Ashoka University Economics
Discussion Paper 68

An Evolutionary Approach to Pollution Control in Competitive Markets

October 2021

Ratul Lahkar, Ashoka University
Vinay Ramani, Indian Institute of Management, Visakhapatnam

https://ashoka.edu.in/economics-discussionpapers

Plot No. 2, Rajiv Gandhi Education City,
National Capital Region,
P.O.Rai, Sonepat, Haryana 131029

An Evolutionary Approach to Pollution Control in Competitive Markets

Ratul Lahkar* Vinay Ramani†

October 11, 2021

Abstract

We consider a large population of firms in a market environment. The firms are divided into a finite set of types, with each type being characterized by a distinct private cost function. Moreover, the firms generate an external cost like pollution in the production process. As a result, the Nash equilibrium outcome is not socially optimal. We propose an evolutionary implementation mechanism to achieve the socially optimal outcome. In contrast to the classical Pigouvian pricing and the VCG mechanism, evolutionarily implementation does not require the planner to know or elicit any private information from firms. Hence, it is informationally parsimonious. By imposing a tax equal to the current external damage being imposed by a firm, the planner can guide the evolution of the society towards the social optimum. The imposition of the tax generates a potential game whose potential function is the social welfare function of the model. Evolutionary dynamics converge to the maximizer of this function thereby evolutionarily implementing the social welfare maximizer.

Keywords: Evolutionary Implementation; Negative Externality; Potential Games; Pigouvian Tax; Dominant Strategy Implementation

*(Corresponding author) Department of Economics, Ashoka University, Rajiv Gandhi Education City, Sonepat, Haryana 131029, India, email: `ratul.lahkar@ashoka.edu.in`

†Indian Institute of Management, Visakhapatnam, AU Campus, Visakhapatnam, Andhra Pradesh 530022, India, email: `v.ramani2710@gmail.com`

1 Introduction

Environmental pollution is a negative externality that impacts all stakeholders, some of whom are beyond the boundaries of the firm (Tietenberg and Lewis [29]). In general, firms do not internalize the negative externality that they impose on other stakeholders. This is because there is a difference between the private marginal cost of production (PMC) and the social marginal cost of production (SMC), which includes the environmental damage. As a result, when the firm chooses its optimal production quantity, in the (Nash) equilibrium outcome by equating the marginal benefit to the PMC, it ends up producing more than the socially optimal (*i.e.,* efficient) level of output. The socially optimal level of output occurs when the marginal benefit is equal to the SMC. This leads to the fundamental question - *what should a regulator or policy maker do to achieve the socially optimal outcome?*

Pigouvian taxation is the most well known policy instrument to combat negative externalities (Pigou [22]). As studied in standard microeconomic theory, such taxation works by making the PMC equal to the SMC at the socially optimal level of output. A fundamental difficulty with classical Pigouvian pricing is that it imposes quite onerous informational requirements on the policy maker (or the social planner/regulator) (Baumol [1]). In particular, the planner needs to have sufficient information about the private characteristics of firms, for example, their cost functions, in order to calculate the socially optimal level of output. In many cases, policy makers may not be privy to such information. It is perhaps due to such difficulties that more market based mechanisms like carbon trading are being favored over direct taxation. At a theoretical level, the informational problem may be resolved by appealing to the mechanism design literature in which the planner may, for example, design a Vickrey-Clarke-Groves (VCG) mechanism which renders truthful revelation of private information a dominant strategy (Vickrey [30], Clarke [6], and Groves [8]). In practice, though, VCG mechanisms may be difficult to apply due to difficulties in collecting information and assigning optimal actions particularly if a large number of agents are involved (Sandholm [26]), and concerns about revealing confidential private information (Rothkopf [23]).

As one possible resolution of such problems with the classical methods, we consider an evolutionary approach towards implementing the socially optimal outcome in a market environment where firms generate a negative externality like pollution during production. The evolutionary approach requires us to consider a large population of firms, each of measure zero. The firms are divided into a finite set of types, with each type being characterized by a distinct private cost function. Although technically, our model is one of Cournot competition, the measure zero characteristic of each firm implies that the solution we obtain is the competitive one with the Nash equilibrium coinciding with the Walrasian equilibrium. The Nash equilibrium is, however, socially inefficient due to the presence of the external cost. The objective of the planner is to implement the socially efficient outcome that maximizes the aggregate welfare of consumers and producers while allowing for the presence of the external cost of pollution.

The method of evolutionary implementation applies the principles of evolutionary game theory to enable the planner to induce the society towards the Pareto efficient state. In evolutionary game theory, society doesn't immediately coordinate on an equilibrium outcome but instead, may converge towards it under the influence of certain standard evolutionary dynamics (Sandholm [27]). The key insight of this approach in our model is that by imposing a tax equal to the external cost being imposed by a firm, the planner can guide the evolution of the society towards the social optimum instead of the sub–optimal Nash equilibrium. At this social optimum, the level of output produced by firms and, therefore, the pollution they generate will be lower than at the original Nash equilibrium because firms would have internalized the external cost they impose upon society.

Crucially and importantly, the informational requirements for such evolutionary implementation are minimal. All that is required is that the planner should observe the output being currently produced by a firm and know the external cost, which is assumed to be common for all firms. Unlike in classical Pigouvian pricing or the VCG mechanism, there is no need for the planner to know the demand function of consumers or the private cost function of firms, or to induce truthful revelation of any private information, which also precludes concerns about confidentiality. This is because evolutionary implementation doesn't require the planner to calculate the social optimum. Instead, updating the tax based on current output levels suffices to nudge the market towards the optimal state. In fact, we also provide an analysis of Pigouvian pricing and the VCG mechanisms in our large population context to highlight the higher informational requirements of these approaches in comparison to evolutionary implementation.

It should be noted that Pigouvian pricing and the VCG mechanism enables instantaneous coordination on the Pareto optimum. In contrast, evolutionary implementation is more gradual. But by greatly economizing on informational requirements as compared to these other classical approaches, evolutionary implementation emerges as a more feasible alternative particularly when information is scarce. Indeed, at a more practical level, market based mechanisms for pollution control like carbon trading that have emerged as alternatives to Pigouvian taxation may be regarded as real world manifestations of these very principles that underlie evolutionary implementation. The regulator seeks to create a market for pollution permits and allows the market to discover the price for pollution that a polluter must pay (Baumol and Oates [2], Tietenberg [28]). Such market based mechanisms also rely on guiding the behavior of participants towards the social optimum rather than immediate enforcement. Nevertheless, despite their gradualist approach, market mechanisms have become popular precisely because policy makers lack the information for ensuring such enforcement through classical taxation methods.

The technical foundations of evolutionary implementation are provided by the class of games called potential games (Monderer and Shapley [18]). These are games in which incentives can be summarised using a real valued function called the potential function. In large population potential games, standard evolutionary dynamics converge to maximizers of the potential function, which are also Nash equilibria of the underlying game (Sandholm [24]). In our model, the planner transforms the original Cournot game into an externality adjusted game by taxing firms for the external cost that they generate on the basis of their current output. This generates a potential game which has the important feature that its potential function is also the social welfare function of the model.

Evolutionary dynamics converge to the unique Nash equilibrium of this game which, given the identity of its potential function, must also be the social welfare maximizer.

Evolutionary implementation was introduced by Sandholm [25, 26] in the context of congestion games. Since then, this method has been extended to other types of games like public goods, public bads and the tragedy of the commons (Lahkar and Mukherjee [14, 16]). An important difference between the present model and the earlier applications is that here, we consider evolutionary implementation in the context of a competitive market that produces a private good but which creates a negative externality in the process of production. The earlier models considered non market products like public or common resources. As far as we know, this is the first application of evolutionary implementation to problems of negative externalities in a market environment. Moreover, unlike the earlier models, this paper provides a comparison of evolutionary implementation with classical approaches like VCG mechanism with respect to their informational requirements.

Methodologically, therefore, this paper extends the scope of applying evolutionary implementation to a greater range of interactions since most economic activity in a modern society is mediated through markets. It also suggests that policies designed to control environmental pollution are feasible even when the planner has little information about the agents causing such pollution. The difference in the nature of the goods being considered also has analytical implications. As we elaborate in Section 4, unlike in the earlier models, we cannot identify social welfare simply with the welfare of the active agents in the model, which are the producers. Instead, we also need to incorporate the welfare of consumers even though they do not have a strategic role in the model. An important technical contribution of this paper is to generalize the procedure of maximizing an abstract measure theoretic potential function in continuous strategy games like the large population Cournot model with multiple types of agents.

There is by now an extensive literature on the application of mechanism design and implementation theory on the regulation of pollution and other environmental hazards. Kim [10] and Wang et al. [31] are recent examples of such papers. As is the case with much of the standard mechanism design literature, such applications also rely on inducing voluntary disclosure of private information. Hence, as with the classical VCG mechanism, there may be practical concerns like collecting and maintaining confidentiality of information when there are a large number of firms. In contrast, in our evolutionary approach, there is no need for firms to disclose any information. Some of these environmental applications are also in the framework of dynamic mechanism design theory where the optimal mechanism unfolds over time. They, therefore, make significant demands on ability of firms to make rational and forward looking calculations about their optimal responses. Like any evolutionary game theoretic model, a dynamic element is also built into our approach. But there is a crucial difference. Firms in our model need not be farsighted or be capable of highly sophisticated strategic thinking. In evolutionary models, agents are assumed to be myopic in the sense that they condition their behavior on current social conditions and their strategy revisions procedures can also be fairly simple.[1] Hence, evolutionary game theory holds even under more realistic norms of

[1]For example, imitating other successful agents generate the replicator dynamic, which is the most well know

human rationality. In our view, this makes policies designed for evolutionary implementation more robust than conventional mechanisms that may require more nuanced strategic thinking.

The remainder of the paper is as follows. In Section 2, we introduce the model. Section 3 defines potential games and solves for the Nash equilibrium. In Section 4, we introduce evolutionary implementation and establish the main result on convergence to the socially efficient state. Section 5 compares our results to the classical methods of Pigouvian taxation and VCG mechanism. We offer concluding remarks in Section 6.

2 Model

We consider the usual set–up of a population game with a continuous strategy set.[2] Let there be a *population* consisting of a continuum of firms of mass 1. Each firm is of measure zero. The population is divided into a finite set $\mathcal{P} = \{1, 2, \cdots, n\}$ of types with each type, as we describe below, having a distinct private cost of production function. We denote the mass of type p of firms as $m_p \in (0, 1)$ with $\sum_{p \in \mathcal{P}} m_p = 1$ and refer to $m = (m_1, m_2, \cdots, m_n)$ as the type distribution. A strategy for a firm is to produce an output $x \in \mathcal{S} = [\underline{x}, \bar{x}] \subseteq \mathbf{R}_+$ with $\underline{x} \geq 0$ and $\bar{x}$ a finite number that can be arbitrarily large.[3]

Due to the continuous strategy nature of the game, we introduce certain measure theoretic notation to represent a state of the game. Let $\mathcal{B}$ be the Borel $\sigma-$algebra on $\mathcal{S}$ and $\mathcal{M}(\mathcal{S})$ be the space of finite signed measures on $(\mathcal{S}, \mathcal{B})$. The subset $\mathcal{M}^+_m(\mathcal{S}) \subset \mathcal{M}(\mathcal{S})$ is then the space of finite measures that impose a total mass of $m > 0$ on $\mathcal{S}$.[4] A state of type p is such a finite signed measure $\mu_p \in \mathcal{M}^+_{m_p}(\mathcal{S})$ with $\mu_p(B) \in (0, m_p)$ denoting the mass of type p firms producing output $x \in B$. If all agents of a particular type p, play the same strategy x, then we refer to that type state as monomorphic and denote it as $m_p \delta_x$. Otherwise, the type state is polymorphic. A state of the entire population of firms is then described by the $n-$tuple $\mu = (\mu_1, \mu_2, \cdots, \mu_n) \in \Delta = \prod_{p=1}^n \mathcal{M}^+_{m_p}(\mathcal{S})$. Thus, Δ is the set of states of the entire population. Suppose $\mathfrak{M}_b(\mathcal{S} \times \mathcal{P})$ is the space of bounded measurable functions on $\mathcal{S} \times \mathcal{P}$ with the supremum norm. A population game is a weakly continuous mapping

$$F : \Delta \to \mathfrak{M}_b(\mathcal{S} \times \mathcal{P}). \tag{1}$$

such that $F_{x,p}(\mu)$ is the payoff of an agent in population p who uses strategy $x \in \mathcal{S}$ at the social state μ.

The particular population game we are interested in is the large population model of Cournot competition, which has also been analyzed in Lahkar [13]. To introduce this game, we define the

evolutionary dynamic. When agents play a perturbed best response to current social conditions, it generates the logit dynamic, another well known evolutionary dynamic. See Sandholm [27] for detailed description of such dynamics and the strategy revision procedures that generate them.

[2]See, for example, Lahkar and Mukherjee [14, 16]

[3]Thus, each firm has a common strategy set irrespective of type.

[4]Thus, $\mathcal{M}^+_1(\mathcal{S})$ is the space of probability measures on $\mathcal{S}$.

aggregate output a population state $\mu = (\mu_1, \mu_2, \cdots, \mu_n)$ as

$$A(\mu) = \sum_{p \in \mathcal{P}} \int_{\mathcal{S}} x \mu_p(dx). \tag{2}$$

Our assumptions about the mass of each population implies $A(\mu) \in [\underline{x}, \bar{x}]$. In our subsequent analysis, we will frequently use the notation α to concisely denote $A(\mu)$. More generally, $A(\mu)$ is the aggregate strategy at the state μ. Let now $\beta : [\underline{x}, \bar{x}] \to \mathbf{R}_+$ be the market inverse demand function such that $\beta(A(\mu))$ describes the market price when aggregate output is $A(\mu)$. Further, let $c_p : \mathcal{S} \to \mathbf{R}_+$ be the private cost function of firms of type p. Thus, each firm of a particular type has a common cost function but firms of different types have different cost functions. Then, the payoff of a firm of type p which produces output x at the state μ in the large population Cournot model F is

$$F_{x,p}(\mu) = x\beta(A(\mu)) - c_p(x). \tag{3}$$

This is of course the standard payoff function of a model of Cournot competition appropriately extended to a large population context. We make the usual assumption that the inverse demand function β is strictly decreasing in $\alpha \in [\underline{x}, \bar{x}]$. Further, we assume that for every $p \in \mathcal{P}$, c_p is strictly increasing and strictly convex.

The payoff in (3) depends entirely upon the agent's individual strategy and the aggregate strategy level. Therefore, F constitutes an aggregative game (Corchón [7]). In our subsequent analysis, we will make a distinction between the private cost of a firm and the social cost. In addition to the private cost $c_p(x)$ that a type p firm producing x faces, we will assume that it also imposes an external cost $s(x)$ on the society as a whole, where $s : \mathcal{S} \to \mathbf{R}_+$ is a convex function. For example, $s(x)$ might be the environmental pollution that a firm generates when it produces output x. Society here refers to the population of firms and the consumers whose behavior is captured by the market demand function β. We note here that following standard practice of partial equilibrium analysis, we are not modeling consumers as active agents in our model. Nevertheless, we will be able to account for their welfare by using the market demand function and the standard concept of consumer surplus.

Thus, the social cost of a type p firm producing output x will be $c_p(x) + s(x)$. For now, we focus only on private costs and, therefore, on the population game defined by (3). The presence of the social cost will, however, introduce a difference between the Nash equilibrium of the model and the socially efficient state, which we will analyze in Section 4.

3 Potential Games and Nash Equilibrium

We now characterize the Nash equilibrium of the large population Cournot competition model (3). The Nash equilibrium of such a large population game is defined as follows.

Definition 3.1 *A Nash equilibrium of a population game F is a state $\mu^* = (\mu_1^*, \mu_2^*, \cdots, \mu_n^*) \in \Delta$*

such that for all $x \in \mathcal{S}$*, all* $p \in \mathcal{P}$*, if* x *lies in the support of* μ_p^**, then* $F_{x,p}(\mu^*) \geq F_{y,p}(\mu^*)$*, for all* $y \in \mathcal{S}$*.*

Lahkar [13] has already found the Nash equilibrium of this model. Here, however, we present an alternative method of deriving that Nash equilibrium by using the fact that the Cournot model is a potential game. The reason for presenting this other method is that it will be relevant later when we consider evolutionary implementation of social efficiency.

Potential games are games in which payoffs may be summarised using a real–valued function called the potential function (Monderer and Shapley [18], Sandholm [24]).[5] The concept has been extended to large population games with continuous strategy sets by, for example, Cheung and Lahkar [5] and Lahkar and Mukherjee [14]. To keep this paper self–contained, we present the definition of such games as well as some essential notation.

Defining potential games in our context requires the notion of the Fréchet derivative, which is a generalization of the usual notion of the derivative to Banach spaces. We first extend the domain of the payoff function (3) from Δ to $\mathscr{M} = \prod_{p=1}^{n} \mathcal{M}(\mathcal{S})$.[6] Consider now a direction $\zeta = (\zeta_1, \zeta_2, \cdots, \zeta_n) \in \mathscr{M}$. The Fréchet derivative $DF_{x,p}(\mu)\zeta$, represents the change in $F_{x,p}(\mu)$ when $\mu \in \mathscr{M}$ changes in the direction ζ. Let $f : \mathscr{M} \to \mathbf{R}$ be a Fréchet differentiable function and suppose there exists a function $\nabla f : \mathcal{S} \times \mathcal{P} \to \mathfrak{M}_b(\mathcal{S} \times \mathcal{P})$, called the gradient of f, such that

$$Df(\mu)\zeta = \sum_{p \in \mathcal{P}} \int_{\mathcal{S}} \nabla f(\mu)(x,p)\zeta_p(dx), \text{ for all } \zeta = (\zeta_1, \cdots, \zeta_n) \in \mathscr{M},$$

where $Df(\mu)\zeta$ is the Fréchet derivative of f at $\mu \in \mathscr{M}$ in the direction $\zeta = (\zeta_1, \cdots, \zeta_n) \in \mathscr{M}$. The definition of a potential game with multiple types and a continuous strategy set is then as follows (Lahkar and Mukherjee [14]).

Definition 3.2 *A population game* F *is a potential game if there exists a Fréchet differentiable (with respect to the variational norm) function* $f : \mathscr{M} \to \mathbf{R}$ *such that*

$$\nabla f(\mu) = F(\mu) \textit{ for all } \mu = (\mu_1, \cdots, \mu_n) \in \Delta$$

or, equivalently, $\nabla f(\mu)(x,p) = F_{x,p}(\mu)$ *for all* $(x,p) \in \mathcal{S} \times \mathcal{P}$*. The function* f *is called the potential function of the game* F*.*

With this definition of potential games, we obtain the following result that characterises the large population Cournot model as a potential game. The part of the result has actually been already established as Proposition 5.3 in Lahkar [13]. In addition, we show that the relevant

[5] Monderer and Shapley [18] introduced the original notion of potential games for finite player games. This concept was extended by Sandholm [24] to large population games.

[6] To ensure that $\mathscr{M}$ is a Banach space, we impose the variational norm on it. See, for example, Appendix A.1.1 in Lahkar and Mukherjee [14] for more details. We also note that this extension of the domain to $\mathscr{M}$ implies that $A(\mu)$ can take any value in $\mathbf{R}$. Therefore, wherever necessary, we will also implicitly extend the domain of the demand function β from $[\underline{x}, \bar{x}]$ to $\mathbf{R}$.

potential function is concave.[7] This significance of concavity is explained later. The proof of the result is in the Appendix.

Proposition 3.3 *The large population model of Cournot competition F defined by (3) is a potential game with potential function $f : \mathcal{M} \to \mathbf{R}$ defined by*

$$f(\mu) = \int_{\underline{x}}^{A(\mu)} \beta(z)dz - \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} c_p(x)\mu_p(dx). \tag{4}$$

Further, f is concave but not strictly concave on Δ.

The fact that F is a potential game with a concave potential function is important for two reasons. First, Nash equilibria of large population potential games is the set of local maximizers and minimizers of the potential function (Sandholm [24], Cheung and Lahkar [5]). In fact, if the potential function is concave, then the set of Nash equilibria is the convex set of global maximizers of the potential function. Thus, concavity of the potential function allows for a convenient characterization of the set of Nash equilibria of a potential game. Indeed, if the potential function is strictly concave, that immediately implies that F has a unique Nash equilibrium. In our case, we don't have a strict concavity. Nevertheless, as we show below, we will still obtain a unique maximizer of the potential function (4) and, therefore, a unique Nash equilibrium in our model. Second, all standard evolutionary dynamics converge to Nash equilibria in potential games (Sandholm [24]). Thus, with a concave potential function, the convex set of Nash equilibria is globally attracting. We will elaborate this point further during our discussion on evolutionary implementation where it will be of immense significance.

Thus, the large population Cournot model F defined by (3) is both an aggregative game and a potential game. Hence, we refer to it as an *aggregative potential game*.[8]. Before calculating the Nash equilibrium of F, we provide an interpretation of the potential function (4) because it will be of relevance to our analysis.[9] For that, we define the aggregate payoff of firms at a state $\mu \in \Delta$ as

$$\begin{aligned}\bar{F}(\mu) &= \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} F_{x,p}(\mu)\mu_p(dx) \\ &= \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} \left(x\beta(A(\mu)) - c_p(x)\right)\mu_p(dx) \\ &= \beta(A(\mu) \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} x\mu_p(dx) - \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} c_p(x)\mu_p(dx) \\ &= A(\mu)\beta(A(\mu)) - \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} c_p(x)\mu_p(dx). \end{aligned} \tag{5}$$

[7]Cheung and Lahkar [5] establish a similar result for a large population Cournot model with a single type of firm. This result is a generalization because it allows for multiple types of firms.

[8]This terminology is from Lahkar [12] and Cheung and Lahkar [5] who analyzed such games with a single type

[9]A version of this interpretation restricted to a single type of firms also appear in Sandholm [27] and Cheung and Lahkar [5].

The aggregate payoff is a measure of the welfare of firms. This is simply the total revenue that firms obtain from consumers minus the total cost of production of firms. Notice from (4) and (5) that $f(\mu) - \bar{F}(\mu) = \int_{\underline{x}}^{A(\mu)} \beta(z)dz - A(\mu)\beta(A(\mu))$. But $\int_{\underline{x}}^{A(\mu)} \beta(z)dz$ is the area below the market demand function till output $A(\mu)$. Since we are subtracting the total revenue obtained by firms, which is also the sum paid by consumers, from this area, this particular difference is nothing but the consumer surplus at the state $\mu \in \Delta$, which we denote as $CS(\mu)$. Hence, we have

$$\begin{aligned} CS(\mu) &= f(\mu) - \bar{F}(\mu) \\ \Rightarrow f(\mu) &= CS(\mu) + \bar{F}(\mu). \end{aligned} \tag{6}$$

Thus, (6) implies that the potential function of a large population Cournot model is the sum of consumers' welfare, measured by consumer surplus, and producers' welfare, measured by the aggregate payoff of the model. We, therefore, interpret the potential function (4) as the total private surplus of all agents in the model. The qualifier private is added because this function does not account for the external cost $s(x)$ that each firm imposes on society. This total surplus is also different from the aggregate payoff of the model because we are not modeling consumers as active agents.

As noted earlier, we need to maximize this potential function to derive the Nash equilibrium of our model. This, however, is not straightforward because (4) is defined on an abstract measure space. To make the task tractable, we introduce another function $g : \prod_{p\in\mathcal{P}}[\underline{x}, \bar{x}] \to \mathbf{R}$ defined as

$$g(\alpha_1, \cdots, \alpha_n) = \int_{\underline{x}}^{\sum_{p\in\mathcal{P}} m_p\alpha_p} \beta(z)dz - \sum_{p\in\mathcal{P}} m_p c_p(\alpha_p). \tag{7}$$

The following lemma establishes the relationship between this function and potential function. Due to such a relationship between the two functions, we follow the terminology of Lahkar [12] and Cheung and Lahkar [5] and refer to (7) as the *quasi–potential function*.[10] The proof is in the Appendix.

Lemma 3.4 *Consider the potential function f defined by (4) and the quasi–potential function g defined by (7). Let μ be the current state of the game. Denote the aggregate strategy (output) of type p firms as $a(\mu_p) = \int_{\mathcal{S}} x\mu_p(dx)$ and define $\alpha_p = \frac{a(\mu_p)}{m_p}$ for all $p \in \mathcal{P}$. Then, the following holds.*

1. *Suppose μ_p is monomorphic for all p. Then, $f(\mu) = g(\alpha_1, \cdots, \alpha_n)$.*
2. *Suppose there exists at least one p such that μ_p is polymorphic. Then, $g(\alpha_1, \alpha_2, \cdots, \alpha_n) > f(\mu)$.*

The first part of this lemma shows that if all type states are monomorphic, then the quasi–potential function equals the potential function. The second part follows from the strict convexity

[10] Lahkar [12] introduced such a function for large population aggregative potential games with a single type and a finite strategy set. Cheung and Lahkar [5] extended that notion to games with continuous strategy sets. The present analysis is a further generalization because it allows for multiple types.

of the cost functions. In addition, we now show that the quasi–potential function is strictly concave on $\prod_{p\in\mathcal{P}}[\underline{x},\bar{x}]$. This will imply that this function has a unique maximizer. We state the result in the following lemma, this proof of which is in the Appendix.

Lemma 3.5 *The quasi–potential function $g:\prod_{p\in\mathcal{P}}[\underline{x},\bar{x}]\to\mathbf{R}$ defined in (7) is a strictly concave function. Hence, it has a unique maximizer $\alpha^*=(\alpha_1^*,\cdots,\alpha_n^*)\in\prod_{p=1}^{n}[\underline{x},\bar{x}]$.*

Using Lemmas 3.4 and 3.5, we can now establish the main result of this section. This result, stated in the following theorem, characterizes the unique Nash equilibrium of the large population Cournot competition model. Its proof is in the Appendix.

Theorem 3.6 *The large population model of Cournot competition F defined by (3) has a unique Nash equilibrium*

$$\mu^*=\left(m_1\delta_{\alpha_1^*},m_2\delta_{\alpha_2^*},\cdots,m_n\delta_{\alpha_n^*}\right), \tag{8}$$

where $\alpha^*=(\alpha_1^*,\cdots,\alpha_n^*)$ is the unique maximizer of the quasi–potential function g as characterized in Lemma 3.5.

Theorem 3.6, therefore, implies that at the unique Nash equilibrium, the state of a type p is $\mu_p^*=m_1\delta_{\alpha_1^*}$. Equivalently, every firm of type p produces output α_p^*, where $(\alpha_1^*,\alpha_2^*,\cdots,\alpha_n^*)$ is the maximizer of the quasi–potential function as characterized in Lemma 3.5. The reason why maximizing the quasi–potential suffices is that due to the strict convexity of the cost functions, the potential function can be maximized only if all type states are monomorphic (Lemma 3.4(2)). But at such states, the potential function equals the quasi–potential function (Lemma 3.4(1)).

The relationship between the potential function and the quasi–potential function, as highlighted in Lemma 3.4, and the fact that the quasi–potential function is defined on real numbers, simplified our task of maximizing the measure theoretic potential function considerably. The reason this technique works is because our Cournot competition model is an aggregative potential game. It is a generalization of the procedure used in Cheung and Lahkar [5] to multiple types of agents and is an important technical contribution of this paper. Lahkar [13] also characterized the Nash equilibrium of the large population Cournot model. But the methodology of that paper was a more direct one of computing best responses and finding the mutual best response. Here, we rely on the indirect method of potential games because this property will be essential for out discussion on evolutionary implementation in the next section.

We conclude this section with a brief discussion of the implications of Theorem 3.6. The Nash equilibrium μ^* maximizes the potential function. But recall our interpretation of the potential function as the total private surplus, i.e. the sum of consumer surplus and aggregate payoff of firms. Thus, the Nash equilibrium also maximizes the total private surplus. But this means that the Nash equilibrium is also the Walrasian equilibrium of the model.

An alternative way of seeing this equivalence is by directly maximizing the quasi–potential function (7). Suppose $(\alpha_1^*,\alpha_2^*,\cdots,\alpha_n^*)$ is an interior maximizer of g. Then, that maximizer is characterized by

$$\beta\left(\sum_q m_q \alpha_q^*\right) = c_p'(\alpha_p^*). \tag{9}$$

Thus the Nash equilibrium derived in Theorem 3.6 equates the market price to the marginal private cost for every type of firm, which is, of course, also the characteristic of the Walrasian equilibrium. Hence, the two equilibria coincide in our model. This is happening due to the fact that every firm is of measure zero, which effectively makes our Cournot model a model of a perfectly competitive market.

4 Social Welfare Maximization

We now bring in an evolutionary perspective into our model. We assume that the firms do not coordinate instantaneously on the Nash equilibrium, an assumption that is certainly plausible when there are a large number of agents. Instead, we assume that the game starts with an arbitrary initial state which then changes as agents revise their strategies according to certain behavioral norms called revision protocols.[11] Such revision protocols generate evolutionary dynamics which take the form of ordinary differential equations and which describe the change in the state of the game. Well known example of such dynamics in evolutionary game theory include the replicator dynamic (Oechssler and Riedel [19, 20], Cheung [4]), the BNN dynamic (Hofbauer et al. [9]), the pairwise comparison dynamic (Cheung [3]), the logit dynamic (Lahkar and Riedel [11], Perkins and Leslie [21]) and the best response dynamic for aggregative games (Lahkar and Mukherjee [14]).[12] Section 4 of Lahkar and Mukherjee [16] provides a description of these dynamics, which, to make the present article self–contained, we summarize in Appendix A.1.

We refer the reader to Section 4 of Lahkar and Mukherjee [16] for a description of these dynamics.

It is known that solution trajectories of all these dynamics converge to Nash equilibria globally in potential games (Sandholm [24], Cheung and Lahkar [5]). Indeed, the potential function of such games acts as a Lyapunov function along which solution trajectories towards a maximizer of the function, which is a Nash equilibrium. Therefore, in our model of Cournot competition, all these dynamics converge to the unique Nash equilibrium μ^* characterized in Theorem 3.6. Equivalently, all such dynamics converge to the maximum of the potential function which we have interpreted as the total private surplus. Hence, even without imposing the strong requirement of instantaneous coordination on equilibrium, we are able to maximize the total private surplus of consumers and producers in an evolutionary sense. Significantly, this happens under a variety of behavioral norms that generates the different evolutionary dynamics we have mentioned, which adds to the robustness

[11] See, for example, Sandholm [27], Cheung [3] and Lahkar and Riedel [11] for discussions of such strategy revision protocols

[12] The references given here study the continuous strategy version of these dynamics. See Sandholm [27] for detailed discussion on the finite strategy antecedents of these dynamics. The best response dynamic cannot be defined for all continuous strategy games due to the possibility that the best response may not exist. Hence, while considering this dynamic, we confine ourselves only to aggregative games where the best response exists and is always unique (Lahkar and Mukherjee [14]).

of our conclusion. This is the main reason why we have emphasised the potential game property of our model.

But although the Nash equilibrium μ^* maximizes total private surplus, it does not maximize the total social surplus. This is because of the presence of the external cost $s(x)$ in our model which, as we discussed in Section 2, is the external cost that a firm producing output x imposes upon society. We now introduce an extension of the potential function (4) and argue that this new function measures the total social surplus. We define this function $\hat{f} : \mathscr{M} \to \mathbf{R}$ as

$$\begin{aligned} \hat{f}(\mu) &= \int_{\underline{x}}^{A(\mu)} \beta(z)dz - \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} c_p(x)\mu_p(dx) - \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} s(x)\mu_p(dx) \\ &= f(\mu) - \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} s(x)\mu_p(dx) \\ &= \int_{\underline{x}}^{A(\mu)} \beta(z)dz - \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} \left(c_p(x) + s(x)\right)\mu_p(dx) \\ &= CS(\mu) + \bar{F}(\mu) - \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} s(x)\mu_p(dx). \end{aligned} \tag{10}$$

The last equality in (10) follows from (4) and (6). For reasons that we will explain later, we call $\hat{f}$ the *social potential function*. The difference between this function and the potential function f defined in (4) is the aggregate external cost $\sum_{p\in\mathcal{P}} \int_{\mathcal{S}} s(x)\mu_p(dx)$. The potential function measures the total private surplus (6) in our model. Once we subtract the aggregate social cost, we are left with (10) which must be the total social surplus.

Before proceeding further, it would be worthwhile to elaborate on the key distinction between (10) and the potential function used in earlier applications of evolutionary implementation in Sandholm [25, 26] and Lahkar and Mukherjee [14, 16]. As alluded to in the Introduction, the earlier applications were in non–market environments like network congestion, public goods and common resources. In such situations, it natural and meaningful to identify the welfare of the entire society with the welfare of the strategically active agents in the model. Hence, the social welfare function or the potential function that the planner sought to maximize in those models was the aggregate payoff function $\bar{F}(\mu)$ itself. Maximizing $\bar{F}$ then automatically implies taking into account the externalities that agents impose upon each other. In the present case, however, the active agents are the firms. Simply maximizing their aggregate payoff or, equivalently, total profits, would result in the monopoly output which cannot be a worthwhile objective for a planner. Instead, to make our problem substantive, we also need to introduce consumers into the model who are not strategic agents but whose welfare, measured using the consumer surplus, must be added to the payoff of firms to derive social welfare. Therefore, in our model, the potential function that measures social welfare needs to be more broadly defined by allowing for both strategic and non–strategic participants in the market. In addition, we need to subtract the aggregate external cost separately giving us the social welfare function (3).

The following proposition characterizes key properties of this function including its maximizer. Given our interpetation of $\hat{f}$ as the total social surplus, this maximizer will also be the social welfare maximizer. In order to establish the result, we follow a similar approach as the quasi–potential function and construct a finite dimensional analog of $\hat{f}$, which we denote $\hat{g}$ (see (11)). The difference between g and $\hat{g}$, as between f and $\hat{f}$, arises due to the inclusion of the social cost $s(x)$. But due to our assumption that s is convex, $c_p(x) + s(x)$ is strictly convex. Hence, the result essentially follows from Proposition 3.3, Lemmas 3.4 and 3.5 and Theorem 3.6. The formal arguments are in the Appendix.

Proposition 4.1 *The social potential function $\hat{f}$ defined in (10) is concave in Δ. Further, it has a unique maximizer $\mu^{**} \in \Delta$. The maximizer takes the for $\mu^{**} = \left(m_1\delta_{\alpha_1^{**}}, m_2\delta_{\alpha_2^{**}}, \cdots, m_n\delta_{\alpha_n^{**}}\right)$, where $(\alpha_1^{**}, \alpha_2^{**}, \cdots, \alpha_n^{**}) \in \prod_{p=1}^n [\underline{x}, \bar{x}]$ is the unique maximizer of the strictly concave function $\hat{g} : \prod_{p=1}^n [\underline{x}, \bar{x}] \to \mathbf{R}$ defined as*

$$\hat{g}(\alpha_1, \alpha_2, \cdots, \alpha_n) = \int_{\underline{x}}^{\sum_{p \in \mathcal{P}} m_p \alpha_p} \beta(z) dz - \sum_{p \in \mathcal{P}} m_p \left(c_p(\alpha_p) + s(\alpha_p)\right). \tag{11}$$

Thus, μ^{**} is the social welfare maximizer in our model of Cournot competition. Further, for every p, $\alpha_p^{**} \leq \alpha_p^*$, where α_p^* is as characterized in Theorem 3.6, with the inequality being strict if $\alpha_p^* \in (\underline{x}, \bar{x})$.

Alternatively, we can interpret this result in a manner analogous to (9). If we assume an interior maximizer, and if $(\alpha_1^{**}, \alpha_2^{**}, \cdots, \alpha_n^{**})$ is the maximizer of (11), then this maximizer is characterized by

$$\beta\left(\sum_q m_q \alpha_q^{**}\right) = c_p'(\alpha_p^{**}) + s'(\alpha_p^{**}) \tag{12}$$

Therefore, the social welfare maximizer μ^{**} equates the market price to the marginal social cost of firms. In contrast, as can be seen in (9), the Nash equilibrium (or the Walrasian equilibrium) equates the market price to the marginal private cost. Once the firms internalize the external cost they impose upon society, the output they produce at the social optimum, α_p^{**}, is less than their Nash equilibrium output level, α_p^*. Hence, aggregate output, as well as the associated level of pollution by each firm, is also less at the social optimum.

4.1 Evolutionary Implementation

Proposition 4.1 characterizes the social welfare maximizer μ^{**} as the maximizer of the social potential function (10). This raises the question of how to implement this social welfare maximizer. Notice that the social welfare maximizer is different from the Nash equilibrium since the Nash equilibrium is the maximizer of a different function, the potential function (Theorem 3.6). Therefore, the natural outcome of our model will not be the social welfare maximizer. This is, in fact, a very robust conclusion in our evolutionary game theoretic context because, as we have noted earlier, all standard evolutionary dynamics converge to the Nash equilibrium.

The problem of implementing the social welfare maximizer is, therefore, a substantial one. For this purpose we introduce a planner who, given any type distribution $m = (m_1, m_2, \cdots, m_n)$ of firms, wishes to implement μ^{**}. In the formal language of mechanism design theory, the mapping $m \mapsto \mu^{**}$ is the *social choice function* the planner wishes to implement. The method of implementation we consider is evolutionary implementation (Sandholm [25, 26]).[13] Unlike conventional mechanisms like the VCG mechanism, this mechanism does not rely on truthful revelation of types to achieve instantaneous coordination on the social welfare maximizer. Instead, as we will see, the planner uses a transfer pricing policy based on current actions of agents and relies on evolutionary processes to gradually direct the behavior of agents towards the social welfare maximizer. In doing so, the planner will not require any knowledge or types or type distribution of the firms.

All we need to assume is that the planner can observe the current output and knows (or can estimate) the external cost function s. The external cost imposed by a firm which produces output x is $s(x)$. The evolutionary implementation mechanism then requires the planner to impose a tax $s(x)$ on any firm which produces output x. The resulting payoff of a type p firm which produces x at the state μ is

$$\hat{F}_{x,p}(\mu) = x\beta(A(\mu)) - c_p(x) - s(x). \tag{13}$$

Intuitively, the imposition of the tax forces a firm to internalize the external cost it imposes upon society. This modification of the original payoff function (3) generates a new population game, which we denote as $\hat{F}$ and call the *externality adjusted game*. The following result follows as a corollary to Proposition 3.3.

Corollary 4.2 *The externality adjusted game $\hat{F}$ defined by (13) is a potential game with potential function $\hat{f}$ defined by (10).*

Proof. Follows from Proposition 3.3 once we replace $c_p(x)$ with $c_p(x) + s(x)$. ■

The fact that $\hat{F}$ is a potential game with potential function $\hat{f}$ is the reason why we called $\hat{f}$ the social potential function. It is the potential function of a game that incorporates the complete social cost of production in our model. Corollary 4.2 is now crucial for our main result on evolutionary implementation. Standard evolutionary dynamics must converge to the unique maximizer of $\hat{f}$. This unique maximizer is not ony the unique Nash equilibrium of $\hat{F}$ but, as established in Proposition 4.1, is also the social welfare maximizer. This is the essence of evolutionary implementation, which we now state formally as follows.

Proposition 4.3 *The potential game $\hat{F}$ has a unique Nash equilibrium μ^{**}, which is the social welfare maximizer in our model. Solution trajectories of the BNN dynamic (27), the pairwise*

[13] See Lahkar and Mukherjee[14, 16] for extension of this methodology to other continuous strategy models like public goods, public bads and tragedy of the commons.

comparison dynamic (28) and the best response dynamic for aggregative games (30) converge globally to μ^{**} in $\hat{F}$. The replicator dynamic (26) converges to μ^{**} from all initial states in the interior of Δ and is, therefore, almost globally converging. The logit dynamic (29) converges globally to a logit equilibrium which, for perturbation parameter η small, approximates the Nash equilibrium μ^{**}.

Proof. By Corollary 4.2, $\hat{F}$ is a potential game with potential function $\hat{f}$. The conclusion that $\hat{F}$ has a unique Nash equilibrium then follows from the fact that $\hat{f}$ is concave in Δ (Proposition 4.1). Hence, every Nash equilibrium of $\hat{F}$ must be a maximizer of $\hat{f}$. But as shown in Proposition 4.1, $\hat{f}$ has a unique maximizer, μ^{**}. Convergence to μ^{**} under various evolutionary dynamics then follows from results about these dynamics established in the papers cited earlier at the beginning of this section. ■

Proposition 4.3 is the main result of our paper. It establishes that the well known dynamics of evolutionary game theory all converge to the unique Nash equilibrium of the externality adjusted game $\hat{F}$ defined by (13) and, therefore, to the social welfare maximizing state of our original model F. This is evolutionary implementation. Convergence is with reference to the weak topology or the topology generated by convergence in distribution. While the BNN, the pairwise convergence and the best response dynamics converge globally, certain caveats are required for the replicator and logit dynamics. Under the replicator dynamic, convergence is assured only from the interior of the state space as strategies which are not initially present in the population cannot reappear under this dynamic. Thus, convergence is almost global. The logit dynamic does converge globally, but to an approximation of the Nash equilibrium, which is called a logit equilibrium. The approximation becomes increasingly precise as the perturbation parameter $\eta \to 0$ (Lahkar and Riedel [11]).

The planner, therefore, succeeds in evolutionarily implementing the socially welfare maximizing state from every, or almost every, initial state. At the social optimum, Proposition 4.1 implies that every firm produces a lower output and, therefore, a lower level of the negative externality than at the Nash equilibrium. It is true that evolutionary implementation is gradual in the sense that firms do not immediately coordinate on the socially efficient state but instead evolve towards it. But it is robust because it happens under a wide variety of evolutionary dynamics or, equivalently, a diversity of revision protocols that generate these dynamics. Moreover, the informational requirements for such implementation are minimal. Recall our claim that the planner does not have to know the type or even the type distribution of firms in this evolutionary implementation mechanism. The validity of that claim is evident now. All the planner requires is to observe the current output x of a firm and then, based on the knowledge of the external cost function, impose the tax $s(x)$. We should note, however, that evolutionary implementation requires the planner to continuously update the tax level $s(x)$. As the firms change output with the evolution of the state variable, the planner also needs to observe and modify the tax being paid by each firm.

In the next section, we compare the informational requirements of our evolutionary implementation approach with that of more classical approaches like Pigouvian pricing and VCG mechanism. We conclude this section by noting the importance of the assumption that the external cost function s is common for all types of firms. It is the feature that enables the planner to impose the tax $s(x)$ without knowing the type of a firm. Had this function been type specific, then evolutionary implementation would not have been feasible without a knowledge of the type of individual firms.

5 Pigouvian Tax and Dominant Strategy Implementation

Evolutionary implementation is not the only way for the planner to achieve the social optimum. Two more classical approaches are Pigouvian taxation and dominant strategy implementation. However, as we argue now, the information that the planner would require to apply Pigouvian taxation or dominant strategy implementation would be significantly more than for evolutionary implementation. In the absence of such information, evolutionary implementation turns out to the more feasible method of implementation.

First, we consider Pigouvian taxation, which is calculated with respect to the social optimum. Notice from Proposition 4.1 that the socially optimal level of output is

$$\alpha^{**} = \sum_p m_p \alpha_p^{**}. \tag{14}$$

We now assume that the planner has sufficient information to calculate this socially optimal output. This would require the planner to know the demand function β, the private cost functions c_p and the external cost function s. In addition, the planner would also need to know the type distribution m although not the type of individual firms. With such information, the planner can construct the function $\hat{g}$ defined by (11) and calculate (14). The planner then announces the market price $\beta(\alpha^{**})$ and allows firms to produce whatever output x they want at that price provided they pay a tax $s(x)$. The firm's problem then reduces to a standard individual maximization exercise where they choose $x \in \mathcal{S}$ so as to maximize

$$x\beta(\alpha^{**}) - c_p(x) - s(x). \tag{15}$$

By (12), this is uniquely maximized at α_p^{**}. With every type p firm choosing α_p^{**}, realized state is the social welfare maximizer μ^{**} as characterized in Proposition 4.1.

Therefore, the planner succeeds in implementing the social welfare maximizer. In this process, the Pigouvian tax on a type p firm is $s(\alpha_p^{**})$, which the planner imposes not directly but by setting the market price at the socially optimal level and allowing the firm to profit maximize. The assumption that the external cost $s(x)$ is the same for all firms allows the planner to implement efficiency without knowing the type of an individual firm. This was also the reason why evolutionary implementation was feasible in our approach. But unlike Pigouvian taxation, the planner could apply evolutionary implementation without knowing the demand function, private cost functions, and the type distribution. Thus, Pigouvian taxation has a greater informational requirement than evolutionary implementation. But unlike evolutionary implementation, it leads to instantaneous coordination on the social optimum.

5.1 Dominant Strategy Implementation

We continue to assume that the planner knows the demand function, private cost functions and the external cost function. But not the type distribution. In this situation, the planner cannot directly calculate the socially optimal output (14) and, therefore, cannot apply Pigouvian taxation as described above. However, even then, as we now show, the planner can implement the social optimum by applying a mechanism designed which is an extension of the classical VCG mechanism to large population games (Lahkar and Mukherjee [15]). We will show that under this mechanism, even if the planner implements Pigouvian taxation on the basis of reported types, every firm will find it strictly dominant to report their true type. Thus, like the classical VCG mechanism, the present approach relies on truthful revelation. This is the key difference with evolutionary implementation where there was no attempt by the planner to induce firms to reveal their type.

Formally, $m \mapsto \mu^{**}$ continues to be the planner's social choice function. To implement this function, the planner invites reports from firms about their type or, equivalently, about their private cost function. Let $\tilde{m} = (\tilde{m}_1, \tilde{m}_2, \cdots, \tilde{m}_n)$ be the reported type distribution. Thus, $\tilde{m}_p$ is the proportion of firms who report their type to be p. It is possible that $\tilde{m}_p \neq m_p$, the true mass of type p firms. The planner now applies the reported type distribution to the function $\hat{g}$ defined in (11) and maximizes that function. This exercise is feasible given our assumption that the planner knows the demand and cost function. Let $(\tilde{\alpha}_1^{**}, \tilde{\alpha}_2^{**}, \cdots, \tilde{\alpha}_n^{**})$ be that unique maximizer. We can then uniquely characterize $\tilde{\alpha}_p^{**}$ as

$$\tilde{\alpha}_p^{**} = \begin{cases} \underline{x} \text{ if } \beta\left(\sum_q \tilde{m}_q \tilde{\alpha}_q^{**}\right) < c_p'(x) + s'(x) \text{ for all } x \in \mathcal{S}. \\ x^{**} \text{ if } \exists x^{**} \in \mathcal{S} \text{ such that } \beta\left(\sum_q \tilde{m}_q \tilde{\alpha}_q^{**}\right) = c_p'(x^{**}) + s'(x^{**}). \\ \bar{x} \text{ if } \beta\left(\sum_q \tilde{m}_q \tilde{\alpha}_q^{**}\right) > c_p'(x) + s'(x) \text{ for all } x \in \mathcal{S}. \end{cases} \tag{16}$$

Define $\tilde{\alpha}^{**} = \sum_p \tilde{m}_p \tilde{\alpha}_p^{**}$. This is the aggregate output that would result when every firm that reports type to be p produces output $\tilde{\alpha}_p^{**}$. The planner now announces the market price to be $\beta(\tilde{\alpha}^{**})$, assigns output $\tilde{\alpha}_p^{**}$ to any firm which reports type to be p and charges tax $s(\tilde{\alpha}_p^{**})$ from that firm. In the conventional terminology of mechanism design theory, the planner constructs a direct mechanism $\phi : (q, \tilde{m}) \to \left(\tilde{\alpha}_q^{**}, \beta(\tilde{\alpha}^{**}), s(\tilde{\alpha}_q^{**})\right)$. This mechanism takes the reported type q of a firm and the reported type distribution $\tilde{m}$ and assigns the output $\tilde{\alpha}_q^{**}$, price $\beta(\tilde{\alpha}^{**})$ and tax $s(\tilde{\alpha}_q^{**})$ to that firm. The payoff to a type p firm in this mechanism ϕ which reports type to be q when the reported type distribution is $\tilde{m}$ is then

$$\phi_p(q; \tilde{m}) = \tilde{\alpha}_q^{**} \beta(\tilde{\alpha}^{**}) - c_p(\tilde{\alpha}_q^{**}) - s(\tilde{\alpha}_q^{**}). \tag{17}$$

Our objective is to show that truthful revelation is incentive compatible in ϕ. Before doing so, we note an important feature of ϕ arising from the fact that each agent is of measure zero. It is that reports by a single firm, or even a finite number of firms, cannot affect aggregate variables. In particular, the reported type distribution $\tilde{m}$ and the aggregate output $\tilde{\alpha}^{**}$ will not be affected by such a finite number of reports. This will have a major bearing on our main result of this section which we now state.

Proposition 5.1 *For any type distribution m, the direct mechanism ϕ defined by (17) implements the social welfare maximizing state μ^{**} characterized in Proposition 4.1 in strictly dominant strategies.*

Proof. We need to show that for any reported type distribution $\tilde{m}$, the unique payoff maximizing strategy of a type p firm in ϕ is p itself. This would imply that truthful revelation is strictly dominant in ϕ. Recall that due to each firm being of measure zero, individual announcements will not have any impact on $\tilde{m}$ and $\tilde{\alpha}^{**}$. Hence, the problem for a type p firm trying to decide its reported type can be equivalently stated as

$$\max_{x \in \mathcal{S}} x\beta(\tilde{\alpha}^{**}) - c_p(x) - s(x). \tag{18}$$

But (16) implies $\tilde{\alpha}_p^{**}$ is the unique solution to (18). Thus, given any reported type distribution $\tilde{m}$, a type p firm finds reporting type to be p and, hence, getting allotted output $\tilde{\alpha}_p^{**}$, the unique best response. Therefore, truthful revelation is strictly dominant. With every firm truthfully reporting type, a type p firm ends up producing α_p^{**}. This implements the social welfare maximizing μ^{**}. ■

The argument in the proof of Proposition 5.1 is akin to that of the classical VCG mechanism. But as noted in Lahkar and Mukherjee [15] in the context of a public goods game, the large population extension generates a stronger result. In the classical finite player mechanism, truthful revelation is weakly dominant whereas in the large population context, truthful revelation emerges as strictly dominant. The intuition behind this result is Pigouvian taxation but with respect to the reported type distribution. No agent can individually affect the reported type distribution or the market price that is a function of aggregate output. Hence, truthful revelation becomes strictly dominant.

Thus, dominant strategy implementation is a variant of the pure Pigouvian taxation approach described earlier but without requiring the planner to know the true type distribution. Like Pigouvian taxation, it also enables instantaneous implementation of efficiency. But the other informational requirements of Pigouvian taxation–knowledge of the demand, private cost and external cost functions–remain unchanged. Hence, dominant strategy implementation also calls upon the planner to have more information than evolutionary implementation, where only knowledge of the external cost function was required. But one advantage of dominant strategy implementation is that it requires the planner's intervention just once unlike evolutionary implementation when the tax level needs to continuously updated.

Recall the importance of the assumption that the external cost function is common for all firms for evolutionary implementation. The same holds true for dominant strategy implementation as well. If the external cost function is different, then the firm may be able to improve its payoff by falsely claiming to be a type with a lower level of external cost and, thereby, reducing its tax burden. False reporting doesn't work in (18) because only private cost functions are type specific and firms cannot avoid their true private cost of production by claiming to be of a different type. The external cost function is qualitatively different because it is not a cost of production but a tax which can be manipulated through strategic reporting.

6 Conclusion

In this paper, we consider a model of evolutionary implementation in a large population game where firms' production activities lead to negative externalities. Our work has useful implications for policy makers operating in environments of regulating negative externalities. Our primary contribution is in proving that a policy maker can implement the socially optimal level of output by using a very simple mechanism ; the policy maker observes the current level of output and sets a transfer price or tax equal to resulting external cost. This generates a potential game with a potential function whose maximizer is the social optimum. Evolutionary dynamics converge to this maximizer resulting in implementation of the social welfare maximizer. Significantly, knowledge of current output levels of firms and the external cost function are all that the planner requires to induce evolution of the society to the social optimum. Unlike the classical Pigouvian tax, there is no need for the planner to have prior knowledge of the socially optimal state. Nor does the planner need to know the type distribution of firms as in classical VCG mechanisms.

We have established convergence to the social optimum under a variety of evolutionary dynamics (Proposition 4.3). Implicitly though, that result has assumed that the behavior of all firms in the market is described by the same evolutionary dynamic. But what if different firms operate under different revision protocols generating different dynamics. Will the result still hold? The answer is yes. This is because the precise structure of the dynamic is not important for convergence in potential games. All that is required are two conditions: (i) Nash equilibria should be rest points of the dynamic (Nash stationarity) and (ii) there should be a positive relationship between the payoff of a strategy and the rate of change of its mass (positive correlation). The underlying dynamic then converges to a Nash equilibrium in a potential game or, equivalently, to a maximizer of the potential function (Sandholm [24]). It is well known that the five canonical dynamics we have considered in Proposition 4.3 satisfy these conditions.[14] Moreover, so does a convex combination of dynamics that satisfy Nash Stationarity and positive correlation (Section 5.7, Sandholm [27]). Therefore, even if different firms operate under different dynamics, the resulting convex combination of these dynamics would still induce evolutionary implementation.

Evolutionary implementation doesn't require the planner to know market demand and private cost functions. But it still requires knowledge of the external cost function. How realistic is the assumption that the planner does know the external cost function? Our argument is that even if the exact function is not known, the planner can still estimate it using well established econometric techniques to measure current levels of externalities (Lin [17]). Taxes based on the estimated

[14] The logit dynamic satisfies these conditions up to an approximation. Hence, we obtain convergence to an approximate Nash equilibrium–the logit equilibrium (Lahkar and Riedel [11]).

external cost function would still generate a potential game with a potential function that would measure not exact social welfare but an estimated level of welfare. Evolutionary dynamics would still converge to the maximizer of the estimated social welfare function implementing, even if not the exact social optimum, but a fairly good approximation of it. Hence, our evolutionary methodology is flexible enough to be effective even if the planner does know the external cost function precisely.

A Appendix

Proof of Proposition 3.3: See Proposition 5.3 in Lahkar [13] for the proof that F defined by (3) is a potential game with potential function 4. For the proof of concavity of the potential function, let $\mu, \nu \in \Delta$, $\mu \neq \nu$, be two states. Fix $\lambda = (0, 1)$. Clearly,

$$\sum_{p\in\mathcal{P}} \int_{\mathcal{S}} c_p(x)(\lambda\mu_p + (1-\lambda)\nu_p)(dx) = \lambda \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} c_p(x)\mu_p(dx) + (1-\lambda) \sum_{p\in\mathcal{P}} \int_{\mathcal{S}} c_p(x)\nu_p(dx). \tag{19}$$

Recall that β is strictly decreasing on $[\underline{x}, \bar{x}]$. Hence, $\int_{\underline{x}}^{\alpha} \beta(z)dz$ is strictly concave on $[\underline{x}, \bar{x}]$. Further, $A(\cdot)$ is linear and since $\mu, \nu \in \Delta$, $A(\lambda\mu + (1-\lambda)\nu) \in \Delta$. Combining these facts, we obtain

$$\begin{aligned}\int_{\underline{x}}^{A(\lambda\mu+(1-\lambda)\nu)} \beta(z)dz &= \int_{\underline{x}}^{\lambda A(\mu)+(1-\lambda)A(\nu)} \beta(z)dz \\ &\geq \lambda \int_{\underline{x}}^{A(\mu)} \beta(z)dz + (1-\lambda) \int_{\underline{x}}^{A(\nu)} \beta(z)dz, \end{aligned} \tag{20}$$

with $\geq$ holding with equality only if μ, ν such that $A(\mu) = A(\nu)$. Combining (19) and (20), we obtain

$$f(\lambda\mu + (1-\lambda)\nu) \geq \lambda f(\mu) + (1-\lambda) f(\nu).$$

Thus, the potential function f defined by (4) is concave but not strictly concave. Strict concavity fails because even if $\mu \neq \nu$, it is possible that $A(\mu) = A(\nu)$. ■

Proof of Lemma 3.4: Notice that the aggregate strategy for all firms is $A(\mu) = \sum_p a(\mu_p) = \sum_p m_p\alpha_p$. Therefore, any difference between f defined by (4) and g defined by (7) can only arise due to the difference between $\sum_{p\in\mathcal{P}} \int_{\mathcal{S}} c_p(x)\mu_p(dx)$ and $\sum_p m_p\alpha_p$.

1. If μ_p is monomorphic for all p, then $\alpha_p = \frac{a(\mu_p)}{m_p}$ implies all firms of type p are playing α_p. Therefore, $\sum_{p\in\mathcal{P}} \int_{\mathcal{S}} c_p(x)\mu_p(dx) = \sum_p m_p\alpha_p$. Hence, $f(\mu) = g(\alpha_1, \alpha_2, \cdots, \alpha_n)$.

2. Suppose μ_q is polymorphic. By definition, $\alpha_p = \frac{a(\mu_p)}{m_p}$ implies $\int_{\mathcal{S}} x\frac{\mu_q}{m_q}(dx) = \alpha_q$. Recall c_q is strictly convex. Hence,

$$\begin{aligned}&\int_{\mathcal{S}} c_q(x)\frac{\mu_q}{m_q}(dx) > c_q\left(\int_{\mathcal{S}} x\frac{\mu_q}{m_q}(dx)\right) = c_q(\alpha_q) \\ \Rightarrow &\int_{\mathcal{S}} c_q(x)\mu_q(dx) > m_q c_q(\alpha_q).\end{aligned}$$

Thus, if μ_q is polymorphic, $\int_{\mathcal{S}} c_q(x)\mu_q(dx) > m_q c_q(\alpha_q)$ while part 1 implies that if μ_p is monomorphic, then $\int_{\mathcal{S}} c_p(x)\mu_p(dx) > m_p c_p(\alpha_p)$. Thus, if the state of at least one type is polymorphic, then $\sum_{p\in\mathcal{P}} \int_{\mathcal{S}} c_p(x)\mu_p(dx) > \sum_p m_p\alpha_p$. Hence, $g(\alpha_1, \alpha_2, \cdots, \alpha_n) > f(\mu)$. ■

Proof of Lemma 3.5: Consider the quasi–potential function (7). Take two points $(\hat{\alpha}_1, \cdots, \hat{\alpha}_n) \neq (\tilde{\alpha}_1, \cdots, \tilde{\alpha}_n) \in [\underline{x}, \bar{x}]$. We need to show that

$$\begin{aligned}&\int_{\underline{x}}^{\sum_q m_q(\lambda\hat{\alpha}_q+(1-\lambda)\tilde{\alpha}_q)} \beta(z)dz - \sum_{q\in\mathcal{P}} m_q c_q(\lambda\hat{\alpha}_q + (1-\lambda)\tilde{\alpha}_q)\\ >&\lambda\left(\int_{\underline{x}}^{\sum_q m_q\hat{\alpha}_q} \beta(z)dz - \sum_{q\in\mathcal{P}} m_q c_q(\hat{\alpha}_q)\right) + (1-\lambda)\left(\int_{\underline{x}}^{\sum_q m_q\tilde{\alpha}_q} \beta(z)dz - \sum_{q\in\mathcal{P}} m_q c_q(\tilde{\alpha}_q)\right). \qquad (21)\end{aligned}$$

Since β is a strictly decreasing function, $\int_{\underline{x}}^{y} \beta(z)dz$ is a strictly concave function, for $y \in \mathbf{R}$. Hence,

$$\int_{\underline{x}}^{\sum_q m_q(\lambda\hat{\alpha}_q+(1-\lambda)\tilde{\alpha}_q)} \beta(z)dz \geq \lambda \int_{\underline{x}}^{\sum_q m_q\hat{\alpha}_q} \beta(z)dz + (1-\lambda)\int_{\underline{x}}^{\sum_q m_q\tilde{\alpha}_q} \beta(z)dz, \qquad (22)$$

with equality holding only if $\sum_q m_q\hat{\alpha}_q = \sum_q m_q\tilde{\alpha}_q$. Further, the strict convexity of c_q implies

$$c_q(\lambda\hat{\alpha}_q + (1-\lambda)\tilde{\alpha}_q) < \lambda c_q(\hat{\alpha}_q) + (1-\lambda)c_q(\tilde{\alpha}_q). \qquad (23)$$

Since (23) holds for all $q \in \mathcal{P}$, it together with (22) implies (21). ■

Proof of Theorem 3.6: Consider the potential function f defined by (4) and the quasi–potential function g defined by (7). Recall the notation from Lemma 3.4 that $a(\mu_p) = \int_{\mathcal{S}} x\mu_p(dx)$. Note that if $\mu^* = \left(m_1\delta_{\alpha_1^*}, m_2\delta_{\alpha_2^*}, \cdots, m_n\delta_{\alpha_n^*}\right)$ is as described in the theorem, then $\frac{a(\mu_p^*)}{m_p} = \alpha_p^*$ for all $p \in \mathcal{P}$. Therefore, by part 1 of Lemma 3.4, $f(\mu^*) = g(\alpha_1^*, \cdots, \alpha_n^*)$.

Now consider $\mu \neq \mu^*$ such that μ_p is monomorphic for every p. Hence, $\mu = (m_1\delta_{\alpha_1}, \cdots, m_n\delta_{\alpha_n})$ for some $(\alpha_1, \cdots, \alpha_n) \in \prod_{p=1}^{n}[\underline{x}, \bar{x}]$, with $\alpha_p \neq \alpha_p^*$ for at least one $p \in \mathcal{P}$. Note that in this case, $\frac{a(\mu_p)}{m_p} = \alpha_p$, $\alpha_p \in [\underline{x}, \bar{x}]$. Therefore, by Lemma 3.4(1), $f(\mu) = g(\alpha_1, \cdots, \alpha_n)$. But then, by applying Lemma 3.5, we obtain

$$f(\mu^*) = g(\alpha_1^*, \cdots, \alpha_n^*) > g(\alpha_1, \cdots, \alpha_n) = f(\mu). \qquad (24)$$

Next, consider μ such that μ_p is polymorphic for at least one $p \in \mathcal{P}$. For any $p \in \mathcal{P}$, define $\alpha_p \in [\underline{x}, \bar{x}]$ such that $\frac{a(\mu_p)}{m_p} = \alpha_p$. Then,

$$f(\mu^*) = g(\alpha_1^*, \cdots, \alpha_n^*) \geq g(\alpha_1, \cdots, \alpha_n) > f(\mu), \qquad (25)$$

where the weak inequality holds if $(\alpha_1^*, \cdots, \alpha_n^*) = (\alpha_1, \cdots, \alpha_n)$ and the strict inequality follows from Lemma 3.4(2).

Combining (24) and (25), we conclude that μ^* is the unique global maximizer of the potential function f. Since f is concave (by Proposition 3.3), this implies μ^* is the unique Nash equilibrium of F. ■

Proof of Proposition 4.1: Recall that $c_p(x)$ is strictly convex for all p while $s(x)$ is convex. Therefore, $(c_p(\alpha_p) + s(\alpha_p))$ is also strictly convex. Hene, the concavity of $\hat{f}$ follows from similar argument as in Proposition 3.3 once we replace $c_p(x)$ in (4) with $(c_p(\alpha_p) + s(\alpha_p))$. To establish the existence of a unique maximizer μ^{**}, we construct (11) and note that it is analogous to the quasi–potential function (7) except for the presence of the $\sum_p m_p s(\alpha_p)$ term. But since $(c_p(\alpha_p) + s(\alpha_p))$, $\hat{g}$ is strictly convex by the same argument as in Lemma 3.5. Hence, $\hat{g}$ has a unique maximizer $(\alpha_1^{**}, \alpha_2^{**}, \cdots, \alpha_n^{**})$.

The strict convexity of $(c_p(\alpha_p) + s(\alpha_p))$ also implies that $\hat{f}$ and $\hat{g}$ share the same relationship as f and g as characterized in Lemma 3.4. The identification of the social welfare maximizer μ^{**} such that each type state $\mu_p^{**} = m_p \delta_{\alpha_p^{**}}$ then follows from Theorem 3.6. The conclusion that $\alpha_p^{**} \leq \alpha_p^*$ follows from the maximization of (7) and (11) and the fact that for every $x \in \mathcal{S}$, $c_p'(x) + s'(x) > c_p'(x)$. ■

A.1 Standard Evolutionary Dynamics

To briefly describe the standard evolutionary dynamics, we consider a population game F with strategy set $\mathcal{S}$ and let the payoff of a population p agent playing strategy x at social state μ be $F_{x,p}(\mu)$. Let $\bar{F}_p(\mu) = \frac{1}{m_p} \int_{\mathcal{S}} F_{x,p}(\mu) \mu_p(dx)$ be the average payoff in population p at μ. The excess payoff of a strategy x in population p at μ is then $F_{x,p}(\mu) - \bar{F}_p(\mu)$. We require this notion of the excess payoff to define the replicator dynamic and the BNN dynamic.

For the logit dynamic, we need to define the probability measure $L_{\eta,p}(\mu)$ on $\mathcal{S}$, known as the logit choice measure, as $L_{\eta,p}(\mu)(B) = \int_B \frac{\exp(\eta^{-1} F_{x,p}(\mu))}{\int_{\mathcal{S}} \exp(\eta^{-1} F_{y,p}(\mu)) dy} dx$, $B \subseteq \mathcal{S}$, $\eta > 0$. The parameter η is a measure of perturbation. The logit choice is is generated when agents best respond to a perturbed version of payoffs, where the perturbation depends upon η (Lahkar and Riedel [11]). For η small, it puts most of the probability mass on the set of best responses to μ. Hence, intuitively, the logit choice measure is an approximation of the best response.

The replicator dynamic, the BNN dynamic, the pairwise comparison dynamic and the logit dynamic respectively in F are now defined as follows.

$$\dot{\mu}_p(B) = \int_B \left(F_{x,p}(\mu) - \bar{F}_p(\mu) \right) \mu_p(dx), \tag{26}$$

$$\dot{\mu}_p(B) = m_p \int_B \left[F_{x,p}(\mu) - \bar{F}_p(\mu) \right]_+ dx - \mu_p(B) \int_{\mathcal{S}} \left[F_{y,p}(\mu) - \bar{F}_p(\mu) \right]_+ dy, \tag{27}$$

$$\dot{\mu}_p(B) = \int_{\mathcal{S}} \int_B \left[F_{x,p}(\mu) - F_{y,p}(\mu) \right]_+ dx \mu_p(dy) - \int_{\mathcal{S}} \int_B \left[F_{y,p}(\mu) - F_{x,p}(\mu) \right]_+ \mu_p(dx) dy, \tag{28}$$

$$\dot{\mu}_p(B) = m_p L_{\eta,p}(\mu)(B) - \mu_p(B), \text{ where } \eta > 0. \tag{29}$$

In each of these dynamics, $\dot{\mu}_p(B)$ is the direction and magnitude of change in the mass of agents in population p who are playing strategies in $B \subseteq \mathcal{S}$. The replicator dynamic (26) and the BNN dynamic (27) depends upon the excess payoff. If the aggregate excess payoff of strategies in B is positive, then the replicator dynamic increases the mass of agents playing strategies in B. Under the BNN dynamic, agents adopt strategy x with probability proportional to the positive part of the excess payoff $F_{x,p}(\mu) - \bar{F}_p(\mu)$ of that strategy (note that $[a-b]_+ = \max(a-b, 0)$). The pairwise comparison dynamic involves agents abandoning strategy y and adopting strategy x with probability proportional to $[F_{x,p}(\mu) - F_{y,p}(\mu)]_+$. The logit dynamic moves the social state μ towards the logit choice measure $L_{\eta,p}(\mu)$.

While the four dynamics (26)–(29) are well defined for all population games, the best response dynamic for aggregative games, as the name suggests, is valid only in such aggregative games where every social state μ generates a unique best response. Denote the aggregate strategy level $A(\mu)$ at μ in such a game as $A(\mu)$. Under the best response dynamic, the change in the population state μ_p is given by

$$\dot{\mu}_p = m_p \delta_{b_p(\alpha)} - \mu_p. \tag{30}$$

The reason why it may be difficult to define this dynamic in more general games is that the best response may not be uniquely defined or, due to the continuous structure of the strategy set, may not even exist.

References

[1] Baumol WJ (1972) On Taxation and the Control of Externalities. Amer. Econ. Rev. 62:307–322.

[2] Baumol WJ, Oates WE (1988) The Theory of Environmental Policy (2nd ed.). Cambridge University Press, NY, USA.

[3] Cheung MW (2014) Pairwise comparison dynamics for games with continuous strategy space. J. Econ. Theory 153:344–375.

[4] Cheung MW (2016) Imitative Dynamics for Games with Continuous Strategy Space. Games Econ. Behav. 99:206–223.

[5] Cheung MW, Lahkar R (2018) Nonatomic Potential Games: the Continuous Strategy Case. Games Econ. Behav. 108:341–362.

[6] Clarke E (1971) Multi-part pricing of public goods. Public Ch. 11:17–23.

[7] Corchón L (1994) Comparative statics for aggregative games the strong concavity case. Math. Soc. Sci. 28:151–165.

[8] Groves T (1973) Incentives in Teams. Econometrica. 41:617–631.

[9] Hofbauer J, Oechssler J, Riedel F (2009) Brown–von Neumann–Nash dynamics: the continuous strategy case. Games Econ. Behav. 65:406–429.

[10] Kim S-H (2015) Disclosure and inspection policies for green production. Oper. Res. 63(1):1–20.

[11] Lahkar R, Riedel F (2015) The logit dynamic for games with continuous strategy sets. Games Econ. Behav. 91:268–282.

[12] Lahkar R (2017) Large population aggregative potential games. Dyn Games Appl. 7:443–467.

[13] Lahkar R (2020) Convergence to Walrasian equilibrium with minimal information. J. Econ. Interact. Coord. 15:553–578.

[14] Lahkar R, Mukherjee S (2019) Evolutionary Implementation in a Public Goods Game. J. Econ. Theory 181:423–460.

[15] Lahkar R, Mukherjee S (2020) Dominant Strategy Implementation in a Large Population Public Goods Game. Econ. Lett. 197:109616.

[16] Lahkar R, Mukherjee S (2021) Evolutionary Implementation in Aggregative Games. Math. Soc. Sci. 109:137–151.

[17] Lin S (Ed.) (1976) Theory and measurement of economic externalities. Academic Press.

[18] Monderer D, Shapley L (1996) Potential games. Games Econ. Behav. 14:124–143.

[19] Oechssler J, Riedel F (2001) Evolutionary dynamics on infinite strategy spaces. Econ. Theory 17:141–162.

[20] Oechssler J, Riedel F (2002) On the dynamic foundation of evolutionary stability in continuous models. J. Econ. Theory 107:223–252.

[21] Perkins S, Leslie D (2014) Stochastic fictitious play with continuous action sets. J. Econ. Theory 152:179–213.

[22] Pigou AC (1920) The Economics of Welfare. Macmillan, London.

[23] Rothkopf MH (2007) Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. Oper. Res. 55:191–197.

[24] Sandholm WH (2001) Potential games with continuous player sets. J. Econ. Theory 97:81–108.

[25] Sandholm WH (2002) Evolutionary Implementation and Congestion Pricing. Rev. Econ. Stud 69:667–689.

[26] Sandholm WH (2005) Negative Externalities and Evolutionary Implementation. Rev. Econ. Stud 72:885–915.

[27] Sandholm WH (2010) Population Games and Evolutionary Dynamics. MIT Press, Cambridge, MA.

[28] Tietenberg TH (2006) Emissions Trading: Principles and Practise (2nd ed.) Resources for the Future. Washington, D.C.

[29] Tietenberg TH, Lewis L (2018) Environmental and Natural Resource Economics (11th ed.) Routledge, United Kingdom.

[30] Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. J. Finance. 16(1):8–37.

[31] Wang S, Sun P,de Vericourt F (2016) Inducing environmental disclosures: A dynamic mechanism design approach. Oper. Res. 64(2):371–389.