

UCLA

Recent Work

Title

Volume and Price Formation in an Asset Trading Model with Asymmetric Information

Permalink

<https://escholarship.org/uc/item/37t357fc>

Authors

Bernardo, Antonio
Judd, Kenneth L.

Publication Date

1996-07-01

Volume and Price Formation in an Asset Trading Model with Asymmetric Information*

Antonio E. Bernardo[†] Kenneth L. Judd[‡]

July, 1996

Abstract

This paper examines how private information affects trading volume, the information content of trading volume data, and if there are any relations between trading volume and price changes which can be explained by informational differences. We develop a model with two trading periods in which asymmetrically informed agents learn the trading volume of the first round of trading prior to trading in the second round. We also develop a general method for computing equilibria in such a model assuming only smooth, concave utility functions, and asset return distributions and signals with smooth densities. The model predicts (i) a positive relation between trading volume and the absolute value of price changes; (ii) a positive relation between trading volume and subsequent stock price volatility; and (iii) that positive price movements on high trading volume lead, on average, to positive future price movements.

*Preliminary and incomplete. We thank seminar participants at the Western Finance Association meetings, Duke University, University of California, Berkeley, Stanford University and especially Deborah Lucas for valuable comments.

[†]Anderson Graduate School of Management at U.C.L.A. e-mail: abernard@agsm.ucla.edu - tel. (310) 825-2198.

[‡]Hoover Institution, Stanford University. e-mail: judd@hoover.stanford.edu.

1 Introduction

There is a vast literature on the formation of prices in asset markets when traders have private information. Much of this literature builds on the seminal work of Grossman (1976, 1981) and Grossman and Stiglitz (1980) in which competitive, risk-averse investors use private information *and* the information in equilibrium prices to make optimal trading decisions.¹ In particular, investors are assumed to use the correct equilibrium statistical relationship between prices and private information to update their beliefs about the value of a risky asset. This literature, however, largely ignores the role of other publicly observable information in this regard, in particular, trading volume.² In Grossman (1976), for example, trading volume plays no informative role since traders learn all relevant information by observing the equilibrium price alone.³ Many subsequent papers include the device of “noise traders” to prevent equilibrium prices from revealing all relevant information. While this device yields interesting and tractable models, trading volume once again has no interesting informational role in these models.⁴

There are many volume regularities which have been noted in empirical studies. Admati and Pfleiderer (1988), and Brock and Kleidon (1992) have studied intraday volume and price movements and found that periods of increased trading volume tend to be periods of increased return variability. In particular, volume and return variability are greatest at the open and close of trading. Interday studies have also turned up interesting facts relat-

¹A vast literature has also built on the work of Kyle (1985) who developed an alternative framework with imperfectly competitive risk-neutral informed traders, a risk-neutral market maker, and liquidity traders.

²A notable exception is Blume, Easley, and O’Hara (1994). Wang (1994) investigates the implications of asymmetric information for trading volume, but does not study the informational role of volume.

³The critical special statistical feature of his model is the existence of a one-dimensional sufficient statistic for the collection of private information, allowing a single price to also be a sufficient statistic for all information.

⁴See Blume, Easley, and O’Hara (1994, pp. 159-165) for an excellent treatment of the trivial informational role played by trading volume in the standard models.

ing volume and price movements. Karpoff (1987) surveys a large empirical literature that finds a positive relation between trading volume and the absolute value of price changes in equity and futures markets. Gallant, Rossi, and Tauchen (1992) and Brock and LeBaron (1993) have noted several statistically significant properties of the auto- and cross-correlation functions between price and volume movements. In particular, large price movements today imply large volume today and tomorrow. These properties are difficult to explain in conventional models.

The empirical facts are only part of the motivation for addressing these issues. When discussing the role of asymmetric information, it is clearly more realistic to use dynamic models wherein information leaks out gradually; this is why the Kyle (1985) and Wang (1994) analyses are so interesting. In these models, however, traders ignore any volume information. Since traders know (or can easily learn) volume in previous trading rounds and, perhaps, nearly contemporaneous volume data, we should include it in our analyses.

In this paper we examine how private information affects trading volume, the information content of volume data, and if there are any relations between trading volume and price changes which can be explained by informational differences. We develop a model with two trading periods in which asymmetrically informed investors learn the first period trading volume prior to trading in the second period. In period one, investors optimally allocate their wealth between the risky asset and a riskless bond conditional on their private information and the information in the equilibrium share price. In period two, however, investors condition their asset demands on their private information, the sequence of prices, *and period-one trading volume*. We compute rational expectations equilibria in which each investor uses the correct statistical relationship between equilibrium prices, trading volume, and private information to make inferences about other investors' private information from observed market data. We will make specific functional form assumptions concerning the elements of the model, but the general method we develop for solving such problems requires only smooth, concave utility functions, and asset return distributions and signals with smooth densities.

Trading volume has information content in our model because we employ a return

and information structure in which investors are uncertain about the mean *and variance* of the risky asset's return. For illustrative purposes, we make distributional assumptions which ensure that two statistics, corresponding to the sample mean and the sample variance of the private information signals, summarize all private information. The period-one price alone does not fully reveal all private information because it is consistent with many possible pairs of the two statistics. However, we demonstrate that trading volume is more highly correlated with the sample variance of the private signals than the sample mean. Intuitively, if all investors observe similar signals, good or bad, trading volume will be relatively low because all investors will share similar posterior beliefs about the value of the stock. However, when investors observe *different* private signals large trading volume is generated by differences in posterior beliefs. Thus, trading volume data allows investors to make more precise inferences because, in equilibrium, it is correlated with the dispersion statistic. Consequently, trading volume has information content beyond that which is in the period-one price.

Our model generates many interesting predictions including (i) a positive relation between trading volume and the absolute value of price changes which has been documented in numerous empirical studies of equity and futures markets; (ii) a positive relation between trading volume and subsequent stock price volatility which has been documented in intraday and interday studies; and (iii) that positive price movements on high trading volume lead, on average, to positive future price movements. These results suggest that "technical analysis" of volume-volatility information may be valuable.⁵

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 describes the economic model and the equilibrium concept. Section 4 describes the numerical methods used to compute equilibria and the accuracy of these methods. In

⁵Brown and Jennings (1989) and Grundy and McNichols (1989) developed multi-trading-period models to demonstrate that a sequence of prices could provide information that a single price could not. In both models, however, investors do not use the information contained in trading volume data. In Brown and Jennings (1989) knowledge of the first-period price and trading volume would lead to a fully revealing equilibrium and thus, there would be no additional information contained in the second-period price. In Grundy and McNichols (1989) trading volume has no informational content.

Pring (1991) and Edwards and Magee (1992) devote considerable attention to the technical analysis of volume data.

Section 5 we present some qualitative results about the nature of the information content in trading volume data and the relation between trading volume and price changes. Section 6 gives concluding remarks.

2 Related Literature

This paper builds on a considerable literature attempting to understand the relation between price movements and trading volume. In a model with competitive, risk-averse investors with private information and liquidity traders, Pflleiderer (1984) demonstrated that (i) successive price changes are negatively correlated; and (ii) trading volume is positively correlated with absolute price changes. However, the first result is a direct consequence of the presence of the liquidity traders and the second result is dominated by the non-speculative component of trade induced by the liquidity traders. Admati and Pflleiderer (1988) argued that many interesting intraday volume-volatility patterns could be explained by investors who strategically time their trades. They extended the Kyle (1985) framework by introducing two types of liquidity traders: discretionary and non-discretionary. Discretionary liquidity traders time their trades to minimize the expected cost of transacting. They showed that trading volume is concentrated because discretionary liquidity traders, and consequently informed traders, prefer to transact when the market is “thick”. Furthermore, if information acquisition is endogenous then price variability is greater in these periods of concentrated trading. Foster and Viswanathan (1993,1995) used a Kyle (1985) framework in which the conditional variance of the underlying value process depends on past information. Trading volume and price movements are correlated in the following way: large price movements occur when there are high realizations, in absolute value, of public information signals. In their model, each informed traders’ demands are increasing in the realization of the public signal; furthermore, high signals imply high conditional variance yielding greater profits to the informed traders *holding the number of informed fixed*. Consequently, in equilibrium, the number of informed traders increases and trading volume is higher.

Campbell, Grossman, and Wang (1993) developed a model in which stock price movements are due to either innovations in fundamentals or innovations in risk preferences. Therefore, a drop in the stock price could be caused by bad news about future cash flows or an increase in the risk aversion of the marginal investor. In the former case, trading volume will be low because the bad news about stock fundamentals is public information while in the latter case trading volume will be high because investors will want to reallocate risk. Consequently, they predicted that price changes accompanied by high volume will tend to be reversed. Wang (1994) developed a multi-period trading model with competitive, risk-averse investors who have asymmetric information and heterogeneous private investment opportunities. The model is structured so that investors have speculative and hedging motives for trade each period. The equilibrium volume-volatility relationship in this model depends on the relative strength of these two motives for trade.

All of these models make predictions consistent with observed volume-volatility relationships, however, they also assume that investors ignore the information content in trading volume. In another class of models, investors are endowed with different models of the economy; consequently, trade occurs because investors interpret *public* information in different ways (e.g. Varian (1989), Harris and Raviv (1993), and Kandel and Pearson (1996)). Since all information in these models is public, the rational expectations equilibrium concept is not applicable because there is no private information that investors can *learn* from equilibrium prices and/or trading volume. Furthermore, it is assumed that investors do not update their beliefs about the true model of the economy when they observe prices and/or trading volume. The no-trade theorems of Milgrom and Stokey (1982) and Tirole (1982) do not apply in these models because traders disagree about the distribution of the payoffs on the security even when they have the same information sets. These models generate many interesting predictions about the relation between price movements and trading volume. Furthermore, speculative trading can persist in these models because investors will never agree on the correct model of the economy. However, these striking results come at the considerable expense of assuming many ad hoc forms of irrationality on the part of investors.

This paper is most closely related to a recent paper by Blume, Easley, and O'Hara (1994) which extended the concept of rational expectations equilibrium to allow investors to learn from observed trading volume data. They developed a Walrasian model in which aggregate supply is fixed but equilibrium prices are "noisy" because the quality of investors' private signals is not known with certainty. Volume data is informative in this setting because prices alone do not fully reveal both the magnitude of the private signals and their precision. However, they assume that investors are myopic in order to compute equilibria in their multi-period model. In this paper, we do not introduce any form of ad hoc irrationality. Verrecchia (1994) also developed a model in which competitive market makers, in a Kyle (1985) framework, use the information contained in trading volume. If the likelihood that a trade is uninformed depends on its size, volume data provides information not contained in net demand data to distinguish informed trades from uninformed trades.

In this paper we propose an alternative model of the information contained in trading volume data. We assume that investors are uncertain about the mean *and the variance* of the risky asset's payoff; consequently, private information is summarized by two statistics: the sample mean and variance of the private signals. Prices are not fully revealing in this setting because a high price could be due to a high mean or low dispersion of the private signals. Holding the dispersion of the signals fixed, a high signal mean implies a high conditional expected payoff since it is more likely that the true asset return was drawn from a distribution with a high mean. Similarly, holding the mean of the signals fixed, high signal dispersion implies a high conditional variance since it is more likely that the true asset return was drawn from a distribution with a high variance. We show that equilibrium trading volume in our model is positively correlated with the sample variance of the private signals; intuitively, trading volume is larger when investors have different information. Thus, trading volume contains information that is not in the equilibrium price alone because it helps to distinguish between assets with high expected cash flows and high risk and assets with low expected cash flows and low risk.

3 A General Gamma-Gaussian Model with Informative Trading Volume

Grossman (1976) presented a model which has been widely used to examine problems of information and asset prices. Its essential features were exponential utility functions, Gaussian returns, common knowledge about return variance, and private information about the mean return. Grossman demonstrated that there existed a fully revealing rational expectations equilibrium. The critical special statistical feature of his model is the existence of a one-dimensional sufficient statistic for the collection of private information, allowing a single price to also be a sufficient statistic for all information. In this section, we examine a generalization of Grossman's model which allows arbitrary tastes and more general return and information structures.

We will assume three types of investors, all with different information and possibly different tastes and endowments. We assume that there are only two assets: one safe asset (a bond) which will be worth R in the final period, and one risky asset (stock) which will be worth \tilde{Z} in the final period. Investors are endowed with shares of stock and some cash and also with some private information. The economy has three periods. In period one, investors observe private information and allocate their initial wealth between stock and the bond. In period two, period-one trading volume is publicly announced and investors once again choose their optimal portfolio; however, in this period investors condition their asset demands on their private information, the sequence of prices, and period-one trading volume.⁶ The price of equity is endogenous in periods one and two; our goal is to determine how these prices evolve and depend on investors' information. In period three, investors liquidate their wealth and consume the proceeds. The total supply of equity is given by the sum of the investors' endowments; we normalize this to unity. The bond will not accumulate any value between

⁶We assume that investors do not observe contemporaneous volume data in period one because the subsequent equilibrium in our setup would be fully revealing and, consequently, there would be no trade and a trivial relationship between price movements and trading volume.

trading periods; that is, we are assuming that the two trading periods are close to each other, whereas date three is some nontrivial time in the future. Therefore, the riskless asset's price is unity in periods one and two.

The realization of \tilde{Z} and investor's information is determined in a three-stage process. First, Nature determines the distribution of $\tilde{Z} \sim N(m, \frac{1}{w})$ in a two-step fashion. It draws the precision parameter, $W = w$, with density proportional to $w^{\alpha-1} \exp(-\beta w)$, which is the Gamma distribution with parameters α and β . Given the precision value w , nature draws the mean parameter, m , from the distribution $N(\mu, \frac{1}{w})$.⁷ Second, information is gathered by investors. Type i investors observe $y_i + \epsilon_i$, where y_i is distributed $N(m, \frac{1}{w})$ and ϵ_i is distributed $N(0, \sigma_{\epsilon_i}^2)$, for $i = 1, 2, 3$. The draws of y_i and ϵ_i are all mutually independent. Thus, Nature chooses a distribution from which it will draw the final return, and investors' information consists of independent experiments modelled as draws from the true distribution chosen by nature, $N(m, \frac{1}{w})$, plus some experimental noise⁸. In the third and final stage, Nature chooses draws a realization of \tilde{Z} with distribution $N(m, \frac{1}{w})$.

This is a more general model of asymmetric information than typically used. First, if the precision w is drawn from a degenerate distribution with a mass at w_0 then $\tilde{Z} \sim N(m, \frac{1}{w_0})$ and we have a model similar to Grossman (1976). In this degenerate case, there exists a one-dimensional statistic for \tilde{Z} , which is perfectly correlated to the full information rational expectations equilibrium price, and, hence, there will be a fully revealing rational expectations equilibrium. Second, the more general model allows us to examine a richer variety of information structures, most of which will not produce fully revealing equilibria. For example, it allows us to consider asymmetric information about the variance of an asset's return. This is particularly important since investors are often uncertain about an asset's riskiness as well as its expected return. If there is asymmetric information about the mean and the variance of the asset's return in our model there does not exist a one-dimensional

⁷Multivariate versions of this model will implement multivariate specifications with a Wishart conjugate prior for the variance-covariance matrix.

⁸The choice of gamma priors for the precision and conditionally normal priors for the mean ensures that there exists a two-dimensional sufficient statistic for all private information.

sufficient statistic for all private information. Thus, a single price will not, in general, be a sufficient statistic for all private information.

We now consider investor i 's problem. The uncertain final-period consumption of a type i investor is given by:

$$\tilde{C}_i = \theta_{2i}\tilde{Z} + [W_i + p_1(\bar{\theta}_i - \theta_{1i}) + p_2(\theta_{1i} - \theta_{2i})]R \quad (1)$$

where p_j is the share price of stock in period j , θ_{ji} is the number of shares held by type i agents after the date j round of trading, $j = 1, 2$, $\bar{\theta}_i$ is the type i endowment of stock, and W_i is the type i cash endowment. The evolution of information implies that θ_{1i} can depend on only p_1 and y_i , and θ_{2i} may depend on only p_1 , y_1 , p_2 , and period one volume. The functions θ_{1i} and θ_{2i} are the trading strategies which type i investors choose.

A type i investor's period one problem is to solve

$$\max_{\theta_{1i}} E[U(\tilde{C}_i) | I_{1i}] \quad (2)$$

where a I_{1i} is type i investor's conditioning information at period one. In this paper, we use constant absolute risk aversion utility functions of the form $U(C) = -\exp^{-a_i C}$ where a_i is investor i 's coefficient of absolute risk aversion. I_{1i} includes investors i 's private signal y_i , and the price, p_1 . Since he sees only p_1 and y_i , a type i investor's first-period demand function will depend only on p_1 and y_i , and can be written $\theta_{1i}(p_1, y_i)$. In every state (p_1, y_i) , differentiation of (2) with respect to $\theta_{1i}(p_1, y_i)$ produces the first-order condition for choice of $\theta_{1i}(p_1, y_i)$

$$0 = E[U'(\tilde{C}_i) (\tilde{Z} - p_1 R) | p_1, y_i].$$

This says that the excess return should be uncorrelated with the final marginal utility consumption when conditioned on an investor's information set.

A type i investor's period two problem is to solve

$$\max_{\theta_{2i}} E[U(\tilde{C}_i) \mid I_{2i}] \quad (4)$$

where a I_{2i} is type i investor's information in period two. In the model of this paper, I_{2i} includes I_{1i} , his private signal, y_i , and the price p_1 , plus the new period two information which is the price p_2 , and first-period trading volume, V_1 . Period one volume is defined by

$$V_1 = \frac{1}{2} \sum_{i=1}^n |\theta_{1i} - \bar{\theta}_i|. \quad (5)$$

The first-order condition, conditional on I_{2i} , for the choice of θ_{2i} will be

$$0 = E[U'(\tilde{C}_i) (\tilde{Z} - p_2 R) \mid p_1, p_2, y_i, V_1]. \quad (6)$$

Again, this is the familiar condition that the excess return of equity over the riskless asset should be uncorrelated with the final marginal utility of consumption when conditioned on an investor's information set.

While this structure is rather simple, it is arbitrary in the number of investors, the distribution of \tilde{Z} , and the information allocation of investors. We can also extend the model to include a variety of assets, including derivatives. This would only involve changing the period one and period two information sets and adding Euler equations for the additional assets. The basic form of the first-order conditions in more general models would remain as in (2) and (4). While these more general models are not totally general, it will be clear that the projection methods used to compute solutions of our simple model are of general applicability.

3.1 Equilibrium Concept

We assume that there is a fixed number of shares available, and that the safe return is fixed exogenously. Equilibrium consists of several functions expressing asset prices as a

function of all the information in the economy at the time of trading, and trading rules as functions of the total information. In particular, the equilibrium concept employed in this paper is the standard concept of rational expectations equilibrium; the only unusual feature is that agents use information about volume when making inferences about asset value. Equilibrium occurs when investors' beliefs about the relationships among trading volume, private information, and asset returns are correct. More specifically, we employ the following definition of equilibrium:

Definition 1: *A rational expectations equilibrium is a collection of price functions at times $t = 1, 2$, denoted $p_t(y_1, y_2, y_3)$, asset demand functions for type i investors in period 1, $\theta_{1i}^*(p_1, y_i)$, and asset demand functions for type i investors in period 2, $\theta_{2i}^*(p_1, p_2, V_1, y_i)$, such that:*

(i) *given $p_1(y)$, $\theta_{1i}^*(p_1(y), y_i)$ solves (2);*

(ii) *given $p_1(y)$, $p_2(y)$, and $V_1(y)$, $\theta_{2i}^*(p_1(y), p_2(y), V_1(y), y_i)$ solves (4);*

(iii) *$\sum_i \theta_{1i}^*(p_1(y), y_i) = \sum_i \bar{\theta}_i$ for all y , and*

(iv) *$\sum_i \theta_{2i}^*(p_1(y), p_2(y), V_1(y), y_i) = \sum_i \bar{\theta}_i$ for all y*

where $V_1(y) = \frac{1}{2} \sum_i |\theta_{1i}^(p_1(y), y_i) - \bar{\theta}_i|$ is period one trading volume.*

Since the utility function is concave, the second-order conditions corresponding to the first-order conditions (2) and (4) are automatically satisfied as well. Note that this is a partial equilibrium model since we take the price of bonds as given, and allow an arbitrary net aggregate position in bonds. It would be trivial to make the bond price endogenous, but would cloud the discussion of some issues since the bond price would then also convey information. Since the bond price is an aggregate price and the risky asset being modelled is intended to be equity in individual firms or other assets based on a small portion of the economy, it would be unnatural to allow the bond price to respond solely to information concerning the risky asset. If our risky asset were intended to represent an aggregate market asset, we would then move to a more general equilibrium analysis. In this paper, we follow the tradition of this literature and make the bond price exogenous.

We do not offer here any existence proof. We expect that one could use the approaches

of Allen (1985a,b), and Anderson and Sonnenschein (1982), to prove existence of equilibrium for the related equilibrium concepts they offer. In this paper, we instead approximate equilibria numerically and discuss the quantitative results.

4 Numerical Methods for Computing Approximate Equilibria

4.1 Computing Conditional Expectations

The first-order conditions in (3) and (6) imply that our equilibrium concept involves a conditional expectation. Numerical implementation of the conditional expectation conditions is the most challenging aspect of this problem. We use Gaussian quadrature methods combined with basic projection ideas to implicitly compute conditional expectations.

To solve this problem, we use the following definition of conditional expectation.

Definition 2: *Assume that Y and X are random variables. The conditional expectation function relation,*

$$Z(X) = E[Y|X]$$

holds if and only if

$$E[(Z(X) - Y)f(X)] = 0$$

for all continuous bounded functions, $f(X)$, of X .

Intuitively, Definition 2 says that the prediction error of the conditional expectation, $E[Y | X]$, is uncorrelated with any continuous function of the conditioning information, X . This definition replaces the conditional expectation with an infinite number of unconditional expectation conditions. In practice, we approximate $Z(X)$ by finitely parameterizing $Z(X)$ and imposing a finite number of the unconditional expectation conditions to identify the free parameters. The details of this will be made explicit below.

We shall now use these ideas to compute asset market equilibrium in our model. Using

the definition of conditional expectation given above, the projection method, as described in Judd (1992), solves for an approximate equilibrium by finitely parameterizing $p_1(y)$, $p_2(y)$, $\theta_{1i}(p_1, y_i)$, and $\theta_{2i}(p_1, p_2, y_i, V_1)$ and imposing a finite number of the conditions implicit in our definition of equilibrium.

More precisely, we approximate the first-period and second-period price laws with the polynomial representations

$$p_1(y_1, y_2, y_3) = \sum_{0 \leq j+k+l \leq N_{p_1}} a_{jkl} H_j(y_1) H_k(y_2) H_l(y_3) \quad (7)$$

and

$$p_2(y_1, y_2, y_3) = \sum_{0 \leq j+k+l \leq N_{p_2}} b_{jkl} H_j(y_1) H_k(y_2) H_l(y_3) \quad (8)$$

where $H_i(\cdot)$ denotes the degree i Hermite polynomial and N_{p_1} represents the total degree of the polynomial approximation. Hermite polynomials are natural in this setting because returns are normally distributed and Hermite polynomials are mutually orthogonal with respect to the normal density with mean zero and variance of one half. See Judd (1992) for a discussion of the advantages of orthogonal bases in projection methods.

Another important aspect of the approximations in (7) and (8) is the restriction to a *complete polynomial* representation. The set of *complete polynomials of degree N* over R^m is defined to be

$$\mathcal{P}_N \equiv \{X_1^{i_1} \cdots X_m^{i_m} \mid \sum_{l=1}^m i_l \leq N, \quad i_l \geq 0, \quad \forall l\}.$$

We see that the polynomials in (7) and (8) are weighted linear combinations of elements in $\mathcal{P}_{N_{p_1}}$ and $\mathcal{P}_{N_{p_2}}$. An alternative is the *tensor product of degree N* over R^m :

$$\mathcal{T}_N \equiv \{X_1^{i_1} \cdots X_m^{i_m} \mid 0 \leq i_l \leq N, \quad \forall l\}.$$

The use of complete polynomials result in little loss of accuracy as compared to the full tensor

product basis but has the advantage of many fewer unknown parameters. For example, if we let $N_p = 3$ in our model we have 64 unknown coefficients in each of our price functions if we use the tensor product basis but only 20 unknown coefficients if we use the complete polynomials. (See Judd (1991) for a more detailed discussion.)

Similarly, we represent the first-period stock demand for a type i investor, $i = 1, 2, 3$, by

$$\theta_{1i}(p_1(y), y_i) = \sum_{0 \leq j+k \leq N_{\theta_1}} c_{jk}^i H_j(p_1(y)) H_k(y_i). \quad (9)$$

and we represent the second-period stock demand for a type i investor, $i = 1, 2, 3$, by

$$\theta_{2i}(p_1(y), p_2(y), y_i, V_1(y)) = \sum_{0 \leq j+k+l+m \leq N_{\theta_2}} d_{jklm}^i H_j(p_1(y)) H_k(p_2(y)) H_l(y_i) H_m(V_1(y)). \quad (10)$$

We will also need to compute the period one volume function, $V_1(y)$. When we need to do so, we will just use the definition in (5).

Our goal then is determine the unknown a_{jkl} , b_{jkl} , c_{jk}^i , d_{jklm}^i coefficients. If we let $N_{\theta_1} = N_{\theta_2} = N_{p_1} = N_{p_2} = 3$ the number of unknown coefficients is 175: twenty for each of the period one and period two price functions, ten for each of the three first-period policy functions, and thirty-five for each of the three second-period policy functions. To determine the unknown coefficients we impose projection conditions on the investors' first-order conditions and market clearing. The total number of conditions will equal the number of unknown coefficients, hoping that they are sufficient to fix the unknown coefficients. This is just one of many possible solution techniques. For example, we could impose more conditions than unknowns and choose coefficient values to minimize some squared error criterion. In this paper, our approach is to attempt to, to use econometric terms, "exactly identify" unknown coefficients by imposing an equal number of "orthogonality conditions."

The first-order-condition for a type i investor in period 1 is given by

$$E\{U'(\tilde{C}_i)(\tilde{Z} - p_1 R) \mid y_i, p_1\} = 0, \quad i = 1, 2, 3.$$

$$E_y\left\{\left(\sum_{i=1}^3 \theta_{2i}(p_1(y), p_2(y), y_i, V_1(y)) - 1\right)H_j(y_1)H_k(y_2)H_l(y_3)\right\} = 0, \quad j + k + l \leq N_{p_2} \quad (16)$$

The alternative is to define $\theta_{t3} = 1 - \theta_{t1} - \theta_{t2}$ for $t = 1, 2$. That approach has the appeal of forcing market clearing in every state of the world and reduces the number of unknowns. However, θ_{t3} , the resulting trading rule for type 3 agents, will almost surely be a full function of y , not measurable in $(p_1(y), y_3)$, their information set. Therefore, since we are focussing on the importance of information, we adopt the approach in (15) and (16). In this formulation, the market will not perfectly clear in every period. While this is not desirable, we must make a choice between market clearing and maintaining the measurability of agents' strategies in their information. We choose to sacrifice market clearing, endeavoring to make the amount of market nonclearing small, possibly motivated (but not here modelled) by inventory holding by a market maker.

The system of projection conditions, (12), (14), (15), and (16), constitute a finite nonlinear system of algebraic equations in the unknown coefficients. We have succeeded in reducing an infinite dimensional functional problem to a finite-dimensional algebraic problem. There are many ways to implement these ideas, choosing alternative bases and fitting criteria. In particular, we have not provided a proof of a solution to this nonlinear system. If there is no solution to this system, we can still use a least-squares solution to the system to compute a good approximation for the unknown coefficients. However, we examine just one approach, the Galerkin method, in detail since the point of this exercise is to demonstrate how to apply projection ideas to a familiar model.

4.2 Computational Details

Before proceeding, we should discuss some of the critical details of the numerical procedure. We first note that the projection conditions are all multidimensional integrals. The greatest computational effort lies in computing these integrals. Therefore, it is important to compute

them accurately and efficiently. To this end we use product Gaussian quadrature methods; in the appendix we present the precise details. The advantage of product Gaussian methods is their well-known high level of accuracy when integrating smooth functions, such as those in this model. Gaussian quadrature formulas are also good in terms of efficiency, but probably not the best. Even with Gaussian quadrature, solving our model strains the most powerful desktop PC's today. Extending our approach to more assets and/or more periods with product Gaussian quadrature would probably require supercomputer methods. However, there are more advanced multidimensional quadrature which are much more efficient; future work will focus on developing the application of these methods to rational expectations models.

The second detail is the method used for solving the nonlinear system. We use HYBRD⁹, a public domain FORTRAN program which implements the Powell hybrid method. The advantage of this method is that it combines a quasi-Newton method for nonlinear equations with an approach for minimizing SSR. Therefore, if the system does not have a solution, HYBRD will switch to a least-squares method and still produce an approximate solution. If one really thought that the least-squares approach was best then it would be more appropriate to immediately use a least-squares method. The advantage of HYBRD is that if it is presented a nonlinear system without a solution, it will not get hung up but instead produce an answer of some value. If HYBRD feels that there is no solution, it will inform the user of that fact. In all of our examples below, HYBRD concluded that it had found a solution to the nonlinear system. We suspect that Powell hybrid method implementations in NAG, IMSL, and other scientific programming libraries would perform as well if not better. We have tried some other nonlinear equation programs but with less satisfactory results.

A very useful trick is to compute the solution to a low-order problem and use the solution as an initial guess for a higher-order problem. For example, one might begin by computing the solution to the model using quadratic price and policy functions and then use these solutions as the initial guesses for the quadratic elements of the cubic price and policy

⁹We actually use a slightly modified version which avoids zero equality comparisons.

functions. These initial guesses will be very accurate because the extra terms that are added in the cubic approximation are, by construction, orthogonal to the quadratic terms.

4.3 A Check on the Method

We would like to have some way to check out this application of the projection method. Fortunately, there are some numerically nontrivial cases where we know the solution. Suppose that the investors have the same utility function and endowment. Then, if the information were common knowledge, each type would hold the same portfolio independent of the information. Since this fact is common knowledge, the uninformed investors will, in the period one rational expectations equilibrium, trade to this point in all states of the world, no matter what the distribution of private information (see Milgrom-Stokey for an elaboration of this fact), and the price would be the full information equilibrium price. Numerical calculation of the full information prices is a trivial calculation, involving only numerical integration which in this case, because of the smooth functions involved, will be very accurate. In the second period, there is no new information; therefore, there will be no change in price and there will be no trade. While we may know these facts, the algorithm does not “know” these facts and instead approaches the problem in the general way. Therefore, we can check our algorithm on these cases.

We used the algorithm above to compute equilibrium for cases covered by the Milgrom-Stokey theorem with a wide variety of utility function (relative risk aversion between one and five) and returns. Table 1 displays a typical example of our general Gamma-Gaussian model with common CARA tastes and normal returns. Each investor has constant absolute risk aversion of 2.5 (which implies a relative risk aversion parameter of approximately 3 given the level of consumption), and each investor begins with $1/3$ shares (total endowment is 1), and 1.0 units in cash. Investors differ only in the noisiness of their signal, $\text{VAR}(e)$. The columns labelled “Signal i ” $i=1,2,3$, denote the value of y seen by type i investors measured in terms of the standard deviation. Thus, if Signal 1 = -1 agent 1 observed a signal which was one standard deviation below his mean signal.

The entries in Table 1 compare the true, full-information, rational expectations equilibrium with the approximation computed by our methods. For example, we see that if the pattern of information was $(0,0,0)$ the approximated equilibrium price (partial-info price) in both periods 1 and 2 is 0.97275 which is correct to within 5 significant digits of the true, full-information price when we use cubic, complete polynomials. By comparing our approximations with the full information calculations, we found that this method generated the correct prices and holding strategies to within at least five significant digits as long as we used degree 3 polynomials for the pricing, and demand functions. We also found that the correct solution was found even if the initial guesses were poor, indicating the stability of the method.

These examples do not provide a proof of the validity of our method in general. However, they do give us confidence to proceed to more interesting problems.

4.4 Calibration

We now turn to examples which use the novel features of our model. Our calculations will use “sensible” values for the basic parameters. The average annual real return on equity is roughly eight per cent and the average annual real return on bonds is roughly one per cent. The standard deviation of stock returns is roughly 40% for individual stocks and 20% for large stock indexes. Since we are discussing the importance of lagged information on volume, the length of a period should be short, such as a day or, at most, a week. We choose values for μ , σ^2 , α , and β consistent with these statistics. Finally, we assume that relative risk aversion is between 0.5 and 5.

4.5 Accuracy Measures

An important issue to address when we use numerical methods is the quality of the resulting approximation and how many polynomial terms do we need to get a good approximation. One way to check an algorithm is to check it out on cases where one knows the solution. The

exercises conducted above for the Milgrom–Stokey cases showed that cubic approximations did very well. Another check is to resolve the problem using another method and compare the procedures' answers. In our case, this could be done by using different basis functions and different integration formulas. Hopefully the solutions are insensitive to these changes; if not, then one must be concerned about any of the solutions.

A second approach is to take the approximation and compute a measure of how much it violates the equilibrium conditions. In this problem, we can ask how much better can an agent do if he used more information than implicit in the computed approximation. Note that the equilibrium conditions has each agent choose decision rules which yield Euler equation residuals which are orthogonal to a restricted set of basis functions. This is, essentially, allowing him to do only a limited regression analysis of the data. We can compute the wealth equivalent of the Euler equation residual when projected in directions not used in computing the approximation. This is the *consumption error*, that is, the difference, in consumption units between following the equilibrium rule versus following a rule which uses more information in making inferences from the price. We operationalize this by taking the equilibrium law and subjecting it to a more refined regression analysis and asking how much an agent will gain if he is allowed to use the better inference rule. The results for a typical case are shown in Table 2. We found that the Euler equation residual errors are very small, approximately one in one hundred thousand parts of wealth, when projected in directions not used to approximate the equilibrium. Moreover, the Euler equation errors give us clues about the optimal number of basis functions and quadrature nodes to use in the approximation. We find that cubic approximation works extremely well, with little to be gained by moving to quartic approximation. We also find that using twice as many quadrature nodes in each dimension as the degree of approximation works very well. Thus, for cubic approximation one should use 6 or 7 quadrature nodes for each dimension of integration.

The second approach is consistent with a costly information interpretation of equilibrium. If these computational methods produce a policy function with small optimization errors, then that approximate policy function is as compelling a description of behavior as

the equilibrium policy function since it is unclear why individuals would bear the nontrivial cost of finding the “true” policy function if the gain is small.¹⁰

5 Results

In this section we report the results of a few simple examples of the types of analyses that can be conducted in this framework. We conducted three sets of experiments for (α, β) pairs shown in the figures. In all of these examples we randomly generated 1000 signal triples (y_1, y_2, y_3) using the true joint probability distribution to determine the equilibrium relation between price movements and trading volume.

Figure 1 was produced by generating 1000 signal triples (y_1, y_2, y_3) and plotting the relation between equilibrium first-period trading volume and the sample variance for each triple. The figure suggests that trading volume is positively correlated with the sample variance of the private signals. Table 3 (I) verifies this by showing the results of a linear least-squares regression of equilibrium trading volume on the sample signal variance. We report the coefficient with the t-statistic in parentheses. The coefficient on the sample signal variance is highly significant. Intuitively, trading volume increases when investors observe disperse signals because they have different beliefs about the future value of the stock. Of course, investors would not “agree to disagree” if there was only a speculative motive for trade; however, investors also have risk-sharing motives for trade in this model. As we mentioned above, private information in this model is summarized by two statistics: one corresponding to the mean and the variance of the private signal observations. Consequently, the period 1 price alone does not reveal all private information because a single price cannot disentangle the two statistics. Figure 1 demonstrates, however, that trading volume data contains information about the variance of the private signals and, thus, the tradeoff between risk and expected return that is not contained in the equilibrium price.

¹⁰There is a strong similarity between this procedure and the approach of Anderson and Sonnenschein (1982). They assume that agents run regressions and use them when making decisions. They prove existence of equilibrium when agents are restricted in the regressors they use.

Figure 2 shows the relation between trading volume and the period 1 price. Again, we randomly generated 1000 signal triples (y_1, y_2, y_3) and plotted the period 1 price against period one trading volume. We conducted three experiments corresponding to different values of (α, β) - the parameters of the distribution from which the variance of the stock payoff is drawn. Smaller values of β correspond to smaller variances of this distribution. The first thing to note is the V-shape relationship between trading volume and the period 1 price. We can also interpret this as the relationship between trading volume and *price changes* if we shift the x-axis by the hypothetical price that would have prevailed in competitive equilibrium *prior* to observing private information. Thus, large price movements tend to occur in periods of high trading volume. This theoretical prediction is also made in Blume, Easley, and O'Hara (1994), Harris and Raviv (1993), Kandel and Pearson (1996), among others, and is consistent with many studies of the volume-volatility relation in equity and futures markets.¹¹ Furthermore, the V-shape pattern is less pronounced when there is less uncertainty about the *variance* of the stock payoff. This is not surprising since in the case where there is *no* uncertainty about the variance of the stock payoff the model collapses into the Grossman (1976) model in which the equilibrium price alone reveals all private information and there is no speculative trading.

Another simple example of the type of analysis that can be conducted in our framework is to examine the relation between sequences of prices and trading volume data and subsequent movements in the stock price. In the first round of trading, investors condition their asset demands on private information and equilibrium prices. Due to the return and information structure in our model, investors do not learn all payoff-relevant information by observing this data. Prior to the second round of trading, first-period trading volume is announced and investors use this information to make better inferences about the payoff of the stock. Thus, second-period prices change to reflect this.

In Figure 3 we compare the movement in the price of the stock in period 2 to period 1 trading volume. Stock prices will move up, on average, from period 1 to period 2 because

¹¹See Karpoff (1987) for a survey of the empirical evidence. See also Gallant, Rossi, and Tauchen (1992).

investors resolve more uncertainty after observing the period 1 trading volume and the period 2 price. There are several interesting things to observe. First, the volatility of prices in period 2 is greater when the period 1 trading volume is high. When period 1 trading volume is low, prices tend to trade in a very narrow range in period 2 but when period 1 trading volume is high, price movements can be much larger. This is consistent with the intraday evidence reported in Admati and Pfleiderer (1988) and Brock and Kleidon (1992). Furthermore, the figure suggests that there is a tendency for prices to rise after periods of high trading volume.

Table 3 (II) gives the results of a linear least-squares regression of price movements from period one to period two on trading volume, period-one price, and an interaction term of volume multiplied by period-one price. The data are those generating Figure 3. The regression results demonstrate that high trading volume in period one is positively correlated with subsequent price movements and that the coefficient on trading volume is highly significant. The coefficient is not very significant economically: an increase in trading volume of 1% of the outstanding shares leads to roughly one-tenth of one percent of an increase in stock prices. Furthermore, the coefficient on the interaction term of trading volume with period-one price is positive and very significant statistically. Thus, high trading volume is a more positive signal when it occurs contemporaneously with a positive price change in period one than when it occurs contemporaneously with negative price change in period one. One note of caution is that the effect is driven in part by outliers in the sample of signal triples. The approximate price and policy functions are less accurate in the tails of the signal distributions, thus, some of this effect may be due to approximation error. Future research will conduct these experiments using a larger number of quadrature nodes to improve accuracy in the tails.

These simple examples give an indication of the potential of our model to explain some of the relations between trading volume and price changes as well as the relation between trading volume and the serial correlation of returns that have been documented in other studies. By computing more examples, sharper intuition can be developed about the precise nature of the economic value of the information contained in trading volume.

6 Conclusions

This paper uses a general numerical approach to compute rational expectations equilibrium in a multi-trading period model in which trading volume data has informational content. The approach is very general - requiring only smooth, concave utility functions, and asset return distributions and signals with smooth densities. We show a simple example of the types of analyses that can be conducted in this framework by demonstrating the relationship between trading volume and stock prices movements.

There are many interesting possibilities for future research. In particular, we could extend the model to multiple rounds of trade. To do so we would have to extend the model to introduce a recurring non-informational motive for trade to circumvent the no-trade theorems of Milgrom and Stokey (1982) and Tirole (1982). One solution is to introduce rational “noise” traders each period as in Judd and Bernardo (1994). These are fully rational traders whose tolerance for risk is unknown to all other traders. This would introduce a risk-sharing motive for trade each period and another source of noise which would prevent investors from perfectly inferring all relevant information from price and volume data.