

TRADABLE DEFICIT PERMITS.
Efficient Implementation of the Stability Pact in the European Monetary Union

Alessandra Casella
Columbia University
New York, NY 10023

May 1999

1. Introduction.

The current provisions of the Pact for Stability and Growth advocate balanced budgets in the longer term and specify a ceiling for deficit spending of 3 per cent of GDP for each member of the European Monetary Union. A violation of the ceiling will trigger warnings and eventually penalties (unless exceptional circumstances can be invoked).¹ In this form, the Pact suffers from several shortcomings that will limit its effectiveness and impose exceptional costs on at least some of the member countries.

This paper will discuss a scheme for achieving the objective of overall fiscal moderation embodied in the Stability Pact at lower costs to countries' growth. It will not question whether the Stability Pact is itself needed, but, taking the intent to regulate the fiscal position of member countries as established, present an efficient mechanism for the implementation of this regulation. In the last two years the political climate in the European Union has changed, while unemployment is still too high. The political will to maintain fiscal discipline will weaken if the costs are too large. If balancing the fiscal accounts is considered a worthwhile goal, reaching it may well depend on our ability to design a better mechanism.

The weaknesses of the Stability Pact have been widely discussed:

- The imposition of the same criterion for each country leaves no room for differences, either in countries' initial fiscal positions (which may intended) or in countries' cyclical phase. For some economies at least the room for maneuver required for the functioning of the automatic stabilizers during the cycle is almost certainly missing.

- The arbitrariness left in the criteria creates uncertainty about the application of the penalties

and is sure to generate difficult negotiations with individual countries. Compliance will be affected. Especially when times are hard (but not hard enough to qualify for exceptions automatically), the incentive to violate the deficit limit and then negotiate will be high.

- There is no reward for virtue. While an exposed fiscal position by any one country is considered a weakness for the whole system, a particularly solid position must be its own reward: "The problem with the Pact as presently framed is that it is all stick and no carrot: rewarding good fiscal behavior in booms .. in addition to punishing bad behavior in slumps would surely make better sense" (Bean, 1998, p.106).

- Politically the Pact seems designed to be unpopular. Combining a draconian criterion with discretion in the application of the penalties, it emphasizes the countries' loss of sovereignty. It is likely to be an obstacle if and when the UK decides to join the monetary union.

The thesis of this paper is that all these limitations can be overcome if we combine the overall objective of fiscal discipline with sufficient flexibility for individual countries. Borrowing from the experience of environmental markets, we should design a system of tradable deficit permits: having set an overall ceiling and an initial distribution of permits, EMU countries could be allowed to trade rights to deficit creation. The scheme need not treat all countries identically and could be designed to penalize countries with higher debt to GDP ratios. Its fundamental virtue is that it exploits countries' incentives to minimize their costs to insure that the final goal is achieved as efficiently as possible. The next section clarifies the logic behind the proposal; section 3 describes the US market for sulfur dioxide emissions, and section 4 discusses options and obstacles in adapting the simplest design to the fiscal concerns of the Monetary Union.

2. The Basic Idea.

The starting point of the Stability Pact must be the belief that markets alone are unable to impose sufficient discipline on the fiscal position of individual EMU members. Regardless of official statements to the contrary, it could be that both markets and governments expect that a profligate country will be bailed out eventually by the rest of the system. This may occur because either budget constraints operating through the banking system or psychological effects tie the robustness of the whole monetary union to that of its weakest member. Or, without getting to the extremes of a crisis, because fiscal exposure translates into inflationary pressure on the European Central Bank (ECB), itself a form of bail-out. In any case, the markets' inability (or unwillingness) to differentiate sufficiently among different countries' debts could lead to excessive spending. All of these motivations have been analyzed elsewhere (for example, Beetsma and Uhlig (1998), Chari and Kehoe (1998), Eichengreen and Wyplosz (1998), Flandreau et al. (1998)), and it is not the purpose of this paper to discuss them further.² Regardless of economists' opinions about the seriousness of these fears, the existence of the Stability Pact makes clear that the problems are seen as important by policy makers. If we accept the will to impose fiscal discipline on the EMU countries as established, we need to ask how the objective can be achieved at the lowest cost to countries' growth.

The different mechanisms described above all amount to stating that high deficit spending by any one country has costs for the other members that the country itself fails to internalize: a high deficit in Italy affects negatively Germany, France and the other members of the Union. We can then think of deficits as a form of pollution, originating in one country's activity, but having

repercussions for all. Once the problem is stated in these terms, we see that the arguments that have been developed for environmental regulations can be adapted to our purposes.

In a world where a benevolent and perfectly informed central planner existed, a centralized solution would be possible. All decisions would be deferred to the center: in the same way as countries have relinquished their monetary policy, they would also lend their fiscal powers to a European-wide body. At least in the short run, neither the institutions nor the political will are in place to make such a scenario feasible or in fact desirable. Alternatively, countries running fiscal deficits could be charged a tax per unit of new debt equivalent to the social cost of their issues, a direct parallel to the Pigouvian tax advocated for environmental problems. But the information required to calculate the tax correctly is very difficult to obtain: a tax scheme imposes a daunting task on the regulator, with the likely result that the realized deficits could be seriously different from the desired objectives. In the presence of uncertainty, fixing the aggregate ceiling to fiscal expansion limits this risk, but must be complemented by an allocation mechanism that distributes the responsibility for fiscal austerity in the most efficient manner. Analogously to pollution permits, a system of tradable deficit permits sets a total limit to fiscal deficits but uses the market to allocate them across the different countries at minimum cost. This is the scheme recommended and discussed in this paper.³

Most traditional forms of pollution control take the form of quantitative limits on pollution sources, typically imposing the same limit on all sources. But because different sources utilize different production technologies, have different access to cleaner inputs and different scope for capital investment in pollution reduction, the same pollution limit imposes widely diverging costs. In a market for pollution permits, instead, while the regulatory authority sets the overall pollution

limit - the total stock of permits available on the market - it is the market itself that ensures that all pollution sources will act to equalize their marginal costs of pollution reduction, achieving the target decrease in pollution at minimum total cost.

The parallel with our fiscal problem is immediate: the scheme currently envisioned by the Stability Pact consists of uniform quantitative constraints on each country's deficit (not more than 3 per cent of GDP, barring a serious recession), and the observance of this limit is likely to be associated with very different costs, depending on the country's structure, debt overhang and cyclical phase. The dispersion in costs is a sign of the scheme's inefficiency: a system of tradable deficit permits would allocate deficits there where their value is higher, making it possible to implement the desired fiscal discipline much more efficiently.

Box on tradable pollution permits.

The idea that externalities could be solved efficiently through the creation of property rights and voluntary trading originates with Coase (1960). If there were no obstacles to information and to contracting, voluntary exchanges would result both in the correct level of total activity and in the optimal allocation among agents. In the presence of real world imperfections, however, the achievement of the first best through private bargaining is unlikely, and advocates of market instruments in environmental regulation make the more modest but more realistic claim of implementing a predetermined target at minimum costs. Accepting the principle that both the total allowed level of pollution and the distribution of the burden necessary to bring it about are chosen politically, they show that a market for pollution permits achieves the target reduction in pollution at lowest total costs. The creation of a market for pollution permits was first proposed

by Dales (1968) and Crocker (1966). The literature is now very large, but a good introduction is provided by the following sources: Montgomery (1972) for a formal treatment; Tietenberg (1985) for a survey of the issues involved; Baumol and Oates (1988) for a thoughtful discussion of the approach's promises and limits.

The mechanism through which the market delivers minimum total costs is shown very clearly in the following example constructed by Tietenberg (1985, ch. 2) and illustrated in Figure 1 (Figure 1 in Tietenberg, 1985, p.20). Consider two sources of pollution, each of which would produce 15 units of pollution if unconstrained. Suppose that the goal is to reduce total pollution by half, i.e. to 15 units in total. The horizontal axis in the diagram measures the reduction in emissions for the two sources, and should be read from left to right for source 1 and from right to left for source 2. The vertical axis measures the marginal costs of pollution reduction for the two sources (thus both curves are zero when pollution is unconstrained). For each source, the total cost of achieving a specified level of pollution control is the area below the marginal cost curve, up to the specified point. Notice that any point on the horizontal axis achieves the desired level of pollution control, and thus the corresponding areas under the marginal costs curves yield the total costs associated with each allocation of pollution reduction between the two sources. There is a unique allocation such that the marginal costs are equal, and that is the unique allocation that minimizes total costs (it is easy to read from the figure that the sum of the two sources' total costs is higher at any other point). This is the first result: the aggregate cost of compliance is minimized when marginal costs are equal across sources. It is also easy to see that a market for pollution permits would take the two sources to the efficient point. Suppose that the initial distribution of permits required source 1 to cut emissions by 7 units and source 2 by 8. As long as the price of a

permit were smaller than G , source 2 would want to buy permits, and as long as it were higher than E source 1 would want to sell. When the price equals P the two sources trade exactly the required number of permits to bring them to the efficient point. But the price cannot be different from P , because at any lower price demand exceeds supply, and at any higher price supply exceeds demand. Notice that any initial distribution of permits leads to the same equilibrium, although to a different distribution of the total costs.

Fig.1

There are of course a number of potential difficulties with applying a system of tradable permits to fiscal discipline, and they are discussed below in Section 4.. For the moment, let us consider how the scheme would work in its simplest realization. Each year each country is allocated a number of deficit permits, equivalent, for example, to 3 per cent of its GDP. In practice, these permits could simply be entries in special accounts maintained by each country at the ECB or at the European Commission. The permits are denominated in Euros and freely tradable. At the time final fiscal statistics are made public, for example by the end of April of the following year, each country must have in its account a sufficient number of permits to cover the year's deficit, and these permits are withdrawn from the system.⁴ If a country is found not in compliance, it faces a steep fee for each of the missing permits and must relinquish a corresponding number of permits from the following year's allocation. Following the example of existing environmental markets, it seems advisable to let countries bank permits for future use while not allowing them to borrow from future allocations. This leaves some room for intertemporal planning and smoothing of anticipated shocks, but limits temptations for governments with too short horizons. In

practice, this means that deficits can be offset by permits carrying a date contemporaneous or preceding the year of the deficit.

An unusual feature of the scheme is its finances. Buying the permit is buying the right to issue a unit of debt; but since the permit must be paid, part of the new debt is in fact devoted to the purchase of the permit itself. If p is the price of the right to issue 1 Euro of debt, a government can devote only a share $1-p$ of the value of each permit to the purchase of new resources; in other words, the right to borrow 1 Euro of new resources for national expenditure requires the purchase of $1/(1-p)$ Euros of permits, i.e. costs the government $p/(1-p)$. This multiplicative factor implies that the price of the right to issue 1 Euro of debt must always be smaller than 1 (as intuition suggests), and creates a wedge between total new debt and those borrowed resources that can be devoted to national expenditure, but has no other implications, and in particular does not affect the theoretical optimality of the scheme. The point is analyzed more formally in the Appendix.⁵

It is true of course that each government would be faced with an extra expense - the purchase of the permits - exactly when it is already surpassing its 3 percent deficit allocation. But this occurs in the current system too, if the penalties are applied. The important difference is that in the scheme suggested here the fee is not fixed arbitrarily, but determined by the market. In particular, consider a country affected by an asymmetric negative shock, the type of idiosyncratic shock that would put the existing system under stress. If the other countries do not anticipate a need for fiscal expansion, the price of the permits will be low (the price is positive only if the total unconstrained deficit of the EMU countries is larger than 3 per cent of their total GDP), and the shock can be countered by fiscal policy at low cost. More generally, the cost of going above the

3 per cent limit at any given time is the market valuation of a fiscal expansion at that time, taking into account the overall ceiling and, possibly, the option of banking permits for the future. If the ceiling has been chosen correctly, this is exactly what the cost should be.

Thus the first important advantage of the scheme is the flexibility it provides to individual countries - the performance of all countries need not be the same, nor needs the performance of each country at different times. Notice, and this is a second benefit, that the increased flexibility works both through imposing the correct costs to fiscal expansion and through creating the correct rewards for fiscal cuts. A country choosing to sell some of its permits collects resources, again according to the value of those permits in their best alternative use. Contrary to the current writing of the Stability Pact, good behavior is rewarded. Finally, flexibility means that a country can intervene *before* experiencing a severe contraction - as opposed to the present plan where exceptions to the fiscal constraint are only triggered when a very serious recession has been experienced, and thus when, by definition, stabilization policy has already failed.

As all regulatory mechanisms, this too will function well if the aggregate ceiling is appropriate. In the case of an unexpected Europe-wide negative shock, the supply of permits may need to be increased to prevent an undesirable and contractionary increase in their price. More generally, the supply of permits should be adjusted to take into account the state of the European economy. This allows the system to overcome two further limitations of the Stability Pact. First, in the case of a symmetric negative shock, the quantitative constraints on fiscal deficits at present leave all responsibility for stabilization policy with the Central Bank. At least until deficits remain close to their upper limits there can be no meaningful discussion of an optimal policy mix. Second, while the Stability Pact advocates the medium term objective of a balanced budget, there is no

mechanism that encourages countries to bring it about. Individual countries are given no inducement to reduce deficits below the limit, and imposing a lower limit would be extremely difficult in the absence of any room for individual deviations. A system of tradable permits, on the contrary, allows enough flexibility to accommodate anticipated declines in supply. Indeed this is one of the reasons that have made the system acceptable to environmental groups in the US.

If changes in permits supply are planned, however, it is important to make them as predictable and as transparent as possible: the scheme functions if the market functions, and the market requires predictability. Thus there is a trade-off between the possibility of fine-tuning the supply of permits and the need to minimize interference with the market. Again as in the case of the environmental markets, there should be an automatic rule that specifies how supply is determined over time, and the rule should be as simple as possible, focusing on indicators that countries' policies cannot manipulate.

One of the fundamental advantages of a system of tradable permits over the current Pact is its transparency, and it should not be compromised. It would be a system regulated by rules as opposed to a system open to political exceptions, and as such much easier to understand, predict and enforce. A standard finding in the literature on environmental regulation is the decline in litigation following the substitution of a system of tradable permits for quantitative constraints. The lack of flexibility that defines the latter makes it necessary to allow for exceptions, and the presence of exceptions makes it often desirable for a firm to violate the constraint first, and litigate later. A significant part of the increase in compliance predicted (and observed) with tradable permits comes from decline in violations whose legitimacy could conceivably be defended in a murkier system (Tietenberg (1985), Stavins (1998)).

Finally, the theoretical superiority of the scheme suggested here is also its main practical advantage: at equal enforcement, it is guaranteed to cost less. No matter how desirable fiscal discipline may appear ideally, it will not be pursued if its costs are too high. This is particularly true at the present moment when other painful reforms must be tackled - to labor markets, to pension systems, to social benefits in general - and people's impatience with further fiscal austerity has been expressed at the polls and is shared by many of their representatives. If a scheme of fiscal discipline is to be imposed at all, then its only chance of implementation comes from being as efficient as possible.

The example of environmental regulation has been mentioned frequently because it is the main field where alternative schemes for controlling externalities were devised, studied and compared. Beyond the theoretical work, what is really exceptional about environmental regulation is that economists' schemes were in fact put into practice as policies. Thus not only the theoretical studies but the practical experience of existing permits markets can be used as benchmark for quite different applications.⁶

3. The Experience of Sulfur Dioxide Emission Trading.

In 1990, Title IV of the Amendments to the Clean Air Act introduced in the US the first national program of environmental regulation based on a system of marketable permits. Smaller local markets for pollution permits had been tried in the past and some are in existence at present, but this program is unique in its scale, its geographical extension, and its unfettered reliance on private markets.⁷ I summarize its main features below; more exhaustive analyses can be found in Ellerman et al. (1996) and (1997), Joskow et al. (1998), Schmalensee et al. (1998), Stavins

(1998), and in the remarkable Web page of the Environmental Protection Agency (EPA) at www.epa.gov/acidrain.

The goal of Title IV is to reduce permanently yearly sulfur dioxide (SO₂) emissions to approximately 50 per cent of their 1980 levels, i.e. to about 9 million tons. SO₂ is the main precursor of acid rain, and is mostly emitted by coal-fired electric utilities.⁸ The program is to be enacted in two phases: the first, from 1995 to 1999, targets the largest and dirtiest generating units, and requires them to cut their aggregate yearly SO₂ emissions by about 40 per cent; phase 2, from 2000 onward, submits effectively all electric-generating units in the US to the aggregate ceiling of 9 million tons of emissions per year. Each unit is required to install and maintain a continuous emissions monitoring system approved by the EPA.⁹

The aggregate targets are implemented through a system of fully tradable allowances. The total amount of emissions allowed for a given year translates into a total number of allowances carrying that year's date, and distributed to all affected units approximately in proportion to their average 1985-87 heat input. An allowance is the right to emit a ton of SO₂ and exists as a record in the accounts system maintained by the EPA (the Allowance Tracking System). Within 30 days of the year's end, each unit must withdraw from its account and deliver to the EPA a number of allowances equal to its SO₂ emissions during the year and carrying either that year's or an earlier date. Thus each unit needs to buy allowances if it emits more than its initial allocation, but can either sell or bank allowances, if it emits less. If a unit is found not in compliance, it must pay \$2,000 for each ton of emissions not accounted for (10 times the current price of an allowance) and withdraw a corresponding number of allowances from its next year allocation.

The most remarkable feature of the program is the lack of government intervention in

influencing how compliance can be achieved. The only intervention in the market is the auction held by the EPA in March of each year, where approximately 3 per cent of allowances are sold, and whose purpose, especially at the beginning of the program, was to generate price data that may help the market identify an equilibrium price (revenues are redistributed to the utilities). Otherwise, the market is open and unregulated. Anybody can enter the market and trade, be it utilities, brokers, businesses, environmental groups, or individuals; traders who are not electric utilities can but are not required to have an account at the EPA; partners can trade in any manner they choose and are free to record or not their transaction with the EPA; new instruments - for example, swaps of current for future date allowances - or new goods - for example bundles of coal deliveries and matching emission allowances - can be (and were) freely created by the market.

The openness of the market is matched by its transparency: the balances of the EPA accounts are public information, available on the EPA Web page, and so are all bids in the EPA auctions. Information on current prices is easily obtained from either Fieldston Publications' market survey or from Cantor Fitzgerald, the largest brokerage firm active in environmental markets; for example, Cantor Fitzgerald's Web page (www.cantor.com/ebs/) reports that the price of one SO₂ allowance for current use was \$212.17 on May 25. Placing an order is equally easy: through Cantor Fitzgerald it can be done by calling toll free (800) 228-2955, ext.5 from a US or Canada phone.¹⁰ Quite naturally, the two brokers act also as consultants, and have played an important role in devising new instruments; as the example of the swaps makes clear, an active market in derivatives is growing.

So far the program has been unanimously judged a success. The first and main concern was

that the market would not develop, but all evidence to date is to the contrary. As described in detail in Joskow et al. (1998), by mid 1994 a single market price had emerged.¹¹ Starting at around that time, the market price indexes reported by all market observers are indistinguishable, and match the prices realized at the EPA auctions. The volume of market transactions has also increased dramatically, from 900,000 allowances transferred between organizations in 1994 to 9.5 millions in 1998, with the percentage of activity generated by the EPA auction falling from 20 per cent in 1994 to 3 per cent in 1998 (EPA, Web page). (See Figure 2). The environmental objectives were easily achieved and in fact surpassed: in 1995 and 1996, years for which final data are available, SO₂ emissions by regulated sources were respectively 39 and 33 per cent below allowed aggregate levels. This very substantial overcompliance took place at allowance prices that were far below what had been anticipated: predicted prices were in the range of \$250-\$350, whereas for all of 1995 prices were below \$130 and for all of 1996 below \$100. The aggregate performance masks large differences in individual responses, exactly as the theory suggests should occur - in both years, approximately one quarter of all affected units had emissions exceeding their allocation of allowances (Schmalensee et al. (1998)).

Whether or not the market that has evolved is efficient is a more difficult question. On one hand, neither concerns with the design of the EPA auction (Cason (1993) and (1995)), nor fears of interference from state regulatory commissions (Fullerton et al. (1996)) have proven justified (Joskow et al. (1998), Bailey (1996)). On the other, the steep initial decline in the market price suggests that utilities were surprised by the low cost of compliance and had overinvested in capital improvements, reducing the demand for allowances below what would have been optimal (Ellerman et al. (1997), Schmalensee et al. (1998)).

A less ambitious but more direct test of the success of the program is provided by an estimate of the savings incurred. How much more would it have costed to achieve the same pollution reduction through an alternative, less flexible regulation? Of course the answer depends on the counterfactual chosen as comparison. For 1995 and 1996, Ellerman et al. (1997) find that the average cost of cutting one ton of pollution was around \$200 (including annualized capital expenses), or about one quarter to one third below the cost that utilities would have incurred if they had been prohibited from trading allowances (and transferring them among units of the same company) at the same initial distribution of rights to pollute. This amounts to total savings of the order of \$300 million per year.

4. Discussion.

The basic system of tradable deficit permits can be amended to capture more faithfully the concerns of European policy makers and insure that it can be implemented. The purpose of this section is to show that potential problems can be faced, not to claim definitive answers. Most of the points raised here will eventually require more detailed analyses.

4.1. Deficits and Debts.

A serious concern at the start of the sulfur dioxide program was the possibility that trade would lead to "hot spots", geographical concentrations of pollutant in specific areas. Theoretically, a system of emission permits is optimal only when it concerns an "assimilative pollutant", that is a substance that mixes perfectly in the atmosphere and whose precise geographical origin is therefore irrelevant. Only in this case are emissions from all sources perfect substitutes and thus

justified to trade at one price. "Hot spots" could develop in the case of sulfur dioxide, but apparently have not, at least to any important extent (Ellerman et al. (1997)).

A very similar problem exists for a system of deficit permits. The system described in section II is the efficient way to implement an aggregate constraint on EMU countries deficits, as percentage of EMU GDP. But it assumes that the distribution of deficits among the different countries is irrelevant, because all individual deficits are perfect substitutes. If the source of the fiscal externality is interest rate spill-overs when the market considers all countries' debts perfect substitutes, then the design of the system is appropriate. But if on the contrary the fear is that a country's excessive exposure can trigger a crisis for the Union as a whole, then Italy's new debt issues, for example, are not perfect substitutes for Germany's, because their outstanding debt volumes are different. Since heavier debt service means higher costs of deficit reduction, purchases of permits could be concentrated among countries with larger stocks of debt, creating "hot spots" of new debt creation exactly where deficits are more costly for the Monetary Union. The distribution of deficits matters.

What is the appropriate design in such a case? Once again the intuition becomes particularly clear when we phrase the problem in terms of pollution. Consider two sources of pollution located at different distances from a town. The policy target is air quality in the town, and air quality is affected by emissions from the two sources in inverse proportion to distance. A system of emission permits is not well suited to the problem because it values emissions from the two sources at the same price, whereas the two sources' impact on the target is different. The correct system should be instead one of *pollution* permits, as distinct from *emission* permits, where the two sources are allocated and trade permits to create pollution as measured at the receptor point,

i.e. in the town (an “ambient permits system”). If one source is twice as distant from the town as the other, then two units of its emissions will have the same effect on air quality as one from the second source. Hence pollution permits trading at 1 to 1 translate into emissions permits trading at 2 to 1: the same level of emissions must cost the source closer to town twice as much. The solution, proposed and studied formally by Montgomery (1972), is simple and elegant: property rights should be created for the targeted externality and not for its origin, because the latter is one-step removed from the policy problem.

In the case of environmental regulation, the simplicity of the theoretical solution translates into an impossibly difficult problem of implementation. Even if the impact of each source of emissions at a target point could be estimated accurately, each source affects environmental quality at a number of target points, and affects each of them differently. Thus each point should correspond to a different market, and each source should be active on all of them. But in fact if a source emits several pollutants, as typically the case, then it should be active on a different set of markets for each pollutant, and all these trades should happen contemporaneously and be interdependent. It seems highly implausible that firms could navigate such a system in practice.

For the purposes of our problem however, the essential intuition can be applied much more simply. The policy objective is to control the impact of each country's fiscal choices on the financial stability of the Monetary Union. This latter variable is equivalent to town air quality in the simple example discussed above; because it is the only relevant objective, a single market will be sufficient to achieve it, and it should be a market in “permits to increase the financial fragility of the Union”. Thus what we need is an index of financial fragility and a measure of the impact that different countries' deficits have on such an index.

Consider the following plausible implementation. The proxy for aggregate financial fragility is given by the average squared deviation of Union countries' debt to GDP ratios from a fixed reference level; the policy objective is maintaining the rate of growth of that index below a specified target. As shown in the Appendix, under reasonable assumptions this amounts to holding the weighted sum of all countries' deficits below the target, using as weight each country's debt to GDP ratio.¹² The Appendix also demonstrates how such a policy can be efficiently implemented through a system of marketable permits: at the end of the year, each country must hold an amount of permits proportional to the year's deficit, using the country's debt to GDP ratio as factor of proportionality. Exactly as in the case of the two sources of pollution at different distances from town, if country A's debt to GDP ratio is double that of country B, then country A will have to hold twice as many permits to issue the same amount of new debt, or equivalently will pay twice the price for each Euro of deficit. And this is just as it should be, since the marginal impact of country A's deficit on the Union's financial fragility is twice as large that of country B's.

The simple example comes close to real policy concerns. But of course different indexes of financial fragility could be constructed, taking other variables into account. As shown by Montgomery (1972) the only important constraint is that the policy target be expressed as a linear function of each country's deficit. Any such policy, choosing freely the factor of proportionality, can be implemented at minimum cost through a system of permits.

Environmental economists, on the other hand, have been forced to consider different approximations to the optimal system, given its forbidding complexity in the applications that are relevant to their concerns (see for example, Krupnick, Oates and Van De Verg (1983), and the

discussion in chapter 4 of Tietenberg (1985)). When undesirable geographical concentrations of pollutants are a possibility, a natural option to consider is that of "zonal permits" where free trading of emission permits is allowed within each zone, but not across zone borders. Such a system guarantees that pollution levels in each zone do not exceed target levels, but does so by raising aggregate costs, because it interferes with the equalization of marginal compliance costs produced by free trade. For our purposes, an analogous system would divide the member countries into different groups, according to their likelihood of triggering a fiscal crisis, fix a maximum deficit to GDP ratio for each group and a corresponding distribution of permits among all countries in the group, and allow trading only within the group. Given the small number of countries, the partition should probably not be finer than two groups, and again a plausible criterion for selecting group members could be debt to GDP ratios.¹³ Individual countries would move between the two groups as their debt to GDP ratios change. To the extent that similar debt to GDP positions do not imply perfect correlation of shocks, some individual disturbances could still be smoothed through trading. It is clear however that the main purpose of a "zonal" system is to reduce trade, and thus the system will be inefficient, unless the ceilings have been chosen exactly right. In addition, the scheme could be politically very unpopular, because it would subject different countries to different constraints, with permits trading at different prices within each group. We have discussed above how a better scheme can be designed for our problem with relative ease; thus the only reason to consider a "zonal" system at all is its simplicity. It would be very transparent, easy to explain to the public, and a possible compromise between the current regime and a more desirable, fully flexible scheme.

It may be important to clarify at this point that our entire discussion follows from the

maintained assumption that the fiscal externality takes the form of global, not local effects. The geographical analogy that allows us to borrow directly from the environmental literature should not be misleading. In our analysis, each country cares only about an overall measure of the Union's fiscal health, be it a straightforward aggregate variable - for example the Union's aggregate deficit to GDP ratio - or a more complex function of member countries' policies - the probability that a financial crisis be triggered, or that the ECB be pressured into an inflationary path. We are not considering the alternative case where each country has preferences over the precise distribution of fiscal spending, independently of any aggregate measure of fiscal stability (as would happen, for example, if a country were affected exclusively by the fiscal stance of its main trading partners, or its geographical neighbors). The logic of our approach generalizes to this scenario, but the appropriate trading mechanism could not be a simple market for permits. With a unique market price and anonymous trading, a country has no instrument with which to impose differential discipline on other member countries. We would need to devise more complex exchange mechanisms where countries could bid over specific distributions of fiscal spending across Union members, a "menu auction" à la Bernheim and Whinston (1986). I am confident that the correct scheme could be designed, at least in theory, but neither the current formulation of the Stability Pact nor the discussions that have accompanied it suggest that local effects are a central concern (see for example, Eichengreen and Wyplosz (1998), and the discussion that follows the paper). All attention is focused on the possible influence of individual countries' fiscal mismanagement on Union-wide indicators: interest rates, inflation, the Euro exchange rate. Thus studying a simple mechanism that can provide a solution to these global effects seems to be of particular importance and is the only issue addressed in this paper. Even in this case the

institution through which countries exchange permits requires some attention, and we turn to it now.

4.2. How should trade be organized?

If the theoretical answer is the creation of a competitive market, how do we make sure that transactions between the governments of 11 countries will indeed be competitive? There are two main difficulties. First, 11 is a small number, and since countries have very different sizes, the impact on the market of a large transaction by one of the big countries could a priori be important. Hahn (1982) has shown that the active presence of a monopsonist in the permits market distorts the equilibrium price and leads to higher compliance costs than in the competitive case. Hahn's result does not apply directly to our case because we have several large countries who are likely to find themselves on opposite sides of the market, not a single monopsonist, and a small number of players is not a guarantee of anticompetitive behavior (think for example of Bertrand's model of price competition in a duopoly). Nevertheless, trusting that a deficit permits market in general will be competitive requires considerable faith.

One particular reason for concern, the second difficulty mentioned above, is that EMU countries' governments interact continuously on a large set of issues, and there would be scope for bundling purchases of permits with other bilateral transactions. Again, bilateral bargaining need not be inefficient in principle, but the terms of trade are bound to reflect the two parties' relative strength. Direct negotiations between any two countries cannot be prevented, but the deviation from competition will be minimized if we can design a parallel market where a country can always buy or sell permits at the competitive price. Is it possible to design realistic trading

rules that would fulfill this function?

Once more, it seems reasonable to begin by studying the environmental markets. As described earlier, the market for sulfur dioxide emission permits was purposely left as unregulated as possible - the EPA does not concern itself with the manner in which a transaction has taken place, and the yearly EPA auction is both very small and increasingly irrelevant in signaling the market clearing price. In practice, in 1997 more than 80 per cent of all trades between organizations involved brokers (EPA, web page), and the latter were overwhelmingly represented by a handful of companies. According to Cantor Fitzgerald's web page, an increasing share of its environmental trades takes place through electronic trading, and more precisely through a proprietary electronic platform that the brokerage house developed for its trades of US government securities (the Cantor Exchange - CX). The technology amounts to a continuous double auction, where the current highest bid and lowest ask prices, and corresponding quantities, are posted on all terminals of connected agents (and with some delay on the Internet).

A continuous double auction (often oral, with increasing frequency computerized) is the trading mechanism followed by most organized exchanges around the world and governing transactions of stocks, bonds, metals, commodities and derivative securities. Because of its practical importance, its intuitive resemblance to the abstract idea of "the market", and the extreme difficulty of characterizing its properties theoretically, it has been the object of a large volume of experimental work. In laboratory experiments, double auctions quickly converge to competitive outcomes with full efficiency even when the number of participants is as small as three players on each side of the market. The conclusion is surprising and very robust, insensitive to the exact details of the mechanism. Box 2 below discusses the experimental literature on

double auctions in more detail.

In experiments, double auctions consistently outperform other exchange mechanisms in thwarting the exercise of market power (Smith et al. (1982), Clauser and Plott (1993)), and yield competitive outcomes even when one side of the market, sellers for example, are given the opportunity to coordinate actions in pre-play communication (Isaac and Plott (1981), Clauser and Plott (1993)). According to these works, participants quickly recognize their common interest, "conspiracies" form and informal agreements to impose cartel prices are easily reached. However, the agreements unravel just as quickly during the play of the game. Clauser and Plott conjecture that two features of double auctions may be responsible: first, the continuous nature of the auction implies that the temptation to defect is continuous too, as opposed to the single act of quoting the cartel price once in a static auction (for example in a sealed bid or posted offer auction). Second, any access to the market requires a seller to underbid the current ask price; thus, unless the cartel has also agreed to a credible *ex post* partition of the profits, the agreement will not be sustainable¹⁴.

Box on double auctions

A double auction is a market mechanism through which multiple buyers and multiple sellers exchange goods. Both bids and asks are permitted and exchange can happen at any time during the trading period; thus trades take place at different prices and net trades are the result of many bilateral transactions. Both in real markets and in experimental settings, double auctions began as oral auctions and have increasingly become computerized. The description below follows the survey by Friedman (1993).

Comparing the double auction to other institutions may help to make its functioning clearer. It differs from one-sided auctions because, in addition to competition among buyers, it allows active competition among sellers through the announcement of ask prices. It is a continuous auction because prices can be quoted and goods exchanged continuously, in contrast to static discrete-time auctions where all bids and asks are collected during the period, but all trades happen at the end, at the single market clearing price (“clearing houses”).

Many variants of double auctions exist in real markets and many have been studied in experiments: all bids and asks may be public (as in an oral auction) or only a subset, possibly only the highest bid and the lowest ask; the good may be divisible or indivisible; it may be an asset, living over time and paying dividends, or it may not; traders may be specialized in buyers and sellers, or be allowed to be both, etc. Real world markets provide the richest data, but as always are difficult to interpret; experiments in laboratories, although inevitably too simple, can be controlled and replicated and have provided a wealth of results.

By far the most important and the most common of these is that experimental double auctions are very efficient market mechanisms, independently of their specific details and, remarkably, of the number of players. Three or four buyers and three or four sellers are sufficient to induce prices and allocations that closely approximate the competitive equilibrium (see for example Smith (1982), Friedman and Ostroy (1995)). In markets with one-period goods or with short-lived assets, the conclusion is so robust that it has led researchers to hypothesize that something in the rules of the exchange is powerful enough to overcome both limits in the players' rationality and opportunities for monopoly profits. The puzzle is compounded by theory's inability to characterize the equilibrium of a double auction convincingly. A continuous double auction is a

very complex game of incomplete information, where not only all price quotes, but also their specific sequence and their time, relative to the end of the period, transmit information and must be chosen and interpreted strategically. When the strategic interactions are studied explicitly (as in Wilson (1987), the problem is so difficult that it seems implausible that real world subjects would be solving it.

Friedman (1984) and Friedman and Ostroy (1995) have conjectured that double auctions induce competitive behavior even in the presence of few large players through a form of Bertrand competition. With the rich exchange of information of a double auction, the argument is intuitively convincing, but the formal result requires that players have knowledge of the market clearing price. In most experiments, players acquire experience with the rules of the game through stationary repetitions of periods of exchange, and thus can use the previous period final price as a good guess of the equilibrium price. The view that these experiments come to approximate games of complete information seems plausible, but raises the obvious concern that the experiments may then not be faithful to the experience of real world markets.

The concern is particularly relevant for asset markets, where assets live for multiple non-stationary periods, information arrives continuously and dividends are state-contingent. In asset markets experiments, the asset yields a random dividend at the end of each period, and the complexity of the real world problem may be captured by letting the asset have a very long life (for example 15 or 30 periods in Smith et al. (1988); or by assigning it a specified probability of "death" each period, with no fixed end (Camerer and Weigelt (1993); or by designing messages that reach the different traders at different times and may convey different information (Copeland and Friedman (1987)). As expected, the increased richness of the environment leads to slower

and more erratic convergence to the efficient rational expectations equilibrium, and bubbles and crashes may appear. However, convergence to the "correct" price is observed in almost the totality of experiments - a result that seems in itself quite striking - and the speed of convergence increases with the experience of the participants, or in the presence of future markets.

For our purposes, these observations suggest that a continuous double auction may be the appropriate trading mechanism for a market in deficit permits. In line with developments in all financial markets, the continuous double auction should be computerized, a feature that reduces transaction costs and to some extent protects anonymity. Thus exchanges would then take place through a two-tier system: direct bilateral negotiations between countries, and a simultaneous electronic and anonymous double auction. A similar dual structure exists in most financial markets: for example the "upstairs" and "downstairs" market of the New York Stock Exchange. The downstairs market is the main market, organized as a (partially) computerized double auction; the upstairs market is reserved for very large trades that could not be concluded without delay through the main market, but are closed through the dealers' personal negotiations. If the downstairs market is sufficiently liquid, it can exercise the necessary disciplinary effect on the upstairs market.¹⁵

To maximize the liquidity of the continuous double auction, we need to consider which role, if any, should be played in it by market makers, and who these market makers might be. Although in theory the auction could function without intermediaries,¹⁶ the presence of market makers actively speculating on their own accounts, and thus acting to maximize the volume of trades, provides liquidity. This is particularly important in a new market, and certainly would be in a

market where the number of traders is expected to be small. In addition, transactions in deficit permits could well be lumpy, with countries entering the market only to purchase or sell relatively large volumes of permits. The continuous nature of the auction, together with the year long horizon over which countries can plan for necessary acquisitions of permits, are meant to alleviate the problem, but we cannot exclude that the market may be required to accommodate large trades, or risk losing all relevance. A hybrid system, where the auction's order book is supplemented by a dealer ready to close large transactions at price quotes that cannot be more favorable to traders than the current best prices in the order book, is a possible solution.

If the market maker comes to play a large role in the market, it is tempting to think that the role should be played by one of the European Union's institutions. If acting directly though, the institution must be constrained so as to refrain from any intervention aimed at affecting the market price of the permits, or the decentralization of the permits scheme loses meaning and credibility. If the function is subcontracted to a private financial intermediary, the choice of the agent and the extent and transparency of the guarantees he would enjoy become critical.

Alternatively, in line with existing securities markets, the market maker's function could be left to the private sector, tapping existing know-how and eliminating concerns about regulatory intervention. An intermediary able to provide the necessary liquidity and absorb potentially very large positions would probably need to be a consortium of existing banks and securities firms. But should banks and securities firms, and private agents in general, be allowed to trade on the deficit permits market?

4.3. Who Should Trade?

We need to face explicitly an issue we have avoided so far: who should be allowed to trade on the permits market? There are two parts to the question: First, should trading be open to all, or should it be restricted to the agents required to hold permits? Second, even if we decide for the latter, should these agents be exclusively the central governments of the member countries, or should we extend the program to state, provincial and local governments? Neither answer is obvious.

In environmental permits markets, anybody can trade. Two arguments support this design: first, the larger the number of traders, the more likely that the outcome of the market will be competitive; second, participation in the market by consumers and environmental groups can lead to a better approximation of the socially desirable total pollution ceiling. In reality the presence of consumers in the permits market remains very limited, but the first argument has proven important: the market would not have developed as it has without the presence of brokers who have acted as market-makers and engineered new contracts.

In the case of fiscal permits, the logic is unchanged. The more dispersed the ownership and trading of the permits, the more likely that the market will be thick and efficient. A fully developed market for deficit permits should not be any more fragile or manipulable than a market for, say, government bonds. If anything, the yearly allocation of permits to the member countries should offer a partial buffer from market instability.¹⁷ In the government bonds market, the presence of large financial intermediaries, private agents, Central Banks and foreign governments is not seen as a threat but as an important contribution to the volume, liquidity and efficiency of the markets.

The problem however is that the market would not be born fully developed. Given the political

sensitivity of the assets traded, opening a thin tentative market to all - and particularly to large foreign players - would probably rely too much on the competitive features of the double auction and would certainly be very controversial. Even within Europe, allowing the ECB to be an active trader would raise the concern that the Bank could control both monetary and fiscal policies for the Union.¹⁸ A more realistic alternative is to begin by restricting entry to the players directly bound by the program and to regulated brokers. Private intermediaries would be monitored, but should be allowed to trade because their objectives are more closely aligned with profit maximization than governments', and thus their presence will improve the functioning of the market. Opening the market further should be considered at a later stage.

Which agents then should be bound by the program? The discussion so far has assumed that only national governments would be required to hold permits. But the Stability Pact specifies ceilings for general government deficit, thus including deficits incurred by state, provincial and local governments. The extent to which compliance with the Pact should be shared among the different levels of government is hotly debated: in Germany, the states have questioned whether the Pact, adhered to by the central government, can constrain their behavior; in Italy, as in other unitary countries, the legislature has issued laws on an "Internal Stability Pact" (see Law no.448, December 23rd, 1998 and the following interpretative directive).

The point is important. In all member countries a sizable share of general government expenditure is channeled through state and local governments (the average share was 24 percent in 1995). If we look at the four largest economies alone - accounting for more than 80 per cent of the Union's GDP in 1997 - not only this remains true, but, what is most significant, state and local governments have consistently borrowed directly on financial markets (Table 1).¹⁹

Table 1

	F	D	I	E
local governments expenditure share (1995)	18.2	24.4	24.8	25.6
net borrowing share (average 1990-1995)	5.6	28.8	4.9	23.4

Source: All calculations are based on OECD National Accounts, vol.2, 1984-1996. The expenditure share includes both current and capital expenditure. For Germany and Spain, the local governments data include state and provincial governments.

The most likely outcome at present is an arbitrary partition of the total allowed deficit between central and local governments. But the deficit permits scheme suggests a superior alternative: if local governments are allowed to borrow, they should also be required to hold permits. After having received its initial allocation of permits, a national government should distribute it among its different jurisdictions. All jurisdictions would then be allowed to intervene in the permits market, so that the final allocation of permits, and the pattern of borrowing, would reflect the different costs of fiscal austerity. Efficiency would be enhanced through two channels: first, with a competitive permits market, the marginal costs of the borrowing constraint would be equalized not only across countries, but also within each country, among its local governments. Second, and equally important, the increase in the number of players and the reduction in the size of each player would improve the functioning of the market and increase the chance that its outcome would indeed be competitive.

Because the initial allocation of permits does not impinge on efficiency, each country would be free to decide on the appropriate distribution across states or regions. Thus without preventing

redistributive schemes within a country, the program could further a policy of fiscal federalism, in line with the often enunciated principle of subsidiarity. Notice however that simply requiring local governments to hold permits would not per se impose a federalist structure on a country: the central government could remain the single decision-maker behind the local governments' trades. The existence of a permits price would make misallocations of spending patterns among local governments more transparent, and thus easier to correct, but the distribution of effective powers is a political decision that the program leaves, as it must, to each country.

If local governments are allowed to trade independently, the market for permits will have its desired effects only if compliance with the program is expected and enforced at the local level. If a local government missing the required permits is able to either acquire them at no cost from the central government or to shift the penalties to the national level, then the program is irrelevant. Indeed, enforcing compliance with permits requirements is not sufficient to ensure that the market will work as desired: it is also necessary that debt bail-outs from the central governments be ruled out. If local governments do not expect to be held responsible for repaying any debt they incur, then they will simply borrow more, so as to buy the permits necessary to support their deficits. Of course, the possibility of debt bail-outs should distort local governments' borrowing behavior independently of whether or not a market for permits is in place. Since local governments already have direct access to financial markets, this cause of distortion presumably is not too severe.

In the end, whether or not local governments can count on bail-outs by central governments must depend on how constrained the latter are. Both at the local level and at the national level, it is clear that the deficit permits program cannot be effective unless the correct enforcement is provided. Some additional observations on the relative ease of enforcement of a market for

permits versus the current quantitative limits of the Stability Pact are in order, and we turn to them now.

4.4. Enforcement and Political Economy.

In the current version of the Stability Pact, penalties for countries violating their deficit ceiling are not automatic, and it is hard to escape the impression that compliance with the Pact may not be enforced. In a market for deficit permits, enforcement will still be at the arbitrium of the collective will of the Union members, but there are some additional elements that should make enforcement easier. First, as emphasized earlier, the costs of compliance would be lower, and extenuating circumstances meriting exceptions would be harder to claim. In addition, suppose a country is found in violation of the scheme. If penalties are not enacted, the price of the permits immediately falls to zero; if the trading price was positive, any country who has saved permits for future use, or who has emitted debt to finance permits purchases suffers a capital loss. Thus not only is the Union affected by the negative externality attached to excessive deficit spending, as in the case of violation of quantitative limits in the Stability Pact, but the market imposes a direct financial penalty on those countries who played by the rules.

As in the previous discussion of local governments, even if penalties are imposed and collected, under both schemes countries will constrain their behavior only if they indeed bear the final responsibility for such penalties. If countries expect that their obligations will eventually be shared among Union members (directly or through inflationary pressures on the ECB), then enactment of the penalties is irrelevant. Although the possible expectation of a debt bail-out is one of the original motivations for the Stability Pact, the Pact does not address the problem

directly, beside taking the important step of making a country's fiscal imbalances very visible. Again, the same general reasoning holds in the case of a market for permits. As in the case of enforcement of penalties for non-compliance, here too the only real differences are the added flexibility enjoyed by all countries and the additional capital loss that a fiscal bail-out imposes on the other Union members.

Within any one country, taking full advantage of the flexibility allowed by the program should mean being able to trade permits freely across time. Rational players choosing their series of deficits over an infinite horizon should be allowed not only to save permits, but also to borrow them so as to smooth the costs of fiscal contraction over time.²⁰ On the other hand, in a world of democratic governments with short horizons and strong electoral pressures, allowing a government to spend with no effective restrictions, while imposing constraints on (remote) successors seems rather unwise. The more pragmatic approach recommended in this paper, such that governments would be allowed to save but not to borrow permits from their future allocations, is still consistent with some intertemporal trade. As discussed in the case of environmental markets, swaps of current permits in exchange for future ones among market participants may still emerge. They are a form of intertemporal borrowing with two main advantages: first, the desired aggregate fiscal ceiling is satisfied each period; second, the terms of the contract are decided by the market, and thus no contract will emerge if future repayments are not credible. If the alternation of different, competing governments makes past obligations unlikely to be honored, then no borrowing will occur. The market would be so new that it is impossible to tell now whether the permits will acquire the status of "standard" financial assets, mostly insensitive to government changes, or not. Thus leaving this matter open for market

participants to resolve seems correct. In addition, preventing governments from borrowing directly from their own future permits limits the mistakes that initial inexperience with the program is sure to produce, an important benefit in its own right.²¹

5. Conclusions.

This paper has discussed the creation of a market for tradable deficit permits as an efficient mechanism for the implementation of fiscal constraints in the European Monetary Union. When compared to the current provisions of the Stability Pact, a system of deficit permits would have a number of advantages. First of all, it would be much more flexible - individual countries could decide to incur larger deficits by purchasing permits on the market. Thus for example a negative idiosyncratic shock could be overcome at low cost, since the market price of the permits would reflect low demand by the other Union members. At the aggregate level, supply could be adjusted in case of Europe-wide recessions.

Second, because of its flexibility a system of tradable permits minimizes the aggregate costs of compliance with the fiscal target. Given the high rates of unemployment in the four largest countries of the Union, insuring that the costs of fiscal discipline are as low as possible is particularly important. And, if such discipline is indeed desired, the probability of enforcing it should be higher the lower the cost of doing so.

Third, by allowing countries to save or sell their unused permits, the present scheme gives them the incentive to reduce their deficit below the fixed 3 per cent ceiling of the Stability Pact. The Pact's recommendation of a balanced budget in the medium run becomes much more likely to be implemented if it is accompanied by appropriate rewards for doing so.

Finally, given the general idea of a system of tradable permits, the design of the market can vary to reflect the specific policy concerns that have inspired the call for fiscal discipline. For example, countries with different debt positions can be treated differently, mirroring the fear that deficits from economies with larger outstanding debts may be particularly destabilizing for the Union as a whole.

Will countries accept a deficit permits market, as alternative to the current provisions of the Stability Pact? In the aggregate the costs of compliance would be lower; individually, countries that remain below the bound of the 3 percent deficit to GDP ratio and have an equivalent initial allocation of permits would gain, by being able to sell their unused permits. Countries that go above their initial allotment would have to purchase permits on the market; we expect that the cost of doing so would be lower than the very high penalties foreseen by the Stability Pact, if the purchase is not too large, but a more precise answer to this question must wait for a quantitative estimate of the permits market equilibrium price.

The real difficulty in comparing the costs that countries would sustain under the two schemes is assessing the probability that the Stability Pact will indeed be enforced. It is quite possible that when the Stability Pact was agreed upon, policy-makers believed that political considerations, more than economic ones, would finally determine whether or not fiscal discipline would be imposed on Union members. The fixed quantitative targets, together with the possible exceptions, create strong temptations for political settlements concluded outside the public view. In an enlightening discussion of political obstacles to environmental regulation, Robert Stavins describes how difficult it was to go beyond fixed ceilings for individual sources of pollution (Stavins, 1998). “Old style” ceilings were popular with the public, because they sounded severe

and unforgiving; they were popular with polluting firms, because firms knew that exceptions could and would be negotiated out of the limelight; they were popular with politicians because politicians maintained final control and benefitted from the exceptions they were able to grant.²²

But international capital markets are not very forgiving of ambiguities. At the time of this writing, five months after the introduction of the Euro, the markets appear to demand compliance with the fiscal targets, demonstrating a loss of confidence in the European currency in response to any indication of a softening fiscal stance. Enacting a market in deficit permits, with its predictability, its transparency and its realistic requirements, could be an important step towards gaining the international role for the Euro that we all expect.

References

- Bailey, Elizabeth M., 1996. "Allowance Trading Activity and State Regulatory Rulings: Evidence from the U.S. Acid Rain Program", W.P. No.96-002, MIT CEEPR, March.
- Baumol, William J. and Wallace E. Oates, 1988. The Theory of Environmental Policy, Cambridge, U.K.: Cambridge University Press.
- Beetsma, Roel and Harald Uhlig, 1998. "An Analysis of the Stability and Growth Pact", mimeo, May.
- Bernheim, Douglas B. and Michael D. Whinston, 1986. "Menu Auctions, Resource Allocation, and Economic Influence", Quarterly Journal of Economics, 101, February, 1-32.
- Brookings Papers on Economic Activity, 1978. Special Issue: Innovative Policies to Slow Inflation, vol. 2., Washington DC: The Brookings Institution.
- Camerer, Colin F. and Keith Weigelt, 1993. "Convergence in Experimental Double Auctions for Stochastically Lived Assets", in Daniel Friedman and John Rust (eds.) The Double Auction Market, 355-396.
- Canzoneri, Matthew B., Robert E. Cumby and Behzad T. Diba, 1998. "Is the Price Level Determined by the Needs for Fiscal Solvency?", CEPR Discussion Paper No. 1772, January.
- Cason, Timothy N., 1993. "Seller Incentive Properties of EPA's Emission Trading Auction", Journal of Environmental Economics and Management, 25, 2, September, 177-95.
- Cason, Tomothy N., 1995. "An Experimental Investigation of the Seller Incentives in EPA's Emission Trading Auction", American Economic Review, 85, 4, September, 905-22.
- Clauser, Laura and Charles R. Plott, 1993. "On the Anatomy of the "Nonfacilitating" Features of the Double Auction Institution in Conspiratorial Markets", in Daniel Friedman and John Rust

(eds.) The Double Auction Market, 333-353.

Copeland, Thomas E. and Daniel Friedman, 1987. "The Effect of Sequential Information Arrival on Asset Prices: An Experimental Study", Journal of Finance, 42, July, 763-797.

Crocker, Thomas D., 1966. "The Structuring of Atmospheric Pollution Control Systems", in Harold Wolozin (ed.) The Economics of Air Pollution, New York, N.Y.: W. W. Norton and Co., Inc..

Dales, John, 1968. Pollution, Property and Prices, Toronto, Canada: University Press.

Eichengreen, Barry and Charles Wyplosz, 1998. "The Stability Pact: More than a Minor Nuisance?", Economic Policy, 26, April, 67-113.

Ellerman, A. Denny, Richard Schmalensee, Paul L. Joskow, Juan Paulo Montero and Elizabeth M. Bailey, 1997. "Emission Trading Under the U.S. Acid Rain Program: Evaluation of Compliance Costs and Allowance Market Performance", Special Report, MIT, CEEPR, October.

Ellerman, A. Denny, Paul L. Joskow and Richard Schmalensee, 1997. "1996 Update on Compliance and Emissions Trading Under the U.S. Acid Rain Program", MIT, CEEPR, November.

Chari V.V. and Patrick J. Kehoe, 1998, "On the Need for Fiscal Constraints in a Monetary Union", Federal Reserve Bank of Minneapolis, Working Paper No. 589, August.

Flandreau, Marc, Jacques Le Cacheux and Zumer, 1998. "La Belle Epoque of Stability Pacts", Economic Policy, 26, April.

Friedman, Daniel, 1984. "On the Efficiency of Experimental Double Auction Markets", American Economic Review, 74, March, 60-72.

Friedman, Daniel, 1993. "The Double Auction Market Institution. A Survey", in Daniel

Friedman and John Rust (eds.) The Double Auction Market, 3-25.

Friedman, Daniel and Joseph Ostroy, 1995. "Competitiveness in Auction Markets: An Experimental and Theoretical Investigation", Economic Journal, 105, January, 22-53.

Friedman, Daniel and John Rust (eds.), 1993. The Double Auction Market. Institutions, Theories and Evidence, Santa Fe Institute Studies in the Science of Complexity, New York, N.Y.: Addison- Wesley Publishing Co.

Fullerton, Don, Shaun P. McDermott and Jonathan P. Caulkins, 1996. "Sulfur Dioxide Compliance of a Regulated Utility", NBER W.P. No.5542, Cambridge, Ma, April.

Hahn, Robert W., 1984. "Market Power and Transferable Property Rights", Quarterly Journal of Economics, 99,4, 753-65.

Hahn, Robert W., 1989. "Economic Prescriptions for Environmental Problems: How the Patient Followed the Doctor's Orders", Journal of Economic Perspectives, 3, 2, Spring, 95-114.

Isaac, R. Mark and Charles R. Plott, 1981. "The Opportunity for Conspiracy in Restraint of Trade" Journal Of Economic Behavior and Organizations, 2, January, 1-30.

Joskow, Paul L. and Richard Schmalensee, 1998. "The Political Economy of Market-Based Environmental Policy: The U.S. Acid Rain Program", Journal of Law and Economics, XLI, April, 37-83.

Joskow, Paul L., Richard Schmalensee and Elizabeth M. Bailey, 1998. "The Market for Sulfur Dioxide Emissions", American Economic Review, 88,4, September, 669-685.

Krupnick, Alan J., Wallace E. Oates and Eric Van De Verg, 1983. "On Marketable Air Pollution Permits: The Case for a System of Pollution Offsets", Journal of Environmental Economics and Management, 10,3, September, 233-47.

- Leeper, Eric M., 1991. "Equilibria Under Active and Passive Monetary and Fiscal Policies", Journal of Monetary Economics, 27, February, 129-147.
- Lerner, Abba P., 1977. "Stagflation - Its Causes and Cure", Challenge, 20, Sept/Oct., 14-19.
- Lerner, Abba P. and David C. Colander, 1980. MAP: A Market Anti-Inflation Plan, Harcourt Brace Jovanovich, Inc.: New York.
- Montgomery, W. David, 1972. "Market Licenses and Efficient Pollution Control Programs", Journal of Economic Theory, 5, 3, 395-418.
- Okun, Arthur M., 1977. "The Great Stagflation Swamp", Challenge, 20, Nov/Dec., 6-13.
- Pagano, Marco, 1998. "The Changing Microstructure of European Equity Markets", in Guido Ferrarini (ed.), The European Securities Markets: The Investment Services Directive and Beyond, Kluwer Law International.
- Riker, William H. and Itai Sened, 1991. "A Political Theory of the Origin of Property Rights: Airport Slots", American Journal of Political Science, 35, 4, 951-69.
- Schmalensee, Richard, Paul L. Joskow, A. Denny Ellerman, Juan Pablo Montero and Elizabeth M. Bailey, 1998. "An Interim Evaluation of Sulfur Dioxide Emissions Trading", Journal of Economic Perspectives, 12, 3, Summer, 53-68.
- Seidman, Laurence S., 1976. "A New Approach to the Control of Inflation", Challenge, 19, July/Aug. 39-43.
- Sims, Christopher A., 1998. "The Precarious Fiscal Foundations of EMU", Yale University, mimeo, December.
- Smith, Vernon L., 1982. "Microeconomic Systems as an Experimental Science", American Economic Review, 72, 923-955.

- Smith, Vernon L., Arlington W. Williams, W. Kenneth Bratton and Michael G. Vannoni, 1982. "Competitive Market Institutions: Double Auctions v/s Sealed Bid-Offer Auctions", American Economic Review, 72, March, 58-77.
- Smith, Vernon L., Gerry L. Suchanek and Arlington W. Williams, 1988. "Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets", Econometrica, 56, 1119-1152.
- Stavins, Robert, 1998. "What Can We Learn from the Grand Policy Experiment? Lessons from SO2 Allowance Trading", Journal of Economic Perspectives, 12, 3, 69-88.
- Tietenberg, Thomas H., 1985. Emissions Trading: An Exercise in Reforming Pollution Policy, Washington D.C.: Resources for the Future, Inc..
- Vickrey, William, 1986. "Design of a Market Anti-Inflation Program" in David C. Colander (ed.) Incentives-Based Incomes Policies, Cambridge, Ma: Ballinger, 149-158.
- Vickrey, William, 1993. "Today's Tasks for Economists", American Economic Review, 83, March, 1-10.
- Wallich, Henry C. and Sidney Weintraub, 1971. "A Tax-Based Income Policy", Journal of Economic Issues, 5, June, 1-17.
- Weitzman, Martin L., 1974. "Prices vs. Quantities", The Review of Economic Studies, 41, October, 479-489.
- Wilson, Robert, 1987. "Equilibrium in Bid-Ask Markets" in G. Feiwel (ed.) Arrow and the Ascent of Economic Theory: Essays in Honor of Kenneth J. Arrow, London: McMillan.
- Woodford, Michael, 1996. "Control of the Public Debt: A Requirement for Price Stability", Princeton University, mimeo, February.

Footnotes

* I thank Kai Konrad, Charles Wyplosz, Ben Lockwood, Peter Birch Sørensen, Michael Bolle, Prajit Dutta, Raquel Fernandez, Louis-Andre Gerard -Varet, Alan Kirman, Dan O'Flaherty and Alex Pfaff for comments and advice. In particular, I thank Peter Kenen, whose observations led me to this paper, Giorgio Basevi, who independently had the same idea, and Jurgen von Hagen, who suggested extending the analysis to local governments, for long and fruitful conversations. The paper was written at GREQAM in Marseilles and at the New York office of the NBER, and I thank both institutions for their hospitality.

1. The Pact, finalized at the June 1997 meeting of the Council of Ministers, specifies the details of the Excessive Deficit Procedure of the Maastricht Treaty. An annual fall in real GDP of 2 per cent or more is considered an “exceptionally” severe downturn, and the deficit to GDP ratio can be above the 3 per cent limit as long as a recession of this magnitude persists (although it must remain “close” to the reference value). If the decline in real GDP is between 0.75 and 2 per cent, an exception can also be invoked if the Council of Ministers concurs.
2. The expectation that a country will not be allowed to default on its debt implies that all debts are seen by the market as perfect substitutes, and thus are equally priced. The resulting interest rate spillovers are a form of pecuniary externality that causes no distortion if each country is too small to affect the equilibrium interest rate, but leads to suboptimally large debts if, acting individually, countries can affect the market and manipulate it strategically. Pressures on the Central Bank to inflate away excessive debts are, as stated in the text, a form of bail-out and thus also create a fiscal externality. The literature has identified a third possible reason for intervention in the theoretical possibility of indeterminacy of the price level (Leeper (1991), Woodford (1996),

Canzoneri, Cumby and Diba (1998), Sims (1998)).

3. It is possible to derive rigorous conditions that determine whether a tax or a quantity ceiling is the superior policy in the presence of uncertainty (Weitzman (1974)). Weitzman concludes generally in favor of quantity constraints, because he believes that small deviations from the optimal quantity are associated with steep reductions in benefits. For a very balanced discussion, see Baumol and Oates (1988), ch.5.

4. As remarked by Peter Birch Sørensen, both the Stability Pact and a system of tradable deficit permits induce countries to underreport their deficits. In the case of tradable permits, the incentive to underreport exists even when a country is below the 3 percent limit, because it benefits from selling or banking its permits. But it seems unlikely that countries could indefinitely resort to accounting legerdemain, or that over time the difference in underreporting between the two schemes would be of significant magnitude.

5. Because the market price p is always smaller than 1, setting the penalty for non-compliance at or above 1 for each missing permit guarantees that the fee is higher than the market price. Alternatively, the fee can be set equal to a stated multiple of the price, again exploiting the ceiling at 1 to ensure that the market price never follows an explosive path.

6. A different proposal for a permits market targeted, like the one discussed in this paper, to macroeconomic policy, was never enacted but is too interesting not to be mentioned. Concern with inflation in the 70's triggered the idea that individual nominal price increases caused an externality through their impact on inflation. Hence they could be curbed through appropriate taxes: TIPS (tax-based income policy schemes) would use penalties or subsidies to induce firms to refrain from granting wage increases and raising prices (Wallich and Weintraub (1971), Seidman

(1976), Okun (1977)). The customary difficulties of a tax scheme led to the suggestion of a market for licenses to increase prices (Lerner (1977), Lerner and Colander (1980), Vickrey (1986)). A special issue of the Brooking Papers on Economic Activity (1978:2) was devoted to the mechanisms, and Vickrey was sufficiently intrigued by the idea to return to it in his presidential address to the American Economic Association (Vickrey (1993)). In principle, the scheme was meant to apply to all firms and makes the proposal in this paper appear very moderate.

7. For a description of experiences preceding Title IV, see for example Tietenberg (1985) or Hahn (1989).

8. When sulfur dioxide and nitrogen oxides react in the atmosphere, they generate sulfuric and nitric acids that fall back on earth in the form of acid rain. Electric utilities were responsible for 70 per cent of U.S. SO₂ emissions in 1985, with the remainder 30 per cent contributed by a wide variety of much smaller sources, from diesel fuel for transportation to residential and industrial boilers. Including these smaller sources into the program was quickly judged unfeasible (Joskow and Schmalensee (1998)).

6. A "unit" is an individual source of emissions; the 263 units whose participation in phase 1 was mandated belong to 110 separate plants in 22 states, managed by 61 operating companies (some of which may be subsidiaries of the same parent company). An additional 182 units were voluntarily brought under the program in phase 1 through incentives created by special provisions.

10. However, if you are thinking of purchasing very small quantities - a birthday present for your environmentalist niece? - you may need to bid at the annual EPA auction, because the brokers will not accept too small orders.

11. This was more than a year and a half before utilities had to demonstrate compliance for the first time.
12. For comparison, the simple case discussed in Section 2 is equivalent to holding the unweighted sum of the Union's countries deficits below a target level.
13. For example, using as threshold a ratio of debt to GDP of 65 percent, the two groups would be: Belgium, Italy, the Netherlands and Spain, as high debt countries; and Austria, Finland, France, Germany, Ireland, Luxembourg and Portugal, as low debt countries.
14. Clauser and Plott find that experiments designed to eliminate this second source of competition still resulted in competitive prices, suggesting that the first feature may be more important.
15. In most European exchanges, this dual structure takes the form of a computerized continuous auction for trades of smaller size, and a quote-driven dealers' market for large trades (Pagano (1998)). A trading system that protects a trader's identity is useful because lack of anonymity in the permits market could create two problems. First, competing countries could collude and artificially raise (or lower) the price at which transactions take place. At least partly, this difficulty is mitigated by those features of a double auction that appear to curb collusion. Second, although governments' baseline demand for permits will be written in the budget plans, and thus be public information, revisions during the fiscal year will become known abroad only with a lag. Thus knowing where sales or purchases of permits originate would transmit information about future demand and could result in large bid-ask spreads for large players, perceived as insiders.
16. Some existing proprietary trading systems for example provide institutional investors with direct access to the order book (see Tradepoint in the London market).
17. It may seem at first that the market could be more volatile than the market for bonds, possibly

subject to bubbles and crashes: contrary to bonds, the permits would not have an expiration date with a fixed redemption value. However the market price of the permits can never exceed 1, and this is sufficient to rule out the possibility of bubbles.

18. An intriguing question is whether agents - political parties for example - could try to influence election outcomes by trading in the permits market. Here too, if the market were thick and competitive there would be no reason why private players should not be allowed to withdraw permits and establish a record as fiscal hawks, exactly as environmental groups can withdraw emission permits. More troublesome however would be the possibility that a group may interfere with a government's fiscal policy by manipulating the price at which the government can trade permits. Once again the problem exists if the market is thin and manipulable.

19. In the smaller countries, local governments appear almost always as net lenders. All data are from the OECD National Accounts statistics; the average of expenditure shares is unweighted; the GDP numbers are at current prices and current exchange rates.

20. The correct dynamic plan depends on whether the target aggregate ceiling for the regulator includes only deficits, or only debts or both. The Appendix discusses in detail the static case only, but the optimal dynamic plan could also be characterized (although it could be quite complex when both deficits and debts are included).

21. Of course, a government can still shift the costs of its fiscal spending on the future administration by not complying with the program. Two factors though may limit this temptation: first, the government must be sure to lose the elections, in which case the political gains from excessive spending would be few; second, the violation would be very visible.

22. The difficulty seems common to all new market-based regulations: Riker and Sened (1991)

discuss identical problems in introducing tradable rights to airport time slots. As in the environmental case, the final adoption of marketable permits occurred only when the inefficiencies of the previous system became too large to bear.

Figure 1

MC of emission reductions




Figure 2

