Environmental and Resource Economics **25:** 101–127, 2003.

Buyer Liability and Voluntary Inspections in International Greenhouse Gas Emissions Trading: A Laboratory Study

TIMOTHY N. CASON
Department of Economics, Krannert School of Management, Purdue University, West Lafayette, IN 47907-1310, USA. E-mail: cason@mgmt.purdue.edu

Accepted 26 November 2002

Abstract. This paper reports a preliminary laboratory experiment in which traders make investments to increase the reliability of tradable instruments that represent greenhouse gas emissions allowances. In one half of the sessions these investments are unobservable, while in the other half traders can invite costless and accurate inspections that make reliability investments public. We implement a buyer liability rule, so that if emissions reductions are unreliable (i.e., sellers default), the buyer of the allowances cannot redeem them to cover emissions. We find that allowing inspections significantly increases the reliability investment rate and overall efficiency. Prices of uninspected allowances usually trade at a substantial discount due to the buyer liability rule, which provides a strong market incentive for sellers to invest in reliability.

Key words: emissions permits, environment, experiments, Kyoto Protocol

JEL classifications: Q25, L51

1. Introduction

The Kyoto Protocol establishes targets and specific commitments to limit greenhouse gas (GHG) emissions and enhance natural sinks for the 24 developed OECD countries as well as to 12 “economies in transition” (Central and Eastern Europe and the former Soviet Union). The Protocol also allows these Parties to use international GHG emissions trading to fulfill these commitments. Emissions trading can significantly reduce the overall cost of implementing emission reductions because, at least in theory, the reductions would occur where they are the cheapest. But before implementing a trading system, Parties must specify numerous design details – such as the extent that countries can use trading to meet emission reduction commitments, as well as the verification, reporting and accountability of emissions reductions and trades.

Another detail concerns liability: if a Party “oversells” their assigned amount and is out of compliance, who is liable for the shortfall in emissions allowances?

The temptation for a firm or governmental agency within a transition economy to oversell allowances could be substantial. For example, according to one estimate, countries that comprise the former Soviet Union would obtain at least $17 billion in allowance sales revenues if emissions trading is widely used to meet the Kyoto Protocol objectives (Edmonds et al. 1999).[1] These countries could obtain hard currency through allowance sales and earn greater short-term profits if they limited their investments in abatement. Sanctions against noncompliant Parties outlined in the 2001 Bonn Agreement include a suspension of emissions trading rights and an obligation in the next compliance period to make up for each excess tonne emitted plus 30 percent extra.

A variety of liability rules to address noncompliance are possible, and the rules provide differing incentives for compliance (Baron 1999a). One option is an "issuer beware" liability rule that sanctions countries who issue and sell emissions allowances (or "permits") when they have not made sufficient reductions in domestic emissions. Unfortunately, the transaction costs necessary to enforce penalties against sovereign nations that have oversold emissions permits could be substantial. For example, Japan insists that the noncompliance penalties established in the Bonn Agreement not be legally binding because that would be an affront its national sovereignty. Resolving legal disputes among governments tends to be slow and expensive, and there are few examples of successful international enforcement procedures in multilateral environmental agreements (Corfee Morlot 1998). In recognition of these constraints, the Bonn Agreement features a reserve requirement for allowances to deter overselling, even though this requirement could dampen market liquidity and increase transactions costs (Baron 2001).[2]

This paper studies an alternative "buyer beware" liability rule that may be effective in using market incentives to help deter seller noncompliance. Under this rule an allowance buyer bears the risk that the issuing Party may not be in compliance. Sold allowances would be returned to the issuer if she is not in compliance, so the buyer would not be able to use the purchased allowances to offset his own emissions. This system encourages buyers to purchase allowances from sellers who are most likely to comply. Moreover, a Party that has acquired allowances from another Party will have a great interest in ensuring the other Party's compliance (Werksman 1999). Most importantly, in a buyer liability setting the allowances issued by different sellers are likely to trade at different prices. In effect, the market prices reflect the noncompliance risk and these prices provide incentives for sellers to meet their emission reduction commitments.[3]

This internal "enforcement" through market prices is complicated, however, by incomplete information. To obtain high prices an allowance seller would like to assure buyers that she will be in compliance – analogous to the used car seller who wishes to assure potential buyers that the car is of high quality.[4] A used car seller could offer the buyer the opportunity to have a professional inspector evaluate the car prior to purchase to partially overcome this incomplete information problem. For the GHG emissions trading case, this buyer liability rule may provide incentives for the selling country to invite international inspectors to verify emission control efforts and report a low potential for noncompliance.[5] The benefit to the seller would be the higher prices she could obtain in the allowance market.

This research reports a laboratory experiment to examine whether a buyer liability rule provides strong incentives for compliance in a highly stylized emissions trading market, and the role of inspections and information dissemination in promoting compliance. All sessions begin with several baseline periods without any noncompliance risk. Noncompliance risk with a buyer liability rule is then introduced about halfway through each laboratory session. In these later periods if traders make a costly investment in emissions control they can substantially reduce the likelihood that they are out of compliance. Although in all sessions the traders can establish reputations for high compliance rates, in only half the sessions buyers have the opportunity to invite inspectors to verify directly their costly investments.

I find that allowance sellers are more likely to invest in abatement reliability when they can allow inspectors to verify their investment, and this increases the allocative efficiency of the emissions and output markets. Prices of uninspected allowances usually trade at a substantial discount due to the buyer liability rule, which provides the incentive for sellers to invest in reliability investment. Consequently, the results suggest how a buyer liability rule for emission allowances can provide incentives for sellers to comply with GHG emission market rules, especially when they have the opportunity to make their compliance choices public and earn a price premium from their reliability investments. The buyer liability rule can increase buyers' compliance costs if some sellers fail to invest in abatement reliability, however, so this rule does not reduce the level of non-compliance for all market participants.

The remainder of the paper is organized as follows. Section 2 describes the experimental environment and design. This section concludes with four research hypotheses derived from the simple model of emissions trading contained in Appendix A. Section 3 collects the experimental results, and Section 4 concludes. The experiment instructions are available for downloading at http:\\www.mgmt.purdue.edu\faculty\cason\papers\liabinst.pdf.

2. Experimental Environment, Design and Hypotheses

RELATION TO THE EXPERIMENTAL LITERATURE

This paper follows in the tradition of laboratory research on the design of allowance markets. The laboratory has been an important source of empirical data on this research topic, in part because more traditional empirical work based on field data is constrained by the limited implementation of these markets. New proposals for allowance market rules usually cannot be evaluated with field data because the rules have not yet been implemented in the field.

Most experimental work has focused on specific features of emissions trading. For example, experiments have evaluated features of the trading institutions implemented or planned for specific emissions trading programs, such as experiments run to testbed and evaluate the sulfur dioxide allowance market in the U.S. (e.g., Franciosi et al. 1993; Cason and Plott 1996). Muller and Mestelman (1994) present experiments that compare the trading rules for the U.S. sulfur dioxide market with rules proposed for trading nitrogen oxide allowances in Southern Ontario, and they find improved efficiency in the proposed Canadian trading institutions. Experiments were also conducted to help design the Regional Clean Air Incentive Market (RECLAIM), a tradable permit program implemented in Los Angeles to reduce the emissions of sulfur and nitrogen oxides (Carlson et al. 1993a, 1993b; Cason and Gangadharan 1998). Following the recent interest in international agreements to reduce carbon emissions, researchers have also used laboratory experiments to investigate the empirical properties of potential GHG trading schemes (e.g., Bohm and Carlén 1999; Søberg 2000a). Godby and Shogren (2002) report an experiment that studies both buyer and seller liability rules in a GHG allowance market, but without investments in abatement. In their laboratory environment subjects interact only in the allowance market and they have very strong overselling incentives in the buyer liability treatment. The authors observe greater noncompliance under buyer liability rules.

To focus on the market features of interest – here the role of a buyer liability rule in promoting seller compliance with and without observable abatement investments – this paper deliberately abstracts from many additional characteristics of emissions trading. For example, I do not include an opportunity for traders to “bank” permits for future use (as studied by Cronshaw and Brown-Kruse 1999), nor do I study the implications of irreversibility in the choice of production process and abatement technologies (as studied by Ben-David et al. 1999). I study the implications of the liability rule in a relatively simple market environment in order to draw clear inferences from the laboratory data. Although this limits the parallelism between the laboratory and the field, my goal was not to capture all of the features of an international allowance trading system. Data from the field are more appropriate for evaluating the overall performance of specific emissions trading systems. But since GHG allowance trading has yet to develop, field data are not currently available. Moreover, even if field data were available, they would not include periods in which investments are or are not observerable, and therefore would not allow one to isolate the impact of observability on market performance.

ENVIRONMENT

The experiment uses the Multiple Unit Double Auction (MUDA) software (Plott 1991) to implement a computerized version of the double auction market trading institution. Although subjects could trade in multiple units, the MUDA program was most useful for this experiment because it permits simultaneous trading in multiple markets. Subjects could trade output (described to them as good X) in one market, as well as emission allowances (described to them as good Y) in another market. Demand for allowances is therefore derived from supply and demand conditions in the output market, and the economic cost of holding insufficient allowances (either due to deliberate noncompliance or allowance default) depends on the traders' individual production costs and output valuations, as would be the case in the field.

Neutral terminology was used throughout the experiment as is standard in experimental economics. The instructions simply described Y as an input that subjects need to produce output X. Once reliability investments are introduced it is important for subjects to be able to differentiate allowances depending on their origin. Consequently, in all periods (even the baseline periods with no reliability investments) each subject with an allowance allocation received this allocation in a separate MUDA market. Subjects therefore always knew the identity of the issuing party and could use this information when deciding on the appropriate price to offer (Baron 2000). Six subjects received allowance allocations, so the MUDA trading screen displayed six Y markets labeled Y0 through Y5. The trading screen displayed the output market for X in market row 7.

Subjects did not receive profits directly from allowances Y, but rather from "redeeming" units of output X. To simplify matters we allow subjects to produce exactly one unit of output for each allowance Y held at the end of the period.[6] They must also pay a non-allowance production cost when producing X, and they can also buy and sell X on the output market. The non-allowance production cost and the output redemption schedules varied across subjects to represent the heterogeneity across countries in technological capabilities, size and income. The production cost, redemption schedules and allowance endowments are shown in Table I. The goal was not to implement parameters that correspond to the details of specific countries that are expected to participate in GHG emissions trading. Rather, the goal was to include some parameters representative of some potential participants – including those with high (excess) allowance endowments but low values for the allowances (subject 0), those with low allowance endowments relative to their value for output (subjects 1–3), and small participants with surplus (subjects 4–5) and insufficient (subjects 6–7) allowances.[7] Consistent with the standard precept of *Privacy* in experimental economics (Smith 1982), the costs and redemption schedules were subjects' private information.[8]

To speed up the profit calculations and eliminate arithmetic errors, subjects used Excel workbooks to calculate profits each period. Each subject sat in front of two computers – one running MUDA for trading Y and X, and one running Excel for calculating profits. An example page from the workbook is shown in Figure 1. The subjects only enter in data for the four gray "shaded in" cells of the workbook, by copying down information shown on their MUDA trading screen. All other cells are password protected and cannot be changed by subjects.

Subjects were recruited from undergraduate economics classes at Purdue University. They first participated in an extensive training session. In this training session they reviewed a self-paced, computerized tutorial teaching them the MUDA

Table I. Subject output redemption values, allowance endowments and production costs.

	Subject 0	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6	Subject 7
Output (X) consumed	Redemption values							
1	175	210	210	210	185	185	195	195
2	135	195	195	195	135	135	120	120
3	120	180	180	180	75	75	90	90
4	90	170	165	160				
5	60	155	150	145				
6	30	140	135	130				
7		125	120	115				
8		105	105	105				
Y endow-ments:	8	2	2	2	3	3	0	0
Output produced	Production costs							
1	30	15	15	15	15	15	45	45
2	60	30	30	30	45	45	75	75
3	120	50	50	50	90	90	135	135
4	180	60	60	60	150	150	165	165

Notes: One allowance is required for each unit of production, in addition to the production costs shown at the bottom of the table. Reliability investments cost 20-francs per unit of Y endowment to reduce the likelihood of Y default from 40 percent to 0 percent.

trading procedures, and then received written instructions that were also read aloud by the experimenter. Subjects were then assigned randomly to trader roles and participated in several actual trading periods of the Baseline treatment with salient monetary rewards. At the end of the training session they received total profits, converted from lab dollars into U.S. dollars at a fixed and known exchange rate, plus a $10 training fee. We do not report results from the training sessions here, because they involved varying numbers of traders (we trained all the subjects who showed up) and because the complexity of this multiple market input, production and output market structure led to considerable learning and noise in the training. Instead we report only "experienced" sessions, consisting of subjects who had completed this preliminary training. In these experienced sessions subjects received their total trading profits in cash, which ranged between $8.25 and $35.25 with an average of $24.76. Sessions lasted about two hours.

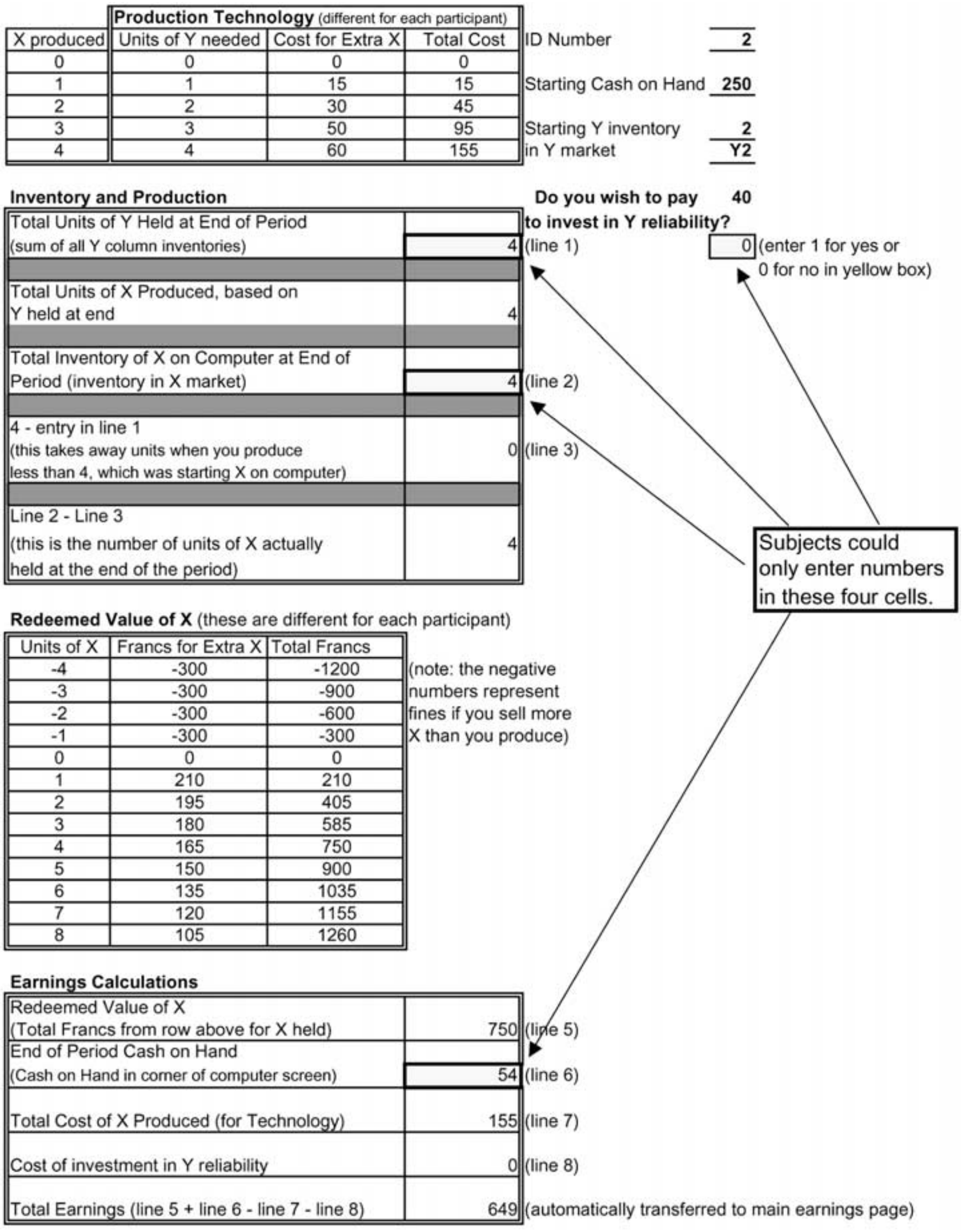

Production Technology (different for each participant)

X produced	Units of Y needed	Cost for Extra X	Total Cost
0	0	0	0
1	1	15	15
2	2	30	45
3	3	50	95
4	4	60	155

ID Number 2

Starting Cash on Hand 250

Starting Y inventory 2
in Y market Y2

Inventory and Production

Total Units of Y Held at End of Period (sum of all Y column inventories)	4	(line 1)
Total Units of X Produced, based on Y held at end	4	
Total Inventory of X on Computer at End of Period (inventory in X market)	4	(line 2)
4 - entry in line 1 (this takes away units when you produce less than 4, which was starting X on computer)	0	(line 3)
Line 2 - Line 3 (this is the number of units of X actually held at the end of the period)	4	

Do you wish to pay 40 to invest in Y reliability? 0 (enter 1 for yes or 0 for no in yellow box)

Subjects could only enter numbers in these four cells.

Redeemed Value of X (these are different for each participant)

Units of X	Francs for Extra X	Total Francs
-4	-300	-1200
-3	-300	-900
-2	-300	-600
-1	-300	-300
0	0	0
1	210	210
2	195	405
3	180	585
4	165	750
5	150	900
6	135	1035
7	120	1155
8	105	1260

(note: the negative numbers represent fines if you sell more X than you produce)

Earnings Calculations

Redeemed Value of X (Total Francs from row above for X held)	750	(line 5)
End of Period Cash on Hand (Cash on Hand in corner of computer screen)	54	(line 6)
Total Cost of X Produced (for Technology)	155	(line 7)
Cost of investment in Y reliability	0	(line 8)
Total Earnings (line 5 + line 6 - line 7 - line 8)	649	(automatically transferred to main earnings page)

Figure 1. Individual period record sheet for subject 2, period 11, session BI20405.

TREATMENTS

The experiment features three treatment conditions: (1) a *Baseline* treatment with no possibility of default risk; (2) a *No Inspections* treatment with default risk but no information dissemination/inspection mechanism for subjects to inform others of their investments in reliability; and (3) an *Inspections* treatment with default risk and an opportunity for subjects to inform others costlessly of their investments in reliability.

All sessions began with 6 or 7 periods of the Baseline treatment. This early baseline phase serves three purposes. It provides some experience beyond the training session with the trading and record-keeping procedures in a comparatively simple environment. It also allows subjects to learn the relationship between total production costs (which include costs of acquiring emission allowances), non-allowance production costs, and output redemption values. Finally, the baseline treatment provides actual market outcomes with no default risk for comparison with the other treatments that include default risk. The total endowment of Y and therefore the total production of X in these baseline periods is fixed at 20 units. Figure 2 displays the aggregate valuation schedule pooled across traders, as well as the aggregate marginal production cost schedule. Because of the discrete trading units used in the experiment the X market can clear at a range of equilibrium prices $p_x \in [140, 145]$, and the allowance (Y) market can clear at a range of equilibrium prices $r \in [70, 85]$.

After these initial baseline periods, without prior warning we introduced the possibility of allowance default – either with or without the opportunity for inspections – by distributing and reading aloud an additional instructions page. Subjects were given the opportunity to invest $t = 20$ experimental francs for each allowance in their endowment to ensure that there is no chance that any of their allowances default. (We framed this default risk to subjects as a chance that the Y units "break down" and cannot be used by anyone to produce X.) If they chose not to make this investment, there was a 40 percent chance that all of a subject's allowances default. In order to keep the environment simple we did not permit partial investment or partial default.[9] Subjects submitted their investment decisions on written forms before trade opened each period. The experimenter randomly generated the default outcomes using a 10-sided die, with independent rolls for each subject.

In the No Inspections treatment, subjects received no public information regarding any reliability investment decisions. Realizations of default, however, were publicly announced and written on the whiteboard at the end of the trading period. [The allowances originating in markets with default were listed (e.g., "Y1, Y4") under a label "breakdowns," along with the period number, and this information remained on the whiteboard for all subsequent periods.] Any subject holding allowances originating in such a default "Y market" could not use those allowances in X production. Note that subjects receive noisy signals regarding reliability investments: a particular Y market may not be in default because the trader with

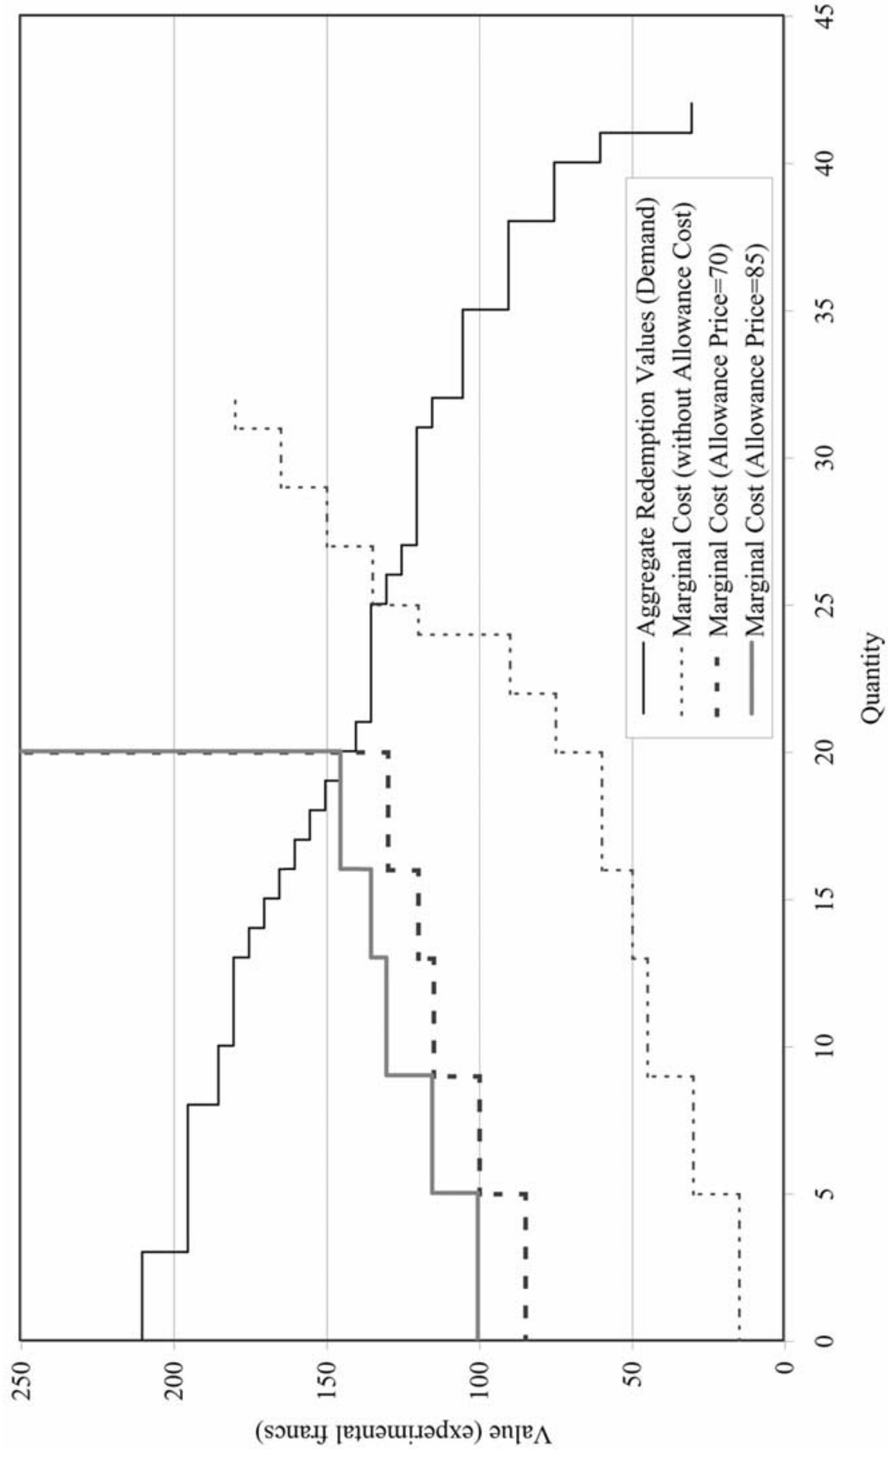

Figure 2. Output redemption values and marginal cost for alternative allowance prices.

Table II. Summary of experimental sessions.

Session name	Baseline periods	No inspections periods	Inspections periods	Experience level
B20401	Periods 1–7	Periods 8–15		Once Experienced
B20402	Periods 1–7	Periods 8–15		Once Experienced
BI20403	Periods 1–7		Periods 8–15	Once Experienced
BI20405	Periods 1–7		Periods 8–15	Once Experienced
B20404x	Periods 1–6	Periods 7–16		Twice Experienced
BI20406x	Periods 1–6		Periods 7–16	Twice Experienced

a Y endowment in that market made the reliability investment, or because that trader received a lucky realization of the random default process. Consequently, this noisy information may make it challenging for subjects to establish reputations to overcome the moral hazard problem of reliability investment.

In the Inspections treatment subjects are given the opportunity to provide information and overcome this moral hazard problem. When submitting their reliability investment decision form they check off a box indicating whether they wish their investment decision to be announced publicly. This is equivalent to inviting inspectors to verify accurately the investment.[10] Subjects could announce that they do not invest, but not surprisingly when announcements were made they almost always were announcements of reliability investment.[11]

Table II lays out the design of the six sessions. Three introduced the No Inspections treatment and three introduced the Inspections treatment in the second half of the session. Two of the six sessions involved twice-experienced subjects. All parameters in Table I and the reliability investment cost were doubled (and consequently the franc-to-dollar exchange rate was cut in half) for the twice-experienced sessions. This creates different equilibrium prices in the once- and twice-experienced sessions. (The results presentation renormalizes prices in the twice-experienced sessions so that they are comparable to the once-experienced sessions.)

HYPOTHESES

Appendix A presents a simple model of emissions trading with default risk in this environment. This model can be summarized with the following testable hypotheses, which are evaluated in the laboratory experiment.

Hypothesis 1: The introduction of default risk (weakly) increases output prices.

Output prices may increase when a positive allowance default probability exists, because in this case the expected output falls below the output in the no default baseline. Consequently, equilibrium output prices rise compared to the no default case because the demand for output is downwardly-sloping. The output price impact of allowance default depends on trader beliefs. For example, if everyone believes that none of the 20 allowances are backed by reliability investments, then the expected output falls from 20 to $0.6 \times 20 = 12$ units. As shown in Figure 2, the marginal valuation of the 12th unit is 180 francs, so prices could reasonably rise this high. On the other hand, if everyone believes that all of the 20 allowances are backed by reliability investments, then prices correspond to the no default risk equilibrium (140–145).

Hypothesis 2: In the presence of default risk, the prices of (inspected) allowances known to be reliable are no lower than the prices of allowances in the no default baseline, and are greater than or equal to the prices of allowances of unknown reliability in the default risk treatment.

The "discount" received on allowances of unknown reliability reflects their default risk. The upper bound on this discount is the 40 percent chance that an allowance without any reliability investment defaults. A discount this great would only be reasonable if subjects believed that an uninspected allowance was certainly unreliable.

Hypothesis 3: Traders will invest in reliability, whether or not reliability inspections are permitted.

One can calculate from Table I at equilibrium production and redemption levels that on the margin the minimum value of a valid allowance varies across subjects and ranges between 85 and 150 francs. An allowance that has a 40 percent chance of default is therefore worth at least $0.6 \times 85 = 51$. The experiment implemented a reliability investment cost of 20 francs, which is substantially below the minimum discount arising from unreliability (85 – 51 = 34 francs). Therefore, in equilibrium the subjects should invest in reliability, regardless of whether they can submit to inspection.

The final Hypothesis 4 follows directly from the observation that if allowances default because some traders fail to make reliability investments, then output falls and final allocations become stochastic. This leads to a decline in system efficiency.

Hypothesis 4: Overall efficiency is positively related to the level of reliability investment.

3. Results

To provide the reader with an overview of the lab behavior, we first present two figures that illustrate individual transaction prices for two typical sessions. The formal evaluation of the research hypotheses follows in later subsections.

EXAMPLE SESSIONS

Figure 3 presents the time series of transaction prices for session B20402, which featured reliability investments and default risk in periods 8–15 but no inspection opportunities. Individual transaction prices are indicated by squares for allowances and by diamonds for output. These prices exhibit considerable variance, and mean output transaction prices remain below the equilibrium range (indicated by solid horizontal lines) throughout the baseline periods.[12] Allowance prices also exhibit high variance in the initial periods, but by period 4 they settle into the equilibrium range (indicated by dashed horizontal lines) for the remainder of the baseline periods.

Output prices drift up slightly and allowance prices temporarily dip down after the default risk and reliability investments are introduced in period 8. As these treatment periods progress, the most salient feature of these transaction price data is the increasing variance of allowance prices. This increased variance may reflect differences in beliefs about the reliability of allowances originating from different subjects. Some subjects may have acquired reputations for selling reliable allowances, so they may be able to obtain a price premium over those allowances originating from less reliable sources.

Figure 4 presents the transaction prices in BI20406x. This session featured inspections in the default treatment periods so it can provide more direct evidence of a price differential between reliable and unreliable allowances. When the reliability investments are first introduced, most traded allowances are inspected (at prices indicated with crosses), although a few uninspected allowances trade at prices indicated by open circles. After several periods the prices of the inspected allowances increase and the prices of the uninspected allowances decrease. The uninspected allowances trade at roughly a 40 percent discount, which suggests that buyers believe that the uninspected allowances are unreliable since the default probability of an unreliable allowance is 40 percent. Also note that in this session output prices and output price variance tends to increase following the introduction of default risk.

TRANSACTION PRICES

Observation 1: The introduction of default risk increases output prices (supports Hypothesis 1).

Support: Table III presents transaction price summary statistics for all individual sessions in the baseline (no default risk) and default risk treatment periods.[13] As shown on the left side of this table, the data support Hypothesis 1 since output prices rise on average in all six sessions after default risk is introduced. This allows us to reject the null hypothesis that the median change in transaction prices after the introduction of default risk is zero, in favor of the alternative that default risk increases transaction prices (Wilcoxon signed rank test $S = 0$, sample size $n =$

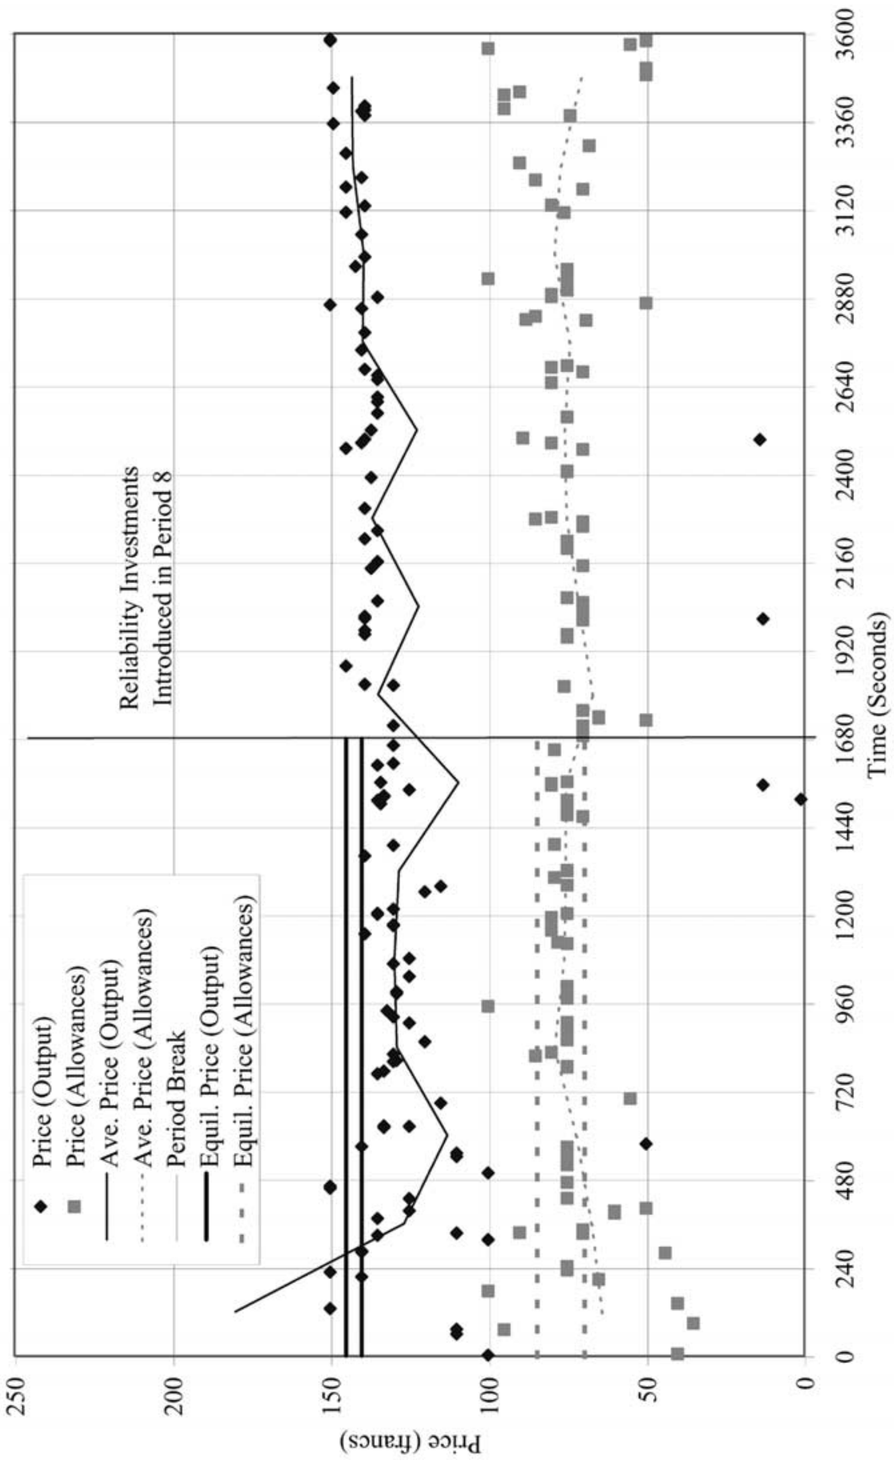

Figure 3. Prices for session B20402 with reliability investments but no inspections (once experienced).

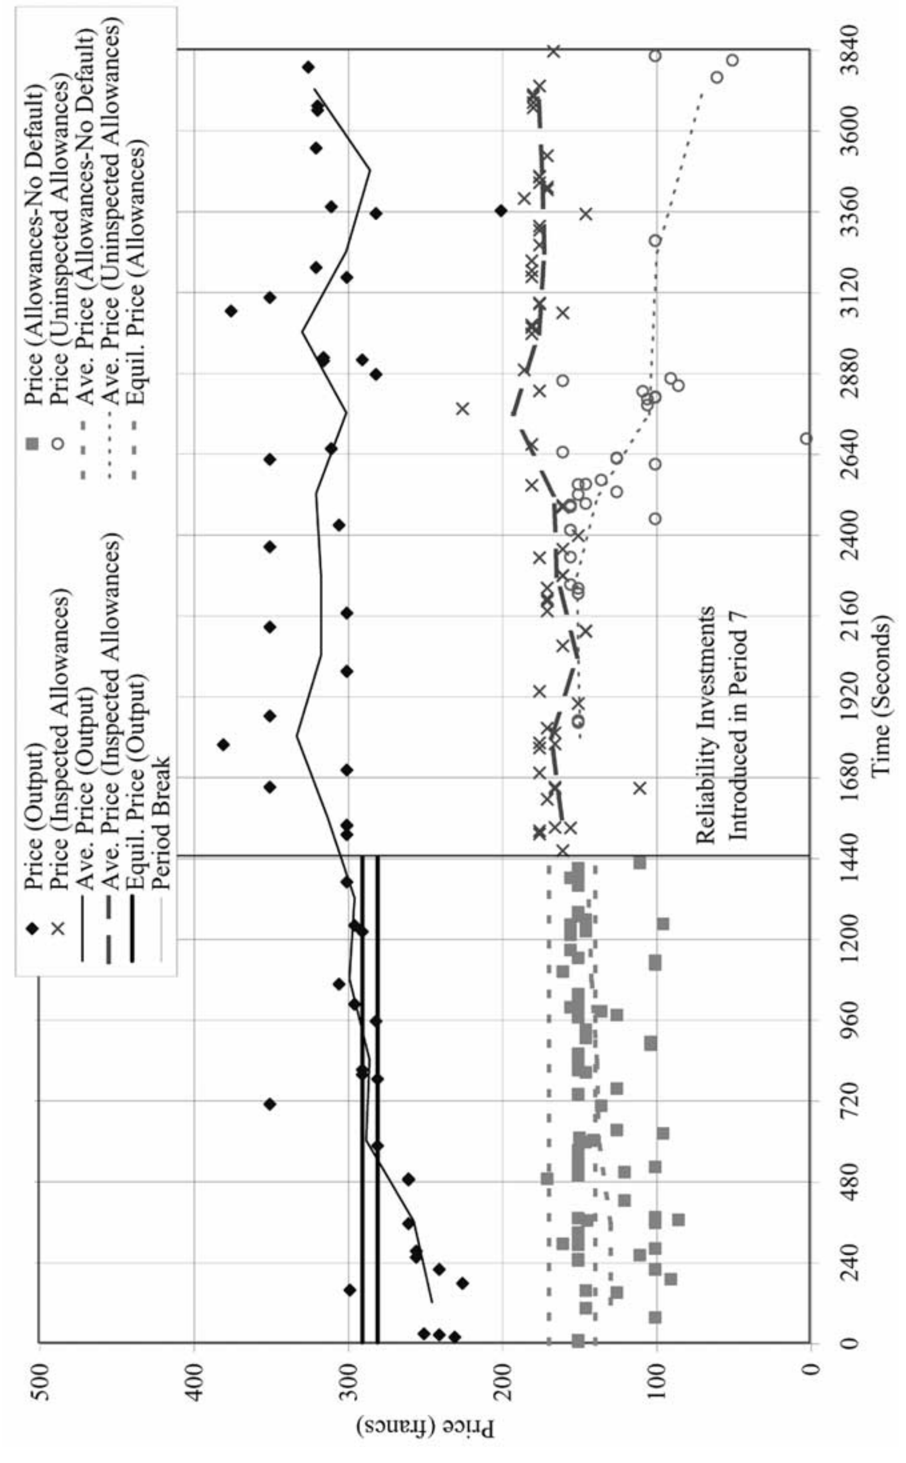

Figure 4. Prices for session BI20406x with reliability investments and inspections (twice experienced).

Table III. Transaction price statistics.

Panel A: No Inspections Treatment

Session	Statistic	Baseline Output (X)	Default Output (X)	Baseline Allowances (Y)	Default Allowances (Y)
B20401	Mean Price	78.3	139.5	60.0	18.9
	Std. Error	3.0	9.3	2.6	1.1
	Variance	363.5	2760.6	638.3	119.1
	Trans./Period	5.7	4.0	13.9	12.8
B20402	Mean Price	128.5	134.6	73.2	73.9
	Std. Error	6.5	3.3	1.5	1.4
	Variance	2813.0	597.0	140.4	125.7
	Trans./Period	9.4	6.8	9.0	8.3
B20404x	Mean Price	146.8	147.0	110.7	91.4
	Std. Error	1.2	2.1	1.1	1.6
	Variance	55.9	291.1	63.1	196.8
	Trans./Period	6.7	6.3	9.3	7.7

Panel B: Inspections Treatment

Session	Statistic	Baseline Output (X)	Default Output (X)	Baseline Allowances (Y)	Default Inspected[a] allowances (Y)	Default Uninspected allowances (Y)
BI20403	Mean Price	146.9	150.7	103.1	114.2	91.8
	Std. Error	3.2	1.5	2.0	1.0	5.0
	Variance	497.9	149.8	292.0	56.0	533.6
	Trans./Period	6.9	7.9	10.6	6.8	2.6
BI20405	Mean Price	114.6	148.3	67.7	95.5	80.6
	Std. Error	4.2	2.2	1.8	0.7	3.2
	Variance	483.5	92.8	214.5	33.5	72.3
	Trans./Period	3.9	2.5	9.4	7.9	0.9
BI20406x	Mean Price	137.0	157.0	68.8	84.6	61.3
	Std. Error	3.0	2.7	1.2	0.9	2.8
	Variance	219.9	250.3	112.8	52.7	320.6
	Trans./Period	4.2	3.5	13.2	7.0	4.0

[a]Inspected and revealed reliability investment only. Does not include the one subject who revealed a lack of reliability investment through inspection.

6; one-tailed p-value < 0.05).[14] The mean price increase in the No Inspections treatment is 22.5 francs and the mean price increase in the Inspections treatment is 19.2 francs; this increase amount is not significantly different in the two treatments (Mann-Whitney $U = 4$, sample sizes $n_1 = n_2 = 3$; p-value > 0.8).

The right side of Table III presents statistics for allowance prices. Subject beliefs about allowance reliability are unobservable, so in the No Inspections treatment the impact of default risk on allowance prices is ambiguous (Panel A). In two of the three No Inspections sessions, average allowance prices decline following the introduction of default risk, while in the third (session B20402) average allowance prices are virtually unchanged after default risk is introduced. As documented below, in the No Inspections treatment subjects make reliability investments for less than one-half of the allowances. The decline in allowance prices often observed following the introduction of default risk is consistent with the observed high default rate.

Observation 2: Under default risk, prices of allowances known to be reliable are higher than (1) allowance prices in the no default baseline, and (2) prices of allowances of unknown reliability (supports Hypothesis 2).

Support: Subject beliefs for allowances that are inspected and found to be reliable are obviously observable. As indicated in Hypothesis 2, since these allowances are certainly reliable their prices should not fall compared to the baseline treatment. In all three Inspections sessions the inspected allowance prices rise compared to the baseline periods, by an average of over 18 francs. Prices of uninspected allowances fall compared to the baseline in two of the three sessions. [In session BI20405 the uninspected allowance prices exceed the baseline periods allowance prices, but virtually no uninspected allowances trade in periods with default risk, as indicated by the average transactions per period of less than one.] This decline in prices for the uninspected allowances probably reflects the (correct) belief that uninspected allowances are often unreliable. In all sessions the uninspected allowances trade at a discount compared to the inspected allowances; on average 114.2 – 91.8 = 22.4 in BI20403 (a 20 percent discount), 95.5 – 80.6 = 14.9 in BI20405 (a 16 percent discount) and 84.6 – 61.3 = 23.3 in BI20406x (a 28 percent discount). As expected, this discount is not greater than the 40 percent default risk for an allowance that is certain to be unreliable.

RELIABILITY INVESTMENT RATES AND OVERALL EFFICIENCY

This subsection documents that reliability investments and overall efficiency both increase in a buyer liability system when traders are able to invite inspections.

Observation 3: Traders do not always invest in reliability (inconsistent with Hypothesis 3), but the investment rate is greater when inspections are allowed. The price premium received for reliable, inspected allowances provides the incentive for reliability investment.

Support: Figure 5 displays the reliability investment rate for each session, as well as the average rate across the three sessions in each treatment. In all but one period the investment rate is higher in all three sessions with inspections (the dashed lines) than in any of the three sessions with no inspections (the solid lines). The investment rate is therefore significantly higher with inspections (Mann-Whitney $U = 0$, sample sizes $n_1 = n_2 = 3$; two-tailed p-value = 0.10).[15] The overall reliability investment rate is 78 percent with inspections, which is more than double the 37 percent rate with no inspections.

As discussed at the end of Section 2, the reliability investment cost was substantially below the minimum value discount arising from unreliability. Therefore, the low reliability investment rate in the No Inspections treatment came as a surprise and is inconsistent with Hypothesis 3. Some subjects value allowances highly, however, and even the unreliable allowances trade on average at prices between 74 and 91 francs in two of the three sessions of the No Inspections treatment (Table III, Panel A). Consequently, in this treatment some subjects apparently determined that their best strategy was to not invest in reliability and try to sell their unreliable allowances.

The price differential between inspected and uninspected allowances documented in the previous subsection, however, provides an incentive for subjects who sell allowances to make reliability investments and allow inspections. The price differential between inspected and uninspected allowances was about 23 francs in the two Inspections sessions that had much trading volume for uninspected allowances (i.e., sessions BI20403 and BI20406x shown in Table III, Panel B). This differential exceeds the reliability investment cost of 20 francs. Subjects allowed most (83 percent) but not all of the allowances with reliability investments to be inspected. Since inspections were free it may seem surprising that this inspection rate is less than 100 percent. Closer examination of the data reveals that 28 of the 32 subjects who invest in reliability but do not invite inspections are in the role of subjects 1, 2 or 3 (see Table I). Because of their low production costs, these subjects typically acquire allowances rather than sell them; consequently, they do not benefit from the higher prices received by inspected allowances because they do not sell their allowances.

Observation 4: Efficiency is greater in the Inspections treatment (which featured high reliability) than in the No Inspections treatment (which featured lower reliability) (supports Hypothesis 4).

Support: Figure 6 displays the time series of mean trading efficiency. Trading efficiency is the actual surplus realized by subjects, divided by the maximum possible surplus they could obtain. This efficiency calculation includes losses due to fines from underproduction that subjects occasionally incurred.[16] In the baseline periods mean efficiency increases from 65 percent in the first period to 80 percent by period 7. Once default risk is introduced, inspections lead to a dramatic divergence between efficiency in the Inspections and No Inspections treatments. Overall effi-

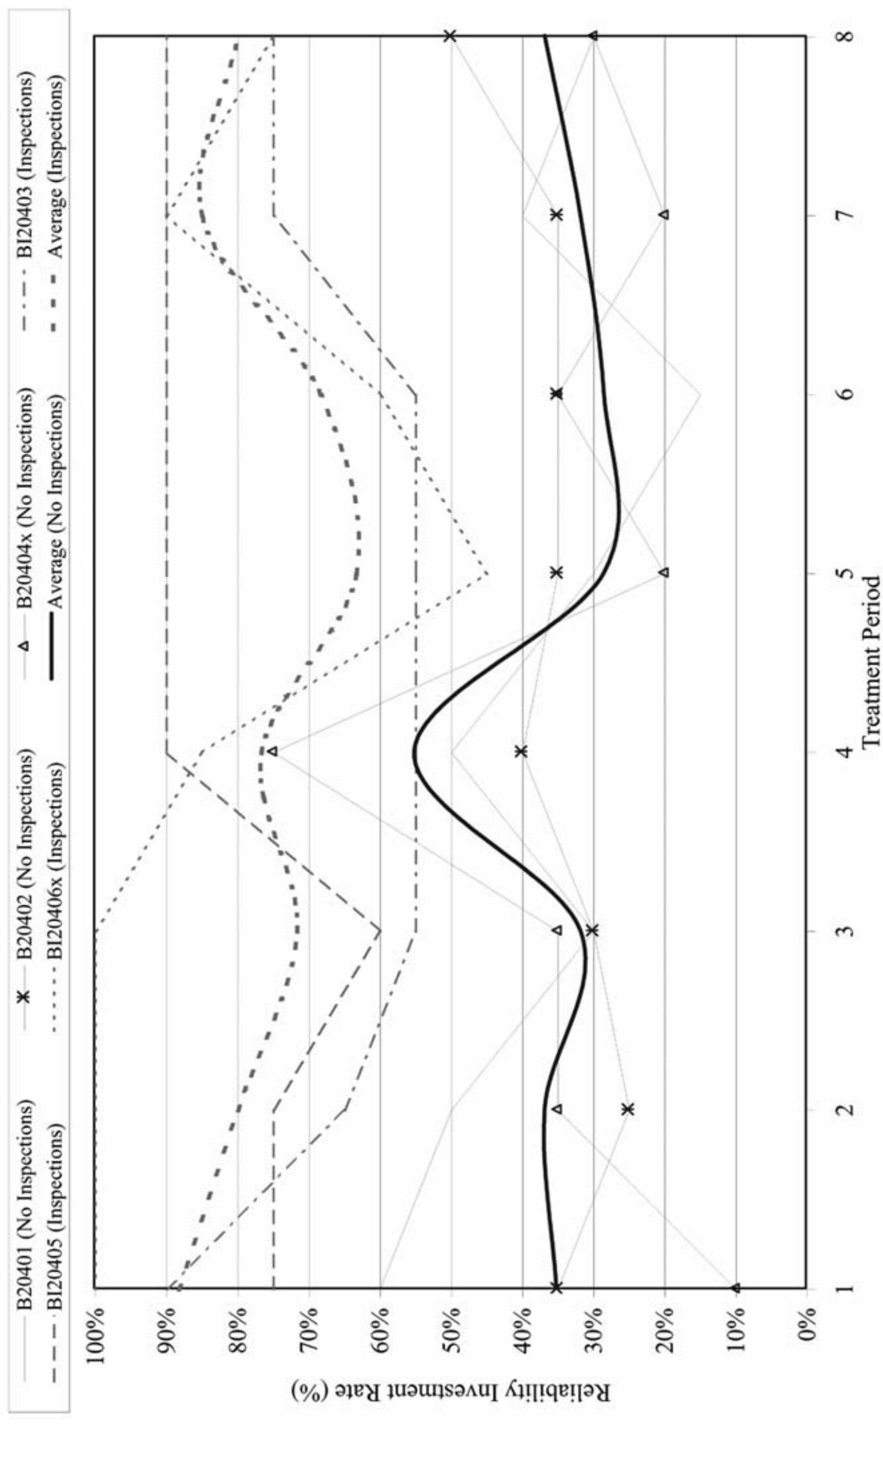

Figure 5. Reliability investment rates, by session and treatment averages.

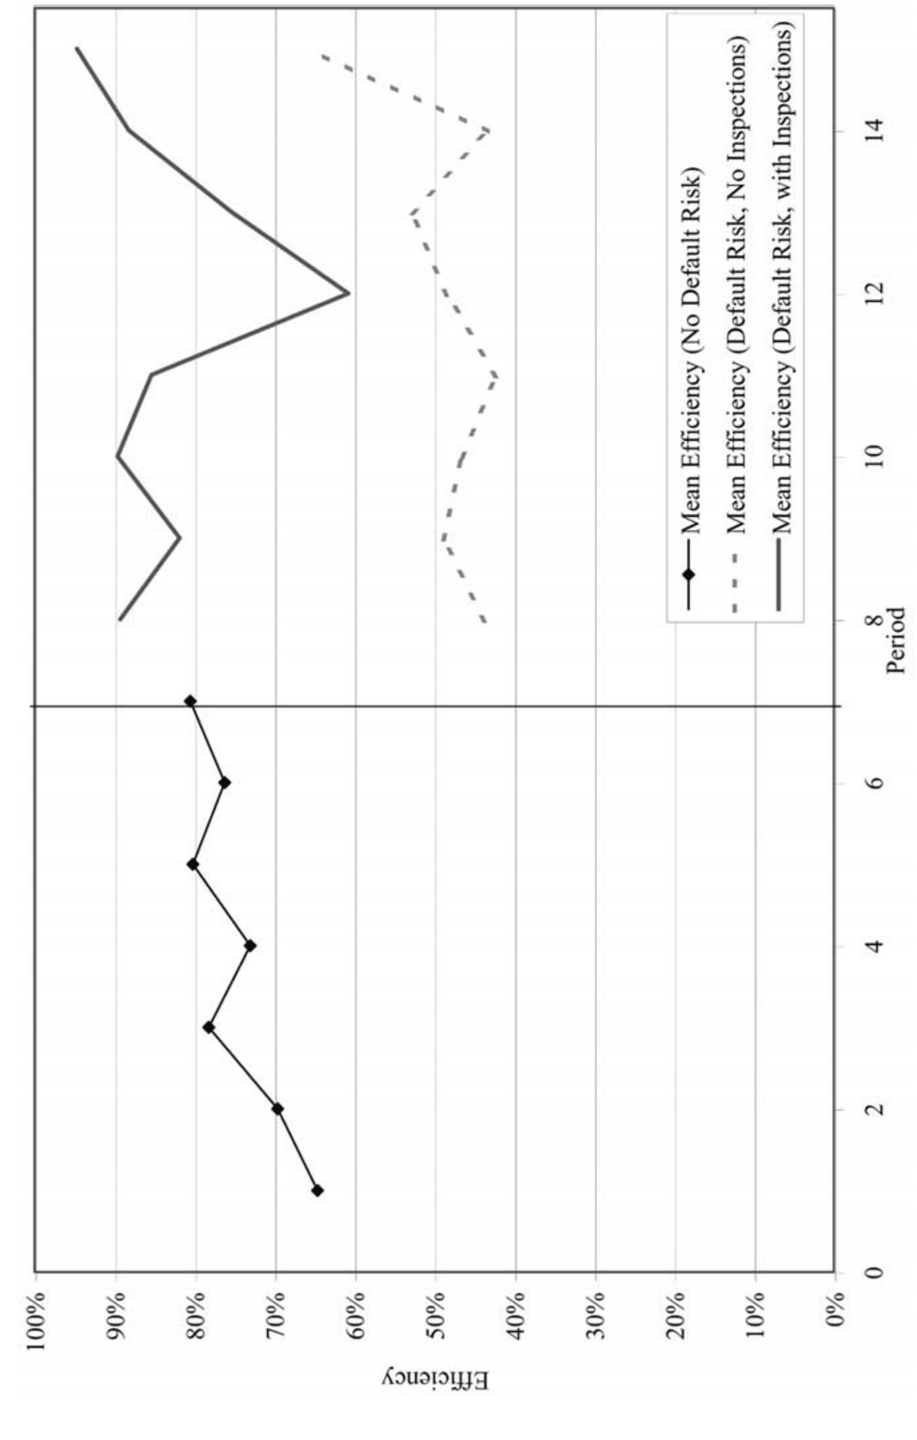

Figure 6. Mean trading efficiency.

ciency averages 83 percent in the Inspections treatment, compared to 51 percent in the No Inspections treatment.[17] Efficiency is higher in all three sessions with inspections (which featured a high reliability investment rate) than in any of the three sessions with no inspections (which featured a low reliability investment rate). Efficiency is therefore statistically greater with a higher reliability investment rate, consistent with Hypothesis 4 (Mann-Whitney $U = 0$, sample sizes $n_1 = n_2 = 3$; one-tailed p-value = 0.05). The low efficiency observed in the No Inspections treatment can thus be traced to its very low reliability investment rate (Figure 5), which led to very frequent allowance default and low profits.

4. Conclusion

This paper demonstrates that when allowance buyers are liable for shortcomings in emission reductions used to generate tradable allowances, they will pay a price premium for "inspected" allowances that are known to be reliable. Sellers of allowances therefore often make investments in emission reduction reliability, and when possible they usually reveal these investments to potential buyers to obtain higher prices. These preliminary results thus suggest how a buyer liability rule can function in an international GHG trading system to promote seller compliance through market incentives that generate appropriate price signals, especially if an inspection process evolves to address problems of incomplete information.

These results, of course, obtain in a laboratory environment that is a simple special case of the more complex emissions markets that will exist in the field. The goal here was not to implement realistic parameters that would approximate an international GHG emissions market, since the details of emissions markets are still being developed. The experiment achieves a more modest goal: to demonstrate how buyer liability rules can create the proper incentives to overcome a key seller moral hazard problem in GHG emissions trading in a simple setting. If the experiment failed to establish that the buyer liability rule promotes seller compliance in this simple case, it seems unlikely that it would generate appropriate incentives in more complex field environments.

The evidence presented here, however, is not sufficient to make strong recommendations for the design of GHG emissions markets. Additional modeling and laboratory experiments in more complex and realistic environments should help determine the robustness of these preliminary findings and provide guidance regarding the relationships between market rules and market performance. For example, it would be valuable to compare buyer liability and issuer liability rules when the international compliance authority has different degrees of enforcement power or different levels of enforcement costs. It would also be useful to relax some of the simplifying conditions used in this initial experiment. Subsequent experiments could include heterogeneous default rates, partial default, and noisy inspections of reliability investments, as well as other extensions to increase the parallelism between these laboratory markets and emissions trading markets that may be introduced in the field.

Acknowledgements

For helpful comments and discussions I would like to thank – without implicating – Chris Anderson, Rob Godby, Sally Kane, Andy Muller, Charlie Plott, Jay Shogren, Catherine Wilson, Tom Wilson, two anonymous referees, and participants in the Krannert behavioral economics brown-bag seminar. The Center for International Business Education and Research at Purdue University provided financial support.

Appendix A: A Simple Model of Emissions Trading with Default Risk

---

The purpose of the simple and highly stylized model in this section is to illustrate the implications of a buyer liability rule in a GHG emissions trading system. It is not intended to capture the rich detail of any particular trading scheme, but rather to highlight the importance of (observable or unobservable) investments in technology that affect the reliability of emissions controls. The main goal is to generate testable hypotheses for the experimental environment described in Section 2.

BASIC ASSUMPTIONS

Consider an environment in which all trades in the competitive output (energy services or products produced from energy) market x occur at the common price p_x, but allowances y^k originated by trader k may trade at potentially different prices denoted r^k, $k = 1, \dots, n$. The n traders could be participating countries, or individual firms that are allowed to issue excess emission allowances for sale on the international market. Denote the net x purchased or sold in the output market by trader i as x_i^m, the allowances that originate in allowance market y^k that trader i obtains or retains to cover emissions used in production as y_i^k, and trader i's initial endowment of emission rights as $\bar{y}_i \geq 0$. To simplify matters assume that each allowance y that is valid for offsetting emissions can be used to produce one and only one unit of output x (that is, the total valid allowances $\sum_{k=1}^{n} y_i^k$ = total allowable production of x by trader i).[18] Differences in output and emissions control technology across traders are captured by non-allowance production costs, which for trader i are summarized by the cost function $c_i(\bullet)$. Since we have assumed the valid allowances y = output x, we shall economize on notation by writing $c_i(y_i)$, where y_i is taken to mean the valid allowances. Output is valued by each trader i according to the output redemption schedule $V_i(x)$. This redemption value schedule summarizes the relevant output demand of this trader. We assume a downwardly sloping demand ($V_i'(\bullet) < 0$) and increasing marginal cost ($c_i''(\bullet) > 0$).

EQUILIBRIUM WITH NO ALLOWANCE DEFAULT RISK

Using this notation we can write the profits of trader i as

$$\pi_i = V_i(y_i + x_i^m) - c_i(y_i) + r^i \bar{y}_i - p_x x_i^m - \sum_{k=1}^{n} r^k y_i^k, \text{ where } y_i = \sum_{k=1}^{n} y_i^k. \tag{1}$$

The output redemption schedule is written $V_i(y_i + x_i^m)$ because trader i redeems the output she produces with the valid allowances y_i, net of the amount of output purchased or sold on the output market x_i^m. Because no default risk exists, the allowances are perfect substitutes. Consequently, in equilibrium the allowance prices do not differ according to their origination; i.e., $r^i = r^j = r\ \forall i, j$. Using this assumption and the definition of y_i given in (1) and simplifying:

$$\pi_i = V_i(y_i + x_i^m) - c_i(y_i) + r\bar{y}_i - p_x x_i^m - r y_i. \tag{2}$$

In the competitive equilibrium (CE) of this environment traders act as price takers and select the output to trade on the market (x_i^m) and allowances to use in production (y_i) to maximize (2). The first order conditions generate the familiar relationships $p_x = V_i'(y_i + x_i^m)$ (the output price equals the marginal benefit of the last unit of output consumed) and $r + c_i'(y_i) = V_i'(y_i + x_i^m)$ (the marginal cost of production equals the marginal benefit of production). Combining these equations we see that the equilibrium difference between the output and allowance prices is the marginal production cost: $c_i'(y_i) = p_x - r$.[19]

STOCHASTIC ALLOWANCE DEFAULT AND RELIABILITY INVESTMENTS

Next consider the case in which traders first decide whether to invest t per unit of emissions endowment to ensure that their emissions control efforts are reliable. These investments can be thought of as technology and maintenance investments so that any emissions control, conservation technologies and generation capacity (e.g., nuclear) used to "back" emissions allowances do not fail. For simplicity this reliability investment is all or nothing: either all of the emission allocation $\bar{y}_i$ is backed up with investments (at total cost $t\bar{y}_i$) or none of it is backed up (at total cost 0). If a trader does not make this reliability investment, with probability $\bar{\alpha}$ all the allowances allocated to him cannot be used to offset emissions. If the trader does make the reliability investment, the default probability is zero. We implement a buyer liability legal environment in which any traders who bought a defaulting seller's allowances also cannot use them to offset emissions in the case of default. Every trader's default realization is independent of the other traders' realizations, and we assume that the default probability $\bar{\alpha}$ is the same for all traders who do not make the reliability investment.[20]

This risk of allowance default leads to the following expected profit expression, where we restrict attention to the case of two traders i and j; the extension to all n traders is straightforward and simply requires more notation.

$$\begin{aligned} \mathrm{E}\pi_i = {} & r^i \bar{y}_i - r^i y_i^i - r^j y_i^j - p_x x_i^m - t\bar{y}_i \\ & + (1-\alpha_i)(1-\alpha_j)[V_i(y_i + x_i^m) - c_i(y_i)] \\ & + (1-\alpha_i)\alpha_j[V_i(y_i^i + x_i^m) - c_i(y_i^i)] \\ & + \alpha_i(1-\alpha_j)[V_i(y_i^j + x_i^m) - c_i(y_i^j)] + \alpha_i\alpha_j[V_i(x_i^m)]. \end{aligned} \tag{3}$$

The first line of this equation represents revenues and expenditures on output, allowances and reliability, which occur prior to and therefore do not depend on the default realization. The other four terms correspond to output redemption and production costs in the four possible default outcomes: (1) neither trader defaults, (2) only trader j defaults, (3) only trader i defaults, and (4) both traders default. The α_i and α_j may equal the $\bar{\alpha}$ or zero default probabilities, depending on information and beliefs. Trader i knows his own investment decision, so $\alpha_i \in \{0, \bar{\alpha}\}$. Trader i's beliefs regarding j's investment lead to $0 \leq \alpha_j \leq \bar{\alpha}$.

Consider now the risk neutral trader who seeks to maximize expected profits by choosing the output to trade x_i^m and the allowances to hold at the end of the period y_i^i and y_i^j, for either reliability decision. The first-order conditions to maximize (3) are:

$$\begin{aligned} x_i^m : p_x &= (1-\alpha_i)(1-\alpha_j)V_i'(y_i + x_i^m) + (1-\alpha_i)\alpha_j V_i'(y_i^i + x_i^m) \\ &\quad + \alpha_i(1-\alpha_j)V_i'(y_i^j + x_i^m) + \alpha_i \alpha_j V_i'(x_i^m). \end{aligned} \tag{4}$$

$$\begin{aligned} y_i^i : r^i &= (1-\alpha_i)(1-\alpha_j)[V_i'(y_i + x_i^m) - c_i'(y_i)] + \\ &\quad (1-\alpha_i)\alpha_j[V_i'(y_i^i + x_i^m) - c_i'(y_i^i)]. \end{aligned} \tag{5}$$

$$\begin{aligned} y_i^j : r^j &= (1-\alpha_i)(1-\alpha_j)[V_i'(y_i + x_i^m) - c_i'(y_i)] + \\ &\quad \alpha_i(1-\alpha_j)[V_i'(y_i^j + x_i^m) - c_i'(y_i^j)]. \end{aligned} \tag{6}$$

As with the baseline case with no default risk, (4) indicates that the trader chooses his expected production and output so that his expected marginal redemption value is equal to the output price p_x. If any allowances have some probability of default then expected output falls below the no default case. Consequently, equilibrium output prices will not fall and may rise compared to the no default case because we have assumed a downwardly-sloping output demand ($V_i'(\bullet) < 0$). The trader beliefs regarding reliability investments determine if and how much output prices rise.

Just as default risk increases expected marginal redemption values, default risk decreases expected marginal production costs. This is because marginal costs increase in output, and output cannot rise when allowance default is introduced. The first-order conditions indicate that the value of an allowance *that is known to be reliable* will not fall and may rise compared to the no default case. To see this, consider (5) in the case where trader i knows that she made the reliability investment so $\alpha_i = 0$. The value of these reliable allowances is

$$r^i = (1-\alpha_j)[V_i'(y_i + x_i^m) - c_i'(y_i)] + \alpha_j[V_i'(y_i^i + x_i^m) - c_i'(y_i^i)]. \tag{7}$$

The first term in this weighted average corresponds to the no default allowance valuation, while the second term corresponds to a higher valuation when $y_i^j > 0$ since $V_i'(\bullet) < 0$ and $c_i'(\bullet) > 0$.

When allowances are subject to default risk and some are inspected and are known to be reliable, the allowances of unknown reliability are likely to trade at a discount compared to the inspected allowances. This discount reflects their default risk. The (risk neutral) upper bound on this price discount is the $\bar{\alpha}$ default risk of an allowance without any reliability investment. This price discount would only be reasonable if subjects believed that an uninspected allowance was certainly unreliable.

If the reliability investment cost is low enough, however, subjects may have an incentive to make reliability investments. The marginal value of an allowance that defaults with probability α is $(1 - \alpha)\ [V'(\bullet) - c'(\bullet)]$. Therefore, the value of incurring the cost t to reduce the default probability of the marginal allowance from $\bar{\alpha}$ to 0 is $\bar{\alpha}[V'(\bullet) - c'(\bullet)]$. We chose parameters so that $t < \bar{\alpha}\ [V_i'(\bullet) - c_i'(\bullet)]$ for all i in the competitive equilibrium. Consequently, for the parameters employed in the experiment all traders have an incentive to invest in reliability.

Notes

1. This is equivalent to three-quarters of Russia's trade surplus in 1997, or all lending by the U.S. to Russia between 1990 and 1996. These figures are in 1992 dollars and are based on a projected price of $106 per ton of carbon, using Pacific Northwest National Laboratory's Second Generation Model. It assumes that both the former Soviet Union and Eastern Europe are treated as though 2010 rather than 1990 were the base year, so there are no excess emissions credits that can be used due to the poor economic growth rates in these regions. Allowing for additional excess credits based on 1990 emission levels (as specified in the Protocol) substantially increases the estimated sales revenues.
2. A system of escrow payments and issuer liability could also help deter overselling, if sellers do not receive allowance sales proceeds until compliance is verified at the end of the compliance period.
3. Buyer liability could be more complicated to administer than issuer liability, however (Mullins 1999). For example, in the case of noncompliance the question of which trades should be invalidated would need to be resolved (e.g., the last allowances sold? or a proportional discounting of all allowances sold?). Like issuer liability, a buyer liability rule also requires high transaction costs to enforce penalties. But if its market incentives are more successful in promoting compliance, a buyer liability rule could have lower total enforcement costs.
4. See Cason and Gangadharan (2002) for a laboratory posted offer study of environmental quality uncertainty.
5. This could include, for example, verification of operation reliability of nuclear or hydro generation facilities, or installation and maintenance of emission control technology on natural gas wellheads.
6. In about five percent of the periods a subject sold more output X than she produced, and in these cases she incurred a fine of 300 francs for each oversold unit of output. This fine exceeds the highest output redemption value for any subject, as shown below in Table I. These occurrences of insufficient production usually occurred because of allowance seller noncompliance (and therefore, allowance default), since production required valid allowances Y. Occasionally, subjects incurred these fines due to accounting errors.
7. This "small numbers" case corresponds roughly to a situation in which only national governments trade on behalf of Parties. Some commentators are concerned that Parties are more likely to exercise market power with small numbers of traders. [Various researchers have used laboratory experiments to explore the impact of market power in emission allowance markets, including Søberg (2000b), Muller et al. (2002), Godby (1999), Carlén (1999) and Ledyard and Szakaly-Moore (1994).] One way to reduce market power and simultaneously increase liquidity is to allow individual firms to trade emissions internationally (Mullins 1999; Barron 1999b). This could raise additional liability and noncompliance concerns, however, due to the poor record of domestic environmental law enforcement in many countries.
8. As a referee correctly points out, letting countries' marginal abatement costs be common information does not contradict the precept of privacy if subjects are paid a fraction of their country's trade gains, where this fraction is private information. This is the procedure use in Søberg

(2000a), for example, who also only provides subjects with the expected and not the exact abatement costs of other subjects.

9. The more complex (but more realistic) case in which only *some* of a subject's allowances default would reduce the risk faced by individual subjects, and would be a useful extension for a subsequent experiment.
10. We have obviously simplified matters by making this inspection perfectly accurate, costless, and publicly announced. Alternatively, we could have the investment observed with some noise, or observed only by some subset of participants who invest real resources to determine other's investments. These are interesting extensions for future research.
11. The only exception was a single subject who announced in three early periods that he did not make the reliability investment.
12. The high mean price in the first period was due to a single transaction at 500 francs (not shown), which may have been an error by a trader. Typographical errors are not uncommon in the first few periods of computerized double auction trading when subjects are not very familiar with the trading software.
13. In an earlier draft we presented these statistics for the later periods of each treatment only, to exclude the adjustment phase of market behavior during the initial periods of each treatment. Table III presents all periods. The results based on only later periods are qualitatively similar; the main difference is that the variance is usually lower for the later periods.
14. The six observations used in this test are generated by six independent groups of subjects, so they are statistically independent as required for a valid Wilcoxon test. Likewise, the Mann-Whitney tests presented below for the comparison between the Inspections and the No Inspections treatments employ only one statistically independent observation for each session, as required for a valid test.
15. This two-tailed p-value is marginally significant, but it cannot be smaller than 0.10 for this nonparametric test with this $n_1 = n_2 = 3$ sample size. Although in equilibrium subjects should invest in reliability regardless of whether they can submit to inspections, inspections should increase the reliability investment rate if they aid at all in helping allowance sellers establish reputations. This suggests that a proper alternative hypothesis is one-tailed, leading to a p-value of 0.05.
16. These fines could be thought of as transfers that allow for consumption elsewhere in the economy, and therefore it also seems reasonable to not subtract the fines from the total surplus realized by subjects in the market when calculating efficiency. Calculating efficiency in this alternative way does not change any qualitative conclusions, however. Efficiency is of course greater when fines are not subtracted from the realized surplus, particularly for the No Inspections treatment with default risk where fines were most common. But the efficiency calculated in this alternative way for all three sessions with no inspections is still less than the efficiency in any of the three sessions with inspections. This provides statistical support for Hypothesis 4 (Mann-Whitney test) as noted below for the version of efficiency that does account for losses due to fines incurred by traders.
17. The unusually low mean efficiency in period 12 in the Inspections treatment is due to the unusually low reliability investment rate for this period in session BI20406x, shown as treatment period 5 in Figure 5. This example highlights the importance of reliability for generating high efficiency in this environment.
18. Occasionally this assumption is violated in the laboratory environment with its small number of units per trader, when subjects sell more output x than they produce due to allowance default or accounting mistakes. These occurrences are heavily fined, and should not arise in equilibrium.
19. In the experiment we induce discrete redemption value and marginal cost schedules V_i and c_i, as well as the endowments $\bar{y}_i$. This permits a direct calculation of the competitive equilibrium, shown in Section 2.

20. A more realistic assumption that default probabilities vary across sellers would change the bounds on the beliefs described below equation (3), but would not change the main results substantively.

References

Baron, Richard (1999a), 'An Assessment of Liability Rules for International GHG Emissions Trading', International Energy Agency Information Paper, Energy & Environment Division, Paris (http://www.iea.org/clim/cop5/pubs/assess.pdf).

Baron, Richard (1999b), 'Market Power and Market Access in International GHG Emissions Trading', International Energy Agency Information Paper, Energy & Environment Division, Paris (http://www.iea.org/clim/cop5/pubs/market.pdf).

Baron, Richard (2000), 'Market Access Issues in International GHG Emissions Trading', International Energy Agency Information Paper, Energy & Environment Division, Paris (http://www.iea.org/clim/clichan/market.pdf).

Baron, Richard (2001), 'The Commitment Period Reserve', OECD Environment Directorate and International Energy Agency, Paris (http://www.oecd.org/pdf/M00020000/M00020141.pdf).

Ben-David, Shaul, David Brookshire, Stuart Burness, Michael McKee and Christian Schmidt (1999), 'Heterogeneity, Irreversible Production Choices, and Efficiency in Emission Permit Markets', *Journal of Environmental Economics and Management* **38**, 176–194.

Bohm, Peter and Björn Carlén (1999), 'Emission Quota Trade Among the Few: Laboratory Evidence of Joint Implementation Among Committed Countries', *Resource and Energy Economics* **21**, 43–66.

Carlén, Björn (1999), 'Large-Country Effects in International Emissions Trading: A Laboratory Test', Department of Economics, Stockholm University, mimeo.

Carlson, Dale, Charles Forman, Nancy Olmstead, John Ledyard, Charles Plott, David Porter and Anne Sholtz (1993a), 'An Analysis of the Information and Reporting Requirements, market Architectures, Operational and Regulatory Issue and Derivatory Instruments for RECLAIM: report Submitted to SCAQMD, Contract No. R-C93074', Technical Report, South Coast Air Quality Management District.

Carlson, Dale, Charles Forman, Nancy Olmstead, John Ledyard, Charles Plott, David Porter and Anne Sholtz (1993b), 'An Analysis and Recommendation for the Terms of the RECLAIM Trading Credit', Technical Report, South Coast Air Quality Management District.

Cason, Timothy and Lata Gangadharan (1998), 'An Experimental Study of Electronic Bulletin Board Trading for Emission Permits', *Journal of Regulatory Economics* **14**, 55–73.

Cason, Timothy and Lata Gangadharan (2002), 'Environmental Labeling and Incomplete Consumer Information in Laboratory Markets', *Journal of Environmental Economics and Management* **43**, 113–134.

Cason, Timothy and Charles Plott (1996), 'EPA's New Emissions Trading Mechanism: A Laboratory Evaluation', *Journal of Environmental Economics and Management* **30**, 133–160.

Corfee Morlot, Jan (1998), 'Ensuring Compliance with a Global Climate Change Agreement', OECD Information Paper, ENV/EPOC(95)5/REV1, Paris (http://www.oecd.org//env/docs/cc/epoc985r1.pdf).

Cronshaw, Mark and Jamie Brown Kruse (1999), 'An Experimental Analysis of Emissions Permits with Banking and the Clean Air Act Amendments of 1990', in R. M. Isaac and C. Holt, eds., *Research in Experimental Economics*, vol. 7. Stamford, Conn.: JAI Press, pp. 1–24.

Edmonds, Jae, Michael Scott, Joseph Roop and Christopher MacCracken (1999), 'International Emissions Trading & Global Climate Change: Impacts on the Costs of Greenhouse Gas Mitigation', Report prepared for the Pew Center on Global Climate Change, Washington, DC.

Franciosi, Robert, R. Mark Isaac, David Pingry and Stanley Reynolds (1993), 'An Experimental Investigation of the Hahn-Noll Revenue Neutral Auction for Emission Licenses', *Journal of Environmental Economics and Management* **24**, 1–24.

Godby, Robert (1999), 'Market Power in Emission Permit Double Auctions', in R. M. Isaac and C. Holt, eds., *Research in Experimental Economics*, vol. 7. Stamford, Conn.: JAI Press, pp. 121–162.

Godby, Robert and Jason Shogren (2002), '*Caveat Emptor* Kyoto', Department of Economics, University of Wyoming, mimeo.

Ledyard, John and Kristen Szakaly-Moore (1994), 'Designing Organizations for Trading Pollution Rights', *Journal of Economic Behavior and Organization* **25**, 167–196.

Muller, R. Andrew and Stuart Mestelman (1994), 'Emissions Trading with Shares and Coupons: A Laboratory Experiment', *The Energy Journal* **15**, 185–211.

Muller, R. Andrew and Stuart Mestelman, John Spraggon and Robert Godby (2002), 'Can Double Auctions Control Monopoly and Monopsony Power in Emissions Trading Markets?', *Journal of Environmental Economics and Management* **44**, 70–92.

Mullins, Fiona (1999), 'International Emissions Trading Under the Kyoto Protocol', OECD Information Paper, ENV/EPOC(99)18/FINAL, Paris (http://www.oecd.org//env/docs/cc/epoc9918.pdf).

Plott, Charles (1991), 'A Computerized Laboratory Market System and Research Support Systems for the Multiple Unit Double Auction', Social Science Working Paper 783, California Institute of Technology.

Smith, Vernon (1982), 'Microeconomic Systems as an Experimental Science', *American Economic Review* **72**, 923–955.

Søberg, Morten (2000a), 'Price Expectations and International Quota Trading: An Experimental Evaluation', *Environmental and Resource Economics* **17**, 259–277.

Søberg, Morten (2000b), 'Imperfect Competition, Sequential Auctions and Emissions Trading: An Experimental Evaluation', Discussion Paper, Statistics Norway, Research Department.

Werksman, Jacob (1999), 'Responding to Non-Compliance under the Climate Change Regime', OECD Information Paper, ENV/EPOC(99)21/FINAL, Paris (http:/www.oecd.org//env/docs/cc/epoc9921.pdf).