulation: multiagent system approach*, Proceedings of the 2008 ACM symposium on Applied computing.

Received December 2011;

Report	**Score**
Design Description & Discussion (50%)	50
Organization (20%)	20
Requirement (10%)	10
Experiments/Results (10%)	10
Grammar & Errors (10%)	10
TOTAL	**100**

Programming	**Score**
Documentation (15%)	10
Software Design (15%)	15
Programming Style (10%)	10
Testing (15%)	13
Program Correctness (45%)	43
TOTAL	**92**

FINAL PROJECT SCORE: 95.5