AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Causality in Futures Markets

by

Henry L. Bryant, David A. Bessler, and Michael S. Haigh

WP 03-07

Department of Agricultural and Resource Economics
The University of Maryland, College Park

Causality in Futures Markets

by

Henry L. Bryant^a, David A. Bessler^b, and Michael S. Haigh^c

^aDoctoral candidate and ^bProfessor, Department of Agricultural Economics, Texas A&M University, College Station, TX 77840, USA

^cAssistant Professor, Department of Agricultural and Resource Economics, University of Maryland, College Park, MD 20742, USA

June 1, 2003

Abstract: This research investigates various unresolved issues regarding futures markets, using formal methods appropriate for inferring causal relationships from observational data when some relevant quantities are hidden. We find no evidence supporting the generalized version of Keynes's theory of normal backwardation. We find no evidence supporting theories that predict that the level of activity of speculators or uninformed traders affects the level of price volatility, either positively or negatively. Our evidence strongly supports the mixture of distribution hypothesis (MDH) that trading volume and price volatility have one or more latent common causes, resulting in their positive correlation.

JEL Classification: C49, G13

Keywords: futures markets, causality

Acknowledgements: The first author gratefully acknowledges the support of the Tom Slick Senior Graduate Research Fellowship from the College of Agriculture and Life Sciences at Texas A&M University. Thanks to Todd Knarr whose public license "Date" class for C++ was used in manipulating the raw futures data.

Correspondence regarding this manuscript should be addressed to:

David A. Bessler
Department of Agricultural Economics
2124 TAMU
College Station, TX 77840
(979) 845-3096
d-bessler@tamu.edu

Causality in Futures Markets

I. Introduction

It has been over seventy years since Keynes wrote his *Treatise on Money*, in which he proposed his theory of “normal backwardation” – the idea that hedgers use futures markets to transfer risk to speculators, causing futures prices to deviate from expected future cash prices so that the speculators might be compensated. Despite decades of empirical investigation, no consensus regarding the validity of Keynes’ conjecture has been reached. Two difficulties have prevented the conclusive confirmation or rejection of the theory. First, the expected future cash price is not observed, and therefore neither is any risk premium. Second, it is not feasible for researchers to seek the answer to this question by experimentation. Systematic manipulation of futures markets is not only impractical; in many cases it is illegal.

Indeed, researchers conducting empirical work in economics and finance generally must work with observational rather than experimental data, and frequently are not able to observe all relevant quantities. This makes the correct inference of causal relationships difficult at best and impossible by some accounts - many assume that the use of controlled experiments is the *only* means by which causal mechanisms can reliably be inferred. Careful empirical researchers in these fields have thus resigned themselves to being able to draw only rather weak conclusions. It is said that evidence consistent with a theory is found, rather than that a theory has been proven. The less cautious investigator, upon finding that two observed quantities A and B covary, might imprudently conclude that A causes B , or vice versa. Consequently, empirical studies in economics and finance rarely unanimously support or reject available theoretical explanations. Such is certainly the case with research into futures markets, the subject of this paper.

This situation is not unique to economics and finance - researchers in numerous fields find themselves operating under such difficult circumstances. This situation has inspired a recent multidisciplinary effort to develop a body of theory concerning the inference of causal relationships using observational data. A subset of this literature further concerns itself with conducting this inference when the observational data are incomplete. Treatments of this subject can be found in Pearl (2000) and Spirtes, Glymour and Scheines (2000). This study uses these causality methods to investigate Keynes's theory, and other unresolved questions of more recent vintage regarding futures markets: we investigate the causes of the well-documented positive correlation between volume and volatility in futures markets, and assess the evidence regarding theories that predict that the activities of certain types of traders affect levels of price volatility. A correct understanding of the causal mechanisms that drive futures markets is obviously important for a variety of parties – hedgers, speculators, exchange officials, and regulators.

The following section extends this introduction by discussing in detail the issues that we investigate and the importance of each. We then describe the specific causal inference procedure that we employ and the data that we use. Finally, we present the analysis and offer some concluding remarks.

II. Issues Investigated

The first issue that we investigate is the Keynes's (1930) theory of normal backwardation, and its extensions. Keynes argues that hedgers enter the futures markets primarily to reduce the risk associated with cash market positions. He further argues that hedgers are generally commodity producers, and are therefore long in the cash market and short in the futures market. This necessarily means that speculators must be long in the futures

market, and he postulates that the current futures price must be below the expected future cash price in order to induce speculators to bear the risk associated with those long positions. A consequence of the theory as stated by Keynes, and supported in Hicks (1939), is that a futures contract's prices are expected to display an upward trend on average. Telser (1958) searches for such trends in the cotton and wheat markets, and reports finding no evidence. Cootner (1960) extends the theory by pointing out that hedgers are not necessarily commodity producers, but may be commodity consumers as well. Thus the net position of hedgers as a whole might be either long or short. As such, he suggests that the current futures price might be either above or below the expected future cash price. The modified theory is sometimes referred to as "net hedging" or "hedging pressure." Cootner reports finding evidence consistent with this hypothesis, namely that speculators appear to be earning profits over his sample period. Cootner thus shifts the empirical focus from searching for trends in futures prices to searching for speculative profits and/or hedging losses in futures markets.

Houthakker (1957) and Rockwell (1967) note however that speculative profits may be due to superior forecasting ability, rather than the collection of a risk premium. Rockwell thus recasts normal backwardation as "the return earned by a hypothetical speculator who follows a naïve strategy of being constantly long when hedgers are net short and constantly short when hedgers are net long." This then implies that speculative profits / hedging losses are a necessary, but not sufficient condition for the net hedging theory to hold. The analysis then must focus on decomposing speculative profits into forecasting ability and naïve components. Both Rockwell and Chang (1985) conduct such analyses, and each finds evidence of speculative profits. Rockwell reports that speculative profits are due to forecasting ability, however Chang reports evidence of naïve profits as well. This approach suffers from the inherent difficulty of dividing

speculators into able and naïve groups, when forecasting ability is unobserved. The reliability with which this task can be performed using aggregate data on trader positions is highly questionable. Further complicating matters, the available data regarding market commitments contain a proportion of traders whose status (either speculator or hedger) is unknown. In contrast to Rockwell and Chang, Hartzmark (1987) uses a very unique, highly disaggregated data set to find evidence that hedgers earn significant positive profits on average, precluding Rockwell's naïve speculator from profiting. Certainly, the evidence from these related empirical approaches to the question is mixed.

Meanwhile, another thread of the literature has developed a somewhat different perspective on the question. The theory of normal backwardation is presented in the context of a single asset, and the hypothesized risk premium should therefore due to the expected futures return and variability of that return. Dusak (1973) and Black (1976) argue that the question should be considered in a portfolio context. The capital asset pricing model (CAPM) states that any risk premium should be due to the relationship between an asset's returns and returns on total wealth. If a futures contract's price changes are correlated with returns on total wealth, then some portion of the risk of holding a futures contract is undiversifiable (a non-zero "beta" in CAPM parlance), and a risk premium should therefore be present (because there is a "systematic" risk). If, on the other hand, futures price changes are independent of returns on total wealth, then the risk of holding a futures contract should be fully diversifiable, and no risk premium should be present. Dusak finds that for the markets that she considers, futures price changes are independent of returns on a proxy for total wealth (the S&P 500 index), and concludes that no risk premiums are present. Carter, Rausser & Schmitz (1983) argue that Dusak's proxy for total wealth is inadequate, and that it should include commodity prices.

Hirschleifer (1988) and Hirschleifer (1990) argue that the assumptions built into the standard CAPM might be inappropriate, however. He argues that there may be a costs associated with futures market participation (perhaps in the form of learning the mechanics of futures market operation), which limit the participation of some types of investors. If this is the case, his theoretical models show that even in the presence of a zero beta, not all risk can be diversified away. He therefore argues that risk premiums in futures markets could be composed of two components – the standard systematic component, and a “residual” component that is a function of hedging pressure. Bessimbinder (1992) finds that futures returns covary with hedger’s net positions, and concludes that this result supports hedging pressure as a determinant of risk premiums in futures markets, consistent with Hirschleifer’s model, and with the generalized concept of normal backwardation.

We now summarize the above discussion regarding existing empirical research on normal backwardation. Risk premiums that may exist in futures markets cannot be observed, because the expected future cash price cannot be observed. The standard empirical practice then is to check for speculative profits, which would be consistent with the existence of risk premiums. If speculative profits exist (the evidence on this is mixed), they must be decomposed into profits due to forecasting ability, which is unobserved, and any residual profits (a dubious proposition). If there are profits that are not due to forecasting ability, it is inferred that risk premiums are present. These premiums may be due to systematic risk if futures price changes are correlated with returns to total wealth (a nebulous concept). After adjusting “observed” risk premiums for systematic risk, it is then inferred that any residual risk premium that is not due to systematic risk may be due to hedging pressure, if measures of these two phenomena are correlated. This path by which a researcher might find evidence consistent with the generalized theory of normal

backwardation is so convoluted, it is little wonder that no consensus has been reached. If such evidence is found and can be believed, it is still only *consistent* with the theory, the real burning issue of causality is never addressed. Does the net position of hedgers *cause* futures price changes?

We believe that the successful evaluation of this question requires that it be reconsidered from scratch, in a framework that explicitly addresses the issue of causality, and simultaneously accounts for the existence of relevant, but unobserved quantities. We provide such an investigation here. Clearly, the correct answer to Keynes' theory is important for market participants. Should a hedger anticipate losing money on average in exchange for enjoying reduced risk? Can speculators expect to be profitable on average by merely taking a position opposite that of hedgers' net position, regardless of the depth of their knowledge of a market and their forecasting ability?

The second issue that we investigate is the cause(s) of positive correlation between the volume of trade and degree of price variability in futures markets. This relationship is well documented; Karpoff (1987) provides a survey of the evidence. There are two theoretical explanations for this phenomenon. First, there is the Mixture of Distributions Hypothesis (MDH), due originally to Clark (1973). He proposes a model in which there is a stochastic number of independent price changes over any time period, due to a non-constant rate of information arrival. This results in the variance of the overall price change for a given period being an increasing function of number of within-period price changes. Volume of trade is also specified as an increasing function of the number of within-period price changes. Thus, in this theory the (unobserved) rate of information arrival is a common cause of trading volume and price change volatility. Epps and Epps (1976) present an alternative formulation of the MDH.

They specify an equilibrium model of intraday price determination in which the level of disagreement among traders causes the magnitude of a day's overall price change. Here volume is also an increasing function of disagreement, and the Epps & Epps model therefore also implies that volume and volatility are both effects of a latent common cause. Tauchen and Pitts (1983) offer a MDH model that incorporates elements of both the Clark and Epps and Epps models. On a related issue, the MDH models also result in a leptokurtotic distribution for observed price changes, which is consistent with empirical evidence. A competing explanation for this phenomenon due to Mandelbrot (1963) is that price changes are drawn from a distribution with infinite variance from the stable Paretian family. Finding evidence supporting the MDH explanation for positive volume and volatility correlation would thus also support one theory of the cause of excess kurtosis in the futures price change distribution.

A competing explanation for the positive correlation between trading volume and price volatility in futures markets is that of noisy rational expectations (NRE). In the NRE model of Shalen (1993), there are two types of traders. Informed traders have private information regarding market values. Uniformed speculators, on the other hand have no private information, and attempt to extract price signals from observed futures price changes. The series of these price changes is noisy, however, due to a random liquidity demand from hedgers (buying or selling due to their activity in the underlying market, not due to information arrival). The uninformed speculators can then misinterpret this liquidity trading as being due to information arrival, causing them to adjust their positions, resulting in increases in volume and volatility. The level of activity of uniformed speculators then is a common cause of volume and volatility.

We investigate the causal mechanisms driving the volume and volatility relationship. In addition to potentially vindicated one of the theoretical explanations given above, understanding

these mechanisms is important for market participants, researchers, and regulators. All market participants are obviously impacted by price volatility and market depth, and clearly should be interested in the underlying causes. Researchers will be interested in the correct specification of empirical models, and results inconsistent with existing theories might inspire new ones. Regulators have displayed an interest in curbing excessive levels of price volatility, and the success of such an endeavor would be greatly aided by a deep understanding of its causes.

The third issue that we investigate is allegations that the activities of specific types of traders are causes for the level of price volatility. This is closely related to the volume - volatility issue; as explained above, the NRE expectation model of Shalen predicts that volatility is an increasing function of the number of uninformed speculators. Similarly, in the model of Stein (1987), rational, but imperfectly informed futures speculators can (but do not necessarily) destabilize prices. These models contrast with the rational expectations model of Danthine (1978), in which imperfectly informed speculators stabilize prices. More recently, the finance literature has become interested in irrational behavior. An example of this is the model of DeLong, et al. (1990). In their model, irrational traders drive an asset's price away from the fundamental value, rational arbitrageurs' fear that the return to fundamental value may be slow coming, and so limit their activity, resulting in increased price volatility. This model is not concerned with futures markets as such, but the underlying principals might still apply.

Empirical evidence compiled regarding this question thus far is limited. Daigler and Wiley (1999) examine various financial futures markets, and report that the activity of futures traders who are on the trading floor is associated with decreased price volatility, while the activity of the "general public" is associated with increased volatility. The on-floor traders can observe the identities of those making large trades, and are therefore in a position to infer the

informational content of those trades. They can therefore be thought of as informed, and the results are thus consistent with the models of Shalen and Stein. Chang, Chou, and Nelling (2000) find that in the S&P 500 futures market, large hedging activity is positively correlated with volatility, and concludes that increased volatility likely results in increased hedging demand. Wang (2002) finds that in exchange rate futures markets measures of speculative activity and volatility are positively related. He suggests then that speculators destabilize markets. Note that these last two studies find very similar empirical evidence, but reach opposite conclusions regarding the likely direction of causality. Neither seems to consider the possibility that the observed relationship might be due to a common cause. This question is important for reasons similar to those given above for our second line of inquiry.

III. Inferring Causal Relationships From Incomplete, Observational Data

As mentioned in the introduction, treatments of the theory of causal inference using observational data can be found in Pearl (2000) and Spirtes, Glymour and Scheines (2000). These methods are just beginning to be adopted in applied economic research, although these efforts to date have largely worked under an assumption of causal sufficiency (i.e. that the researcher has collected observations for all variables present in the unknown causal structure). Swanson and Granger (1997) search for causal relationships among the variables in a vector autoregression to guide an appropriate Bernanke decomposition of the innovation covariance matrix and Demiralp and Hoover (2003) investigate the reliability of such a procedure. Haigh and Bessler (2003) investigate price discovery in cash grain markets and a related transportation market. Akleman, et al (1999) investigate causal relationships among corn exports and exchange rates using causal methods, both with and without the assumption of causal sufficiency.

We now provide a description of the algorithm that we employ to inferring causal relationships, the Fast Causal Inference (FCI) algorithm. The FCI was developed to be appropriate for inferring causal relationships from observational data (to the extent possible), even in the presence of latent variables. This section is adapted from Chapter 6 of Spirtes, Glymour and Scheines (2000); see that work for a more thorough description. The causal literature has developed the directed graph as a tool for visually representing a group of related causal relationships. A graph is a set of variables ($V_1, V_2, \dots, V_n$) that are connected by lines called edges, which may represent causal flows. If two variables are connected by an edge, they are said to be adjacent. Directed edges have arrowheads on the ends indicating the direction of causal flow between two adjacent variables. For example, $V_1 \rightarrow V_2$ indicates that V_1 is a cause of V_2 . V_1 is a parent of V_2 if there is a directed edge from V_1 to V_2 . A path is a sequence of variables such that each pair of variables that are adjacent in the sequence are also adjacent in the graph. A directed path is a path containing only directed edges in which causal flow runs from the first endpoint on the path to the last. An undirected path is a path in which causal flow is not required to run from the first endpoint on the path to the last. If there is a directed path from V_1 to V_2 , we say that V_1 is an ancestor of V_2 (e.g. as is the case in the graph $V_1 \rightarrow V_3 \rightarrow V_2$) and that V_2 is a descendant of V_1 . Note that parents are always ancestors, but the reverse is not true. A cyclic path is one in which causal flow begins at a variable and eventually returns to that variable (e.g. $V_1 \rightarrow V_2 \rightarrow V_3 \rightarrow V_1$). If a variable is caused by two other variables on a path, it is said to be a collider. For example, in the graph $V_1 \rightarrow V_2 \leftarrow V_3$, V_2 is a collider on the paths $\langle V_1, V_2, V_3 \rangle$ and $\langle V_3, V_2, V_1 \rangle$. A graph that contains directed edges, and no cyclic paths is a directed acyclic graph (DAG). The set of variables in a DAG is assumed to be causally sufficient – there are no latent common causes for any pair of variables in V .

Nature may choose to hide some variables, however. Suppose there is a DAG G over a set of variables V , and that O is a subset of the variables in V that are observed. A path U is an inducing path relative to O if and only if a) every member of O on U , except for the endpoints, is a collider on U , and b) every collider on U is an ancestor of either one of the endpoints. For example, in the graph G (see figure 1), the path $U = \langle V_1, V_2, V_3, V_4, V_5, V_6 \rangle$ is an inducing path from V_1 to V_6 over $O = \{V_1, V_2, V_4, V_6\}$. As required, each of the colliders on U , V_2 and V_4 , is an ancestor one of the endpoints, V_1 and V_6 , and the variables on U that are in O (other than the endpoints) are each colliders on U . Inducing paths provide the critical connection between statistical independence relations and causal mechanisms represented in graphs over observable variables. The existence of an inducing path between V_1 and V_2 is implied if V_1 and V_2 are statistically dependent conditional on every subset of $O \setminus \{V_1, V_2\}$ (although this fact alone does not imply the direction of causal flow). This implies that for our example, we would be able to find no subset of $\{V_2, V_4\}$ (including the empty set) that could be conditioned on to render V_1 and V_6 independent.

A graph is an inducing path graph (IPG) over O if there is an edge between two variables V_1 and V_2 with an arrowhead at V_2 if and only if there is an inducing path in G from V_1 to V_2 relative to O . To continue the example from the above paragraph, suppose we observe only the variables in $O = \{V_1, V_2, V_4, V_6\}$. As previously established, there is an inducing path running from V_1 to V_6 over O . The inducing path graph G' over O (shown in figure 2) thus features a directed edge running from V_1 and V_6 . Note, however, that in G there was no edge running from V_1 to V_6 . This illustrates an important point – the existence of a directed edge between two variables in an IPG implies that one variable is an ancestor of the other in the underlying DAG, but not necessarily a parent. Note also that some edges in G' have arrowheads on both ends.

These result from the existence of inducing paths running from each variable to the other. For example, $\langle V_2, V_3, V_4 \rangle$ is an inducing path over O , as is $\langle V_4, V_3, V_2 \rangle$. Hence the definition of IPG above requires arrowheads at both ends of the edge between V_2 and V_4 in G' . Such edges are referred to as bidirected edges, and they imply that the two adjacent variables have a latent common cause (in this case V_3).

Unfortunately, the statistical conditional independence relations over a set of observed variables will not necessarily identify a unique IPG. For an IPG G' , the set of all IPGs that entail equivalent sets of statistical independence relations over a given set of observed variables O is denoted $\text{Equiv}(G')$. A partially-oriented inducing path graph (POIPG) is a pattern that represents set of IPGs in $\text{Equiv}(G')$, where G' is the true IPG over O for the DAG G . The ends of the edges in a POIPG can have any one of three types of marks: no mark, and arrowhead, or an “o”. We use the symbol “*” to denote any one of these three types of end marks. We say that π is a POIPG of DAG G with IPG G' over O if and only if: a) π and G' have the same variables and adjacencies; b) if $V_1 o \rightarrow V_2$ is in π , then either $V_1 \rightarrow V_2$ or $V_1 \leftrightarrow V_2$ is in every IPG in $\text{Equiv}(G')$; c) if $V_1 \rightarrow V_2$ is in π , then $V_1 \rightarrow V_2$ is in every IPG in $\text{Equiv}(G')$; d) if $V_1 * - \underline{V_2} * - V_3$ is in π , then V_2 is a non-collider in every IPG in $\text{Equiv}(G')$; e) if $V_1 \leftrightarrow V_2$ is in π , then $V_1 \leftrightarrow V_2$ is in every IPG in $\text{Equiv}(G')$; and f) if $V_1 o - o V_2$ is in π , then either $V_1 \rightarrow V_2$, $V_1 \leftarrow V_2$, or $V_1 \leftrightarrow V_2$ is in every IPG in $\text{Equiv}(G')$. The adjacencies that exist in a POIPG then convey information about the conditional independence relations among the observed variables, and end marks on the edges other than “o” convey information about the direction of causal flow in the underlying DAG. The output of the FCI algorithm that we describe below is a POIPG.

One special type of path that may be found in a POIPG is a definite discriminating path, the existence of which may be used when orienting the edges in the FCI algorithm. A path U is a

definite discriminating path for V_1 if and only if U is an undirected path between V_2 and V_3 containing V_1 , every variable on U , except the endpoints, is either a collider or definite non-collider on U and the following conditions also hold:

- A) If V_4 and V_5 are adjacent on U and V_5 is between V_4 and V_1 on U , then $V_4^* \rightarrow V_5$
- B) If V_4 is between V_3 and V_1 on U and V_4 is a collider on U then either $V_4^* \rightarrow V_3$ or $V_4 \leftarrow V_3$.
- C) If V_4 is between V_2 and V_1 on U and V_4 is a collider on U then either $V_4^* \rightarrow V_2$ or $V_4 \leftarrow V_2$.
- D) V_2 and V_3 are not adjacent.

Some conditions regarding the underlying DAG G are necessary for inferring the set of IPGs over O that are in $\text{Equiv}(G')$ using a set of conditional independence relationships. First, the Markov condition assumes that in the probability distribution over the variables V in the underlying DAG G , a variable V_1 is independent of every set of variables that does not contain V_1 or its decedents, conditional on V_1 's parents. This essentially states that it is possible to represent a set of conditional independence relations graphically, using the definitions of a DAG and the related terminology that we laid out in the first paragraph of this section. Second, the faithfulness or stability condition requires that the conditional independence relations among the variables V in the underlying DAG G are due to the topology of G , rather than peculiar, offsetting parameter values in the causal relationships. Pearl (2000) gives the following example. Suppose we have DAG H (see figure 3), and that the causal relationships are represented by the structural equations

$$V_2 = cV_1 + u_2$$

and

$$V_3 = aV_1 + bV_2 + u_3$$

where u_2 and u_3 are independent stochastic errors. Note that generally we would expect V_2 and V_3 to be dependent, but if the parameter a just happened to take the value $-bc$, then V_2 and V_3 would be independent. The faithfulness or stability condition states that one is unlikely to encounter this type of independence relation in practice. If, in examining a hypothetical data set associated with figure 3, the only conditional independency we find is that V_2 and V_3 are independent conditioned on the null set, the faithfulness condition allows us to infer that V_2 and V_3 should not be adjacent, and that neither V_2 nor V_3 is caused by V_1 .

We now describe the Causal Inference (CI) algorithm, the basic functioning of which underlies the FCI algorithm that we use, and then describe how the two algorithms differ. The input of either algorithm is observations over a set of possibly causally insufficient variables, and the output of either algorithm is a POIPG. Both algorithms consist of two phases, determining the adjacencies in the POIPG using a statistical test of conditional independence relationships (Fisher's z test), and then deducing the maximally informative orientation of the resulting edges that is consistent with the faithfulness condition and the assumption that the underlying graph is a DAG (i.e. there are no cycles). The CI algorithm involves the following steps:

- A) Form a complete undirected graph on the set of variables O , in which every variable is connected to every other variable by an undirected edge.
- B) If two variables V_1 and V_2 are independent conditional on any subset S of $O \setminus \{V_1, V_2\}$, remove the edge between V_1 and V_2 , and record S in the separating set for V_1 and V_2 , denoted $\text{Sepset}(V_1, V_2)$.
- C) Let F be the graph that results from step B). Orient each edge as $o \rightarrow o$. For each triple of variables (V_1, V_2, V_3) such that the pairs (V_1, V_2) and (V_2, V_3) are adjacent in F but

the pair (V_1, V_3) is not, orient $V_1^* -^* V_2^* -^* V_3$ as $V_1^* \rightarrow V_2 \leftarrow^* V_3$ if and only if V_2 is not in $\text{Sepset}(V_1, V_3)$ and arrange $V_1^* -^* V_2^* -^* V_3$ as $V_1^* -^* \underline{V_2}^* -^* V_3$ if and only if V_2 is in $\text{Sepset}(V_1, V_3)$.

D) Repeat the following sequence of instructions until no more edges can be oriented:

- i) If there is a directed path from V_1 to V_2 , and there is an edge $V_1^* -^* V_2$, orient $V_1^* -^* V_2$ as $V_1^* \rightarrow V_2$.
- ii) Else, if V_1, V_2, V_3 is a collider along $\langle V_1, V_2, V_3 \rangle$, V_1, V_2, V_3 is adjacent to V_1, V_2, V_3 , and V_1, V_2, V_3 is not in $\text{Sepset}(V_1, V_3)$, then orient $V_1^* -^* V_2$ as $V_1 \leftarrow^* V_2$.
- iii) Else, if U is a definite discriminating path between V_1 and V_2 for V_3 , and V_4 and V_5 are adjacent to V_3 on U , and V_3, V_4 , and V_5 form a triangle, then
 - a) If V_3 is in $\text{Sepset}(V_1, V_2)$, then mark $V_3^* -^* V_4^* -^* V_5$ as $V_3^* -^* \underline{V_4}^* -^* V_5$
 - b) Else, $V_3^* -^* V_4^* -^* V_5$ as $V_3^* \rightarrow V_4 \leftarrow^* V_5$
- iv) Else, if $V_1^* \rightarrow \underline{V_2}^* -^* V_3$ then orient as $V_1^* \rightarrow V_2 \rightarrow V_3$.

Step B above is computationally infeasible, as the number of possible subsets of O grows very rapidly with the cardinality of O . Checking for conditional dependence of two variables V_1 and V_2 over all possible subsets of $O \setminus \{V_1, V_2\}$ then becomes very difficult. The FCI and CI algorithms differ in the way that step B is performed. The FCI uses an intermediate step to infer that some variables cannot be in $\text{Sepset}(V_1, V_2)$, thereby reducing the number of conditional independence tests that must be performed. This procedure is relatively complicated, and does not offer any additional understanding of the means by which the causal structure is inferred, and we therefore do not describe it. The important fact is that the FCI algorithm is essentially a computationally feasible version of the CI algorithm. The FCI algorithm is implemented in the Tetrad 3 computer program, which we use in our analysis.

IV. Data

We analyze eight futures markets: Chicago Board of Trade (CBOT) corn, New York Mercantile Exchange (NYMEX) crude oil, Chicago Mercantile Exchange (CME) Eurodollar deposits, New York Commodity Exchange (COMEX) gold, CME Japanese Yen, New York Board of Trade (NYBOT) coffee, CME live cattle, and the CME S&P 500. Observations for all data over the interval March 21, 1995, through January 8, 2003, are used. We construct three types of data series for use in the analysis: i) those related to trader activity and positions, ii) those related to futures prices and trading volume, and iii) trend and seasonal series. We now discuss each category of data in turn.

The Commodity Futures Trading Commission (CFTC) requires certain exchange members and futures commission merchants (i.e. brokers) to file daily reports with the Commission. Those reports show the futures positions of traders that hold positions above specific reporting levels set by CFTC regulations (these are referred to as “reportable positions”). Each trader is classified as being either commercial or non-commercial, with commercial traders being those engaged in hedging activity. Ederington and Lee (2002) caution that this distinction is not always entirely accurate, and our data regarding trader type are thus noisy. Henceforth we refer to reportable commercial positions as being those of “large hedgers”, to reportable non-commercial positions as being those of “large speculators”, and to non-reportable positions as being those of “small traders”. The data collected as of a markets close on each Tuesday are released to the public in the CFTC’s Commitments of Traders (COT) report, generally on the following Friday. We use this data in two ways. First, we calculate the net position of large hedgers (*LH Net Position*) as the number of open long futures positions minus the number of open short futures positions held by large hedgers. Second, we calculate the aggregate level of

activity of each trader type (*LH Activity*, *LS Activity*, and *ST Activity* for large hedgers, large speculators and small traders, respectively) as the sum of their open long and short futures positions. These three variables, at any point in time, sum to twice the level open interest in the market. Some adjustments to the COT data are necessary. Before 1998, corn futures positions are measured in numbers of 1,000 bushels, rather than number of contracts (each calling for delivery of 5,000 bushels). We therefore divide all corn COT data prior to 1998 by five so that the related data series we use are measured in consistent units over the sample period. The size of the cash settlement called for by the S&P 500 futures contract was halved in late 1997, and we therefore multiply all S&P 500 COT data prior to the change by two, to make our measures of trader positions consistent with the current contract specification. In the crude oil and coffee markets, observations for the COT series are missing for September 11, 2001, and are linearly interpolated.

Daily price data for individual deliveries for each market are provided by Commodity Research Bureau, and the corresponding volume data are provided by Primark Datastream. We construct a continuous futures price level series (*Nearby*) for each market using week-ending observations of the futures contract nearest to expiration. We use weeks that run Wednesday through Tuesday in constructing all price, volume, and volatility series, so as to correspond with the COT data. We also construct a nearby weekly returns series (*Return*) using weekly returns series for each individual delivery. Thus no observations in our *Return* series are constructed using price level observations from two different deliveries (as would be the case if one simply constructed a return series using a previously constructed nearby levels series). A weekly total volume series (*Volume*) was constructed for each market by summing the total trading volume for all deliveries for each day. A measure of futures price volatility (*Volatility*) was constructed

by taking the log difference between the high and low prices for each week for the nearby contract.

A linear time trend series (*Time*) is used in the analysis, as are weekly observations of two annual seasonal harmonic variables. These are defined as

$$Annual\ Sin = \sin\left(\frac{2\pi Time}{52}\right)$$

and

$$Annual\ Cos = \cos\left(\frac{2\pi Time}{52}\right).$$

These harmonic series account for the possibility of seasonal influences on the volume, volatility, and activity variables which we expect in the agricultural commodity futures markets.

V. Analysis

We begin by investigating the possibility of non-stationary behavior in the series. *A priori*, the Efficient Market Hypothesis gives us strong reason to suspect that the *Nearby* series may contain a unit root, however, we have no such grounds for suspicion with respect to the remaining series. Indeed, it would seem rather implausible to believe that *LH Net Position*, for example, might drift off toward infinity. All data series save the trend and seasonal harmonic series are subjected to Augmented Dickey-Fuller (ADF) tests for non-stationary, with the results given in table 1. We find that for seven of the eight *Nearby* series we cannot reject the null hypothesis of no unit-root, confirming our initial suspicion. We therefore use the *Return* series for all markets in the causal analysis that follows (we use the *Return* series even the live cattle market, as we wish to keep the interpretation of the results consistent across markets). For the remaining series (*Volatility*, *Volume*, *LH Activity*, *LS Activity*, *ST Activity*, and *LH Net Position*),

we generally reject the null hypothesis that each series contains a unit root. Given that a) the burden of proof was on proving that there is no unit-root, b) our prior expectations, c) it is a well-established fact that ADF tests have low power against plausible alternatives (see, for example, DeJong, et al, 1992), and d) our desire to use consistent types of series (i.e. levels or differences) across markets, we proceed to use the levels series for all variables other than *Return*.

We apply the FCI algorithm to the 10 data series for each market. In all cases, the algorithm is restricted from allowing inducing paths running from any variable to *Time*, and from allowing a latent common cause for any variable and *Time*. Similar restrictions are placed on the allowed orientations of edges attached to the seasonal harmonic series, although the possibility of inducing paths running from *Time* to either of the seasonal harmonic series is not prohibited. The resulting POIPGs are presented in figures 4 through 11.

We first consider the evidence with regard to the generalized theory of normal backwardation. Our analysis considers the relationship between week-ending level of *LH Net Position* and the *Return* that was realized over those weeks. We interpret this as follows. Suppose the futures price begins at exactly the unobserved spot price that is expected to prevail at the time of expiration. A move by *LH Net Position* to a higher level should, if normal backwardation holds, then cause the futures price to move higher, to some price above the expected future spot price so that speculators who are now more short can be compensated. We then expect a positive relationship between the week-ending level of *LH Net Position* and *Return*.¹ The *prima facie* evidence in this regard is not generally supportive of the hypothesis that hedging pressure causes risk premiums, as we find a negative correlation between these two

¹ The *LH Net Position* variable is constructed using open interest data that are pooled across contract maturities, while the *Return* series is constructed from price data for the nearby contract. While not ideal, we are confident in this approach, as open interest tends to be heavily concentrated in the nearby contract. The *LH Net Position* variables should thus provide a reasonable approximation to the net position of large hedgers in the nearby contract, slightly scaled upward.

variables in all markets except the S&P 500. Examining the POIPGs for the eight markets, we find that *LH Net Position* and *Return* have a latent common cause in three markets (gold, japanese yen, and coffee), no causal connection in three markets (Eurodollar deposits, live cattle, and the S&P 500), and that in the remaining two markets (corn and crude oil) either there is a latent common cause or there is an inducing path from running *Return* to *LH Net Position* in all IPGs consistent with the observed set of conditional independencies. In no case do we find the possibility that causal flow might run from *LH Net Position* to *Return*, and we firmly conclude that hedging pressure does not cause returns, and we thus find no support for the generalized theory of normal backwardation. We can conclude, then, that it does not appear that hedgers need not expect to automatically pay a risk premium to speculators. Note, however, that this is not the same as concluding that risk premiums do not exist in these markets, only that there are not risk premiums *caused* by hedging pressure. The speculative profits sometimes found in other research could then be due to speculators collecting risk premiums that are due to other causes, or due to superior forecasting ability.

We now describe how the algorithm arrived at this conclusion. In the cases where the two variables are adjacent in the POIPG, a sufficient condition to conclude that causal flow does not run from *LH Net Position* to *Return* is the existence of an arrowhead on the *LH Net Position* end of the edge. We explain the existence of such an arrowhead using the corn market as an example. After the adjacencies are determined for the corn market POIPG (step B in the FCI algorithm), the following sub-graph is present: *Volume*—o*LH Net Position*—o*Return*. The adjacency between *Volume* and *Return* is removed because the unconditional correlation between the two is not significantly different from zero. Finding that *Volume* and *Return* are unconditionally uncorrelated prevents us from believing that we could have either *Volume*—o*LH*

Net Position→*Return* or *Volume*←*LH Net Position*←*Return*. Furthermore, the faithfulness condition prevents us from believing that *Volume*←*LH Net Position*→*Return* (if this were the case, it would be very unusual to find that *Volume* and *Return* were unconditionally uncorrelated). We therefore must accept the only remaining possibility, that *Volume*→*LH Net Position*←*Return* is the appropriate orientation (*LH Net Position* is a collider). We thus have an arrowhead at the *LH Net Position* end of the edge between it and *Return*. This type of edge orientation is due to step C of the FCI algorithm, and this rule can be used to orient all of the edge end marks that are critical to our analysis in this paper.

We next discuss the evidence regarding relationships between trader type and volatility levels, and afterwards discuss the related issue of the *Volume* and *Volatility* relationship. The theories discussed earlier in the paper make predictions regarding causal relationships between speculators and/or uninformed traders. *LS Activity* obviously represents speculative activity, and some would argue that this category represents uninformed traders to some extent as well. It does not seem unreasonable to interpret *ST Activity* as representing uninformed traders. Such distinctions turn out not to be necessary, however. We find no evidence of causal flow running from either *LS Activity* or *ST Activity* to *Volatility* in any of the eight markets. The edges directly connecting *ST Activity* and *Volatility* are removed by conditioning on the empty set in five markets, by conditioning on *Volume* in crude oil and eurodollars, and by conditioning on the *Time* trend for the S&P 500 market. The edge between *LS Activity* and *Volatility* is removed in the following markets by conditioning on the variables given in parentheses: corn (*Annual Sin*), crude oil (*LH Activity*), Eurodollars (*Return*), gold (*Volume* and *Time*), japanese yen and coffee (the empty set), live cattle (*LH Net Position*), and S&P 500 (*Time*). This information is summarized in table 2. Two variables need not be connected directly by an inducing path in the

POIPG for causal flow to run between them – we may find a roundabout directed path from one variable to the other. We find no evidence of such indirect causal flow in this case however. We thus find no evidence supporting theories that predict that the activity levels of speculators and/or uninformed traders affects volatility (either positively or negatively).

With regard to the *Volume* and *Volatility* relation, we find that in six of the eight markets there is a latent common cause for the two variables. In the coffee market, there is either a latent common cause, or that there is an inducing path from *Volatility* to *Volume* in all observationally equivalent IPGs. The evidence from these markets is then consistent with the MDH, which predicts that either the rate of information arrival (the Clark version) or the level of disagreement among traders (the Epps and Epps version) or some combination of these causes the positive correlation. In the crude oil market, we find an inducing path running from *Volume* to *Volatility*, which is consistent with neither the MDH nor Shalen's prediction that the level of activity of uninformed speculators causes the positive correlation between *Volume* and *Volatility*. If Shalen's theory is true, we expect to find causal flow running from either *LS Activity* or *ST Activity* to both *Volume* and *Volatility*. We find no such evidence in any of the markets that we analyze. The edges between the activity levels and *Volatility* were removed for the reasons discussed previously. The edges between *Volume* and *ST Activity* generally are not removed (crude oil and Eurodollars being the exceptions), but are bidirected, implying a latent common cause. Edges between *LS Activity* and *Volume* are removed in all markets save one (live cattle). This was accomplished by conditioning on *LH Activity* (Eurodollars, Japanese yen, and coffee), *Volatility* (gold and S&P 500), and the empty set (corn). Thus most of the necessary edges to support Shalen's theory are removed, and even though edges connecting *ST Activity* and *Volume* are generally not removed, they are bidirected implying a latent common cause. We also again find

no evidence of indirect causal paths that would support Shalen's theory. The practical implication of our findings is that those attempting to model time-varying volatility may indeed find volume to be a useful proxy for some unobserved cause or causes of volatility. It is not necessarily prudent, however, to assume that any event that will affect an increase in volume will also result in an increase in volatility. Contract expiration, for example, generally results in increased volume as traders roll positions out of the expiring contract. This event has nothing to do with either of the unobservable common causes of volume and volatility that have been suggested in the theoretical literature, and should not therefore be expected to cause an increase in volatility.

VI. Conclusions

In this article, we examine various unresolved issues regarding causal relationships in futures markets. To this end, we apply the Fast Causal Inference (FCI) algorithm, which has been developed in the formal causality literature as an appropriate tool for inferring causal relationships using observational data, even in the presence of relevant unobserved quantities. Such an approach is highly attractive, considering that most research in empirical economics and finance is conducted in such an environment. We find no support for the generalized theory of normal backwardation, and thus no reason to believe that hedgers will generally transfer a risk premium to speculators in exchange for risk-bearing services. We find no support for theories predicting that particular types of traders affect the level of price volatility, either positively or negatively, in futures markets. We find evidence that supports the mixture of distributions hypotheses (MDH), which posit the existence of one or more unobservable common causes of trading volume and price volatility. This suggests that models of time-varying volatility can

benefit from the information about the latent variable(s) contained in volume, but caution in the interpretation of such a model is necessary as volume does not actually cause volatility.

There are abundant opportunities for the further application of causal inference methods to empirical research into derivatives markets. Other open questions need to be addressed, some of which are: is the level of futures trading activity a cause of price volatility in the underlying cash market? What are the causes and/or effects of changes in the shape of the forward curve? What are the causes of basis movements? Does the size of the margin deposit required to trade futures impact any of the quantities that we have considered? What are the causal relationships that exist across related markets (e.g. the soy complex or the crude oil complex)? These issues offer a fertile ground for future research.

References

- Akleman, Derya G.; Bessler, David A.; and Burton, Diane M. (1999): "Modeling Corn Exports and Exchange Rates with Directed Graphs and Statistical Loss Functions," in Clark Glymour and Gregory F. Cooper (eds.) *Computation, Causation, and Discovery*, Menlo Park: American Association for Artificial Intelligence and Cambridge: MIT Press.
- Bessimbinder, Hendrick (1992): "Systematic Risk, Hedging Pressure, and Risk Premiums in Futures Markets," *Review of Financial Studies*, 5(4):637-667.
- Black, Fischer (1976): "The Pricing of Commodity Contracts," *Journal of Financial Economics*, 3:167-179.
- Carter, Rauser & Schmitz (1983): "Efficient Asset Portfolios and the Theory of Normal Backwardation," *Journal of Political Economy*, 91(2):319-331.
- Chang, Eric C. (1985): "Returns to Speculators and the Theory of Normal Backwardation," *Journal of Finance*, 40(1):193-208.
- Chang, Eric C.; Chou, Ray Y.; Nelling, Edward F. (2000): "Market Volatility and the Demand for Hedging in Stock Index Futures," *Journal of Futures Markets*, 20(2):105-125.
- Clark, Peter K. (1973): "A Subordinated Stochastic Process Model With Finite Variance for Speculative Prices," *Econometrica*, 41(1):135-155.
- Cootner, Paul H. (1960): "Returns To Speculators: Telser vs. Keynes," *Journal of Political Economy*, 68(4): 396-404.
- Daigler, Robert T. and Wiley, Marilyn K. (1999): "The Impact of Trader Type on the Futures Volatility-Volume Relation," *Journal of Finance*, 54(6):2297-2316.
- Danthine, Jean-Pierre (1978): "Information, Futures Prices, and Stabilizing Speculation," *Journal of Economic Theory*, 17:79-98.
- DeJong, David N.; Nankervis, John C.; Savin, N. E.; and Whiteman, Charles H. (1992): "Integration Versus Trend Stationary in Time Series," *Econometrica*, 60(2):423-433.
- DeLong, J. Bradford; Shleifer, Andrei; Summers, Lawrence; Waldman, Robert J. (1990): "Noise Trader Risk in Financial Markets," *Journal of Political Economy*,
- Demiralp, Selva; and Hoover, Kevin D. (2003): "Searching for the Causal Structure of a Vector Autoregression," University of California, Davis Working Paper No. 03-03.
- Dusak, Katherine (1973): "Futures Trading and Investor Returns: An Investigation of Commodity Market Risk Premiums," *Journal of Political Economy*, 81:1387-1406.

Ederington, Louis; and Lee, Jae Ha (2002): "Who Trades Futures and How: Evidence from the Heating Oil Futures Market," *H\Journal of Business*, 75(2):353-373.

Epps, Thomas W. and Epps, Mary Lee (1976): "The Stochastic Dependence of Security Price Changes and Transaction Volumes: Implications for the Mixture of Distributions Hypothesis," *Econometrica*, 44(2):305-321.

Fuller, Wayne A. (1976): *Introduction to Statistical Time Series*, New York: Wiley.

Haigh, Michael S.; and Bessler, David A. (2003): "Causality and Price Discovery: An Application of Directed Acyclic Graphs," *Journal of Business*, forthcoming.

Hartzmark, Michael L. (1987): "Returns to Individual Traders of Futures: Aggregate Results," *Journal of Political Economy*, 95(6):1292-1306.

Hicks, J. R. (1939): *Value and Capital*, Oxford: Clarendon Press.

Hirschleifer, David (1988): "Residual Risk, Trading Costs, and Commodity Futures Risk Premia," *Review of Financial Studies*, 1(2):173-193.

Hirschleifer, David (1990): "Hedging Pressure and Futures Price Movements in a General Equilibrium Model," *Econometrica*, 58(2):411-428.

Houthakker, H. S. (1957): "Can Speculators Forecast Prices?" *Review of Economics and Statistics*, 39(2):143-151.

Karpoff, Jonathan M. (1987): "The Relation Between Price Changes and Trading Volume: A Survey," *Journal of Financial and Quantitative Analysis*, 22(1):109-126.

Keynes, John Maynard (1930): *Treatise on Money*, London: MacMillan & Co.

Mandelbrot, Benoit (1963): "The Variation of Certain Speculative Prices," *Journal of Business*, 36:394-419.

Pearl, Judea (2000): *Causality*, Cambridge: Cambridge University Press.

Reichenbach, H. (1956): *The Direction of Time*, Berkeley: University of California Press.

Rockwell, Charles S. (1967): "Normal Backwardation, Forecasting, and the Returns to Commodity Futures Traders," *Food Institute Research Studies*, 7:107-130.

Schwarz, G. (1978): "Estimating the Dimensions of a Model," *Annals of Statistics*, 6:461-464.

Shalen, Catherine T. (1993): "Volume, Volatility, and the Dispersion of Beliefs," *Review of Financial Studies*, 6(2):405-434.

Spirtes, Peter; Glymour, Clark; and Scheines, Richard (2000): *Causation, Prediction, and Search*, Cambridge: The MIT Press.

Stein, Jeremy C. (1987): "Informational Externalities and Welfare-Reducing Speculation," *Journal of Political Economy*, 95(6):1123-1145.

Swanson, Norman R.; and Granger, Clive W. J. (1997): "Impulse Response Functions Based on a Causal Approach to Residual Orthogonalization in Vector Autoregressions," *Journal of the American Statistical Association*, 92:357-367.

Tauchen, George E. and Pitts, Mark (1983): "The Price Variability-Volume Relationship on Speculative Markets," *Econometrica*, 51(2):485-506.

Telser, Lester (1958): "Futures Trading and the Storage of Cotton and Wheat," *Journal of Political Economy*, 66(3):233-255.

Wang, Changyun (2002): "The Effect of Net Positions By Trader Type on Volatility in Foreign Currency Futures Markets," *Journal of Futures Markets*, 22(5):427-450.

Table 1 Results of Augmented Dickey-Fuller Tests^a

Commodity	<i>Nearby</i>	<i>Return</i>	<i>Volatility</i>	<i>Volume</i>	<i>LH Activity</i>	<i>LS Activity</i>	<i>ST Activity</i>	<i>LH Net Position</i>
Corn	-2.73 (1)	-19.59 (0)	-6.53 (2)	-11.49 (0)	-3.70 (1)	-4.10 (0)	-3.42 (1)	-3.94 (1)
Crude Oil	-2.03 (1)	-21.66 (0)	-15.53 (0)	-6.65 (3)	-2.30 (13)	-3.54 (0)	-1.76 (9)	-4.78 (1)
Eurodollars	-0.24 (0)	-19.07 (0)	-6.67 (2)	-5.44 (3)	-2.60 (13)	-2.81 (16)	-3.79 (1)	-4.01 (0)
Gold	0.02 (0)	-19.20 (0)	-14.17 (0)	-12.91 (0)	-3.47 (0)	-4.23 (1)	-3.38 (2)	-4.33 (1)
Japanese Yen	-2.25 (0)	-18.49 (0)	-5.97 (3)	-2.08 (12)	-4.17 (13)	-5.48 (0)	-7.73 (0)	-5.81 (1)
Coffee	-2.69 (3)	-21.39 (0)	-9.50 (1)	-14.11 (0)	-3.72 (0)	-4.55 (0)	-4.89 (1)	-5.90 (1)
Live Cattle	-3.82 (0)	-17.95 (1)	-6.78 (2)	-14.25 (0)	-1.96 (0)	-4.96 (1)	-4.77 (10)	-3.30 (1)
S&P 500	-0.24 (1)	-24.26 (0)	-8.91 (1)	-1.71 (12)	-1.97 (14)	-3.18 (13)	-1.92 (14)	-3.45 (0)

^a The null hypothesis is that the series listed in the row and column intersection has a unit root. We reject this hypothesis if the ADF test statistic is less than the critical value -3.13 (10%) given in Fuller (1976). Both an intercept and a time trend were included in the tests. The optimal lag length given in parenthesis was chosen using the Schwarz (1978) information criterion.

Table 2 Conditioning Sets that Result in Vanishing Correlations^b

Market	between <i>LS Activity</i> and <i>Volatility</i>	between <i>ST Activity</i> and <i>Volatility</i>	between <i>LS Activity</i> and <i>Volume</i>	between <i>ST Activity</i> and <i>Volume</i>
Corn	<i>Annual Sin</i>	Empty Set	Empty Set	(None)
Crude Oil	<i>LH Activity</i>	<i>Volume</i>	<i>ST Activity, LH Net Position, Time</i>	<i>LS Activity, Time</i>
Eurodollars	<i>Return</i>	<i>Volume</i>	<i>LH Activity</i>	<i>LS Activity</i>
Gold	<i>Volume, Time</i>	Empty Set	<i>Volatility</i>	(None)
Japanese Yen	Empty Set	Empty Set	<i>LH Activity</i>	(None)
Coffee	Empty Set	Empty Set	<i>LH Activity</i>	(None)
Live Cattle	<i>LH Net Position</i>	Empty Set	(None)	(None)
S&P 500	<i>Time</i>	<i>Time</i>	<i>Volatility</i>	(None)

^b For the market listed in a row, the correlation between the pair of variables listed in the column is not significantly different from zero, conditional on the variables given in the row and column intersection. “(None)” indicates that no set of variables is found that results in a correlation not significantly different from zero. “Empty Set” indicates that the unconditional correlation between the two variables is not significantly different from zero. *Return* is the log change in the price of the nearby futures contract, *Volume* is the total volume of trade, *Volatility* is the log difference between the high and low nearby futures prices for a week, *LH Net Position* is net futures position of large hedgers, *LH Activity*, *LS Activity*, and *ST Activity* are the total number of open futures positions of large hedger, large speculators, and small traders, respectively. *Time* is a linear time trend, and *Annual Sin* is an annual seasonal harmonic variable.

Figure 1 Directed Acyclic Graph G

Figure 2 Inducing Path Graph G' over $O = \{V_1, V_2, V_4, V_6\}$ Associated with Directed Acyclic Graph G

Figure 3 Directed Acyclic Graph H

Figure 4 Partially Oriented Inducing Path Graph for Corn

Figure 5 Partially Oriented Inducing Path Graph for Crude Oil

Figure 6 Partially Oriented Inducing Path Graph for Eurodollars

Figure 7 Partially Oriented Inducing Path Graph for Gold

Figure 8 Partially Oriented Inducing Path Graph for Japanese Yen

Figure 9 Partially Oriented Inducing Path Graph for Coffee

Figure 10 Partially Oriented Inducing Path Graph for Live Cattle

Figure 11 Partially Oriented Inducing Path Graph for S&P 500