



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.*

Simulation of factors impeding water quality trading market performance

Craig M. Smith
Dept. of Agricultural
Economics
Kansas State University
Graduate Research Assistant
304C Waters Hall
Manhattan, KS 66506
craigsmith@agecon.ksu.edu

Jeffrey M. Peterson
Dept. of Agricultural
Economics
Kansas State University
Associate Professor
216 Waters Hall
Manhattan, KS 66506
jpeters@ksu.edu

John C. Leatherman
Dept. of Agricultural
Economics
Kansas State University
Professor
331G Waters Hall
Manhattan, KS 66506
jleather@k-state.edu

***Selected Paper prepared for presentation at the Agricultural & Applied Economics
Association 2010 AAEA, CAES, & WAEA Joint Annual Meeting, Denver, Colorado, July
25-27, 2010***

*Copyright 2010 by Smith, Peterson, and Leatherman. All rights reserved. Readers may make verbatim
copies of this document for non-commercial purposes by any means, provided that this copyright notice
appears on all such copies.*

Introduction

Environmental economists have argued that pollution trading programs are an efficient means of improving environmental quality, as they give firms with the lowest pollution control costs the largest incentive to reduce pollution. Such low-cost firms are able to sell pollution credits to firms with higher control costs. Aligning incentives with control costs is the condition needed to ensure minimum-cost control of pollution overall. Such incentives typically do not arise from traditional, uniform regulations.

Following on the highly successful trading programs for air emissions such as sulfur dioxide (NCEE 2001), many states have recently adopted trading programs to improve water quality. There are at least 37 WQT programs currently active or under development in over 20 states (Faeth 2000; Woodward 2002; Breetz 2004). In principle, such programs could be applied to any water-borne pollutant and allow trading among point sources, among nonpoint sources, or between point and nonpoint sources (the latter is known as 'point-nonpoint trading'). Most of the existing programs are designed with point-nonpoint trading to limit nutrient loading: point sources are allowed to meet their nutrient emission limits by purchasing water quality credits from agricultural producers in the surrounding watershed. These producers are then obligated to implement a best management practice (BMP) that reduces expected nutrient loading by an amount commensurate with the number of credits sold.

Substantial evidence exists that nonpoint sources can reduce nutrient loading at a much lower cost than point source polluters in many watersheds, suggesting substantial scope and gains from point-nonpoint trading (Faeth 2000). Despite the potential gains, perhaps the most commonly noted feature of existing programs is low trading volume; none of the programs have had extensive trading activity and many have had no trading at all (Hoag 1997). A widely cited and vivid example is the Fox River program in Wisconsin (Hahn 1989), which had only one trade after its inception in 1981 even though an early study (ONEil 1983) found substantial gains from trading among all participating firms.

These outcomes suggest the presence of obstacles to trading that were not recognized in the design of existing programs. While these obstacles have not been researched in a systematic way, individual studies have identified various trading barriers in different contexts. These barriers can be divided into four groups. First, potential participants in a WQT market may be deterred by hidden or intangible costs, such as the perceived risks of future regulation (King and Kuch 2003), or public relations problems from "buying the right to pollute" (Hoag 1997). Trades also may be prohibited by high transactions costs, which would include the expenses for searching for trading partners and bargaining over price and other trading terms (Stavins 1995). Transactions costs are likely to be high in WQT programs that rely on a bilateral negotiation to generate trades instead of a centralized exchange (Woodward 2002). Limited or low information levels among market participants regarding the knowledge of each others' bid prices is another potential inefficiency discussed in the literature (e.g., Ermoliev et al. 2000). Finally, many existing programs set high "trading ratios" between nonpoint and point sources to adjust for the greater uncertainty in nonpoint loading reduction (EPA 1996). However, such ratios operate like a tax to dampen the benefits from trading, hence theoretically reducing trading volume (Malik 1993, Horan 2001).

To examine the ways that various market imperfections may impact the performance of a WQT market, an agent-based model was constructed which simulated a hypothetical point-nonpoint market. In particular, the market was modeled using a variant of the sequential, bilateral trading algorithm proposed by Atkinson and Tietenberg (1991). This paper first presents an overview of the simulation modeling technique and then analyzes the effects of two prominent market impediments identified in the WQT literature: information levels and trading ratios.

Relevant literature

While WQT has been promoted by economists as a way to cost-effectively achieve water quality goals, experience with actual WQT programs has yet to produce these results. To better understand the factors impeding trading, much theoretical work has been done. But very few articles have actually simulated an environmental trading market in action and only a small number of these have focused on water quality trading. Two articles in particular have performed trading simulations and are very applicable to this paper. An often cited article in the environmental markets literature is that of Atkinson and Tietenberg (1991) where a sulfur dioxide trading market was simulated. Netusil and Braden (2001) followed later with a simulation of a water quality market with varying transaction costs. There also are several relevant articles which addressed the effects of a trading ratio including Horan (2001) and Horan and Shortle (2005).

There are two notable studies in particular which simulated environmental trading markets using sequential, bilateral trading processes.

Atkinson and Tietenberg examined the bubble policy of the Emissions Trading Program. This well-known and often cited article attempted to explain the divergence in costs between the least cost solution and incentive based emissions trading approaches in air quality. More specifically, the article examined the hypothesis that a sequential, bilateral process cannot achieve a cost effective equilibrium in markets dealing with non-uniformly mixed pollutants.

The data used in this study came from 27 point sources in the St. Louis Air Quality Control Region. The pollution transfer coefficients were a function of its location, stack height, average mixing height, stack exit conditions, stability wind rose and pollutant decay rates. Each source's control costs were a function of the temperature and volume of the effluent gas stream, type and efficiency of existing controls, fuel requirements and maximum process rate. Seven different control technologies were examined: wet scrubbers, mechanical collectors, electrostatic precipitators, mist eliminators, fabric filters, afterburners and fuel substitution. The total annual costs were a summation of the annual capital and installation costs and the annual operating and maintenance costs. The potential cost savings were calculated as the difference between the pre-trade equilibrium costs and the post-trade costs.

Four trading scenarios were simulated. Simultaneous trading with full information was the first scenario modeled. This most closely mimicked the least cost solution. The second scenario consisted of sequential trading with full information. That is, firms were assumed to have complete knowledge of each other's control costs, so that trades occurred in the order of gains from exchange – i.e., the first

trade was between the two traders that had the most to gain from a transaction. The last two scenarios were both sequential, but they operated under differing forms of partial information. The first partial information scenario began by finding the source with the lowest marginal control cost. The best possible trading partner was found and the trade was consummated. Then the next lowest marginal control cost firm was identified and the “best” trading partner was found and the trade was executed. This process continued until no trades remained. In the second partial information scenario, a firm was randomly selected and a “best” trading partner was found. This random selection process was continued until no trades remained. This algorithm was run 500 times for each air quality standard. In all of the scenarios the air quality standards had to be met.

The results showed that more stringent standards resulted in greater divergences from the least cost benchmark, for all of the scenarios. The authors also concluded that the amount of information available and the sequencing of trades played a large role in the amount of cost savings realized. They thought that the most realistic scenario should be found somewhere in between the full information, sequential trading scenario and the random partial information scenario (thus achieving anywhere between 7% to 88% of the least cost benchmark). They do admit, however, that their cost savings results may be too optimistic due to not accounting for transaction costs. They also suggested that a market for uniformly mixed pollutants may come closer to achieving the least cost benchmark.

Netusil and Braden built upon Atkinson and Tietenberg (1991) and extended their previous work in the area of transferable discharge permits. This is one of only a few studies that examined markets for water quality. The authors examined the effects of sequential bilateral trading under imperfect information in a hypothetical sediment loading market. Their model allowed market participants to make multiple trades as opposed to a single trade.

The paper also incorporated different levels of transaction costs into each trade. This is one of a few studies that included these costs. Another unique issue that the authors addressed is lumpy abatement technologies. Trading does not always result in perfectly divisible transactions. This effectively simulated how a market would function in the real world, since the quantity supplied by a given trader does not always equal the quantity demanded by a given trading partner.

The data used in this analysis came from a 1,064 acre watershed area in Macon County, Illinois. Modeling was performed using a gain-ranked (high information) and a random (low information) contracting scenario. This is consistent with Atkinson and Tietenberg (1991). In each scenario, the number of internal and external contracts was computed. Internal contracts were defined as trades between two sites under common ownership and external contracts were trades between two separate entities.

The results showed that under the gain-ranked scenario, the sediment load under all transaction costs levels were lower than the regulatory policy’s requirement. Another important finding in this scenario was that the distribution of internal and external contracts changed as the transaction costs levels changed. High transaction costs resulted in a decrease in overall trading and caused a shift towards internal contracting.

A very interesting finding was that as transaction costs increased, the overall spending on abatement activities can sometimes decrease. The reasoning is that high transaction costs block low value contracts from occurring and allow the higher value trades to happen. Under random contracting (low information), however, an increase in transaction costs always resulted in an increase abatement and total costs.

The conclusions drawn were that neither trading scenario matched the least cost solution. This is because the least cost solution allows for simultaneous multilateral reallocations. Nevertheless, it is important to note that both trading scenarios resulted in substantial cost savings relative to the regulatory approach even at the highest transaction costs level and lowest information level.

Horan (2001) and Horan and Shortle (2005) analyzed different levels of trading ratios in the context of water quality trading. Horan (2001) presented trading ratios utilized in several existing, pilot, and planned point- and nonpoint-source trading markets. These ranged from 1.3:1 to 3:1. Horan and Shortle (2005) performed a numerical example of trading in Susquehanna River Basin and arrived at “optimal” trading ratios in the range of 0.89:1 to 3.3:1

Horan (2001) argued that from an economically optimal standpoint, the optimal trading ratio would necessarily be less than one when a WQT model is specified to have uniformly mixed pollutant loads, stochastic nonpoint loads, convex damages, and no transactions costs. This is because the variability in nonpoint loadings creates stochastic ambient pollution concentrations and stochastic damages from pollution. This leads to more social risk if damages are convex in ambient pollution and if increases in nonpoint loadings increase the variability of ambient pollution. Social risk is costly, so there are more benefits to reducing the variable nonpoint source pollution. Thus, smaller trading ratios are more economically efficient.

But, Horan (2001) suggested that it is realistic to assume policies are designed to allocate resources with the context of policy makers’ preferences, not to maximize aggregate economic surplus. Thus, trading ratios are designed to be politically optimal. He further argued that trading ratios in excess of one may be the rational public sector response to the risk associated with stochastic nonpoint pollution. Thus, trading ratios must be greater than one for most trading programs to be politically viable.

Conceptual model


As with most markets, WQT markets can suffer from various imperfections and frictions which tend to hinder trading and/or reduce the overall gains from trading. This section will utilize classic demand and supply diagrams to assess the impact of different market imperfections.

Frictionless market

As a point of comparison, the equilibrium of a frictionless market with no imperfections is discussed first, represented in Figure 1. The demand curve in this figure represents treatment plants’ willingness to pay (WTP) for purchasing credits, reflecting the cost of controlling pollution through technology upgrades. The supply curve represents farmers’ cost of pollution control through best management practices, which is their willingness to accept (WTA) to sell credits. When $Q = 0$ credits, treatment plants

are meeting their limits by controlling all of their pollution through their own facility upgrades or technological improvements. As Q increases, plants are buying credits to allow more of the pollution to be controlled by the nonpoint sources. Thus, at any point on the diagram the total amount of pollution control does not change; however, the sources responsible for the pollution control does change. Stated differently, the quantity of trades has zero effect on expected water quality.

Figure 1 Frictionless water quality trading market



This frictionless market condition assumes there are no intangible or transaction costs, and also that the trading ratio is 1:1. In the equilibrium of this market, point sources purchase Q^* credits from nonpoint sources at a price of P^* . Area A represents the market gains to point sources, reflecting the difference between the potential cost of technology upgrades (points along the demand curve) and the actual cost of purchased credits (the price P^*). Area B is the gain to nonpoint sources, or the price received for the credits sold (P^*) less the cost of generating those credits (points along S). The sum of these two areas is equal to total benefits or total cost savings from the program. Cost savings are maximized under these frictionless or “perfect” market conditions.


It is important to note that the areas delineating the gains to point source and nonpoint sources in the figure assume that every contract is traded at the equilibrium price, P^* . This would only occur under a simultaneous trading scenario. However, the way water quality markets are designed, trading must occur in a sequential and bilateral fashion. So, each contract results in a potentially unique price. Acknowledging this would change the individual values of the point- and nonpoint-source gains, but the

total cost savings (sum of the two gains) would not vary. This limitation is true for all of the following market scenarios.

Trading ratios

Figure 2 displays the impact of imposing a trading ratio on an otherwise frictionless market. As explained above, maximum efficiency is achieved by a 1:1 trading ratio. Imposing a 2:1 trading ratio affects the nonpoint sources or the suppliers in the market. They must reduce nutrient loading by two pounds in order to receive one tradeable credit. This essentially doubles the price of all credits sold, resulting in the steeper supply curve shown in Figure 2.

Figure 2 Effects of the trading ratio on market performance



The quantity of credits traded reduces to Q' and the equilibrium price of credits increases to P' . The gains to point sources with the 2:1 trading ratio is area A , compared to area $A + B + C + D$ in the efficient market. Thus, raising the trading ratio to 2:1 induces a loss to point sources of area $B + C + D$. The gains to nonpoint sources with a 2:1 trading ratio is area $B + C + E + F$, compared to area $E + F + G$ in the efficient market. Thus, the net effect of the 2:1 trading ratio to the nonpoint sources is equal to $B + C - G$. If $B + C$ is bigger than G , then the nonpoint sources benefit from the higher trading ratio. The change in total cost savings from the higher trading ratio is equal to a loss of area $D + G$.

Unlike the frictionless market, expected loading in this case does respond to changes in the volume of credit trades. Because nonpoint traders must reduce loading by 2 pounds for every 1 pound emitted by point source traders, there will be a net reduction of 1 pound of expected loading for each trade. For the equilibrium depicted in the diagram, point sources increase their expected loading by Q' pounds, while

nonpoint sources reduce expected loading by $2Q'$ pounds. This implies a net reduction in expected loading equal to Q' pounds.


Information levels

Another type of imperfection considered is limited information, which impacts the sequencing of trades. A frictionless market presumes full information, where every participant in the market knows precisely the willingness-to-pay and willingness-to-accept of all potential traders. In this situation, the trades would be executed in order of their market gains: the first trade would be between the buyer with the highest WTP and the seller with the lowest WTA, with successive trades yielding progressively narrower gaps between WTP and WTA until all gains have been exhausted at the equilibrium point. In a low information scenario, participants have little or no information on other traders' WTP or WTA values. In this limiting case, buyers and sellers would be paired together in a random order uncorrelated to the gains from trading.

Ermoliev et al. (2000) actually proved that random-ordered, sequential trading can lead to an efficient outcome (Q^* in Figure 3). However, this can only occur when every participant has the ability to be a buyer and a seller and there are no transaction costs. That is, traders can back out of earlier trades at no penalty if they find a new trading partner that is more advantageous. This assumption is unlikely to hold for water quality trading programs in practice, where each trade usually involves a binding contract that can only be breached at some financial penalty. All of the models in this paper operate under the assumptions that only point sources are able to buy credits, only nonpoint sources can sell credits, and that the penalties for breaching trade contracts are prohibitively large. Since Ermoliev et al.'s (2000) assumptions are not met in these models, different information levels should result in different levels of cost savings.

Figure 3 shows the effects of different information levels in the market. For this example, the focus will only be on the point sources located at points 1 and 3 along the demand curve (hereafter plant 1 and plant 3), and the nonpoint sources located at points 2 and 4 along the supply curve (hereafter farm 2 and farm 4). For simplicity, let us assume that all four of these entities would trade at most one credit. As in any market, the net gain from a given trade is equal to the difference between the price along the demand curve and the price along the supply curve. In a full information and frictionless market, the first transaction involving any of these traders would be between plant 1 and farm 4. Plant 3 and farm 2 will not engage in trading because there would be a negative net gain from doing so. So, for the four traders combined, the net gain from trading under full information is $P_1 - P_4$.

Figure 3 Effects of information on market performance




A low information scenario, on the other hand, has the potential to result in different net gains (theoretically, it also has the potential to result in the same net gains). Suppose plant 1 trades with farm 2. The resulting net gain from this transaction is $P_1 - P_2$. Suppose also that plant 3 trades with farm 4 for a net gain of $P_3 - P_4$. The combined net gain from this sequence of trading is $(P_1 - P_2) + (P_3 - P_4) = (P_1 - P_4) - (P_2 - P_3)$. So, assuming that all other traders are paired the same as the full information scenario, this “ill-ordering” of trades would reduce the overall market gains by $P_2 - P_3$. This suggests that lower information is likely to increase trading volume while reducing the total gains from trading. However, whether point sources or nonpoint sources gain or lose from less information depends on the order of trading that is realized and cannot be unambiguously predicted.

Co-effects of trading ratios and information levels

The last type of market imperfection covered in this paper are the co-effects of trading ratios with low information levels. The combinatorial effects of these two imperfections can be seen graphically in Figure 4. The overall gains from trading should decrease when a trading ratio greater than one is introduced to marketplace, regardless of the information levels. However, the market will not necessarily become less cost-effective when there is a larger trading ratio and lower information. This is because as the trading ratio increases to 2:1, the potential for “ill-ordered” trading decreases as the number of potential sellers declines. The relative lengths of line segments F_1 and F_2 demonstrate this decline. Thus, in the case of higher trading ratios the information levels are not as crucial because the potential for “ill-ordered” trading decreases.

Figure 4 Co-effects of trading ratio and information levels



Conceptual model summary

The main predictions of theory are that the lack of information increase the quantity of credits traded and reduce overall cost savings. The trading ratio will also reduce overall cost savings but will be negatively related to the quantity of credits traded. However, several relationships cannot be resolved from theory alone, nor can the magnitude of any of the impacts be assessed. These items will be determined in the empirical simulation analysis in the following sections. Table 1 summarizes the predictions made.

Table 1 Summary of predictions from market analysis

Type of Market Imperfection	Trading Volume	Producer Surplus	Cost Savings	Expected Loading
Trading ratio	-	?	-	-
Lack of information	+	?	-	0
Trading ratio & low information	?	?	-	-

Simulation Model

An agent-based model (ABM) was created to simulate a hypothetical point-nonpoint source market. ABM's have been increasingly applied to study the micro-level decision making that occurs in complex systems (Tsfatsion 2006). This simulation was one in which all wastewater treatment plants were required to meet proposed stricter limit of nutrient concentrations in their discharge stream. Plants

could either upgrade their technology to meet this limit, or else keep their old technology and buy water quality credits to offset their excess discharges. Such a regulatory driver is necessary for the market to function.

The model relies on user inputs of willingness-to-pay (WTP) for purchasing credits by each of the wastewater treatment plants and willingness-to-accept (WTA) for selling credits by each farmer who is a potential trader. A sequential, bilateral trading algorithm (Atkinson and Tietenberg 1991) then simulates market outcomes from these base data. As described in the subsections below, the effects of interest were captured either by varying the input data or by altering the assumptions in the trading algorithm that govern how buyers and sellers are paired together.

ABMs require specification of two types of computational objects: the agents themselves and the environment in which they operate (Parker, Berger, and Manson 2002).

Agents

The agents in this model were point- and nonpoint-sources of a water contaminant (e.g., nutrients). To create model agents, costs and quantities were generated for each of $I = 10$ point sources and $J = 500$ nonpoint sources using random draws from independent lognormal distributions. The lognormal distribution was chosen to allow for the well-documented skewness in the distribution of costs and environmental impacts across the population of polluters (Nowak, Bowen, and Cabot 2006). The parameter values of the lognormal distributions for both buyers and sellers are shown in Table 2. The distributional parameters and the population sizes were chosen to be roughly in line with the data used by Smith (2004) to model phosphorous trading in the Middle Kansas subbasin.¹

Table 2 Lognormal distribution parameters for buyers and sellers

Item	Mean	Standard deviation
Buyer quantities (lbs)	5,000	1,250
Buyer costs (\$/lbs)	20	15
Seller quantities (lbs)	200	50
Seller costs (\$/lbs)	15	8

Environment

The environment is the trading mechanism that determines how buyers and sellers are paired together in the water quality trading market. The trading mechanism used here is a variant of the sequential, bilateral trading algorithm proposed by Atkinson and Tietenberg (1991). Two possible market impediments were modeled: trading ratios and lack of information.

¹ To ensure that the final results were not sensitive to a particular set of draws from the lognormal distributions, all simulations were repeated 10,000 times in Monte Carlo fashion, with a new set of prices and quantities assigned to all agents each time. Using the same model calibrated to data from the Lower Kansas River watershed, Peterson et al. (2009) found that 10,000 repetitions was sufficient to ensure that the mean market performance measures computed across the 10,000 repetitions was a stable statistic.

The marginal gains matrix and the trading ratio

The WTP and WTA data were derived from the cost information presented above. Assuming no transactions or intangible costs, the WTP and WTA values were exactly equal to the costs. These values were used to form the core element of the simulation model, the marginal gains matrix. This matrix contains the potential gains from each possible pairing of the farmers and treatment plants. The rows of this matrix correspond to treatment plants while its columns correspond to farmers. In simulation s , the cell in row i and column j of this matrix is

$$MarGain_{s,i,j} = WTP_{s,i} - t_s WTA_{s,j} \quad (11)$$

where t_s is the assumed trading ratio in simulation s (expressed as the number of credits a farmer must sell to offset one unit of plant discharge) and $MarGains$ is the mutual gain if plant i buys one more credit from farm j under the assumptions embedded in simulation s .

A related matrix, Q has the same dimensions and tracks the quantity of credits available for trade between each trading partner. If traders i and j agree on a trade of q credits, the entry in that cell is reduced by q . The quantity data consisted of the amount of credits demanded from the point sources and the number of credits supplied by the nonpoint sources. The individual sources were removed from the market when their quantity demanded (for point sources) or quantity supplied (for nonpoint sources) of credits equaled zero.

The trading algorithm and information levels

The effect of marketplace information was captured by varying the assumptions in the sequential, bilateral trading algorithm that pairs buyers and sellers together in a specific order. Four possible information scenarios were modeled, described in turn below.

1. Full information or marginal-gains-ranked trading

This scenario assumed that every point source and every nonpoint source in the watershed knew precisely all the WTP and WTA values of all traders. In this situation, the most advantageous trades would be executed first. Action began by the plant with the highest WTP trading with the farmer having the lowest WTA. This was determined by the element in the marginal gains matrix that exhibited the greatest positive value.

The point source would purchase as many credits as it needed or until it bought out the nonpoint source, whichever occurred first. The quantity data and the marginal gains matrix were both updated accordingly when the trade was consummated.

The second trade began by finding the greatest positive number in the updated marginal gains matrix. This determined the next two trading partners. The aforementioned process was then run again. This marginal gain-ranked process continued until there were no more gains to be made by trading.

2. Low information trading

The second scenario presumes low information, in which none of the stakeholders knew their own or anyone else's nutrient control costs. Therefore, the trades occurred in a completely random order. The only restriction was that only trades resulting in positive gains were eligible to be chosen. A single element from the marginal gains matrix was chosen at random and this determined the trading partners. The trade was then consummated and the marginal gains matrix and quantity data were updated. Subsequent trades operated in the same random fashion. Trading continued until no potential positive gains remained.

3. Partial information trading: WTP known

The third scenario modeled the case where the treatment plants' WTP values were known to all traders, but farmers' WTA values were unknown. This depicts a situation in which the treatment plants drive the market. Trading began by choosing the plant with the highest WTP. The nonpoint source trading partner was then chosen at random. The prices were determined by the same method as before, and after the trade was consummated, the marginal gains matrix and quantity data were updated. This process continued until no potential positive gains remained from trading.

4. Partial information trading: WTA known

The fourth scenario modeled the case where the treatment plants' WTP values were unknown, but farmers' WTA were known to all traders. This scenario assumed that farmers drive the market. Trading began by first choosing the nonpoint source with the lowest WTA. The point source trading partner was then chosen at random. The prices were determined by the same method as before, and after the trade was consummated, the marginal gains matrix and quantity data were updated. This process continued until no potential positive gains remained from trading.

The simulation experiments

Table 3 displays the varying assumptions in each of the 14 simulation experiments conducted. Simulations 1a-1f assumed that all market participants had perfect, full information regarding others' WTP and WTA values. On the other hand, simulations 2a-2f assumed that there was no information known regarding these values. A comparison of simulation sets 1 and 2 will reveal the effect of full information on market performance. If market performance changes significantly between these two cases, simulations 3 and 4 will illuminate the separate effects of marketplace information on WTP or WTA under "partial" information scenarios.

For both the full and zero information level simulations, differing levels of the trading ratio were included. The trading ratio varied from a low of 0.5:1 in the 'a' scenarios up to 3:1 in 'f' scenarios. The incremental step length was 0.5. The 'a' scenarios, with a trading ratio of less than one, were included based on the work of Horan (2001). Comparing simulations 1b, 2b, 3, and 4 will give insight into the impact of the trading ratio on the market along with its interaction effects with respect to information levels.

Table 3 The simulation experiments

Simulation	Trading Ratio	Information Level
1a	0.5:1	Full
1b	1:1	Full
1c	1.5:1	Full
1d	2:1	Full
1e	2.5:1	Full
1f	3:1	Full
2a	0.5:1	Zero
2b	1:1	Zero
2c	1.5:1	Zero
2d	2:1	Zero
2e	2.5:1	Zero
2f	3:1	Zero
3	1:1	WTP known
4	1:1	WTA known

Simulation Results

Before presenting the simulation results, two matters of interpretation should be noted. First, in evaluating the performance of the WQT market, comparisons are made back to a baseline situation in which treatment plants would be required to meet the nutrient reduction limit by upgrading technology. Based on the information on the plants in this hypothetical watershed, the limits would require the plants to reduce their annual nutrient load by a combined (expected value) 50,000 pounds annually. The expected total annual cost of these technology upgrades would be \$1.0 million. These two values form a baseline for comparing market outcomes. As trades occur in a WQT market, the same loading reduction is achieved but an increasing share of loading reduction is obtained from farmers instead of treatment plants. Trading also will reduce the overall cost of achieving the target. Therefore, cost savings can be expressed both in dollar terms and as a percentage of the baseline costs. Likewise, trading volume can be expressed as the number of credits traded (measured in the pounds of loading reduction borne by farmers) or as a percentage of the loading reduction target.

Second, the gains from trading (geometrically, the area between the WTA and WTP curves) is equivalent to the cost savings to *society* from trading. A portion of these cost savings would be a gain to the treatment plants, to the extent that their credit purchases are less costly than the technology upgrades would have been. The remaining portion would be a benefit to farmers, to the extent that credit revenue is larger than their costs of the BMP. However, these simulations make no attempt at partitioning the total cost savings into the benefits to the two groups. The relative sizes of the gains would depend on the actual credit prices, which would vary across transactions and would depend on the relative negotiating power of the two groups. Lacking any reliable means to estimate the relative

bargaining power and contract prices, estimates of the gains to the two groups could only be obtained by making arbitrary assumptions.

Table 4 summarizes the results of the simulation experiments. The first column serves as a cross reference for the simulation inputs and assumptions delineated in Table 3. The second column reports the number of trades, which is approximately equal to the number of farmers who participate; an individual farmer occasionally participates in more than one trade but this is rare. Thus, on the whole, the simulations indicated that 120-300 farmers would participate depending on how the market is structured. From the 500 potential sellers of credits, this would reflect participation rates ranging from about 24% to 60%, which on the top end, is slightly higher than observed participation rates in new conservation programs (Peterson et al. 2009; Smith et al. 2007).

The third through seventh columns report trading volume, both in terms of the number of credits traded and the loading reductions by type and source. Simulated trading volume varied widely across simulations, ranging from about 6,000 credits to just over 48,000 credits. But, of most importance are the loading reductions which took into account the trading ratios. The next section will cover these results in more detail, but overall the outcomes of WQT are sensitive to the structuring of the market.

The next two columns report the cost savings, as a total in dollars and as a percentage of the baseline total costs (\$1 million). Simulated cost savings also varied widely, ranging from about \$60,000 to \$700,000 or from about 6% to 70% of baseline costs. Again, the potential cost savings varied substantially under the various market structures.

Table 4 Simulation results

Simulation	Number of Trades	Credits Traded	Volume Traded				Cost Savings		Final Costs	
			Base Loading Reduction by NPS (lbs)	Loading Reduction by PS (lbs)	Additional Loading Reduction by NPS (lbs)	Total Loading Reduction (lbs)	Total (\$)	Percent (%)	Total (\$)	Avg. (\$/lb)
1a	134	48,177	24,088	1,823	0	25,912	701,703	70.2	298,297	11.51
1b	227	40,816	40,816	9,184	0	50,000	497,161	49.7	502,839	10.06
1c	224	26,396	26,396	23,604	13,198	63,198	319,788	32.0	680,212	10.76
1d	188	16,210	16,210	33,790	16,210	66,210	205,138	20.5	794,862	12.01
1e	151	9,993	9,993	40,007	14,990	64,990	134,540	13.5	865,460	13.32
1f	121	6,386	6,386	43,614	12,772	62,772	90,824	9.1	909,176	14.48
2a	134	48,223	24,112	1,777	0	25,888	597,226	59.7	402,774	15.56
2b	255	47,156	47,156	2,844	0	50,000	392,259	39.2	607,741	12.15
2c	301	36,035	36,035	13,965	18,017	68,017	234,861	23.5	765,139	11.25
2d	258	22,515	22,515	27,485	22,515	72,515	142,475	14.2	857,525	11.83
2e	203	13,613	13,613	36,387	20,419	70,419	90,259	9.0	909,741	12.92
2f	154	8,356	8,356	41,644	16,712	66,712	59,869	6.0	940,131	14.09
3	252	46,330	46,330	3,670	0	50,000	406,054	40.6	593,946	11.88

4	243	43,916	43,916	6,084	0	50,000	479,847	48.0	520,153	10.40
---	-----	--------	--------	-------	---	--------	---------	------	---------	-------

The last two columns report the final or post-trading costs. Due to the different trading ratios, some of the simulations exactly achieved the loading reduction target while others were either below or above the target level. The next-to-last column was computed simply as the baseline (or pre-trading) costs less the cost savings from trading (e.g., in simulation 1a, \$1,000,000 - \$701,703 = \$298,297), while the last column expresses the final cost in average or per-unit terms (in simulation 1a, \$298,297/25,912 lbs = \$11.51/lbs of loading reduction). The last column provides a useful comparison across the simulations in terms of cost effectiveness. With no trading, the cost per unit of loading reduction is \$1,000,000/50,000 lbs = \$20.00/lbs. With trading, this cost ranged from \$10.06/lbs to \$15.56/lbs, so as expected, trading should only occur at a lower average cost than the no-trade option.


Information levels

The effect of marketplace information on overall cost savings was unambiguously positive. This can be illustrated by comparing simulation 1a (full information), which resulted in net cost savings of \$701,703 to simulation 2a (zero information), which resulted in savings of only \$597,226. This relationship between full and zero information held for every scenario modeled regardless of trading ratios. These results were expected and are similar to the findings of Atkinson and Tietenberg (1991).

When the gains per trade are depicted graphically, the effects of information levels on market performance become more pronounced. Figure 5 graphs the gains per trade under different information levels with a 1:1 trading ratio. Simulation 1b ended at \$497,161 of total gains. This level of gains was only reached after 227 trades had been consummated. Simulation 2b, on the other hand, reached a maximum of \$392,259, but did so after 255 trades. In simulation 1b, trading could have ceased after 118 trades and more gains would have been realized (\$393,724) than the total for simulation 2b. Similarly, if trading was halted after 118 trades in simulation 2b, only \$188,623 (48% of its final value) of gains would be realized.

Figure 5 also reveals the different effects of one side having full information and the other side having low information. When traders are informed of buyers' prices (WTP values - simulation 3), the cumulative cost savings curve behaves very similarly to the full information case (simulation 1b) across the early trades. Whereas simulation 4 (sellers' WTA prices are known) behaves similarly to the low information case. Simulation 3 resulted in more cost savings than simulation 4 across the first 65% of trades. These results imply that if market designers feel that only a limited number of trades will be consummated, creating an institution that provides accessible information about buyers' prices is preferred to providing information about sellers' prices. Of course, this still depends on other factors such as the cost of providing this information to market participants.

Figure 5 Effects of market place information on cost savings with a 1:1 trading ratio



The effect of marketplace information on cost-effectiveness, however, was not unambiguously positive or negative - it depended on the trading ratio. When the trading ratio was below 1.5:1, the effect of information on average cost-effectiveness was positive. On the other hand, when the trading ratio was 2:1 or greater, information levels were negatively related to average cost-effectiveness. That is, the more information available to market participants, the less cost-effective the trading. Table 5 shows this ambiguous effect.

Table 5 Effects of information levels on cost-effectiveness across different trading ratios


Simulations for Comparison	Trading Ratio	Difference in Average Cost-Effectiveness (\$/lb)	More Cost-Effective Simulation?	Conclusions
1a & 2a	0.5:1	-4.05	1a	← More information is more cost-effective
1b & 2b	1:1	-2.10	1b	
1c & 2c	1.5:1	-0.49	1c	
1d & 2d	2:1	0.18	2d	← Less information is more cost-effective
1e & 2e	2.5:1	0.40	2e	
1f & 2f	3:1	0.39	2f	

Trading ratio

As expected, there was a negative relationship between the trading ratio and potential gains from trading. Focusing on the '1' simulations, the cost savings ranged from just over \$700,000 down to \$90,000 when the trading ratio increased from 0.5:1 to 3:1. But, this can be somewhat misleading because each of these simulations resulted in a different amount of nutrient loading reduction. In the case of a 0.5:1 trading ratio, the nutrient target was not met. And in the case of a 2:1 trading ratio (simulation 1d), there were an additional 16,210 pounds of nutrient reduction beyond the target.


As stated earlier, the purpose of a trading ratio greater than one is to account for nutrient reduction uncertainty and ensure that there is an overall increase in water quality (beyond that which would occur in the absence of WQT and reliance on only technology upgrades). According to the simulation results, this was generally the case. Figure 6 illustrates the trading volume and net environmental gains in the different simulations. The height of the red bars represents the amount of loading reduction transferred from point sources to nonpoint sources through trading. The green bars represent the amount of loading reduction achieved from necessary upgrading of wastewater treatment plants. In cases of a trading ratio greater than 1:1 (i.e., simulations 1c-1f and 2c-2f), there were additional loading reductions achieved beyond the target. With a 2:1 trading ratio for example, each unit of increased plant loadings is offset by a 2-pound reduction in expected loading by farmers, resulting in net environmental gains equal to the height of the blue bars. Simulation 1d resulted in 16,210 credits traded. Because of the 2:1 trading ratio, nonpoint sources potentially reduced loading by a total of 32,420 pounds (2*16,210), so combining this with the 33,790 pounds of reduction achieved from WWTP upgrades the total expected loading reduction amounted to 66,210 pounds. So, the introduction of a trading ratio greater than 1:1 did result in a net environmental improvement (more than 50,000 lbs reduced) as predicted by the theory.

Figure 6 Trading volume and additional loading reduction by simulation



Because of reasons stated earlier, the most useful metric for evaluation and comparison across scenarios may be the average cost of nutrient reductions. In regards to cost-effectiveness, the effect of the trading ratio was not independent of the information levels. This can be seen graphically by analyzing Figure 7 and Figure 8. Under full information, the trading ratio had an immediate and persistent damaging effect on cost-effectiveness from trading. Increasing the trading ratio from 1:1 to 2:1 (Simulations 1b and 1d) increased final average costs from \$10.06 to \$12.01 an increase of 19.3%. Going from 1:1 to 3:1 increased final average costs by 43.9%. Further, with full information, reducing the trading ratio from 1:1 to 0.5:1 resulted in an increase in final average costs. But if Figure 7 is analyzed more closely, one can see that a trading ratio of 0.5:1 actually resulted in the lower average costs through the first 138 trades. But, feasible trading ended after 134 trades in simulation 1a. Thus, a 1:1 trading ratio resulted in the most cost-effective loading reductions under full information.


Figure 7 Effects of a trading ratio under full information



A slightly different story played out under low information. Here, the 0.5:1 trading ratio was nearly the least cost-effective (highest average costs) from the first to last trade (Figure 8). Initially, there was no difference in average costs between simulations 2d, 2e, and 2f. If trading ceased after just 75 trades, simulation 2f would have had the lowest average costs of all of the low information simulations. But, the higher trading ratio prevented some trades from occurring, so the final average costs were lower for simulation 2f than for the other two. A 1.5:1 trading ratio (simulation 2c) actually resulted in the lowest final average costs at \$11.25/lb - even lower than a the 1:1 trading ratio simulation (2b).

Based on these results, the determination of an “optimal” trading ratio should necessarily depend on the amount of information available to market participants. The next section discusses points to consider in the characterization of an “optimal” trading ratio.

Figure 8 Effects of a trading ratio under low information



Characteristics of an “optimal” trading ratio

Since real world WQT markets most likely operate somewhere in between full and no information, a “partial” information scenario, where there is some information known about both WTP and WTA values, may be more realistic. If the partial information scenario is defined as halfway in between full and low information, then averages can be calculated across appropriate simulations’ output data. For example, the output from simulation 1a (full information, 0.5:1 trading ratio) and simulation 2a (low information, 0.5:1 trading ratio) can be averaged. The averaged data is reported in

Table 6.

Table 6 Averaged simulation results → "partial" information scenarios

Scenarios	Number of Trades	Credits Traded	Volume Traded				Cost Savings		Final Costs	
			Base Loading Reduction by NPS (lbs)	Loading Reduction by PS (lbs)	Additional Loading Reduction by NPS (lbs)	Total Loading Reduction (lbs)	Total (\$)	Percent (%)	Total (\$)	Avg. (\$/lb)
a	134	48,200	24,100	1,800	-	25,900	649,464	64.9	350,536	13.54
b	241	43,986	43,986	6,014	-	50,000	444,710	44.5	555,290	11.11

c	263	31,215	31,215	18,785	15,608	65,608	277,325	27.7	722,675	11.01
d	223	19,362	19,362	30,638	19,362	69,362	173,806	17.4	826,194	11.92
e	177	11,803	11,803	38,197	17,705	67,705	112,400	11.2	887,600	13.12
f	138	7,371	7,371	42,629	14,742	64,742	75,347	7.5	924,653	14.29

Based on the information in

Table 6, the poorest performing scenarios appear to be the cases where there are extremely low or high trading ratios. The 'a' (0.5:1 trading ratio) scenario had one of the highest average costs of all scenarios at \$13.54/lb of reduction. Most importantly, however, is the fact that the 'a' scenario resulted in only 25,900 lbs of loading reduction - 24,100 lbs short of the goal. This was the only scenario that failed to meet the goal of 50,000 lbs of loading reduction. For this reason, the 0.5:1 trading ratio was ranked as the overall poorest performing scenario.

The 'f' scenario (3:1 trading ratio) exhibited the highest average costs of nutrient reduction at \$14.29/lb. The overall cost savings were minimal in this case at 7.5%. On the positive side, however, was the fact that this scenario did result in 14,742 lbs of additional loading reduction. The 'e' scenario (2.5:1 trading ratio) performed only slightly better than the 'f' scenario in terms of additional loading reduction, total cost savings, and average reduction costs.

It is important to note that each of the remaining scenarios were superior in at least one of the evaluation criteria. For that reason, it is difficult to rank these because it would be dependent upon normative judgments about the relative weighting of the criteria.

The most typical trading ratio used by real-world WQT programs is a 2:1 ratio. This was represented by the 'd' scenario. This scenario actually produced the greatest amount of total loading reduction at 69,362 lbs. However, this scenario did not result in the most cost savings or the lowest average cost of loading reduction.

Scenario 'c' (1.5:1 trading ratio) resulted in the lowest average cost of reduction at \$11.01/lb and the greatest number of transactions (263 trades). This scenario also resulted in 15,608 additional pounds of loading reduction.

Scenario 'b' (1:1 trading ratio) resulted in the greatest amount of cost savings (\$444,710) among the scenarios which actually achieved the nutrient reduction goal. By definition, the 1:1 trading ratio resulted in no additional loading reduction. The average costs of the 'b' scenario were just slightly higher than 'c' scenario.

A few conclusions regarding an "optimal" can be made based on these results. Earlier it was shown that scenarios 'a', 'b', and 'c' became more cost-effective as information became more readily available. On the other hand, scenarios 'd', 'e', and 'f' exhibited a negative relationship between information and cost-effectiveness. Based on this, if adequate information availability is a concern in a certain WQT market,

then it may be reasonable to implement a 2:1 trading ratio. Also, if the goal of the WQT market is to maximize total loading reduction, a 2:1 trading ratio may be attractive.

On the other hand, if it is presumed that information will be readily disseminated and understood by market participants, a lower trading ratio may be more economical. If the goal is to simply maximize cost savings while achieving the reduction goal, a 1:1 trading ratio shows the greatest potential. However, if WQT market designers wish to account for nonpoint source loading uncertainty and increase the probability of having additional loading reduction, a 1.5:1 trading ratio may be most appropriate. Based on the simulation results, a 1.5:1 trading ratio also results in the lowest average costs of reduction. Again, the cost of disseminating information to market participants would need to be considered during the decision-making process.

Simulation of real-world WQT markets

While the model developed here utilized hypothetical data, there is no reason why these same market simulation algorithms could not be used to simulate a real-world hypothetical market. There is, however, because WQT programs, by nature, involve complex interactions between economics and the biophysical world, accurately simulating a real-world WQT market requires superior cost and watershed modeling data.

There are two types of cost data needed. On the point source side, facility upgrade costs and annual operation maintenance costs of meeting a more stringent nutrient standard are needed for wastewater treatment plants in the study watershed. This data can either be attained from surveys or by using general industry cost functions (e.g., Greenhalgh and Sauer 2003). In either case, the one-time and annual costs along with the appropriate time horizon should be used to calculate the annualized costs which considers the time value of money by including a discount rate.

On the nonpoint source side, the expected costs for BMPs are needed. These costs can come from surveys or from previous research. University Extension fact sheets can often times provide general estimates for this type of data (e.g., Devlin et al. 2003). Again, all one-time and annual costs should be converted to an annualized basis in analogous fashion to the point source data.

Related to the pollution control costs, is the idea that traders may perceive ‘intangible’ costs that are weighed against any potential gains. That is, the assumption that only monetary trading gains enter traders’ utility functions may not hold. A growing literature documents that the behavior of participants in an institution is influenced by institutional processes and rules, independent of the participants’ fiscal outcomes (Berg, Dickhaut, and McCabe, 2005; Johnston and Duke, 2007). Obtaining the information necessary to estimate intangible costs that may exist is crucial for simulating a real-world WQT market. Since this data is subjective by nature, it only can be obtained accurately through interaction with potential market participants via meetings and/or surveys.

Along with the economic data, biophysical watershed data is needed. Watershed models play a central role in the simulation and execution of real-world WQT markets. Watershed models represent a scientific understanding of how land characteristics, BMPs, and other factors relate to pollutant loading

into surface waterbodies (Nejadhashemi et al. 2009). There are many types of models ranging from very simple to very advanced (see Nejadhashemi et al. 2009 for guidance in choosing a model). Regardless of the type of model used, the necessary output from the model should be, at minimum, the following: the baseline nutrient loading from each subwatershed, reduction in loading from each subwatershed after a given BMP(s) is implemented, and relevant delivery ratios.

Combining all of this information will allow the researcher to generate the necessary WTP and WTA curves discussed previously in this paper. The steps laid out in the “Simulation model” section should be followed to simulate sequential, bilateral trading in the real-world watershed.

Conclusion

While there is substantial evidence that nonpoint sources have lower nutrient reduction costs than point sources, experience with WQT reveals a common theme: little or no trading activity. The success of WQT seems, in part, to depend on the structure of the market created to bring buyers and sellers together to transact exchanges. These outcomes suggest the presence of obstacles to trading that were not recognized in the design of existing programs.

To examine the ways that various market imperfections may impact the performance of a WQT market, an agent-based model was constructed which simulated a hypothetical point-nonpoint market. In particular, the market was modeled using a variant of the sequential, bilateral trading algorithm proposed by Atkinson and Tietenberg (1991). This paper presented an overview of the simulation modeling technique and then analyzed the effects of two prominent market impediments identified in the WQT literature: information levels and trading ratios.

Information levels refer to buyers’ and sellers’ knowledge of each others’ bid prices. A frictionless WQT market would be one where all of the potential buyers (i.e., point sources) would know all of the sellers’ (i.e., nonpoint sources) offer prices and vice versa. In this full information environment, we can expect that trades would be consummated in the order of their gains. That is, first buyers and sellers to be paired together for trading would be the buyers with the highest offer prices and the sellers with the lowest bid prices. Successive trades will have successively smaller gains until the gap between bid and offer prices reaches zero. This is the textbook Walrasian market and would closely approximate a double auction institution, where all buyers and sellers submit their offers and bids, which are then sorted and matched by a centralized market manager.

While the full information scenario serves as a useful benchmark, most existing WQT markets are decentralized in nature, so that limited information causes traders to be matched in a less efficient sequence. A variety of information levels are possible. One side of the market may have more information than the other (limited information) or neither side having any knowledge of the other side’s bid or offer prices (low information). Each of these scenarios leads to a different sequencing of trades. Market performance was measured in terms of cost savings, the number of credits traded, and the average reduction costs under different information scenarios.

Several notable results were found in regards to information levels. The results implied that if market designers feel that only a limited number of trades will be consummated, creating an institution that provides accessible information about buyers' prices is preferred to providing information about sellers' prices. In terms of average cost-effectiveness, when the trading ratio was below 1.5:1, the effect of information on cost-effectiveness was positive. On the other hand, when the trading ratio was 2:1 or greater, information levels were negatively related to average cost-effectiveness. That is, the more information available to market participants, the less cost-effective the trading.

Trading ratios are a common component of many existing WQT programs. A typical trading ratio of 2:1 requires a nonpoint source to reduce two pounds of expected nutrient loading in order to receive one pound of trading credit. These ratios serve as a "safety factor" and are incorporated to account for the uncertainty in the measurement and monitoring of nonpoint source loading. Because nonpoint traders must reduce loading by 2 pounds for every 1 pound emitted by point source traders, there will be a net reduction of 1 pound of expected loading for each trade. So, while inhibiting some trades from ever occurring, trading ratios also have the potential to improve water quality beyond trading with a 1:1 trading ratio. This paper examined these tradeoffs in terms of effects on market performance and then described procedures that can be used to characterize an optimal trading ratio if one exists. Based on the findings of this study, an "optimal" trading ratio should be in the range of 1:1 to 2:1. Further distinction between values in this range depends on the market designers' goals and the amount of information available and the cost of disseminating this information.

There are several limitations to this study. One is that these simulations did not consider the risk and variability associated with NPS loading. Mean loading values were used. In the real world, there will most definitely be some years in which the BMPs put in place by farmers will over-perform and significantly reduce nutrient runoff and in other years the BMPs may significantly under-perform. Incorporating this stochastic process into the model would most likely not change any of the main relationships found in this study, but it would build in the realism necessary for determining the percentage of time nutrient reduction targets would be exceeded and by how much.

Two other important market impediments not addressed in this study were transactions costs and intangible costs. These two factors may play a major role in determining the success or failure of a market of water quality.

References

- Atkinson, S. and T. Tietenberg (1991). Market failure in incentive-based regulation: The case of emissions trading. *Journal of Environmental Economics and Management* 21 (1), 17–31.
- Berg, J., J. Dickhaut, and K. McCabe. (2005). Risk preference stability across institutions: A dilemma. *Proceedings of the National Academy of Sciences* 102(11), 4209-4214.
- Breetz, H.L., e. a. (2004). Water quality trading and offset initiatives in the united states: A comprehensive survey. Technical Report Report to the EPA, Dartmouth College Rockefeller Center.

- Devlin, D., K. Dhuyvetter, K. McVay, T. Kastens, C. Rice, K. Janssen, & G. Pierzynski (2003). Water Quality Best Management Practices, Effectiveness, and Cost for Reducing Contaminant Losses from Cropland, Department of Agronomy and Department of Agricultural Economics. MF-2572, Kansas State University.
- Ermoliev, Y., M. Michalevich, and A. Nentjes (2000). Markets for tradeable emissions and ambient permits: A dynamic approach. *Environmental and Resource Economics* 15 (1), 39–56.
- Faeth, P. (2000). Fertile ground: Nutrient trading's potential to cost-effectively improve water quality.
- Greenhalgh, S. and A. Sauer. (2003). Awakening the Dead Zone: An Investment for Agriculture, Water Quality, and Climate Change. *World Resources Institute Issue Brief*.
- Hahn, R. (1989). Economic prescriptions for environmental problems: How the patient followed the doctor's orders. *Journal of Economic Perspectives* 3 (1), 95–114.
- Hoag, D. and J. Hughes-Popp (1997). Theory and practice of pollution credit trading. *Review of Agricultural Economics* 19 (1), 252–262.
- Horan, R. (2001). Differences in social and public costs and public risk perceptions and conflicting impacts on point/nonpoint trading ratios. *American Journal of Agricultural Economics* 83 (1), 934–941.
- Horan, R. and J. Shortle (2005). When two wrongs make a right: Second-best point-nonpoint trading ratios. *American Journal of Agricultural Economics* 87 (2), 340-352.
- Johnston, R.J. and J.M. Duke (2007). Willingness to pay for agricultural land preservation and policy process attributes: Does the method matter?" *American Journal of Agricultural Economics* 89(4), 1098-1115.
- King, D. and P. Kuch (2003). Will nutrient credit trading ever work? An assessment of supply, problems, demand problems, and institutional obstacles. *Environmental Law Institute News & Analysis* 33 ELR, 10352-10368.
- Malik, A., D. Letson, and S. Crutchfield (1993). Point/nonpoint source trading of pollution abatement: Choosing the right trading ratio. *American Journal of Agricultural Economics* 75 (1), 959–967.
- NCEE (2001). The United States experience with economic incentives for protecting the environment. Technical Report Publication EPA-240-R-01-001, Washington, D.C.: Environmental Protection Agency.
- Nejadhashemi, A.P., C.M. Smith, and W.L. Hargrove (2009). Adaptive Watershed Modeling and Economic Analysis for Agricultural Watersheds, Kansas State University Agricultural Experiment Station and Cooperative Extension Service. MF-2847.

- Netusil, N. and J. Braden (2001). Transactions costs and sequential bargaining in transferable discharge permit markets. *Journal of Environmental Management* 61 (1), 253–262.
- O’Neil, W. (1983). Transferable discharge permit trading under barying stream conditions: A simulation of multiperiod permit market performance on the fox river, wisconsin. *Water Resources Research* 19 (1), 608–612.
- Parker, D.C., T. Berger, and S. Manson, eds. (2002). *Agent-based Models of Land-Use and Land-Cover Change*. Report and Review of an International Workshop. Indiana University.
- Peterson, J.M., C.M. Smith, J.C. Leatherman, K.R. Douglas-Mankin, T.L. Marsh, J.A. Fox, M. Lee, and M. Henry (2009). Integrating economic and biophysical models to assess the impacts of water quality trading.” Final Technical Report to the USEPA-NCER.
- Smith, C., J. Peterson, and J. Leatherman (2007). Attitudes of great plains producers about best management practices, conservation programs, and water quality. *Journal of Soil and Water Conservation* 62 (September/October), 97A–103A.
- Stavins, R. (1995). Transaction costs and tradable permits. *Journal of Environmental Economics and Management* 29 (1), 133–148.
- Tesfatsion, L. (2006). Agent-based computational economics: A constructive approach to economic theory. In L. Tesfatsion and K. L. Judd (editors), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, Handbooks in Economics Series, North-Holland.
- USEPA (1996). Draft framework for watershed-based trading. Technical Report EPA800-R-96-001, Office of Water.
- Woodward, R. and R. Kaiser (2002). Market structures for U.S. water quality trading. *Review of Agricultural Economics* 24 (1), 366–383.