NBER WORKING PAPER SERIES

PRICES VS. QUANTITIES VS. TRADABLE QUANTITIES

Roberton C. Williams III

Working Paper 9283
http://www.nber.org/papers/w9283

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2002

The author would like to thank Don Fullerton, Larry Goulder, Preston McAfee, Robert Mohr, Jonathan Rubin, Tom Tietenberg, Marty Weitzman, and seminar participants at the University of Texas, the Lonestars Conference, and the NBER Summer Institute for helpful comments. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research

Prices vs. Quantities vs. Tradable Quantities
Roberton C. Williams III
NBER Working Paper No. 9283
October 2002
JEL No. L51, D81, Q2

ABSTRACT

This paper extends Weitzman's (1974) seminal paper comparing price and quantity instruments for regulation to consider a third option: tradable quantity regulations, such as tradable permits. Contrary to what prior work has suggested, fixed quantities may be more efficient than tradable quantities if the regulated goods are not perfect substitutes, even when trading ratios are based on the ratio of expected marginal benefits between goods, not simply one-for-one. Indeed, when benefits are independent across goods, or when the goods are complements, tradable quantities are never the most efficient instrument. This theory is applied to dynamic pollution problems, and suggests that permit banking should be allowed for stock pollutants, but not for flow pollutants. These results indicate that many regulations, including the current sulfur dioxide trading program and proposed greenhouse gas regulations, are inefficient.

Roberton Williams
Department of Economics
University of Texas
Austin, TX 78712
and NBER
rwilliam@eco.utexas.edu

I. Introduction

This paper examines the efficiency of tradable quantity regulations: regulations that fix the aggregate quantity of a set of goods, but provide flexibility in how to divide that aggregate quantity among the goods in the set. Such policies are typically implemented via a system of tradable permits; in the case of pollution regulation, for example, the goods in question are the pollution emissions from each of a group of firms, and a system of tradable emissions permits would cap total emissions. Tradable quantities are already widely used in environmental regulation, and are increasingly considered as an option in other regulatory settings.[1]

This increasing popularity is due in part to efficiency considerations. Economists typically view tradable quantities as being more efficient than fixing the quantity of each good. However, this paper shows that is not generally true; unless the goods are perfect substitutes, fixed quantities may be more efficient than tradable quantities, even if trading occurs at ratios that reflect the expected ratios of marginal benefits between goods, not at a one-to-one ratio. These results indicate that many existing and proposed regulations are inefficient, including the current sulfur dioxide trading program and proposed regulations for nitrous oxides and greenhouse gases.

The key difference between this paper and prior research on the efficiency of tradable permits is that this paper explicitly models asymmetric information; firms know their costs, but the regulator does not. To do so, it uses a framework based on Weitzman's (1974) seminal paper, which compared two archetypal regulations: setting a price for each good versus fixing the quantity of each good. This paper extends that comparison to include tradable quantities. Asymmetric information is an essential part of this comparison; without it, price, quantity, and tradable quantity instruments yield identical outcomes.

[1] Uses of tradable permits include limits on foreign imports to New Zealand (see McAfee *et al*., 1999), sulfur dioxide under the 1990 Clean Air Act Amendments, and smog under the Los-Angeles-area RECLAIM program. Such permits have also been proposed for a wide range of other pollutants, including greenhouse gases and nitrous oxides, as well as for non-environmental problems, including price increases leading to inflation (Vickrey, 1992 and 1993) and national budget deficits within the European Union (Casella, 1999). Furthermore, some other policies are functionally equivalent, in that they fix an aggregate quantity but allow the market to determine the allocation across particular goods. One example would be a purchase contract that would be satisfied by providing a certain combined total of any of several different goods; this could be done explicitly, or simply by writing a contract with very loose specifications.

Weitzman showed that the relative efficiency of price regulation versus quantity regulation depends on the relative slopes of the marginal benefit and marginal cost curves. If the marginal cost curve is steeper than the marginal benefit curve, price regulation will be more efficient, whereas if the marginal benefit curve is steeper, quantity regulation will be more efficient. The intuition for this result is that price regulation provides firms with more flexibility. This flexibility reduces expected costs, but also reduces expected benefits. The steeper the marginal cost curve, the greater the expected cost savings, but the steeper the marginal benefit curve, the greater the reduction in expected benefits.[2]

Weitzman mentions tradable quantities in a footnote (p. 490), stating that they would be better than fixed quantities in the case in which all goods are identical, and that in this case the question reduces to one of whether the aggregate quantity should be controlled via price or quantity regulation. However, the paper does not explicitly show this, and does not discuss the efficiency of tradable quantities when the goods are not identical.

Subsequent work either ignores tradable quantities altogether, or it assumes that tradable quantities will be superior to fixed quantities, and then applies a Weitzman-style analysis at the aggregate level to compare prices and tradable quantities.[3] Neither approach is satisfactory; the former ignores a potentially attractive option, while the latter–as this paper will show–is not valid except in the special case in which all the goods are perfect substitutes.

This paper explicitly compares tradable quantities to prices and to fixed quantities in a setting with asymmetric information. It shows that the choice between tradable quantities and prices or fixed quantities depends on the relative slopes of the marginal cost and marginal benefit curves–as was true for the choice between prices and fixed quantities–and also on the degree of substitutability or complementarity between the goods. If the goods are perfect substitutes, tradable quantities are always

[2] Adar and Griffin (1976) provide a clear graphical analysis of this result.

[3] For example, Newell and Pizer (2002) and Hoel and Karp (1999) take the former approach, while Oates, Portney, and McGartland (1989) take the latter. An exception is the simultaneous and independent paper by Yates (2000), which explicitly compares tradable and non-tradable pollution permits in a framework similar to the one used in this paper. However, this paper considers a broader set of instruments and a broader range of cases, which allows it to develop an intuition that is absent in Yates.

more efficient than fixed quantities. In this case, the decision between tradable quantities and prices depends on the relative slopes of the aggregate marginal cost and marginal benefit curves.

When the goods are substitutes, but not perfect substitutes, then each of the three instruments could be the most efficient. If the marginal cost curve is substantially steeper than the marginal benefit curve, prices will be most efficient. If the marginal benefit curve is substantially steeper, then fixed quantities will be most efficient. In an intermediate range, where the slopes are similar, tradable quantities will be most efficient. The less substitutable the goods are, the smaller that range will be. Finally, if marginal benefits are independent across goods, or if the goods are complements, then tradable quantities are never the most efficient of the three instruments.

These results represent a significant theoretical contribution, and have substantial practical importance. To provide one example of the latter, this paper briefly considers the question of whether pollution permits should be bankable (whether firms should be allowed to emit less pollution in one year in return for being allowed to emit more in future years). It shows that several current and proposed policies (including the sulfur dioxide permit program and the Kyoto global warming agreement) are inefficient, in that they allow banking in cases when it shouldn't be allowed, or don't allow it in cases when it should be allowed.

The next section of the paper develops a model of regulation of production of a set of goods, derives expressions for the relative efficiency of the three instruments, and then considers the implications of these expressions. The third section demonstrates an application of these results to the problem of whether pollution permits should be bankable. The final section offers conclusions and suggestions for future research.

II. The Model

This section develops a simple model of regulation of production of a set of goods, and uses that model to investigate the relative efficiency of three types of regulatory instruments: prices, fixed quantities, and tradable quantities. The structure of this model is similar to the model used in

Weitzman (1974), though it considers tradable quantities in addition to prices and fixed quantities.

A. Assumptions

A set of *N* goods is assumed. The cost of producing good *i* is given by $C_i(q_i, \theta_i)$, where q is the vector of quantities of each good and θ is a vector of random variables. The total benefit is given by $B(q)$.[4] The functions $C_i(\cdot)$ and $B(\cdot)$ are assumed to be continuous and twice-differentiable. To assure a unique internal solution, I adopt the standard assumptions on C_i and *B*; specifically, C_i is increasing and strictly convex in q_i, *B* is increasing and concave, and for any θ_i, marginal costs exceed marginal benefits for q_i sufficiently large, and are less than marginal benefits for $q_i = 0$.[5] Firms are assumed to minimize costs, subject to any regulatory constraints. In the absence of regulation, then, production will be set such that the marginal cost is zero.

The regulator can choose one of three regulatory instruments–a set of prices, a set of quantities, or a system of tradable quantities.[6] Firms will set production to minimize costs, subject to meeting the regulatory constraint. While firms have perfect information about their costs, the regulator does not know the realization of θ when setting the regulation. For the price instrument, this implies the first-order condition

---

[4] For simplicity, I ignore uncertainty in benefits. Weitzman (1974) showed that if cost uncertainty and benefit uncertainty are not correlated, then such uncertainty has no effect on the choice between price and quantity regulation. Stavins (1996) further analyzed the case in which the uncertainty is correlated. Analogous results hold here, so this paper's results would be unchanged if benefits were also uncertain, as long as the benefit uncertainty is not correlated with the cost uncertainty.

[5] For consistency with prior work, especially Weitzman (1974), this model assumes that each good has a positive marginal cost and benefit; thus, for pollution regulation, each good is not pollution from a particular firm, but the reduction in pollution by that firm.

[6] Several papers (for example, Roberts and Spence, 1976, Weitzman (1978), and Kaplow and Shavell, 1997), have suggested more complex regulatory instruments that will generally be more efficient than the simple instruments considered here. Greenwood and McAfee (1991) derive the optimum over all possible regulatory mechanisms, though they consider only the special case in which the benefit function is additively separable. In practice, those more complex instruments are rarely used, and so this paper focuses only on simple regulatory instruments.

(1) $$\frac{\partial C_i}{\partial q_i} = p_i$$

where p is the vector of prices set by the regulator. Under the quantity instrument, firms set production equal to the required quantity.

(2) $$q_i = \overline{q}_i$$

where $\overline{q}$ is the vector of quantities set by the regulator.

Under the system of tradable quantities, the first-order condition is

(3) $$\frac{\partial C_i}{\partial q_i} = r_i \lambda$$

and the overall quantity constraint is

(4) $$\sum_i r_i q_i = Q$$

where *r* is a vector of trading ratios[7] (with r_i units of good *i* being tradable for r_j units of good *j*), *Q* is the total quantity required, and λ is the shadow price for that total quantity constraint.[8] Note that *r* and *Q* are set by the regulator, while λ will be determined by *r*, *Q* and the cost functions.

The goal of the regulator is to maximize the expected value of benefits minus costs

$$E\left[B(q) - \sum_i C_i(q_i, \theta_i)\right].$$[9]

In the absence of uncertainty on the part of the regulator, all three instruments would be set to

[7] Many studies have shown that when the regulated goods are not identical–for example, when pollution damages differ by location or time period–tradable quantity regulations should not allow one-to-one trades, but instead should allow trades at a ratio that reflects the ratio of marginal benefits between goods. See, for example, McGartland and Oates (1985) in the context of emissions permit trading, Kling and Rubin (1997) or Leiby and Rubin (2000) in the context of emissions permit banking, or Casella (1999) in the context of tradable limits on EU budget deficits. All of these studies assume that the regulator has perfect information.

[8] In the case of a tradable emissions permit program, *Q* is the total number of permits issued, λ is the market price of a permit, *r* is the number of permits required for one unit of emissions, and equation (4) is the market-clearing condition for the permit market. In other contexts, such as a contract with loose specifications, λ would be a shadow price within the firm.

[9] Note that distributional or other considerations could be incorporated into the benefit and cost functions, so the assumption that the regulator maximizes benefits minus costs does not necessarily imply that the regulator is concerned only with economic efficiency.

achieve the same optimal vector of quantities. For the price instrument, this would require the regulator to set the price for each good such that the price equals the marginal benefit from that good.

(5) $$p_i = \left.\frac{\partial B}{\partial q_i}\right|_{q=q^*} = \left.\frac{\partial C_i}{\partial q_i}\right|_{q=q^*}$$

Under the system of tradable quantities, the trading ratios would be set proportional to the marginal benefit from each good; thus, without loss of generality, they can be set equal to the marginal benefit

(6) $$r_i = \left.\frac{\partial B}{\partial q_i}\right|_{q=\bar{q}}$$

The optimal total quantity would be given by

(7) $$Q = \sum_i r_i q_i^*$$

Without uncertainty, then, each instrument would achieve a first-best outcome.[10] With uncertainty, however, the three instruments will have different effects. Let $\bar{q}$ denote the *ex ante* optimal vector of quantities–the vector of quantities that maximizes the expected value of benefits minus costs–defined by

(8) $$E\left.\left(\frac{\partial C_i}{\partial q_i}\right)\right|_{q=\bar{q}} = \left.\frac{\partial B}{\partial q_i}\right|_{q=\bar{q}}$$

In order to proceed further, assume that the uncertainty is sufficiently small to justify a second-order approximation for the cost and benefit functions in the neighborhood of $\bar{q}$. Without loss of generality, quantities will be normalized such that at $\bar{q}$, the marginal benefit and the expected marginal cost each equal 1. This implies that

(9) $$C(q,\theta) \approx C(\bar{q},\theta) + \sum_i (1 + \alpha_i(\theta))\hat{q}_i + \frac{1}{2}\sum_i \gamma_{ii}\hat{q}_i^{\,2}$$

and

[10] The price instrument could also achieve a first-best outcome if prices were not fixed, but could vary based on the quantity of production of each good. Note, however, that this would be more complex than just allowing a non-linear price schedule, because the price for a particular good would have to depend not just on the quantity of that good, but also on the quantities of all other goods as well.

(10) $$B(q) \approx B(\bar{q}) + \sum_i \hat{q}_i + \frac{1}{2}\sum_i \sum_j \beta_{ij}\hat{q}_i\hat{q}_j$$

where $\hat{q}$ is the deviation in quantity from $\bar{q}$, given by

(11) $$\hat{q}_i = q_i - \bar{q}_i$$

β and γ are the matrices of the second derivative of the benefit function and of the expected value of the second derivative of the cost function, respectively, each evaluated at $\bar{q}$.

(12) $$\gamma_{ii} = E\left(\frac{\partial^2 C_i}{\partial q_i^2}\right)\Bigg|_{q=\bar{q}}$$

(13) $$\beta_{ij} = \frac{\partial^2 B}{\partial q_i \partial q_j}\Bigg|_{q=\bar{q}}$$

and $\alpha_i(\theta)$ is a function that translates the vector of random variables θ into a vertical shift in the marginal cost curve for good *i*. The unit normalization implies that $\alpha_i(\theta)$ has an expected value of zero.

(14) $$E(\alpha_i(\theta)) = 0$$

Finally, the model assumes that that the distribution of $\alpha_i(\theta)$ is independent across goods.

B. Quantities Under Different Instruments

Under fixed quantities, production is simply equal to $\bar{q}$. But production will generally deviate from $\bar{q}$ under each of the other two instruments. Taking a derivative of the approximation to the cost function (9) gives an approximation to marginal cost for good *i*.

(15) $$\frac{\partial C_i}{\partial q_i} \approx 1 + \alpha_i(\theta) + \gamma_{ii}\hat{q}_i$$

Combining the approximation to marginal cost (15) with the first-order condition for the price instrument and rearranging give an approximation for the deviation from $\bar{q}$ under the price instrument

(16) $$\hat{q}_i \approx \frac{p_i - 1 - \alpha_i(\theta)}{\gamma_{ii}}$$

The optimal price is equal to the marginal benefit at quantity $\bar{q}$, which, as noted earlier, is normalized to 1. Thus,

(17) $$p_i \approx 1$$

Substituting (17) into (16) yields

(18) $$\hat{q}_i^p \approx \frac{-\alpha_i(\theta)}{\gamma_{ii}}$$

where $\hat{q}_i^p$ is the deviation from $\bar{q}$ that will result under the *ex ante* optimal price.

A similar process yields an expression for quantity under a system of tradable quantities. Substituting the expression for marginal cost (15) into the first-order condition for tradable quantities (3) and rearranging yield

(19) $$\hat{q}_i \approx \frac{r_i\lambda - 1 - \alpha_i(\theta)}{\gamma_{ii}}$$

The tradable quantity instrument will set r (the vector of trading ratios) to be proportional to the marginal benefit at quantity $\bar{q}$, which was normalized to 1. Without loss of generality, then, the trading ratios are also normalized to 1.

(20) $$r_i \approx 1$$

The optimal total quantity will then equal the sum over all goods of $\bar{q}$.

(21) $$Q = \sum_i \bar{q}_i$$

Substituting (19), (20), and (21) into (4) and rearranging (using (11)) yield an expression for the shadow price of the quantity constraint under the tradable quantity system (which will be the equilibrium permit price in the case of tradable permits)

(22) $$\lambda \approx 1 + \psi \sum_i \frac{\alpha_i(\theta)}{\gamma_{ii}}$$

where ψ is the slope of the aggregate marginal cost curve, which is equal to the horizontal sum of the marginal cost curves over all goods (in the case of tradable permits, this will be the same as the market demand curve for emissions permits).

(23) $$\psi = 1/\sum_i \frac{1}{\gamma_{ii}}$$

Substituting (20) and (22) into (19) yields

(24) $$\hat{q}_i^t \approx \frac{\psi}{\gamma_{ii}} \sum_j \frac{\alpha_j(\theta)}{\gamma_{jj}} - \frac{\alpha_i(\theta)}{\gamma_{ii}}$$

where $\hat{q}_i^t$ is the deviation from $\overline{q}$ that will result under the system of tradable quantities. Note that if γ_{ii} is equal for all *i* (an assumption that will be maintained for much of the analysis that follows), this expression reduces to $\hat{q}_i^t \approx -(\alpha_i(\theta) - s)/\gamma_{ii}$, where *s* is the sample mean of the α's, given by $s = \sum_j \alpha_j(\theta)/N$. This is similar to the quantity deviation under the price instrument (18). However, under the price instrument, the quantity deviation for good *i* depends on the marginal cost deviation for good *i*, whereas under the tradable quantity instrument, the quantity deviation for good *i* depends on the difference between the marginal cost deviation for good *i* and the average of the marginal cost deviations over all goods.

C. Comparative Advantages of Different Instruments

Following Weitzman (1974), define the comparative advantage of one instrument over another as the difference between the two instruments of the expected value of benefits minus costs.[11] The comparative advantage of one of the other two instruments relative to fixed quantities is thus given by

(25) $$\Delta = E\left[\left(B(q) - C(q,\theta)\right) - \left(B(\overline{q}) - C(\overline{q},\theta)\right)\right]$$

Substituting the approximations to the cost and benefit functions (12) and (13) into (25) and simplifying, using (8) and (11) give

[11] In calculating the comparative advantage of one instrument over another, the model assumes that the regulatory parameters–quota levels, tax rates, trading ratios, and the number of permits allocated–are set optimally *ex ante*. If the parameters for one or more of the instruments are not set optimally, this may change the relative efficiency of the various instruments.

(26) $$\Delta \approx E\left[-\sum_i \alpha_i(\theta)\hat{q}_i + \frac{1}{2}\sum_i\sum_j \beta_{ij}\hat{q}_i\hat{q}_j - \frac{1}{2}\sum_i \gamma_{ii}\hat{q}_i^2\right]$$

For now, assume that β and γ are symmetric across goods. This assumption will be relaxed later, but it is useful now in yielding clear, simple results. Specifically, assume that $\gamma_{ii} = \gamma \,\forall\, i$, that $\beta_{ii} = \beta \,\forall\, i$, and that $\beta_{ij} = \phi\beta \,\forall\, i \neq j$.[12] The parameter ϕ determines whether each good is a complement or a substitute for each of the other goods. When ϕ is positive, the goods are substitutes. For $\phi = 1$, they are perfect substitutes. When $\phi = 0$, the marginal benefit of a particular good is independent of the quantities of the other goods; thus, the goods are neither complements nor substitutes. Finally, if ϕ is negative, the goods are complements.

Using these assumptions, substituting the expression for the quantities under the ex ante optimal prices (18) into (26) and simplifying give an expression for the comparative advantage of prices relative to quantities.

(27) $$\Delta^{PQ} \approx \frac{\beta + \gamma}{2\gamma^2}\sum_i \sigma_i^2$$

where σ_i^2 is the variance of the marginal cost of good *i* at quantity $\bar{q}$, which is approximated by the variance of $\alpha_i(\theta)$

(28) $$\sigma_i^2 = E\left[\left(\left.\frac{\partial C_i}{\partial q_i}\right|_{q=\bar{q}} - E\left(\left.\frac{\partial C_i}{\partial q_i}\right|_{q=\bar{q}}\right)\right)^2\right] \approx E\left[\alpha_i(\theta_i)^2\right]$$

Equation (27) corresponds to the expression for the comparative advantage of prices over quantities from Weitzman (1974). This expression implies prices will be more efficient than quantities when the slope of the marginal cost curve is greater than the slope of the marginal benefit curve. Quantities will be more efficient when the marginal benefit curve is steeper. Prices allow firms more flexibility, which reduces expected costs. This is especially important when marginal cost is very sensitive to quantity. However, to the extent that marginal benefit is sensitive to quantity, this flexibility

[12] Note that, because units are normalized such that the first derivatives of the benefit and cost functions are equal across goods, this assumption that the second derivatives are equal is less restrictive than it might at first appear.

reduces expected benefits.

Substituting (24) into (26) and simplifying, using the assumption that costs and benefits are symmetric across goods, give an expression for the comparative advantage of tradable quantities relative to fixed quantities.

$$(29) \qquad \Delta^{TQ} \approx \frac{\beta(1-\phi)+\gamma}{2\gamma^2}\left(\frac{N-1}{N}\right)\sum_i \sigma_i^2$$

This expression is somewhat similar to the expression for the comparative advantage of prices relative to quantities (27); when the marginal benefit curve is relatively steep, fixed quantities are more efficient, whereas when the marginal cost curve is relatively steep, tradable quantities are more efficient. This reflects the fact that, like prices, tradable quantities offer firms more flexibility than do fixed quantities.

Unlike in expression (27), however, the degree of complementarity or substitutability between different goods matters. The more substitutable the goods are for each other (the closer ϕ is to 1), the more efficient tradable quantities will be relative to fixed quantities. This occurs because under tradable quantities, producing more of one good allows less production of other goods. No such linkage occurs under prices or fixed quantities. If the goods are close substitutes for each other, this is efficient; the more of one good is produced, the lower the efficient quantity will be for each of the other goods. If the marginal benefit of a particular good is independent of the quantities of the other goods (ϕ is 0) this linkage provides no advantage. And if the goods are complements, then this linkage is a disadvantage; in that case, the more of a particular good is produced, the higher will be the efficient quantities of the other goods.

Finally, the comparative advantage of tradable quantities relative to prices is equal to the difference between the comparative advantage of tradable quantities relative to fixed quantities and the comparative advantage of prices relative to fixed quantities.

$$(30) \qquad \Delta^{TP} = E\left[\left(B(q^t) - C(q^t,\theta)\right) - \left(B(q^p) - C(q^p,\theta)\right)\right] = \Delta^{TQ} - \Delta^{PQ}$$

Substituting the expressions for Δ^{PQ} (27) and Δ^{TQ} (29) into (30) and simplifying give

$$(31) \qquad \Delta^{TP} \approx -\frac{1}{N}\frac{\beta\left[1+\phi(N-1)\right]+\gamma}{2\gamma^2}\sum_i \sigma_i^2$$

Again, this expression depends on the relative slopes of the marginal benefit and marginal cost curves. Tradable quantities provide less flexibility than do prices (because while the amount of each good can vary, the total amount of all goods cannot), and thus tradable quantities are relatively efficient when the marginal benefit curve is steep relative to the marginal cost curve. And again, the more substitutable the goods are for each other, the more efficient tradable quantities will be.

These three expressions for comparative advantage (expressions (27), (29), and (31)) represent the key findings of this paper. Comparing these three expressions shows which instrument will be the most efficient in any particular situation. First, consider the case in which the goods are substitutes for each other, so $\phi > 0$. In this case, each of the three instruments could be the most efficient, depending on the relative slopes of the marginal benefit and marginal cost curves. When the marginal cost curve is steep relative to the marginal benefit curve–when $\gamma > -\beta\left[1+\phi(N-1)\right]$–prices will be the most efficient instrument. When the marginal benefit curve is relatively steep–when $\gamma < -\beta(1-\phi)$–fixed quantities will be the most efficient instrument. And for some middle range, when the slopes of the two curves are similar–when $-\beta\left[1+\phi(N-1)\right] > \gamma > -\beta(1-\phi)$–tradable quantities will be most efficient. Note that the size of this range depends on how substitutable the goods are for each other. When they are very close substitutes, there will be a wide range of slopes for which tradable quantities would be most efficient.

In the extreme case in which the goods are perfect substitutes, $\phi = 1$. In this case, tradable quantities will always be more efficient than fixed quantities, regardless of the slopes of the marginal benefit and marginal cost curves. The intuition behind this result is simple. When benefits take this form, only the expected cost differs between the two instruments; the benefits are the same. Tradable quantities have a lower expected cost, because they equalize marginal cost across goods, whereas fixed quantities do not allow such flexibility.[13]

In this perfect substitutes case, tradable quantities will be more efficient than prices if $\frac{\gamma}{N} < -\beta$.

The left-hand term, $\frac{\gamma}{N}$, is the slope of the market-wide marginal cost curve, which is the horizontal sum of the marginal cost curves over all goods. Given the symmetry assumption, that slope is equal to the slope of the marginal cost curve for each good, divided by the number of goods. Thus, in this case, tradable quantities are more efficient than prices if the marginal benefit curve is steeper than the market-wide marginal cost curve. If the marginal cost curve is steeper, then prices are more efficient.

This is an analogue of Weitzman's (1974) result comparing price and quantity instruments; quantity regulation is more efficient if the marginal benefit curve is steeper than the marginal cost curve, while price regulation is more efficient if the marginal cost curve is steeper. The difference is that for multiple goods regulated through a system of tradable quantities, the appropriate marginal cost curve is the market-wide marginal cost curve, whereas when using a fixed quantity instrument to regulate a single good, it is the marginal cost curve for that good.

In contrast, when goods are only very weakly substitutable for each other, there will be only a very small range of slopes for which tradable quantities would be the most efficient instrument. Indeed, when marginal benefits are independent of the quantities of other goods ($\phi = 0$) or the goods are complements ($\phi < 0$) tradable quantities will never be the most efficient instrument, regardless of the relative slopes of the marginal benefit and marginal cost curves. In this case, when the marginal benefit curve is steeper than the marginal cost curve, fixed quantities will be the most efficient instrument, and when the marginal cost curve is steeper, prices will be the most efficient instrument.

In such cases, the way that tradable quantities link the production of the different goods provides

[13] An alternative intuition is that under tradable quantities, the goods are perfect substitutes to the firms in meeting the regulation. Thus, when the goods are perfect substitutes in the benefit function, tradable quantities match firms' incentives with marginal benefits.

no advantage (and is actually a disadvantage when the goods are complements).[14] Thus, when flexibility is beneficial (when the marginal cost curve is relatively steep), the most efficient instrument will be the one that provides the most flexibility–prices. When flexibility is harmful (when the marginal benefit curve is relatively steep), the most efficient instrument will be the one that provides the least flexibility–fixed quantities. Of course, this assumes that all three types of regulation are feasible. If, for example, price instruments are infeasible (due to political considerations, perhaps), then tradable quantities may be the best remaining option even when the goods are complements.

These results differ substantially from prior work in the context of tradable permits.[15] That literature has suggested that as long as trading ratios are set based on the ratio of marginal benefits between the goods–not simply allowing one-for-one trades–tradable quantities will always dominate fixed quantities, even when the goods are not perfect substitutes. The problem is that unless the goods are perfect substitutes, the ratio of marginal benefits between any two goods will depend on the quantities of those two (and perhaps other) goods. Thus, even if the trading ratios are set optimally *ex ante*, they will in general not be optimal *ex post*. Because none of that prior work explicitly incorporated asymmetric information, it did not recognize this issue.

These results have far-reaching implications. They suggest, for example, that tradable pollution permits will be very efficient for a pollutant with global effects, such as greenhouse gas emissions, where emissions from different locations are perfect substitutes, but generally should not be used to regulate

[14] When the goods are complements, it might be efficient to link the production of different goods, but not in the way that tradable quantities link production. This could be accomplished by setting a price for production of a bundle consisting of a certain amount of each good, rather than a separate price for each good. Thus, to the regulated firms, the goods would be perfect complements, just as tradable quantities cause the goods to be perfect substitutes to the regulated firms. Such a "bundle price" policy logically completes the set of archetypal instruments. Fixed quantities fix both the total amount of production and the composition of that production, prices allow both to vary, tradable quantities fix the total amount, but allow the composition to vary, and this "bundle price" instrument fixes the composition but allows the total amount to vary. Such a policy would be more efficient the more complementary the goods are to each other. When the goods are perfect complements, it would dominate prices, in the same way that tradable quantities dominate fixed quantities when the goods are perfect substitutes. This instrument is rare in environmental policy, but is common in some other contexts, such as purchase contracts for complementary goods.

[15] See, for example, McGartland and Oates (1985), Kling and Rubin (1997), Leiby and Rubin (2000), or Casella (1999).

pollutants with very localized effects. Similarly, firms should be allowed to trade reduced emissions of one pollutant for increased emissions of a different pollutant only if the two are close substitutes.[16] In the context of purchasing contracts, these results suggest that firms should write very flexible contracts when purchasing a set of closely substitutable goods, but that such contracts would be highly inefficient when purchasing complementary goods.

D. Comparative Advantages of Different Instruments in the General Case

Removing the assumption that benefits and costs are symmetric across goods causes the expressions for the comparative advantages of the different instruments to become somewhat more complex. Following the same steps used to derive equation (27) yields an expression for the comparative advantage of prices over fixed quantities in this general case

$$(32) \qquad \Delta^{PQ} \approx \sum_i \frac{\sigma_i^2}{2\gamma_{ii}^2}\left(\beta_{ii} + \gamma_{ii}\right)$$

This expression is very similar to (27), but now, because the slopes of the marginal cost and benefit functions differ across goods, the comparative advantage of prices relative to fixed quantities depends on the weighted averages of those slopes. The basic intuition behind the result is unchanged; prices provide more flexibility than fixed quantities, and when the average marginal cost curve is steeper than the average marginal benefit curve, that flexibility is an advantage.

Following the same steps used to derive equation (29) yields an expression for the comparative advantage of tradable quantities over fixed quantities in this general case

$$(33) \qquad \Delta^{TQ} \approx \sum_i \frac{\sigma_i^2}{2\gamma_{ii}^2}\left[\left(\beta_{ii} + \gamma_{ii}\right)\left(1 - \frac{\psi}{\gamma_{ii}}\right) + \psi\left(\frac{\beta_{ii}}{\gamma_{ii}} - \sum_j \frac{\beta_{ij}}{\gamma_{jj}}\right) - \psi^2 \sum_j \sum_k \frac{\beta_{ij} - \beta_{jk}}{\gamma_{jj}\gamma_{kk}}\right]$$

The first term in this expression is the value of the extra flexibility afforded by tradable quantities relative to fixed quantities, and is generally similar to expression (32), depending on the weighted

---

[16] Montero (2001) analyzes regulation of multiple pollutants through tradable permits. Its model and results correspond to those in this paper, for the special case in which ϕ=0, and it does not consider pollution taxes.

averages across goods of the slopes of the marginal cost and marginal benefit curves. When the average marginal cost curve is steeper than the average marginal benefit curve, that flexibility is an advantage, and so tradable quantities are relatively efficient.

The second term reflects whether the goods are complements or substitutes for each other. When the goods are substitutes, this term will be positive, reflecting the fact that tradable quantities are relatively efficient in this case. When the marginal benefit from a particular good is independent of the quantities of other goods, this term will equal zero. And when the goods are complements, this term will be negative, because in this case, the way that tradable quantities link production across goods is a disadvantage.

The third term reflects deviations from symmetry across goods in benefits, costs, and the degree of uncertainty. This term is difficult to interpret except in special cases, but will go to zero if the ratio of benefits to costs equal for all goods, if the degree of uncertainty is equal for all goods, or if those two parameters are uncorrelated across goods. For the special case in which benefits are independent across goods, this term is proportional to the weighted covariance across goods between σ_i^2/γ_{ii} and β_{ii}/γ_{ii}, weighted by $1/\gamma_{ii}$.

Finally, substituting (32) and (33) into (30) yields an expression for the comparative advantage of tradable quantities relative to prices.

$$(34) \qquad \Delta^{TP} \approx \sum_i \frac{\sigma_i^2}{2\gamma_{ii}^2}\left[(\beta_{ii}+\gamma_{ii})\left(-\frac{\psi}{\gamma_{ii}}\right)+\psi\left(\frac{\beta_{ii}}{\gamma_{ii}}-\sum_j \frac{\beta_{ij}}{\gamma_{jj}}\right)-\psi^2\sum_j\sum_k \frac{\beta_{ij}-\beta_{jk}}{\gamma_{jj}\gamma_{kk}}\right]$$

This expression and its interpretation are very similar to expression (33), except that prices provide more flexibility than tradable quantities, and thus the first term will be negative when flexibility is an advantage–when the average marginal cost curve is steeper than the average marginal benefit curve. The second and third terms are the same as in expression (33).

When the goods are perfect substitutes, the third term in (33) and (34) equals zero, and the second term in each expression reduces to $\beta_{ii}\left(\frac{\psi}{\gamma_{ii}}-1\right)$. Thus, expression (33) will reduce to

$\Delta^{TQ} \approx \sum_i \frac{\sigma_i^2}{2\gamma_{ii}^2}\left[\gamma_{ii}\left(1-\frac{\psi}{\gamma_{ii}}\right)\right]$, which is always positive. Thus, just as in the simpler case considered earlier, when the goods are perfect substitutes, tradable quantities are more efficient than fixed quantities, regardless of the slopes of the marginal cost and marginal benefit curves. And expression (34) will reduce to $\Delta^{TP} \approx -\sum_i \frac{\sigma_i^2}{2\gamma_{ii}^2}\left[\beta_{ii}+\psi\right]$. Again, the sign of this expression depends on whether the marginal benefit curve is steeper than the market-wide marginal cost curve.

E. Caveats

A few caveats are in order, due to the simplifying assumptions that have been made. First, this analysis assumes that the marginal cost of each good does not depend on quantity of any other good. If this assumption is violated, then this will affect the efficiency of tradable quantities. If increasing the quantity of one good were to lower the marginal cost of other goods, then the optimal quantity of those goods would rise. However, under tradable quantities, producing more of one good will decrease the required quantity of the other goods, and so tradable quantities would be less efficient than this model indicates. In the opposite case–if producing more of one good increases the marginal cost of the other goods–then tradable quantities will be more efficient than this model indicates.

Second, the analysis assumes that the cost shocks are uncorrelated across goods. Positively correlated shocks would tend to make tradable quantities behave more like fixed quantities. If cost shocks are perfectly correlated, then they will cause the shadow price on the overall quantity constraint (the permit price under tradable permits) to change, but will have no effect on the quantity of each good; thus, the efficiency of tradable quantities will be the same as that of fixed quantities. Negatively correlated shocks would have the opposite effect, magnifying the difference between fixed and tradable quantities.

III. An Application to the Problem of Banking and Borrowing of Pollution Permits

This section applies the earlier results to the problem of whether pollution permit programs should allow permit banking and borrowing (emitting less pollution now in exchange for being allowed to emit more in the future, or vice-versa). These are, in effect, permit trades between different time periods. Thus, the different goods in this case are reductions in pollution in different time periods. This section presents a simple model of pollution regulation in a dynamic setting, first for a flow pollutant and then for a stock pollutant, and uses the tools developed in the previous section to examine the efficiency of bankable/borrowable permits (tradable quantities) relative to pollution taxes and fixed quotas.

A. Regulation of a Flow Pollutant

First, consider regulation of a flow pollutant–that is, a pollutant that causes damage only in the period in which it is emitted. The discounted benefits and costs of pollution abatement are given by

$$B = \sum_{i=0}^{\infty} \delta^i B(q_i) \tag{35}$$

$$C = \sum_{i=0}^{\infty} \delta^i C(q_i, \theta_i) \tag{36}$$

where δ is the discount factor, which is assumed to be the same for firms and for the regulator. The cost and benefit functions are assumed to be time-stationary. Substituting the second derivatives of (35) and (36) into the definitions of γ (12) and β (13) yields

$$\frac{\beta_{ii}}{\gamma_{ii}} = \frac{B''(\overline{q}_i)}{C''(\overline{q}_i)} \tag{37}$$

Because the stationarity of the cost and benefit functions in each period implies that $\overline{q}$ will be constant over time, the ratio of γ_{ii} to β_{ii} will also remain constant, so $\frac{\beta_{ii}}{\gamma_{ii}} = \frac{\beta_{jj}}{\gamma_{jj}} \forall i, j$. Together with the fact that in this case $\beta_{ij} = 0 \; \forall \; i \neq j$, that implies that the second and third terms in (33) and (34) will equal zero.

Then a comparison of the first terms of (33) and (34) shows that either taxes or quotas will be more efficient than tradable permits. If $\frac{-\beta_{ii}}{\gamma_{ii}} < 1$–if the marginal cost curve for abatement in each period is steeper than the marginal benefit curve–then taxes will be more efficient than bankable permits. If the marginal benefit curve is steeper, then quotas will be more efficient.

Thus, as a general rule, permit banking should not be allowed for flow pollutants; if it is better to use a quantity instrument than a price instrument, then separate quotas for each period will be more efficient than bankable permits. In practice, nearly all emissions permit programs allow firms to bank permits. For example, the 1990 Clear Air Act Amendments, California's Low-Emission Vehicle Program, the CAFE vehicle fuel-economy standards, and the leaded gasoline phaseout program have all allowed permit banking in some form.

There are some exceptions. If the benefit and cost functions are sufficiently non-stationary that the ratio $\frac{\beta_{ii}}{\gamma_{ii}}$ varies significantly across periods, and that ratio is higher in periods for which the regulator has more uncertainty about costs, then the third term in (33) and (34) will be positive and thus it is possible for bankable permits to be the most efficient of the three instruments. A more promising argument for the use of bankable permits to regulate a flow pollutant would be if $\frac{-\beta_{ii}}{\gamma_{ii}} < 1$, and thus emissions taxes are the most efficient instrument, but taxes are not politically feasible. In this case, bankable permits would be the next best choice.[17]

B. Regulation of a Stock Pollutant

Consider instead the case of regulation of a stock pollutant–one for which pollution damages in a

[17] Adjustment costs are sometimes raised as another argument for allowing banking. A full analysis of the impact of adjustment costs is beyond the scope of this model, but the intuition for this case is relatively simple. If adjustment costs are significant, the marginal cost of abatement in one period will depend on abatement in other periods; abatement in a given period lowers the marginal cost of abatement in later periods. As discussed in section II.E, however, this non-separability of the cost function will actually work against tradable permits. It implies that abatement in any given period will raise the optimal amount of abatement in later periods. However, with permit banking, more abatement in one period implies just the opposite; less abatement is required in later periods.

given period depend on the amount of emissions in all previous periods. In this case, the benefit function will take the form.

$$B = \sum_{i=0}^{\infty} \delta^i B_i \left(\sum_{j=0}^{i} s^{i-j} q_j \right) \tag{38}$$

where s is the stock decay factor. The term in parentheses is the stock of pollution at time i. Unlike in the case of the flow pollutant, benefits are not independent across time periods. Nor is abatement in any given period a perfect substitute for abatement in any other period, though it is close. To see this, consider the ratio of the marginal benefits of abatement between two periods.

$$\frac{\partial B}{\partial q_i} = \sum_{k=i}^{\infty} \delta^k s^{k-i} B_i{}' \Leftrightarrow \frac{\partial B / \partial q_i}{\partial B / \partial q_i} = s^{j-i} + \frac{\sum_{k=\min(i,j)}^{\max(i,j)} \delta^k s^{k-i} B'}{\sum_{k=i}^{\infty} \delta^k s^{k-i} B'} \tag{39}$$

The ratio of the marginal benefit from abatement in period i to that in period j is a constant plus the discounted value of avoided damages between the two periods divided by the marginal benefit in period i. This ratio is the marginal rate of substitution of abatement between the two periods. If it is constant, then abatement is perfectly substitutable between the two periods. Thus, if the second term in (39) were zero, then abatement would be perfectly substitutable between the two periods. This is the case in the limit as both the stock decay factor and the discount factor go to one. In that case the damages that occur after period j–for which abatement is perfectly substitutable between the two periods–swamp the effects that occur in the interval between period i and period j.

Thus, in the limit as the stock decay factor and discount factor go to one, the results from section II.C for the case of perfect substitutes will apply. While neither factor will actually equal one in practice, as long as both factors are close to one, abatement in one period will be a close substitute for abatement in other periods. In that case, bankable permits will typically be more efficient than emissions quotas. This suggests that banking and trading should be allowed in implementing the Kyoto protocol on greenhouse gas emissions.

Similarly, the choice between bankable permits and taxes would depend on the relative slopes of the marginal benefit and marginal cost curves, where the appropriate marginal cost curve is the marginal cost curve for emissions reductions in all periods. This result is also important for policy decisions. Newell and Pizer (2002) and Hoel and Karp (1999) each found that taxes are generally superior to emissions quotas for the regulation of long-lived stock externalities, because the marginal benefit curve is relatively flat compared to the marginal cost curve for a single period. However, neither study considered bankable permits. The marginal cost curve for abatement over all periods will be much flatter than the curve for any single period, since it is the horizontal sum (weighted by discounted marginal damages) of the marginal cost curves across all periods. As a result, while taxes will generally be superior to emissions quotas, there is a significant range of cases in which bankable emissions permits will be more efficient than taxes. Thus, the argument for using taxes to regulate long-lived stock pollutants is not so clear-cut once bankable permits are considered.

IV. Conclusions

This paper develops a framework to evaluate the relative efficiency of three different types of instruments–prices, fixed quantities, and tradable quantities–in controlling production of a set of goods. It then illustrates these results by considering the problem of whether pollution permits should be bankable.

The paper shows that if the goods are substitutes, each of the three instruments could be the most efficient, depending on the relative slopes of the marginal benefit and marginal cost curves. When the marginal cost curve is steep relative to the marginal benefit curve, prices will be the most efficient instrument. When the marginal benefit curve is relatively steep, fixed quantities will be the most efficient instrument. And for some middle range, when the slopes of the two curves are similar, tradable quantities will be the most efficient instrument. The more substitutable the goods are for each other, the larger is the range in which permits are the most efficient instrument.

When the goods are perfect substitutes, tradable quantities will always be more efficient than fixed quantities. The choice between tradable quantities and prices in this case depends on the relative slopes of the aggregate marginal cost and marginal benefit curves. If the aggregate marginal cost curve is steeper, then prices are more efficient, while if the aggregate marginal benefit curve is steeper, tradable quantities are more efficient. This is the aggregate-level analogue of Weitzman's (1974) result.

In contrast, when the marginal benefit from a particular good is independent of the quantities of the other goods, or the goods are complements, either fixed quantities or prices are typically more efficient than tradable quantities. Only if the goods are sufficiently asymmetric, both in the amount of uncertainty about costs and in the slopes of the marginal benefit and marginal cost curves, and these asymmetries are strongly correlated, can tradable quantities be the most efficient instrument. These results hold even when trading ratios are set based on the ratio of expected marginal benefits between the goods, rather than merely allowing one-for-one trades. Thus, this result differs sharply from past work, which has suggested that tradable quantities will dominate fixed quantities even when goods are not perfect substitutes, as long as trading ratios are proportional to marginal benefits.

The question of whether firms should be allowed to bank and borrow emissions permits is used to illustrate the implications of these results. The case of a flow pollutant–where pollution damage depends only on the current flow of pollution emitted, not on past or future emissions–corresponds to the case in which the marginal benefit from a particular good is independent of the quantities of other goods produced. Thus, in this case, either non-bankable permits (fixed quantities) or pollution taxes (prices) will be more efficient than bankable permits (tradable quantities).

For a stock pollutant–where pollution damage comes from the stock of accumulated past emissions–pollution reduction in one time period is a close substitute for reductions in other periods. Thus, bankable permits will generally be more efficient than non-bankable permits, though this will not always be true; non-bankable permits will be more efficient if the stock decay rate and discount rate are sufficiently low, and the marginal benefit curve is substantially steeper than the marginal cost curve. The choice between bankable permits and pollution taxes will depend on the relative slopes of the intertemporal marginal cost curve (which will be substantially flatter than the marginal cost curve in a single period) and the marginal benefit curve.

These results could have tremendous importance for environmental policy. They show that many current and proposed environmental regulations are not optimal. Emissions permit programs for flow pollutants typically allow banking (for example, the lead phase-out and sulfur dioxide permit programs), while the Kyoto agreement to reduce carbon dioxide emissions–a stock pollutant–does not yet allow for such intertemporal flexibility.

These results also show that emissions permit programs should allow trading for pollutants with global effects, such as greenhouse gases, where pollution in one location is a perfect substitute for pollution elsewhere, but not for pollutants with very localized effects. But most current proposals for implementing the Kyoto agreement limit emissions trading between countries, while some existing and proposed US regulations allow trading even for relatively localized pollutants (such as the proposed use of tradable permits to regulate emissions of nitrogen oxides).[18]

There are a number of promising directions for future research on this topic. The tradable quantity regulations analyzed in this paper allow unlimited trading, but the framework developed here could also be used to consider policies that limit trading in some way. Such policies are common in practice. For example, some emissions trading programs allow trading only within a local area, rather than across the entire universe of pollution sources. The question of how large these trading areas should be is quite interesting.

Similarly, Roberts and Spence (1976) and Weitzman (1978) showed that hybrid instruments–instruments that combine elements of both prices and quantities, such as emissions permits with a price cap–will always be more efficient than a pure price or pure quantity instrument. A similar result should hold for hybrids between fixed and tradable quantities. And fixed/tradable hybrid instruments are relatively common compared to price/quantity hybrids; most emissions permit programs allow permit banking, but not permit borrowing, and Los Angeles's RECLAIM smog-trading program allows permits to be traded from coastal locations to inland locations, but not vice-versa. These programs behave like

[18] Note, however, that if pollution taxes and other price instruments are politically infeasible, tradable permits may be the best available option even for a localized pollutant.

tradable quantities when firms want to trade one direction (banking, or trading inland), but behave like fixed quantities when firms would like to trade the opposite direction (borrowing, or trading to the coast). Further research could determine whether such hybrid instruments are more efficient, as is the case for price/quantity hybrids.

Finally, empirical applications of this framework could be very useful for policy. Given the increasingly rich data on environmental regulation that is becoming available, it should not be too difficult to estimate the parameters of interest for a particular case.

References

Adar, Zvi and James M. Griffin, 1976. "Uncertainty and the Choice of Pollution Control Instruments," *Journal of Environmental Economics and Management*, 3:178-88

Casella, Alessandra, 1999. "Tradable Deficit Permits: Efficient Implementation of the Stability Pact in the European Monetary Union" *Economic Policy*, 29:323-61

Greenwood, Jeremy and R. Preston McAfee, 1991. "Externalitites and Asymmetric Information," *Quarterly Journal of Economics*, 106:103-21

Hoel, Michael and Larry Karp, 1999. "Taxes Versus Quotas for a Stock Pollutant," working paper, Universities of Olso and California-Berkeley

Kaplow, Louis and Steven Shavell, 1997. "On the Superiority of Corrective Taxes to Quantity Regulation," NBER Working Paper #6251

Kling, Catherine and Jonathan Rubin, 1997, "Bankable Permits for the Control of Environmental Pollution," *Journal of Public Economics*, 64:101-115

Leiby, Paul and Jonathan Rubin, 2000, "Bankable Permits for the Control of Stock and Flow Pollutants: Optional Intertemporal Greenhouse Gas Emission Trading," *Environmental and Resource Economics,* forthcoming

McAfee, R. Preston, Wendy Takacs and Daniel Vincent, 1999, "Tariffying Auctions," *RAND Journal of Economics* 30:158-79

McGartland, Albert and Wallace Oates, 1985, "Marketable Permits for the Prevention of Environmental Deterioration," *Journal of Environmental Economics and Management* 12:207-228

Montero, Juan-Pablo, 2001, "Multipollutant Markets," *RAND Journal of Economics*, 32:762-774

Newell, Richard and William Pizer, 2002, "Regulating Stock Externalities Under Uncertainty," *Journal of Environmental Economics and Management*, forthcoming

Roberts, Marc and Michael Spence, 1976, "Effluent Charges and Licenses Under Uncertainty," *Journal of Public Economics*5:193-208.

Rubin, Jonathan, 1996. "A Model of Intertemporal Emission Trading, Banking, and Borrowing," *Journal of Environmental Economics and Management*, 31:269-286

Stavins, Robert, 1996, "Correlated Uncertainty and Policy Instrument Choice," *Journal of Environmental Economics and Management*, January 1996

Vickrey, William, 1992. "Chock-Full Employment without Increased Inflation: A Proposal for Marketable Markup Warrants" *American Economic Review* 82:341-45

Vickrey, William, 1993. "Today's Task for Economists" *American Economic Review*: 83:1-10

Weitzman, Martin, 1974, "Prices vs. Quantities," *Review of Economic Studies* 41:477-491

Weitzman, Martin, 1978, "Optimal Rewards for Economic Regulation," *American Economic Review* 68:683-91

Yates, Andrew, 2000. "Decentralization in Pollution Permit Markets," working paper, Tulane University