



King's Research Portal

Document Version
Other version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Hargreaves Heap, S., & Ismail, M. (in press). No-harm principle, rationality, and Pareto optimality in games. *SYNTHESE*.

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

No-harm principle, rationality, and Pareto optimality in games

Shaun Hargreaves Heap* Mehmet S. Ismail†

Sunday 28th November, 2021

Abstract

Mill's classic argument for liberty requires that people's exercise of freedom should be governed by a no-harm principle (NHP). In this paper, we develop the concept of a no-harm equilibrium in games where players maximize utility subject to the constraint of the NHP. Our main result is in the spirit of the fundamental theorems of welfare economics. We show that for every initial 'reference point' in a game the associated no-harm equilibrium is Pareto efficient and, conversely, every Pareto efficient point can be supported as a no-harm equilibrium for some initial reference point.

Keywords: Pareto optimality, rationality, classical liberalism, no-harm principle, non-cooperative games

*Department of Political Economy, King's College London, London, UK.
E-mail: s.hargreavesheap@kcl.ac.uk

†Department of Political Economy, King's College London, London, UK.
E-mail: mehmet.s.ismail@gmail.com

“It’s a pretty long-standing principle that goes back to the 19th century that you are free to do things but not if you inflict harm on others.” Nick Clegg (Facebook, Vice-President Global Affairs) explaining Facebook’s decision to ban Donald Trump

Nick Clegg reminds us that the sense of liberty which animates the liberal democracies of North America, Europe and elsewhere is not a free-for-all. At least since John Stuart Mill’s *On Liberty*, the writ of individual freedom in a society founded on the principle of liberty is not without constraint: people’s freedom does not extend to harming others. For Mill, the failure to satisfy this condition supplied the only reason for interfering in what people wish to do:

“...the only purpose for which power can be rightfully exercised over any member of a civilised community, against his will, is to prevent harm to others.” (Mill, 1859)

This may be the only condition for constraining individual action according to Mill, but, as Nick Clegg reminds us, it is, nevertheless, a constraint. In this paper, we examine how the introduction of such a no-harm principle (NHP) as a constraint on action in normal form games affects the analysis of rational play in these interactions. In particular, we develop the concept of a no-harm equilibrium where players maximize utility subject to the constraint of the NHP. Our main result is in the spirit of the fundamental theorems of welfare economics, where consumers maximize utility subject to their budget constraints. We show that for every initial ‘reference point’ in a game the associated no-harm equilibrium is Pareto efficient and, conversely, every Pareto efficient point can be supported as a no-harm equilibrium for some initial reference point.

This is a striking theoretical result because it is in marked contrast to the well known insight, for example in the conventional analysis of the prisoners’ dilemma, that the Nash equilibrium of normal form games need not be Pareto efficient. Our result is not only potentially important for this reason, it also has

policy implications. In the conventional wisdom, situations like a prisoners' dilemma provide a *prima facie* case for policy intervention to secure Pareto improvements. Our contrary result suggests something different. It is in the spirit of Coase's theorem in the sense that we agree that it is not the fact that there are, say, externalities in some settings that supplies the *prima facie* grounds for policy intervention. For Coase (1960), it is the presence of non-negligible transaction costs that supplies the grounds. In our case, it would be the failure, in practice, of agents to be guided by the NHP that establishes the *prima facie* grounds for intervention. The NHP might not be operational in practice for a variety of reasons. For example, people may subscribe to a free-for-all sense of liberty and not the NHP version; or the NHP may be too cognitively demanding to guide citizens in a liberal society. Alternatively, the NHP may be rendered inoperable because there are disagreements over the relevant 'reference point' or how to define 'harm'. It is these conditions, our result suggests, that should trigger a policy interest.

Of course, our contrary result may be theoretically arresting and point to a different policy agenda, but it is only really important in so far as the NHP is thought plausibly to be relevant to the analysis of action in normal form games. The premise of the NHP as a constraint on action, in other words, has to be plausible. Nick Clegg's appeal to the principle is an illustration of its possible relevance in this respect. More compelling, perhaps, is the fact that the business of the judiciary in contemporary liberal societies frequently involves deciding when one person's exercise of liberty may or may not be reasonably said to have conferred a 'harm' upon another. What counts as a 'harm' is, of course, naturally controversial and this is why the courts get involved. The point, however, about the courts' involvement in such matters is that it arises because they are upholding the NHP. In other words, the NHP, and not a free-for-all, is constitutive, notably through the activities of the courts, of what freedom means in liberal democratic societies. As a result, the NHP might plausibly be thought to apply to all who voluntarily live in such societies.

This constitutive role of NHP in the liberal democratic understanding of

freedom might be granted but it does not necessarily mean that game theory should adopt the principle when analysing what players should rationally do in games. It might be argued, for instance, that the judiciary identifies actions that satisfy this principle and so, when game theory's available actions are legal, all actions satisfy the principle and there is no need to apply a further NHP test in the analysis of what players will freely and rationally do. Game theory, though, does not typically require that actions available to players be legal. If game theory only considered legal actions, it would be unable to analyse when people follow the law (or more generally uphold a contract) or indeed when they engage in civil disobedience. Furthermore, what is or is not legal is only infrequently tested by courts with the result that the distinction is not apparent in many settings. People in liberal societies, therefore, may need to anticipate and apply the NHP themselves: either because they need to anticipate what the courts might say or because they simply value freedom and so are guided by its entailment in liberal society, the NHP. Nick Clegg, for example, was not executing the court's judgement when banning Donald Trump. He was either saying something like 'Facebook values freedom and this means being guided by the NHP'; or he was anticipating what a court might decide when applying this principle, as it does when determining what freedom requires in a liberal society.

For these reasons, the NHP appears relevant to the analysis of action in games. Our paper, therefore, proceeds, as follows. In the next section, we provide an informal introduction to the challenges posed by introducing the NHP and sketch how we respond to them. We then set out our approach formally in section 2; and section 3 gives the results associated with the inclusion of NHP in three theorems. We reflect on these results in section 4 with some further illustrations and then discuss their relation to the literature in section 5. Section 6 concludes the paper.

1 An informal sketch of the challenges and our approach to introducing NHP in a strategic context

The NHP does not permit a person to take an action that causes ‘harm’ to others. Several challenges arise when deciding how to represent this principle as a constraint on actions in games. In this section we offer an informal sketch of how we respond to these challenges by making four key assumptions.

Game theory has one advantage over the world that the judiciary addresses when trying to decide whether someone’s actions cause a ‘harm’ to another: game theory deals with interactions where the pay-offs to each player in each outcome are given. With players’ interests captured in this way by their pay-offs, it seems natural and uncontentious to say that a person suffers a ‘harm’ when their pay-offs are reduced. But, reduced relative to what? This is the first modelling challenge. There has to be a reference set of pay-offs for players of the game so that one can judge whether any outcome in the game produces a ‘harm’ to some players. Since games contain all the relevant available actions for players in the setting captured by the game, we assume that the reference pay-offs must be given by one of the outcomes in the game. We make no argument over which outcome should be used. Any outcome might serve as the reference. Instead, we seek to characterise in general terms the equilibria that result when players take any of the outcomes in the game as the shared reference point. The specific attributes of an equilibrium that satisfies the NHP may depend on the actual reference point, but we are interested whether there are any general properties of such equilibria.

The next challenge arises because the NHP requires that an individual’s action should not cause a harm to any other person and outcomes in normal form games typically result from the joint actions of several players. We need, therefore, some way of making sense of how an individual produces an outcome in a game and so judge whether that individual’s actions cause a harm. Our approach follows Steven Brams (1994). It takes one outcome in the game as

the reference point (one might think of this as the status quo) and builds an extensive form game on the basis of the possible sequential player deviations from this reference point. That is, the first player in this sequence must decide between ‘staying’ at the reference outcome and ‘moving’ to an alternative outcome by changing their action so as to produce that outcome; and each player thereafter decides in this extensive form game between ‘staying’ at the outcome they have inherited from the previous decisions of others in this branch of the extensive form game and ‘moving’ to an alternative one by changing their action. To fix ideas here consider the Prisoners dilemma game below.

	C	D
C	3, 3	1, 4
D	4, 1	2, 2

Suppose DD is the reference outcome and Row is the first player to consider a deviation. The extensive form game is constructed as follows. Row can either ‘stay’ at DD or ‘move’ to CD by changing their action to C. If Row chooses to ‘stay’ at DD, then, Column faces the same decision at the second node of this extensive form game as the one we have just considered for Row. If Row chooses to ‘move’ to CD, Column at the second decision node in the extensive form game must decide between ‘staying’ at CD and ‘moving’ to CC by changing their action from D to C; if columns stays at CD then Row decides between ‘staying’ at CD and ‘moving’ to DD; and so on. Figure 1 captures these early decision nodes in the extensive form game based on DD as the reference outcome.

The virtue for our purpose of adopting Brams’s Theory of Moves in this way is that each outcome in this extensive form game now occurs through an individual action at some decision node and this allows us to identify how an individual’s action can be said to cause a particular outcome. Thus, we can say in our illustration that, if Row begins by ‘moving’ to CD through a ‘move’ to C from the reference point of DD, then Row has caused CD at this point in the extensive form game.




Figure 1: Illustration of early decision nodes in the extensive form game based on DD as the reference outcome

Brams (1994) assumes farsighted rationality and solves by backward induction the extensive form game created by his procedure of sequential deviations. He calls the outcome of his proposed sequential procedure a non-myopic equilibrium (see also Brams and Wittman, 1981, and Kilgour, 1984). We adopt the same approach of applying the subgame perfect equilibrium solution concept but we introduce the NHP as a constraint on play in the extensive form game and so call the outcome a no-harm equilibrium. Thus, to return to the prisoners' dilemma with DD as the initial reference outcome, we ask players whether the sub-game perfect equilibrium of the NHP constrained extensive form game based on DD is DD or some other outcome. If it is DD, then DD is the no-harm equilibrium outcome of the game. If it is not, then the no-harm equilibrium with DD as the reference outcome is whatever the sub-game perfect equilibrium is in this NHP constrained extensive form game. The next challenge is, therefore to represent the NHP in this extensive form game. This requires three further assumptions.

One is innocuous in the sense that a well-defined finite extensive form game requires a set of terminal nodes. One part of how we do this is by saying that when both players decide to 'stay' at an outcome, then this outcome is implemented. The first decision to 'stay' is like a 'proposal' to implement this outcome and the second decision to 'stay' amounts to an acceptance of




Figure 2: An illustration of the extensive form game where DD is the reference outcome and $k = 1$.

this ‘proposal’. This naturally produces a terminal node for some branches of the extensive form game. But we also need to prevent the ‘moving’ branches of the tree creating what are infinite cycles through the possible outcomes in the game. We do this by preventing a player repeating the same move a certain number of times (say, k times); though, our results hold irrespective of k . Suppose, for example, we follow the branch in the extensive form game that begins with Row ‘moving’ to CD, Column next ‘moves’ to CC, Row then ‘moves’ to DC and column ‘moves’ to DD. When $k = 1$, Row can now only ‘stay’ at DD because to move to CD would be to repeat Row’s initial decision at the beginning DD on this branch (see Figure 2).

The next assumption embodies the NHP. We say that in so far as someone’s choice of action causes an outcome that is implemented, then they are only permitted to take that action if the resulting outcome does not harm other players. An individual can only cause an outcome that is implemented by ‘staying’ at an outcome. Of course, it takes more than one decision to ‘stay’ for an outcome actually to be implemented. But when a player ‘stays’ at




Figure 3: The beginning of the NHP constrained extensive form game where DD is the reference outcome.

an outcome, they cause that outcome and the only contribution that they can make as an individual to turning this outcome into the one that is implemented is by deciding to ‘stay’. Thus we apply the NHP to an individual’s decision to ‘stay’. Hence, in the extensive form game that begins with DD as the reference outcome in the prisoners’ dilemma, Row ‘staying’ at DD satisfies the NHP and so does Row ‘moving’ to CD (because NHP only applies to ‘stay’ decisions). However, Column’s option to ‘stay’ at CD would not satisfy the NHP, but moving to CC does. Thus, the beginning of the NHP constrained extensive form game looks like Figure 3 (as compared with the unconstrained version in Figure 1).

This illustration gives an immediate insight into how individual (farsighted) rationality and the NHP might contribute to making CC the no-harm equilibrium in this extensive form game. The pursuit of the best option for column at the second CD node is precluded by the NHP. So, the initial choice by Row of CD can only lead to CC at the third decision node because column can only move to CC from CD. For similar reasons, at any later decision node on this branch of the game, Row will not be able to ‘stay’ at CD and so the only possible terminal options will be DD and CC along this branch and individual rationality will secure CC. The details are, of course, a bit more complicated. One comment is worth making, nevertheless.

We require that two stay decisions at a particular outcome be made for it to be implemented in an ‘ n -person’ game. This is because we wish to avoid building in Pareto improvements through some version of unanimity rule that requires everyone to agree on some outcome before it is implemented. Our condition for implementation is, therefore, in general, much weaker than unanimity. However, in a two-person game it does amount to unanimity and this supplies another part of the intuition behind why CC emerges as the no-harm equilibrium in the prisoners’ dilemma illustration.

Our final key assumption is for more complicated games than the two person prisoners’ dilemma. We assume that if say player i decides to stay at an outcome but the next player j decides to reject this proposal by moving to another outcome, then the reference point is updated to the outcome proposed by i through their decision to ‘stay’. In effect, this outcome has been endorsed by i , it satisfies the NHP for i and it could have been implemented by j ; so, it is natural to use this as the (new) reference point for judging future deviations. Thus, in general, the decision to ‘stay’ is also a decision to change the reference point and this can only be done by a player when to do so satisfies the NHP.

To conclude, it is important to note that the assumptions behind our operationalization of the no-harm principle in non-cooperative games do not imply Pareto optimality; nor do they even secure Pareto improvements from a starting reference outcome (for examples and counterexamples, see section 3). It is the combination of individual rationality with the no-harm principle that characterizes Pareto optimality in n -person games.

2 The no-harm principle in non-cooperative games

2.1 The setup

Let $G = (A, u_i)_{i \in N}$ denote a normal form game in which $N = 1, 2, \dots, n$ denotes a society whose members are called players, $u_i : A \rightarrow R$ player i ’s Bernoulli utility function representing a strict ranking over the consequences,

A_i finite pure action set of player i , and $A = \times_{i \in N} A_i$ set of pure action profiles. Let $a = (a_1, a_2, \dots, a_n) \in A$ denote a pure action profile in game G .¹ As is standard in normal form games, every action profile is associated with an outcome (i.e., a pay-off profile) and vice versa. We use the terms action profile and outcome interchangeably.

A profile a Pareto dominates a' if for all i , $u_i(a) \geq u_i(a')$ with at least one strict inequality. A profile is called Pareto optimal or efficient if there is no other profile that Pareto dominates it. A profile is called weakly Pareto optimal if there is no other profile in which everyone is strictly better off.

For a given action profile $a_0 \in A$, a natural number $k \in \mathbb{N}^+ = \{1, 2, \dots\}$, and an ordering of players π in game $G = (A, u)$, we define an associated extensive form game with perfect information denoted by $\Gamma_{a_0, k, \pi} = (N, X, I, u, S, H)$, where players take turns in a given order to decide whether either to change their action and so ‘move’ to another available profile in G or not change their action by deciding to ‘stay’.² We refer to this extensive form game as simply “ Γ ” when its parameters $a_0, k \in \mathbb{N}^+$, and π can be inferred from the context, or when they are generic.

We next formally define Γ . Let X denote a game tree, $x \in X$ a node in the tree, x_0 the root of the game tree, and $z \in Z$ a terminal node, which is a node that is not a predecessor of any other node.

Let $I : X \rightarrow N$ denote the player function, where $I(x)$ gives the “active” player who moves at node x . Permutation function $\pi : N \rightarrow N$ denotes a permutation of players, according to which players take turns. Player $\pi(1)$ moves first at node x_0 , player $\pi(2)$ moves second, \dots , player $\pi(n)$ moves n th, and then by player $\pi(1)$ moves again, and so on.

Let $h \in H$ denote an information set, which is a singleton. With a slight abuse of notation, an information set h at node x is denoted by x , i.e., $h = x$.

¹Our definitions can be extended to the games with cardinal von Neumann-Morgenstern utilities in a straightforward way. We keep the current framework for its simplicity.

²For a standard textbook on extensive form games, see, e.g., Osborne and Rubinstein (1994). For a seminal bargaining game with alternating-offers and a discount rate, see Rubinstein (1982).

A *subgame* $\Gamma|x$ of a game Γ is the game Γ restricted to an information set $h = x$ and all of its successors in Γ .

Next, we define $O : X \rightarrow A$, called the *outcome function*, that maps each node x to action profile a which would be the outcome if x were a terminal node. Note that $a_0 = O(x_0)$ —i.e., the outcome at the root of the game would be a_0 , which is the action profile in G where the extensive form game Γ starts.

Let $A_i(x)$ denote the set of pure actions of player i at x . For each $x \in X$, $A_i(x)$ is defined as follows. First, define $X'(a_i, x) = \{x' \in X | a_i \text{ is chosen at } x', O(x') = O(x), \text{ and } x' \text{ is a predecessor of } x\}$. Then, $A_i(x) = \{a_i \in A_i | |X'(a_i, x)| < k\}$. Simply put, the same action by the same player at the same outcome can be chosen at most k times. For example, suppose that $k = 1$, $i = I(x_0) = I(x'')$, and $O(x_0) = O(x'')$. If i chooses a_i at x_0 , then $a_i \notin A_i(x'')$ because $k = 1$ implies that a_i can be chosen at outcome $O(x_0)$ only once.

Let $A'_i = \bigcup_{x \in X_i} A_i(x)$ denote player i 's set of all pure actions where X_i is player i 's set of all information sets. Let $\Sigma_i = \times_{x \in X_i} A_i(x)$ denote the set of all pure strategies of i where a pure strategy $\sigma_i \in \Sigma_i$ is a function $\sigma_i : X_i \rightarrow A'_i$ satisfying $\sigma_i(x) \in A_i(x)$ for all $x \in X_i$. Let $\sigma \in \Sigma$ denote a pure strategy profile and $u_i(\sigma)$ its (Bernoulli) utility for player i .

Let σ be a strategy profile and a be its outcome. With slight abuse of notation we use the same utility function for $u_i(\sigma)$ and $u_i(a)$ because their outcomes, and hence their utilities are the same.

Reference points and terminal nodes

Let $R(\sigma) = \{a \in A | \sigma_i(x) = a_i \text{ where } a = O(x), i = I(x), A_i(x) \setminus \{a_i\} \neq \emptyset\} \cup \{a_0\}$ be the set of all *reference points* under strategy profile σ . In other words, an outcome is called a reference point if the player who acts at the associated node “stays” at it: that is deliberately chooses not to change it. Set $R(\sigma)$ includes the initial outcome a_0 . Note that given a profile σ , for every decision node $x \neq x_0$ there is a unique reference point $y_x \in R(\sigma)$, which is a predecessor of $x \neq y_x$, such that there is no other reference point between x and y_x . Thus, we can refer to *the* reference point at x .

Game Γ comes to an end under two situations. First, let x and x' be two

nodes such that the reference point at x is x' where $O(x') = O(x) = b$. If some players i and then j stay at b , then, x is called a *terminal node*. Second, $x \in X$ is called a *terminal node* if $A_j(x) = \emptyset$, where $j = I(x)$. In plain words, the game terminates if either (i) two players choose to stay at an outcome (the first one is like a ‘proposal’ to implement this outcome and the second one amounts to an acceptance of this ‘proposal’), or (ii) there is no available action left in the game.³

An illustrative example

To illustrate our notation, we return to the prisoners’ dilemma (PD). Suppose that $a_0 = (C,C)$ is the reference point. Let $\Gamma_{a_0,k,\pi}$ be the extensive form game that begins this time with the reference point (C,C) , where $k = 1$, $\pi(1) = \text{Row}$ and $\pi(2) = \text{Column}$ —i.e., Row goes first and Column second. Starting from (C,C) players might end up at (D,D) if they play as follows. Row unilaterally switches their action to D, hence ‘moving’ to (D,C) . Column then moves to (D,D) , where Row chooses D to ‘stay’ which makes (D,D) the new reference point. Column also stays at (D,D) , where both players receive pay-offs of $(2, 2)$. Note that $k = 1$ means that if Row chooses C at (D,D) rather than stay, and then Column also chooses C at (C,D) , moving back to (C,C) , then Row cannot choose D again at (C,C) . We have not so far imposed any restrictions on the players’ choices such as ‘rationality’ or ‘no-harm principle’. We explore these notions in the following subsections.

2.2 The no-harm principle

Our specification of the no-harm principle applies when a player stays. We make this assumption for two reasons. First, in a dynamic strategic setting, the classical liberal has no reason to be concerned with the properties of any transitional (i.e., non-reference point) states in the extensive form game, par-

³Note that the cardinalities of Z and A are the same, $|Z| = |A|$. In other words, for every action profile $a \in A$ there is a terminal node $z \in Z$, and vice versa.

ticularly if they are purely mental constructs. But if a player chooses to stay, then this matters for everyone because any further change and eventually the final outcome will (perhaps indirectly) depend on it. Second, an individual can only contribute through their own decision to making an outcome the final outcome by deciding to stay and so change the reference point; it is this decision where the classical liberal should apply the no-harm principle—i.e., their decision should not harm other players with respect to the previous reference point.

Definition 1. Let $x \in X$ be a node, $b \in A$ the reference point at x , $O(x) = a$, $I(x) = i$, $A_i(x) \setminus \{a_i\}$ nonempty, and σ a strategy profile where $\sigma_i(x) = a_i$. Player i 's action $a_i \in A_i(x)$ is said to satisfy the *no-harm principle* (NHP) if $u_j(a) \geq u_j(b)$ whenever $j \neq i$. We say that strategy profile σ satisfies the NHP if every player i 's every action $a_i \in A_i(x)$ satisfies the NHP.

In plain words, a player's stay action satisfies the NHP if their decision does not harm others with respect to the current reference point. Accordingly, a strategy profile satisfies the NHP if every player's every stay action under that strategy profile satisfies the NHP. Of note, the no-harm principle implies neither Pareto optimality nor even Pareto improvement from a reference point. In section 3, we illustrate that assuming the no-harm principle may lead a society to a Pareto inferior outcome compared to the initial reference point.

2.3 Solution concept

We assume that players are individually rational and farsighted in the usual sense of subgame perfection and are additionally constrained by the no-harm principle (NHP) in their action choices in Γ . Moreover, we assume that G , Γ , and the previous sentence is common knowledge (Lewis, 1969; Aumann, 1976). First, we define subgame perfect equilibrium (Nash, 1951; Selten, 1965).

A pure strategy profile $\sigma \in \Sigma$ in game Γ is called a *subgame perfect equilibrium* (SPE) if for every player i and for every non-terminal $x \in X$ where $i = I(x)$, $u_i(\sigma|x) \geq u_i(\sigma'_i, \sigma_{-i}|x)$ for every $\sigma'_i|x \in \Sigma_i|x$. Put differently, σ is

a subgame perfect equilibrium if it constitutes a Nash equilibrium in every subgame of Γ .

Definition 2. A pure strategy profile $\sigma^* \in \Sigma$ that satisfies the no-harm principle is called a *no-harm equilibrium* (NHE) in G if for every player i and for every non-terminal $x \in X$ where $i = I(x)$

$$u_i(\sigma^*|x) \geq u_i(\sigma'_i, \sigma^*_{-i}|x)$$

for every $\sigma'_i|x \in \Sigma_i|x$ such that $(\sigma'_i, \sigma^*_{-i}|x) \in \Sigma|x$ satisfies the no-harm principle.

In plain words, a strategy profile is a no-harm equilibrium if at every node the active player plays a best response under the constraint of the no-harm principle. Like subgame perfect equilibrium, in finite games no-harm equilibrium can be computed using backward induction with the condition that players maximize their utility under the constraint of the NHP. Note that an NHE is *not* equivalent to a strategy profile that is both a subgame perfect equilibrium and satisfies the no-harm principle, in part because in general there may be no SPE that satisfies the NHP, but as we show in section 3 an NHE always exists.

The NHE depends, of course, on the initial reference point, a_0 . But for now it is important to note that the NHP *per se* does not require Pareto optimality of the final outcome. Players simply act independently and maximize their individual utility; they do not act to maximise the pay-offs of others. They can stay wherever they want as long as the outcome does not harm others with respect to the reference point and there could always be other outcomes that are as good for the individual who decides to ‘stay’ and which would be better for the other players. We illustrate this point in section 3 with an example where the NHP by itself does not produce a Pareto efficient outcome.

3 Existence, uniqueness, and efficiency

In this section, we first show that the no-harm equilibrium exists under very general conditions in normal form games.

Theorem 1. *For every initial reference point a_0 in a normal form game G , for every $k \in \mathbb{N}^+$, and every permutation of players π , there exists a no-harm equilibrium in pure strategies.*

Proof. We fix an initial reference point a_0 , a permutation of players π , and some $k \in \mathbb{N}^+$.

Notice that for every a_0 , the game Γ always possesses a pure subgame perfect equilibrium. This is true because Γ is a well-defined finite extensive form game with perfect information. To see this, notice that the root of the game is a_0 and that π gives a unique order of players who take their turns sequentially by construction of Γ . Because there are finitely many players and that k is finite, the game Γ ends after finitely many steps. This implies that there is always a subgame perfect equilibrium in pure strategies.

Next, we assume that players act according to the NHP, which essentially puts a constraint on their choices in Γ . This implies that they have fewer (finitely many) choices under the NHP than they have under Γ . Because the NHP is common knowledge, the constrained game—i.e., the game in which all strategy profiles satisfy the NHP—is still of perfect information. Let σ^* be a subgame perfect equilibrium in the constrained game, which exists by the same arguments as above. We note that σ^* is a no-harm equilibrium in Γ because σ^* satisfies the NHP and at every node every active player plays a best response among the profiles that satisfy the NHP, since by construction all those profiles satisfy the NHP. This concludes the proof that σ^* is a no-harm equilibrium. \square

We next show under what conditions the uniqueness of the NHE outcome is guaranteed from an initial reference point.

Theorem 2. *For every initial reference point a_0 in a normal form game G , for every $k \in \mathbb{N}^+$, and every permutation of players π , the no-harm equilibrium outcome is unique.*

Proof. Given an initial reference point a_0 , a finite k , and a permutation π , the associated Γ possesses a pure subgame perfect equilibrium as shown in

the proof of Theorem 1. We next show that this subgame perfect equilibrium outcome is unique. The reason is that no matter which player moves on a non-terminal node either (i) the player has a unique pure best response or (ii) the pure best responses all lead to the same outcome because the preferences of the players are strict in G . Thus, the subgame perfect equilibrium outcome in Γ must be unique. Analogously, the subgame perfect equilibrium outcome in Γ which is constrained by the NHP must also have a unique outcome. Together with Theorem 1, this implies that the NHE outcome must be unique. \square

Finally, we illustrate the relationship between the no-harm principle, rationality, and efficiency.

Theorem 3. *Let G be a game, π be a permutation of players, and $k \in \mathbb{N}^+$. A strategy profile is a no-harm equilibrium if and only if it is Pareto optimal in G .*

Proof. Given π and $k \in \mathbb{N}^+$, we first show that if an initial reference point a_0 is Pareto optimal then it is the no-harm equilibrium outcome. We know that the NHE is unique because the subgame perfect equilibrium is unique due to the fact that the preferences in G are strict (see Theorem 2). To reach a contradiction, suppose that a_0 is not the no-harm equilibrium outcome, and the outcome of the no-harm equilibrium from a_0 is given by some $a' \neq a_0$. Then, a' must Pareto dominate a_0 , which is a contradiction to the assumption that a_0 is Pareto optimal. This is because if some player j chooses to stay at an outcome that is different than a_0 , which is Pareto optimal, then either (i) the NHP would be violated or (ii) player j would receive a strictly lower pay-off, violating farsighted rationality.

Next, we show that for an initial reference point a_0 that is not Pareto optimal, its no-harm equilibrium must be Pareto optimal. Let σ^* be the no-harm equilibrium from a_0 , which exists by Theorem 1 for any $k \in \mathbb{N}^+$ and any π . To reach a contradiction, suppose that the outcome of σ^* is a , and a is Pareto dominated by some action profile b . Because σ^* is a no-harm equilibrium and satisfies the no-harm principle, there exists a player j who chooses to stay (i.e., a_j) at reference point a . However, player j 's choice of a_j

cannot be a best response because (i) there is a path from a to b , (ii) no other player can stay at an action profile which harms player j along the path because a is the reference point, and (iii) every player including j strictly prefers b to a . Statement (i) is correct for any $k \geq 1$. To reach a contradiction, suppose that (i) is false. Then, player j must have played suboptimally prior to a' because in the beginning of the game there exists a path between any two action profiles by construction of Γ . Statement (ii) holds by definition of the NHP and statement (iii) holds by our supposition that a is Pareto dominated.

As a result, the outcome of no-harm equilibrium σ^* cannot be Pareto dominated. Thus, it must be Pareto optimal. \square

3.1 Discussion of the assumptions

We next discuss how different assumptions in the definition of Γ and the no-harm equilibrium affect the results we have presented so far.

We begin with three brief observations. First, while the three theorems hold for any k and any π , the associated no-harm equilibrium would potentially be different for different k and π (see, e.g., the example in subsection 4.3). However, this does not change the conclusion of, e.g., Theorem 3 that any such no-harm equilibrium is Pareto efficient. Second, one might wonder what happens to the NHEs when there are indifferences between the outcomes in G . In that case, Theorem 1 would still be valid, though Theorem 2 would no longer hold. This is because subgame perfect equilibrium outcomes in Γ need not be unique, which implies that NHE outcomes need not be unique either. For analogous reasons as in the proof of Theorem 3 we can conclude that every Pareto optimal profile must be an NHE, and an NHE cannot be strictly Pareto dominated. Moreover, for every initial reference point, for every k , and for every π there would always be an NHE that is Pareto optimal. Third, one might also wonder what would happen if we give each player the opportunity to stop the process as well as stay whenever it is their turn. In that case, *mutatis mutandis*, our results would still be valid as long as NHP applies to ‘stop’ decisions as well.

We next discuss the roles of our two main assumptions.

The role of the no-harm principle

Next, we illustrate that the NHP is essential for Theorem 3. First, notice that the no-harm equilibrium definition would reduce to subgame perfect equilibrium if the no-harm principle were not assumed. Consider the following simple example and suppose that the NHP is not assumed.

	L	R
L	4, 3	1, 4
R	2, 1	3, 2

Let the initial reference point be $(1,4)$, $k = 1$, $\pi(1) = \text{Row}$, and $\pi(2) = \text{Column}$. On grounds of individual rationality alone, Row would move to $(3,2)$, where Column as well as Row would stay, making it the final outcome in Γ . This is because (i) Column would not gain by moving to $(2,1)$ from $(3,2)$ because Row would not move to $(4,3)$ as Row anticipates that Column would then go back to $(1,4)$ in which case Row would have to stay because $k = 1$, and (ii) if Row moves back to $(1,4)$ from $(3,2)$, Column would stay there as well, which makes Row worse off. Thus, without the no-harm principle and starting at $(1,4)$, players would end up at $(3,2)$, and this is Pareto dominated by $(4,3)$.

On a different note, if we strengthen the no-harm principle, e.g., by requiring many players to stay at the reference point to implement it as the outcome of the game, then Theorem 3 would still hold; though, the associated the no-harm equilibrium outcome would be a potentially different Pareto efficient outcome.

The role of farsighted rationality

Farsighted rationality is also a necessary assumption for Theorem 3 because a strategy profile might satisfy the no-harm principle alone and yield a Pareto inferior outcome with respect to the initial reference point. The following 2×2 game provides a simple example.

	L	R
L	2, 2	0, 3
R	1, 0	4, 4

Suppose that $a_0 = (2,2)$, $k = 1$, $\pi(1) = \text{Column}$, and $\pi(2) = \text{Row}$. Consider the strategy profile in which Column moves from $(2,2)$ to $(0,3)$ where Row stays at L, making $(0,3)$ the updated reference point. Row's choice of L satisfies the no-harm principle since it does not harm Column player. Next, Column moves back to $(2,2)$ and Row moves to $(1,0)$ where first Column stays and then Row stays, making $(1,0)$ the outcome of Γ . Column's decision to stay at $(1,0)$ satisfies the no-harm principle since it does not harm Row player with respect to the updated reference point $(0,3)$. Anticipating this and if the players were farsightedly rational, Column would *not* stay at $(1,0)$. But in the absence of the assumption of rationality, the aforementioned moves cannot be ruled out and it results in an outcome, $(1,0)$, that is strictly Pareto dominated by $(2,2)$.

4 Illustrations

4.1 The Prisoners' Dilemma

We first go back to the PD. It follows CC, CD and DC are Pareto efficient for some reference points and so all are NHEs, but DD is not Pareto efficient and is not an NHE. Nevertheless, although it is clear DD is not Pareto efficient, it is perhaps not immediately obvious why deviation from DD satisfies both (1) NHP and (2) farsighted rationality.

Let $\Gamma_{a_0, k, \pi}$ be the extensive form game that begins with the reference point $a_0 = (D,D)$, $k = 1$, $\pi(1) = \text{Row}$ and $\pi(2) = \text{Column}$. Consider the deviation from DD by Row to C. This deviation may seem to be precluded because it does not immediately satisfy Row's individual rationality—since Row is worse off—at CD. Nevertheless, to see whether it might satisfy Row's farsighted individual rationality (2), we need to consider what Column does at CD because CD may not be stopping point. Indeed, Column cannot stay at CD because CD harms Row relative to the reference point of DD and so will not satisfy the no-harm

principle (1). Hence if Column were to find themselves at CD, they would have to move and CC is the only option. Will Row stay at CC? CC satisfies (1) the NHP. It also satisfies Row's farsighted rationality (2) because a move to DC would produce a 'cycle' back to DD from which no further deviation would be permitted because $k = 1$. Thus since CC is better for Row than DD, it is also farsightedly rational for Row to stay at CC, which becomes the outcome of the NHE. In other words, you have to trace through what happens with a deviation by Row using DD as the reference point before you can see that (2) is also satisfied by the deviation of Row to C from reference point DD; and DD is not an NHE. Instead, CC is the NHE associated with the reference point of DD.

4.2 Stag-Hunt and Hawk and Dove

It is well-known that Pareto optimality and the Nash equilibrium are logically distinct concepts in the sense that neither concept is a refinement of the other. As we show in Theorem 3 the NHEs coincide with Pareto optimal profiles. Thus, there is no logical relationship between the set of NHEs and the set of Nash equilibria. Two further illustrations bring this out. In the Stag-Hunt game, NHE is a case of Nash refinement, and in the Hawk-Dove game, NHE expands the Nash equilibria; whereas, as we have just seen, in the PD the Nash equilibrium is not an NHE.

Consider, first, the Stag-Hunt game:

	Stag	Hare
Stag	4, 4	1, 3
Hare	3, 1	2, 2

Clearly, irrespective of the reference point the players will end up at (Stag, Stag), which is the Pareto dominant profile and also a Nash equilibrium. (Hare, Hare) is a Nash equilibrium but not an NHE.

Next, in the Hawk and Dove (Chicken) game, (Dove, Dove) is a NHE as well as the two Nash equilibria (H,D) and (D,H):

L	C	D	R	C	D
A	3, 5, 1	6, 1, 2	A	8, 4, 5	4, 7, 7
B	2, 2, 3	1, 3, 4	B	7, 8, 8	5, 6, 6

Figure 4: No-harm equilibria in a three-person illustrative game

	Hawk	Dove
Hawk	1, 1	4, 2
Dove	2, 4	3, 3

This is an interesting game in that both non-myopic equilibrium and no-harm equilibrium predictions coincide. The two concepts in general give different predictions mainly due to the no-harm principle. In general, not every non-myopic equilibrium is a no-harm equilibrium such as (D,D) in the PD. Conversely, not every no-harm equilibrium is a non-myopic equilibrium because not every Pareto optimal outcome is a non-myopic equilibrium such as (A,D) in game 22 (Brams, 1994), which is given below.

	C	D
A	2, 4	3, 3
B	1, 2	4, 1

The Pareto optimal profiles in this game are (A,D), (B,D), and (A,C), which is the non-myopic equilibrium.

4.3 A three-person illustrative example

We next illustrate the no-harm equilibria in a three-person game presented in Figure 4. Throughout this example, we assume that $k = 1$ and that $a_0 = (A,D,L)$ is the initial reference point whose associated outcome is (6, 1, 2).

First, assume that $\pi(1) = \text{Row}$, $\pi(2) = \text{Column}$, and $\pi(3) = \text{Matrix}$ —i.e., Row moves first, Column second, Matrix third, and so on. Figure 5 illustrates part of the game tree where the arrows show the on-path moves of the no-harm equilibrium, which can be described as follows. Row moves to (1, 3, 4),

Column moves to $(2, 2, 3)$, Matrix moves to $(7, 8, 8)$, Row moves to $(8, 4, 5)$, Column stays at $(8, 4, 5)$, and finally Matrix also stays, making $(8, 4, 5)$ the outcome of the no-harm equilibrium, which is Pareto optimal. The reason why Column stays at $(8, 4, 5)$ is that Column cannot gain any pay-off by moving to $(4, 7, 7)$. This is because Matrix cannot stay at $(4, 7, 7)$ because it does not satisfy the no-harm principle with respect to $(6, 1, 2)$. If Matrix moves to $(6, 1, 2)$ from $(4, 7, 7)$, then Row would have to stay because the no-cyclicity condition precludes Row from moving again to $(1, 3, 4)$. Then, Column's best response would be to move to $(3, 5, 1)$, followed by Matrix's move to $(8, 4, 5)$, where Row would stay. Column would also have to stay at $(8, 4, 5)$ because it is not possible to implement another outcome that satisfies the no-harm principle with respect to the updated reference point $(8, 4, 5)$.

Second, assume that Matrix moves first, Row second, Column third, and so on. The on-path actions of the no-harm equilibrium is given as follows. Matrix moves to $(4, 7, 7)$, Row moves to $(5, 6, 6)$, Column moves to $(7, 8, 8)$, where Matrix stays. Then Row also stays because there is no other profile that satisfies the no-harm principle with respect to $(7, 8, 8)$. Third, if Column moves first, Matrix moves second, and Row third, then the outcome of the no-harm equilibrium would be $(7, 8, 8)$. This is because Column would stay at $(6, 1, 2)$, and then Matrix would move to $(4, 7, 7)$. The rest is analogous to the second case.

At the outset, it looks like Column and Matrix should be able to implement their most preferred outcome $(7, 8, 8)$ in the game. However, as shown above this is not possible if Row is the first-mover at the initial reference point. This three-person example illustrates that the player ordering can affect the NHE associated with any reference point, but the ordering does not affect the conclusion that NHEs are Pareto optimal.

5 Discussion

5.1 Rule-like constraints on individual action

The no-harm principle directs players to take account of other's interest in a very particular way and it is well known from models of altruism and other regarding preferences that the orientation to the interests of others in such models can turn the Pareto superior action profile of CC in the PD into a Nash equilibrium. Since the no-harm principle has the same effect in the PD, it is worth commenting on how the no-harm principle differs from such models of altruism and other regarding preferences.

The no-harm principle is a rule-like constraint on individual action that someone who believes in or subscribes to classic liberalism will wish to follow. In contrast, in models of altruism and more generally models where individuals have pro-social preferences, an individual personally values the pay-offs enjoyed by others and this enjoyment typically grows with the size of the pay-off to others. There is no such monotonic relationship between another's pay-offs and an individual's likelihood of taking an action with the no-harm principle. An action either satisfies the no-harm principle or it does not and in the one case the action is permissible and in the other it is not. In this respect the no-harm principle is akin to a version of rule rationality: that is, if the language of preference satisfaction is retained individuals have a lexicographic preference for following a rule(s) and so when they act to satisfy their preferences, they act in accordance with the rule(s). Another example of rule rationality is provided by Kant's categorical imperative: to 'act only according to that maxim whereby you can at the same will that it should become a universal law'. The rule-like constraint in this instance is that the action is universalized (and is evaluated under this constraint independently of whether others actually take the same action). It is well-known that Kantian rule rationality can produce similar results to those we have derived for the no-harm principle in prisoners' dilemma interactions (e.g. see Roemer, 2010, in the economics literature and for a more general discussion, O'Neill, 1990). The difference is that the Kantian rule has a more controversial connec-

tion to classical liberalism than the no-harm principle (e.g. see Berlin, 1969); and, to our knowledge, the no harm principle has not been studied before in non-cooperative game theory, whereas the Kantian one has (e.g. see Roemer, 2010).

5.2 No-harm principle in social choice theory

The no-harm principle is related in classical liberal political philosophy to the presumption that the State should not intervene in individual decision making when the consequences of those decisions apply only to the individual(s) making the decisions. This non-intervention principle has, of course, since Sen (1970) featured prominently in the social choice literature. The no-harm principle has also been used more recently in this social choice literature (see, e.g., Lombardi et al., 2016). Mariotti and Veneziani (2009; 2013; 2020) introduce a notion called “Non-Interference” principle which roughly says that society’s preferences should not change following a change in circumstances that affect only one individual and for which everyone else is indifferent. Recently, Mariotti and Veneziani (2020) show that there is inconsistency between their “Non-Interference” principle and the Pareto principle (i.e., if everyone in a society prefers an alternative x to y , then society should prefer x to y) in a non-dictatorship.

Our formalization of the NHP differs from Mariotti and Veneziani’s in two main respects, the framework and the conceptual definition. The most obvious difference between the two principles is that ours applies to actions within a game theoretical framework whereas theirs applies to the preferences within a social choice context. Conceptually, under our no-harm principle a player is allowed to choose any action as long as this action does not eventually harm (and may benefit) other players with respect to the reference point. However, the “Non-Interference” principle does not apply to a change in social situations that leave some members of the society better off.

Although we share the interest in the implications of subscribing to the tenets of classical liberalism with the social choice literature, the approach

here is very different. We are not interested in the implications of classical liberalism for a social planner—as is the case in the social choice literature. Instead, we are concerned with how the introduction of the no-harm principle as a constraint on individual decision making in games affects the equilibrium outcomes of those games.

5.3 Pareto efficiency

In a recent and related development, Che, Kim, Kojima, and Ryan (2020) provided a characterization of the Pareto optima via utilitarian welfare maximization. While both our and their approaches are sequential in nature, the main difference between the two papers is that their framework is non-strategic whereas we provide a non-cooperative foundation for Pareto optimality via the no-harm principle.

Ray and Vohra (2020) recently introduced a general class of games called “games of love and hate” to describe strategic situations with pay-off-based externalities. They show that every Nash equilibrium in this class is Pareto optimal under some regularity conditions; though, not every Pareto optimal profile is a Nash equilibrium. Note that both their and our approaches ensure that all equilibria are Pareto efficient. Ray and Vohra (2020) make certain assumptions on the pay-off functions of the players (e.g., prisoners’ dilemma falls outside of that class), whereas we assume that players maximize utility subject to the rule-like constraint of the no-harm principle.

In modeling the no-harm principle, we have followed the approach of Brams (1994). In this same tradition, Brams and Ismail (2021) show that there is always a non-myopic equilibrium that is Pareto optimal; though, not all Pareto optimal profiles are non-myopic equilibria, and not all non-myopic equilibria are Pareto optimal, as we discussed in section 4. The main difference in our results comes from the no-harm principle that we assume, which restricts the actions of the players to ones that satisfy the well-known principle of classical liberalism.

5.4 Welfare economics

There is an analogy between Theorem 3 and the first and the second welfare theorems. The first fundamental theorem of welfare economics states that, roughly speaking, irrespective of the set of initial endowments the competitive equilibrium is Pareto efficient. In our setting, Theorem 3 says that for every initial reference point in the society the associated no-harm equilibrium is always Pareto efficient. The second fundamental theorem of welfare economics states that any Pareto efficient allocation can be achieved as a competitive equilibrium allocation for some set of initial endowments. Analogously, by Theorem 3 for every Pareto efficient profile in a strategic game there is an initial reference point for which this is the no-harm equilibrium.

While these theorems point to a similar outcome, the processes and assumptions behind them are quite different. There is no account of how a competitive equilibrium is reached in general equilibrium theory and, critically, a competitive general equilibrium assumes that agents are price takers. There is no strategic interaction in the sense of a game theory.

6 Conclusion

Game theory standardly makes no assumption about what motivates individuals to act other than they have preferences they seek to satisfy. While this is an admirably parsimonious assumption, it is also misleading when people either subscribe to the political philosophy of classical liberalism or live in a society that is legally founded on the principles of classical liberalism. Such people are additionally constrained, either legally or by their own beliefs, by the no-harm principle. This is because the principle is the key constraint placed on the exercise of individual freedom by J. S. Mill in his classic manifesto for individual liberty: *On Liberty*. Thus, for those who live in a classically liberal society and/or who believe in classical liberalism, a question naturally arises: how is behaviour in games affected by the additional individual constraint on action supplied by the no-harm principle? We offer part of an answer to this

question.

We show with our operationalization of the no-harm principle, that this addition dramatically alters the predictions regarding what happens in games. In particular, the no-harm equilibria are always Pareto optimal. This stands in marked contrast to standard game theory where there is no necessary connection between Nash equilibria and Pareto optimality. It is important in the derivation of this result to note that our operationalization of the no-harm principle does not require Pareto optimality; nor does it even secure Pareto improvements from a starting position. It is the combination of individual rationality with the no-harm principle that secures Pareto optimality.

Our paper opens up two main directions for future theoretical research. First, what are the other frameworks in which strategic foundations of Pareto optimality can be studied? Second, while the definitions of no-harm principle and the no-harm equilibrium can be extended to games under incomplete and imperfect information in a straightforward way as subgame perfect equilibrium is well-defined under these settings, can it be extended to games with infinite horizons?

It also suggests an important new direction for empirical research that has possible implications for public policy. It is well known from experiments, for example, that some people behave selfishly in public goods/PD interactions and others behave pro-socially by contributing to the public good. The pro-social contributions are typically understood through the prism of social preferences and selfishness is understood through the absence of such preferences. To what extent, then, might they be better understood through the differing sway or influence that the no-harm principle has on individuals? In particular, while the puzzle from these experiments from the perspective of standard game theory has centred on why subjects contributed anything to the public good, it changes with the result of this paper. Rather, the puzzling question becomes: why do so many subjects in these experiments, when they come from liberal societies, behave selfishly?⁴ This, in turn, connects to the

⁴Amadae (2016) has an answer to this question: the rise and influence of Game Theory. In fact, the seeds of the analysis in this paper were sown

new policy agenda suggested by this paper: we need, when considering policy interventions, to understand better and focus on the circumstances under which people are not guided by the no-harm principle in liberal democratic societies.

To put this last point slightly differently, it is often argued that the prisoners' dilemma helps explain why people decide to restrict their freedoms (e.g. Hobbes, 1651). From the perspective of this paper, it is not the occurrence of prisoners' dilemmas in social and economic life, a sort of brute fact about some types of interactions, that occasions this retreat from liberty. Rather, it is a retreat from the liberal conception of liberty that is responsible for making prisoners' dilemma interactions problematic; and from a policy perspective, it is important to get the source of the problem right.

References

Amadae, S. M. (2016). *Prisoners of Reason: Game Theory and Neoliberal Political Economy*. Cambridge University Press.

Aumann, R. J. (1976). Agreeing to disagree. *The Annals of Statistics* 4(6), 1236–1239.

Berlin, I. (1969). *Four Essays on Liberty*. Oxford University Press.

Brams, S. (1994). *Theory of Moves*. Cambridge, UK: Cambridge University Press.

Brams, S. J. and M. S. Ismail (2021). Every normal-form game has a Pareto-optimal nonmyopic equilibrium. *Theory and Decision*, 1–14.

by Amadae (2016) and Hargreaves Heap (2016). Amadae (2016) argues that game theory has encouraged a form of neo-liberalism that is distinct from classical liberalism precisely because game theory dispenses with the no-harm principle. She conjectures that the no-harm principle would dramatically alter the prediction of what rational individuals would do in a prisoners' dilemma. Hargreaves Heap (2016) reviewed this book and found this conjecture intuitively plausible and so, in effect, reproduced it in the review.

- Brams, S. J. and D. Wittman (1981). Nonmyopic equilibria in 2×2 games. *Conflict Management and Peace Science* 6(1), 39–62.
- Che, Y.-K., J. Kim, F. Kojima, and C. T. Ryan (2020). Characterizing Pareto Optima: Sequential Utilitarian Welfare Maximization. Available at <https://arxiv.org/abs/2008.10819>.
- Coase, R. H. (1960). The problem of social cost. *The Journal of Law and Economics* 3, 1–44.
- Hargreaves Heap, S. (2016). Review: Prisoners of Reason: Game Theory and Neoliberal Political Economy, by S. M. Amadae. *Journal of Economic Literature* 54(4), 1392–1394.
- Hobbes, T. (1651). *Leviathan Or The Matter, Forme, Amp Power Of A Common-wealth Ecclesiasticall And Civill*. London.
- Lewis, D. (1969). *Convention: A Philosophical Study*. Harvard University Press.
- Lombardi, M., K. Miyagishima, and R. Veneziani (2016). Liberal egalitarianism and the harm principle. *The Economic Journal* 126(597), 2173–2196.
- Mariotti, M. and R. Veneziani (2009). ‘Non-interference’ implies equality. *Social Choice and Welfare* 32(1), 123–128.
- Mariotti, M. and R. Veneziani (2013). On the impossibility of complete non-interference in Paretian social judgements. *Journal of Economic Theory* 148(4), 1689–1699.
- Mariotti, M. and R. Veneziani (2020). The Liberal Ethics of Non-Interference. *British Journal of Political Science* 50(2), 567–584.
- Mill, J. S. (1859). *On Liberty*. John W. Parker and Son: London.
- Nash, J. (1951). Non-Cooperative Games. *The Annals of Mathematics* 54(2), 286–295.

- O'Neill, O. (1989). *Constructions of Reason: Explorations of Kant's Practical Philosophy*. Cambridge University Press.
- Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory*. MIT Press.
- Roemer, J. E. (2010). Kantian equilibrium. *Scandinavian Journal of Economics* 112(1), 1–24.
- Rubinstein, A. (1982). Perfect Equilibrium in a Bargaining Model. *Econometrica* 50(1), 97–109.
- Selten, R. (1965). Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft*.
- Sen, A. (1970). The impossibility of a Paretian liberal. *Journal of Political Economy* 78(1), 152–157.