arXiv:2412.11984v4 [econ.TH] 20 Aug 2026

Quantifying Inefficiency*

Yannai A. Gonczarowski Ella Segev

August 20, 2026

Abstract

We axiomatically define a *cardinal* social inefficiency function, which, given alternatives and individuals' vNM preferences, assigns a number—the *social inefficiency*—to each alternative. These numbers induce not only an ordinal comparison but also a cardinal comparison between alternatives that is uniquely defined by our axioms despite no exogenously given interpersonal comparison, outside option, or disagreement point. We interpret these numbers as per-capita losses in endogenously normalized utility. We apply our social inefficiency function to a setting where interpersonal comparison is hard to justify—object allocation without money—leveraging techniques from computer science to prove an approximate-efficiency result for the Random Serial Dictatorship mechanism.

*Gonczarowski: Department of Economics and Department of Computer Science, Harvard University (*e-mail*: yannai@gonch.name); Segev: The Hebrew University Business School and the Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem (*e-mail*: ella.segev@mail.huji.ac.il). We thank Cassidy Shubatt for excellent research assistance, and thank Gabriel Carroll, Francesco Fabbri, Loren Fryxell, Jerry Green, Sergiu Hart, Scott Kominers, Eric Maskin, Debraj Ray, Tim Roughgarden, Assaf Romm, Itai Sher, Ran Shorrer, Tomasz Strzalecki, Omer Tamuz, Rakesh Vohra, and Frank Yang for insightful comments and discussions. We thank participants at the Georgetown Microeconomic Theory Seminar, the Normative Economics and Economic Policy Webinar, the University of Notre Dame Economic Theory Mini-Conference, the Harvard–MIT Economic Theory Seminar, the Caltech Social Sciences Seminar, the Hebrew University Economic Theory Seminar, the Warwick Economic Theory Workshop, the A16z Crypto research seminar series, the Harvard Theory of Computation Seminar, the Northeastern Computer Science Theory Seminar, the Penn State Microeconomic Theory Seminar, and the Stony Brook Workshop on New Perspectives on Algorithmic Game Theory, for helpful comments. We gratefully acknowledge research support by the following sources. Gonczarowski: National Science Foundation (NSF-BSF grant No. 2343922), Harvard FAS Inequality in America Initiative, and Harvard FAS Dean's Competitive Fund for Promising Scholarship. Segev: Israel Science Foundation (ISF grant No. 301/23).

Contents

1 Introduction **1**

2 Model **7**

3 Axioms **9**

4 Characterization **13**

4.1 Explicit Construction 13

4.2 Cardinal Interpretation via Nonmonetary Compensating Differentials 17

4.3 Logical Independence of Our Axioms 20

4.4 Proofs 21

4.4.1 General Properties 22

4.4.2 Proof of Theorem 1 23

5 Application to Object Allocation **25**

5.1 Model 25

5.2 Robust FLAW Guarantees 26

5.2.1 Preliminaries 28

5.2.2 A Lower Bound on the Robust FLAW of Truthful Ordinal Mechanisms 28

5.2.3 An Upper Bound on the Robust FLAW of Random Serial Dictatorship 29

5.3 FLAW Computation 32

6 Related Work **34**

7 Discussion **39**

References **40**

A Proofs Omitted from Section 4 **A.1**

B Proofs Omitted from Section 5 **A.6**

1 Introduction

The question of how to aggregate the preferences of individual members of a society over a set of possible alternatives into a social order is at the heart of social choice theory, dating back to Arrow (1951) and Harsanyi (1955). Such ordinal aggregation allows us to compare any two alternatives to determine which of the two is socially better, yet does not generally allow us to quantify the *magnitude* by which one alternative is socially better than another. For example, such ordinal aggregation allows us to identify whether a given alternative is socially optimal; however, in case a given alternative is not socially optimal, such ordinal aggregation does not generally allow us to quantify *how far* this alternative is from being optimal—i.e., *how socially inefficient* this alternative is. In this paper, we fill this lacuna.

Consider an individual's ordinal preferences (over lotteries over alternatives). Under the standard von Neumann–Morgenstern (vNM) axioms, these preferences have an expected-utility representation, that is, utility values for each pure alternative, such that a lottery over alternatives is preferred over another if and only if its expected utility is higher. However, these utilities are only unique up to a monotone affine transformation: Multiplying all the utilities of any specific individual by a positive constant, or shifting them by a constant, does not change the preferences that they represent.[1] This underscores the problematic nature of interpersonal comparability of utilities, which is precluded by the mainstream view within social choice theory (Arrow, 1951). Beyond the challenges that the preclusion of interpersonal comparability poses for the ordinal aggregation of preferences, it is even less clear how one might go about defining a notion of *how socially inefficient* some alternative is while depending only on individuals' ordinal preferences and not on any arbitrary choice of affine transformation/normalization of utilities.

In this paper, despite the aforementioned challenges, we axiomatically define not an ordinal but a *cardinal* measure of the social (in)efficiency of any given alternative. This measure provides a unique cardinal comparison between alternatives—even across different settings, e.g., allowing one to numerically compare the social (in)efficiency of the different equilibria that arise in a game for different preference

[1] While some papers use the term *cardinal preferences* for such preferences (whose utility representation is unique up to a monotone affine transformation), in this paper we call these *ordinal preferences* to emphasize that any utility representation of them captures no more than an order over lotteries. We reserve the term *cardinal* for measures that capture richer information beyond only orders.

profiles. And yet, this measure is based solely on the (ordinal) von Neumann–Morgenstern (vNM) preferences of the individuals.

One of the early major contributions of theoretical computer science to economic theory is in its abundant use of (and techniques for deriving) *inefficiency-bounding theorems*.[2] Such theorems, rather than ascertaining the efficiency of some equilibrium, mechanism, or strategy, ascertain that it is guaranteed to not be far from efficient, in a precise, quantifiable sense. Such theorems are most straightforward to state (but not necessarily to prove, of course) in transferable-utility settings. For example, Jin and Lu (2023) recently proved that in a Bayes–Nash equilibrium of a first-price auction with independent bidders' values, the expected welfare (value of the bidder who wins the item) is no less than 86% of the expected optimal welfare (highest value).[3] This inefficiency-bounding theorem allows for a more nuanced understanding of the well-known inefficiency of first-price auctions due to information asymmetries. Other notable examples of inefficiency-bounding theorems include McAfee (2002); Rogerson (2003); Neeman (2003); Akbarpour et al. (2023).

Inefficiency-bounding theorems can also be useful, albeit more elusive to state, in non-transferable-utility settings. For example, an early celebrated contribution of Roughgarden and Tardos (2002) is that the cumulative driving time in every equilibrium of any of a large class of congestion games is no greater than $4/3$ times the optimal cumulative driving time. This inefficiency-bounding theorem allows for further insights beyond the well-known fact (Braess, 1968) that equilibria in many congestion games are inefficient due to unpriced externalities. We emphasize that social (in)efficiency is viewed here not as a binary (Pareto efficient vs. not Pareto efficient) but as a continuum. In fact, the absence of transferable utility allows some Pareto efficient alternatives to be considered less socially efficient than others (e.g., one driver driving alone on one road and all other drivers driving on the other available road can result in a very high cumulative driving time). We observe that three key properties of the cardinal measure of social inefficiency underlying this inefficiency-bounding theorem—the ratio between the cumulative driving time and the optimal cumulative

[2]These are many times referred to as *approximation theorems* within computer science.

[3]In computer-science nomenclature, this result proves that the "Price of Anarchy" (Koutsoupias and Papadimitriou, 1999) in this setting is at least 0.86. The precise fraction is $1 - 1/e^2$, and the same paper proves that it cannot be improved upon. This tight bound follows a sequence of papers that gradually proved tighter and tighter bounds (Syrgkanis and Tardos, 2013; Hartline et al., 2014; Hoy et al., 2018). Inefficiency-bounding theorems indeed allow for fast-paced progress via different sets of authors building upon each other's insights through publication of gradually tighter bounds.

driving time—contribute to its interpretability: *ordinal aggregation* (over the individuals), *intra-context cardinal comparability*, and *inter-context cardinal comparability*.

Ordinal aggregation (over the individuals). Within a congestion game, if one alternative has higher cumulative driving time than another, then the former is less socially efficient than the latter. This is driven by the disutility of every individual being exogenously cast in the same interpersonally comparable unit of measurement (minutes of travel). Ordinal aggregation might be more challenging when the (dis)utility function of each individual merely represents vNM preferences (and is thus defined only up to a monotone affine transformation). This has been the subject of research and debate within social choice theory for decades.

Intra-context cardinal comparability. Within a specific context (specific congestion game), the social inefficiency loss between two alternatives whose cumulative driving times differ by a specific amount (e.g., five minutes) is the same as that between any other two alternatives in the same context whose cumulative driving times differ by the same amount. Intra-context cardinal comparability might be more challenging if the (dis)utility functions of the individuals are not cast in the same unit of measurement.

Inter-context cardinal comparability. Considering two separate contexts (two congestion games), figures such as "4/3 of the optimum" are comparably interpreted in each context. This is due to social inefficiency being not merely an aggregate utility using a common unit, but furthermore (1) measured relative to a meaningful reference point, and (2) viewed as a fraction of the optimal aggregate utility. Indeed, if one accepts that a cumulative driving time of zero is a meaningful reference point, then 4/3 of the optimal cumulative driving time is in a precise sense "4/3 times as bad" (relative to the reference point) as the optimal cumulative driving time, and one might make the point that two equilibria, in two different games, that are each "4/3 times as bad" as the optimum for its respective game are comparable in their social inefficiency. Inter-context cardinal comparability might be more challenging in the absence of a natural reference point. In such settings, even if we have a way to cardinally aggregate over the individuals using a common unit of measurement, it might not be clear how to compare across contexts (e.g., it might be hard to normalize by dividing by the optimum, since shifting any aggregate numeric function by any constant changes the

ratio of function values between any alternative and the optimum).[4]

It is unclear how to state economically meaningful inefficiency-bounding theorems in settings such as matching and voting, in which a measure of social inefficiency that satisfies the three properties just discussed is not readily available. In such settings, if we were to have at our disposal a meaningful cardinal way—that satisfies the above three properties—for measuring the social (in)efficiency of an alternative, this would open the door to using techniques, fine-tuned over several decades in the computer-science literature, to bound this (in)efficiency at alternatives of interest (such as equilibria) and derive new insights in settings in which economically interpretable inefficiency-bounding theorems might currently be lacking.

In this paper, we initiate the study of axiomatically founding a *cardinal* social welfare function—or more specifically, a *cardinal social inefficiency* function—satisfying the above three properties, with the aim of demonstrating a new direction for building sound economic foundations for inefficiency-bounding theorems. Our social inefficiency function is uniquely defined (up to the unit in which it is globally measured, across all contexts) and depends only on the vNM preferences of each individual over the various possible alternatives. The definition of our social inefficiency function does not assume any exogenously given way to compare utility across individuals, nor does it assume any exogenously given reference point or disagreement point (such as having special meaning for a utility value of zero). We demonstrate the use of this social inefficiency function by applying it to a setting in which interpersonal comparison is notoriously challenging to justify and a reference point might not be available: Object allocation without money. In this setting, we use our social inefficiency function to leverage existing mathematical constructions and techniques from the computer-science literature to prove an approximate-efficiency result for the popular Random Serial Dictatorship mechanism.

We define a social inefficiency function as a function that for every *context*—set of possible alternatives, number of individuals, and vNM preferences for the individuals over (lotteries over) these alternatives—and for every specific alternative (or lottery over alternatives), specifies a nonnegative number, called the *social inefficiency* of

[4]Luce and Raiffa (1957, p. 33) write that when restricting to ordinal comparisons, "the non-uniqueness of the zero point is of no real concern in any of the applications of utility theory, but the arbitrary unit of measurement gives trouble." For inefficiency-bounding theorems, the non-uniqueness of a zero point is of real concern as well.

this alternative given the context.[5] After formalizing these definitions in Section 2, in Section 3 we define six axioms that such a social inefficiency function should satisfy.

Our first two axioms relate to ordinal aggregation (the first of the three properties discussed above): *Pareto monotonicity* is a natural desideratum on the ordinal nature of social inefficiency functions, while *anonymity* is a natural desideratum on the aggregative nature of social inefficiency functions. Our third axiom relates to intra-context cardinal comparability, and requires that the social inefficiency of lotteries be measured ex-ante, i.e., that the social inefficiency of a lottery equal the expected social inefficiency of the realized alternative.

The next two axioms relate to *relative* inter-context cardinal comparability, i.e., inter-context cardinal comparability of differences between social inefficiencies. The first of these two axioms is a cardinal version of *independence of irrelevant alternatives (IIA)*. The second, *population-size stability*, requires that relative social inefficiencies in a context remain the same, and do not arbitrarily change, upon self-composition of multiple copies of that context; this axiom pushes our social inefficiency function to have a per-capita flavor, facilitating the comparison on equal footing of social inefficiencies in contexts with different numbers of individuals.

The final, sixth axiom relates to absolute (rather than relative) inter-context cardinal comparability: *feasibility* requires that at each context, *some* alternative has zero social inefficiency.

In Section 4, we state the main result of the first part of this paper: That these six, logically independent axioms characterize a unique social inefficiency function up to a global choice of unit of measure, that is, up to a positive multiplicative factor applied *simultaneously across all contexts and alternatives*. As we show, the characterized social inefficiency function takes a particularly natural form: It measures the per-capita additive loss in welfare relative to the maximum-welfare alternative, where welfare is defined as the sum of *endogenously* normalized vNM utility functions, each utility function normalized with respect to the magnitude of the distributional conflict, that is, normalized so that its unit is the diameter of its range over the Pareto frontier.[6] We therefore call this social inefficiency function the *Frontier-normalized*

[5] Considering the social inefficiency of randomized and not only deterministic alternatives allows us to quantify the social inefficiency of the outcomes of *randomized* mechanisms. More generally, as Arrow (1951, p. 20) writes: "if conceptually we imagine a choice being made between two alternatives, we cannot exclude any probability distribution over those two choices as a possible alternative."

[6] The *kindness function* of Rabin (1993) also uses (albeit for a very different purpose) a normal-

Loss in Average Welfare (FLAW). FLAW has additional appealing properties beyond those explicitly required by our axioms, such as that the social inefficiency of an alternative is not changed if dominated alternatives are added to, or removed from, the context in question.[7]

In Section 5, we apply FLAW to the study of a setting in which interpersonal comparison is notoriously challenging to justify due to the absence of monetary transfers and the repugnance of measuring using money, and any normalization is challenging due to the absence of a reference point such as an outside option: Object allocation (without money).

Random Serial Dictatorship (RSD) is a widely used allocation mechanism that is known to *not* be ex-ante efficient (Zhou, 1990; see also Bogomolnaia and Moulin, 2001). We apply FLAW (which measures ex-ante inefficiency) to prove that nonetheless, RSD is *approximately* ex-ante efficient. We consider the robust FLAW (i.e., worst FLAW over all object allocation problems) of truthful ordinal mechanisms at their truthtelling equilibrium (ex-ante over the randomness of the mechanism). We utilize tools that were originally developed within computer science, for the study of the "Price of Anarchy" (Filos-Ratsikas et al., 2014), to prove that no such mechanism has a robust FLAW that is lower than (i.e., improves upon) that of RSD by more than 28%. We emphasize two crucial features of FLAW that give meaning to the statement of this result: First, we can define the robust FLAW of a mechanism (recall that this is the worst FLAW over all object allocation problems) because FLAW values are cardinally comparable across contexts/object allocation problems. Second, no other mechanism being able to improve by more than 28% is interpretable because our axioms pin down FLAW as a unique *cardinal* social inefficiency function up to a *global* unit of measure (and hence one mechanism's FLAW being lower by 28% than another's is well defined independently of this unit of measure). Finally, we consider the task of computing the FLAW of any matching with respect to specific given preferences, and construct a computationally feasible algorithm for doing so.

To our knowledge, our paper is the first that, for arbitrary sets of alternatives,

ization that depends on utility ranges over the Pareto frontier. The choice of normalization there is an explicit part of the definition rather than a consequence of an axiomatic characterization.

[7]For not satisfying (the ordinal analogue of) this property, Arrow (1951, p. 32) criticizes the social choice function that normalizes each utility function to the range of $[0, 1]$ (over all alternatives rather than over the Pareto frontier like FLAW does) and sums up the normalized utilities, calling this social choice function "extremely unsatisfactory."

axiomatically microfounds a social welfare function that depends *only* on the individuals' ordinal (vNM) preferences over the alternatives, whose *numeric value* (rather than only its induced order) is unique (whether or not up to a global choice of unit of measure) and cannot be arbitrarily additively shifted, and which can be meaningfully compared *across contexts*. We are furthermore not aware of any paper that uses or characterizes the *order* induced by FLAW. In Section 6, we survey related work. We conclude in Section 7.

There is a vast body of mathematical techniques within computer science for proving what computer scientists call *approximation theorems*: theorems that ascertain that some quantity is at least a specific fraction of some reference quantity. This paper aims to make this vast body of work more applicable to economic settings by providing a microfoundation for a cardinal—rather than ordinal—aggregation of individuals' (vNM) preferences. While we focus on cardinally measuring (in)efficiency, the program presented in this paper has the potential to be applicable far more broadly. For example, one might equally imagine axiomatically microfounding a cardinal social (un)fairness function and using techniques from computer science (and beyond) to prove approximate fairness guarantees. We indeed hope that *cardinal* social choice theory could facilitate bringing economic theory and theoretical computer science even closer together.

2 Model

A *context* is a pair $\left(X, (\succeq_i)_{i=1}^n\right)$ where X is a nonempty finite set called the set of *alternatives*, $n \geq 2$ is called the number of *individuals*, and for each $i = 1, \ldots, n$, $\succeq_i$ is a preference order (of individual i) over $\Delta(X)$ that satisfies the von Neumann–Morgenstern (vNM) rationality axioms (i.e., that has a vNM expected-utility representation).[8] We use the notation $\text{Pareto}(C) \subseteq \Delta(X)$ to denote the *Pareto frontier* of C: the set of all $x \in \Delta(X)$ such that there does not exist $x' \in \Delta(X)$ such that $x' \succeq_i x$ for every i and $x' \succ_i x$ (i.e., not $x \succeq_i x'$) for some i. We restrict attention to contexts in which no individual is indifferent along the entire Pareto frontier. (Such

[8] As is standard, $\Delta(X)$ denotes the set of lotteries over elements of X. For $x, y \in \Delta(X)$ and $0 \leq \alpha \leq 1$, we write $\alpha \cdot x + (1-\alpha) \cdot y$ to denote the lottery whose outcome is with probability α drawn from x and otherwise drawn from y.

an individual does not play a real part in the distributional conflict.)[9] A main example of a context is derived from an n-player game, where X is the set of strategy profiles of the game and $\succeq_i$ is the preference order of player i (over $\Delta(X)$) for every i.

Definition 1 (Social Inefficiency Function)**.** A *social inefficiency function* specifies for each context $C = \left(X, (\succeq_i)_{i=1}^n\right)$ and each $x \in \Delta(X)$ a number $I(C, x) \geq 0$, called the *social inefficiency* of x with respect to C.[10,11]

We emphasize that the higher the social inefficiency, the *less socially efficient* the alternative is viewed to be.

While the richness of the class of social preference orders induced by various social welfare functions has been repeatedly demonstrated in social choice theory, we conclude this section by demonstrating the further richness of the class of social inefficiency functions that is due to their cardinal nature.

To demonstrate this richness, we focus on *relative utilitarianism* (Dhillon and Mertens, 1999), which is the social preference order induced by normalizing each individual's (vNM) utility function to be between 0 and 1 over all alternatives, summing the individuals' normalized utilities to define the social welfare of a given alternative, and comparing alternatives by their social welfare.[12] Formally, using our notation, for every context $C = \left(X, (\succeq_i)_{i=1}^n\right)$ let $u_C^1, \ldots, u_C^n$ be vNM utility representations of $\succeq_1, \ldots, \succeq_n$, respectively, such that for every $i = 1, \ldots, n$, $\min_{x \in X} u_C^i(x) = 0$ and $\max_{x \in X} u_C^i(x) = 1$. Let $U_C(x) \triangleq \sum_{i=1}^n u_C^i(x)$. For every context $C = \left(X, (\succeq_i)_{i=1}^n\right)$ and $x, y \in \Delta(X)$, the relative utilitarianism order is defined as $x \succeq_C y$ if and only if $U_C(x) \geq U_C(y)$. This social preference order received several axiomatic foundations (Karni, 1998; Dhillon and Mertens, 1999; Segal, 2000; see Section 6 for a discussion).

One social inefficiency function that induces the relative utilitarianism order is

[9] A previous version of this manuscript contained a generalized analysis that does not make this assumption.

[10] A main example is analyzing the context C derived from a game and taking x to be an equilibrium of the game. One might even take the supremum (or infimum) of $I(C, x)$ over all equilibria x of the game as a measure of the inefficiency of the game. This would serve a similar purpose to the Price of Anarchy from computer science, but would be invariant to strategic equivalence (i.e., not depend on the choice of vNM utility function for each player).

[11] We use the term social inefficiency *function* even though, strictly speaking, I is not a function because the class of all contexts is not a set. For clarity of exposition, we gloss over this point, which is immaterial for this paper because no set-theoretic issues arise here that affect our analysis.

[12] In contrast, recall from the introduction that our social inefficiency function normalizes each utility function so that its unit is the diameter of its range over the Pareto frontier (rather than over the set of all alternatives).

$I(C,x) = n - U_C(x)$. However, the relative utilitarianism order is also induced by many other social inefficiency functions, such as $I(C,x) \triangleq \frac{n-U_C(x)}{|X|}$ and $I(C,x) \triangleq n - U_C(x) + |X|$, and even $I(C,x) \triangleq n^n - \big(U_C(x)\big)^n$. Despite all of these social inefficiency functions being equally well motivated by all of the existing axiomatic foundations of relative utilitarianism that we are aware of (which all axiomatize only the induced order), the distinction between them is crucial to our applications (see Section 5): $n - U_C(x) + |X|$ and $\frac{n-U_C(x)}{|X|}$ differ from $n - U_C(x)$ on inter-context cardinal comparability (e.g., when calculating the maximum of the social inefficiencies of several alternatives—each in a different context—using one of these social inefficiency functions rather than the other can change the set of worst alternatives); $n - U_C(x) + |X|$ and $n^n - \big(U_C(x)\big)^n$ differ from $n - U_C(x)$ on intra-context cardinal comparability (e.g., using one of these social inefficiency functions rather than the other can change the set of alternatives that are at least twice as inefficient as some given alternative). More generally, the relative utilitarianism order is induced by any social inefficiency function of the form $I(C,x) = f_C\big(-U_C(x)\big)$ where f_C is any context-dependent monotone transformation (so long as the resulting social inefficiencies are nonnegative for all alternatives). The fact that f_C can vary arbitrarily across contexts drives the richness of this class of social inefficiency functions. Our goal in this paper is to define a social inefficiency function that provides a unique cardinal comparison between alternatives, leaving open no such degrees of freedom. For more examples of social inefficiency functions (which induce different social preference orders), see Section 4.3 below.

3 Axioms

In this section, we state our axioms for a social inefficiency function. The first axiom is a standard (ordinal) Pareto monotonicity axiom—i.e., if y is weakly (strictly) Pareto dominated by x, then the social inefficiency at y is weakly (strictly) greater than at x.

Axiom 1 (Pareto Monotonicity)**.** Let $C = \big(X, (\succeq_i)_{i=1}^n\big)$ be a context. For every $x, y \in \Delta(X)$, if $x \succeq_i y$ for every $i = 1, \ldots, n$, then $I(C,x) \leq I(C,y)$, and if furthermore $x \succ_i y$ for some such i, then $I(C,x) < I(C,y)$.

The second axiom stipulates that the social inefficiency is unchanged if the order

of individuals is permuted.

Axiom 2 (Anonymity[13])**.** Let $C = \big(X, (\succeq_i)_{i=1}^n\big)$ be a context, let π be a permutation on $\{1, \ldots, n\}$ and set $C_\pi \triangleq (X, (\succeq_{\pi(i)})_{i=1}^n)$. For every $x \in \Delta(X)$, it holds that $I\big(C_\pi, x\big) = I(C, x)$.

One of our main use cases for social inefficiency (and our application in this paper, in Section 5) is quantifying the social inefficiency of the outcome of *randomized* mechanisms. In such a case, it is natural to measure the ex-ante social inefficiency of the outcome. Accordingly, our third axiom stipulates that the social inefficiency of lotteries be measured ex-ante. That is, this axiom stipulates that the social inefficiency of a lottery over two alternatives is the expected social inefficiency of the realized alternative. This axiom is hence named *expected inefficiency*, analogously to *expected utility*.[14]

Axiom 3 (Expected Inefficiency)**.** Let $C = \big(X, (\succeq_i)_{i=1}^n\big)$ be a context. For every $x, y \in \Delta(X)$ and $0 \leq \alpha \leq 1$, it holds that $I\big(C, \alpha \cdot x + (1-\alpha) \cdot y\big) = \alpha \cdot I(C, x) + (1-\alpha) \cdot I(C, y)$.

The fourth axiom is the familiar *independence of irrelevant alternatives (IIA)*, with two adjustments. The first adjustment to IIA is required since we are not given any exogenous reference point such as a disagreement point (or worst-case alternative

[13]The following, weaker yet bulkier variant of anonymity suffices just as well for our characterization. Let $C = \big(X, (\succeq_i)_{i=1}^n\big)$ be a context, let π be a permutation on $\{1, \ldots, n\}$, and set $C_\pi \triangleq \big(X, (\succeq_{\pi(i)})_{i=1}^n\big)$. For every $x, y \in \Delta(X)$, it holds that $I(C_\pi, x) - I(C_\pi, y) = I(C, x) - I(C, y)$. That is, the *relative* social inefficiency between two alternatives is unchanged if the order of individuals is permuted.

[14]One interpretation of this axiom (in tandem with Pareto monotonicity) is that the social inefficiency at a lottery between two alternatives is the same as at an imaginary alternative that every individual prefers as much as her certainty equivalent for the lottery. This axiom therefore allows the risk preferences *of individuals* to be taken into account in the social inefficiency function, and is one of the drivers of our social inefficiency function taking a utilitarian form (Harsanyi, 1955). While the dependence of utilitarian forms on risk preferences has drawn criticism in some settings, we refer readers to the excellent, eloquent discussion in the first five pages of the introduction of Dhillon and Mertens (1999) regarding why such criticism might be less applicable to social choice functions not intended as voting rules but rather for measurement and policy discussions. Our social inefficiency function is indeed intended for the latter, which drives our focus on cardinal comparability, both intra-context and inter-context. Another criticism of this axiom that is worth mentioning is that of Diamond (1967), who, in a nutshell, argues that it abstracts from distributional fairness considerations. We agree, but view this property as a desirable one: our social inefficiency function is indeed intended to measure the (overall) social efficiency loss rather than how the effects of this loss are distributed among the various individuals.

for society). Recall that IIA traditionally requires independence of irrelevant alternatives as long as the disagreement point is unchanged (Nash, 1953). In an elegant paper, Roth (1977) explores variants of IIA with respect to various reference points, and in particular considers two alternative, *endogenous* reference points: the *ideal point*—an imaginary alternative that every individual prefers as much as her ideal alternative—and the *point of minimal expectations*—an imaginary alternative that every individual prefers as much as her least-preferred *undominated* alternative. Our IIA axiom requires independence of irrelevant alternatives as long as both of these points are unchanged.[15]

The second adjustment to IIA is in what precisely is invariant to the removal of irrelevant alternatives. When the only information that an axiomatized object conveys about any two alternatives is their relative order, then IIA naturally stipulates that this order be invariant to the removal of irrelevant alternatives. In contrast, social inefficiency functions are cardinal, i.e., they are intended to convey more information than just an order. As such, the distances between the values (social inefficiencies) assigned to different alternatives must also contain salient information. (If these distances were not to contain salient information, then any monotone transformation of the social inefficiency function would convey the same information, i.e., the function would not be cardinal but rather ordinal.) Therefore, these distances should be retained when irrelevant alternatives are removed (otherwise, these alternatives are not truly irrelevant). Hence, instead of IIA requiring that if the social inefficiency of one alternative is lower than that of another then this order remains the same upon removal of irrelevant alternatives, IIA here requires that if the social inefficiency of one alternative is lower *by any specific difference (distance)* than that of another then this *difference* remains the same upon removal of irrelevant alternatives.

For a nonempty closed set $A \subseteq \Delta(X)$, we use $\max_{\succeq_i} A$ (resp. $\min_{\succeq_i} A$) to denote some $x \in A$ such that $x \succeq_i y$ (resp. $x \preceq_i y$) for every $y \in A$. These are well-defined since vNM preferences are continuous and since $|X| < \infty$ and hence $\Delta(X)$ is compact.

Definition 2 (Ideal Point and Point of Minimal Expectations (Roth, 1977))**.** Let $C = \big(X, (\succeq_i)_{i=1}^n\big)$ be a context, let $\emptyset \neq X' \subset X$, and set $C' \triangleq \big(X', (\succeq_i|_{\Delta(X')})_{i=1}^n\big)$.

- C and C' are said to have *the same ideal point* if $\max_{\succeq_i} X' \sim_i \max_{\succeq_i} X$ for every $i = 1, \ldots, n$.

[15] We emphasize that the more points we require to be unchanged, the *weaker* IIA becomes.

- C and C' are said to have *the same point of minimal expectations* if $\min_{\succeq_i} \text{Pareto}(C') \sim_i \min_{\succeq_i} \text{Pareto}(C)$ for every $i = 1, \ldots, n$.

Axiom 4 (Independence of Irrelevant Alternatives (IIA))**.** Let $C = \left(X, (\succeq_i)_{i=1}^n\right)$ be a context, let $\emptyset \neq X' \subset X$, and set $C' \triangleq \left(X', (\succeq_i|_{\Delta(X')})_{i=1}^n\right)$. If C and C' have the same ideal point and the same point of minimal expectations, then for every $x, y \in \Delta(X')$ it holds that $I(C', x) - I(C', y) = I(C, x) - I(C, y)$.

Before introducing the fifth axiom, following the definition by von Neumann and Morgenstern (1953) of the composition of two games, we define an analogous notion of composition of several contexts into one, composed context that models the original contexts as coexisting alongside each other, with their sets of individuals made disjoint, and without any interaction between these contexts. Let $C^1 = \left(X^1, (\succeq_i^1)_{i=1}^{n^1}\right)$ and $C^2 = \left(X^2, (\succeq_i^2)_{i=1}^{n^2}\right)$ be two contexts (with possibly different numbers of individuals). We define the *composed context* $C^1 \otimes C^2$ as the context that captures (respective copies of) these two contexts existing concurrently, without affecting each other. Formally, $C^1 \otimes C^2 \triangleq \left(X^1 \times X^2, (\succeq_i)_{i=1}^{n^1+n^2}\right)$ where for every $x, y \in \Delta(X^1 \times X^2)$, we have that:

1. For every $i = 1, \ldots, n^1$, it holds that $x \succeq_i y$ if and only if $x^1 \succeq_i^1 y^1$, where $x^1, y^1 \in \Delta(X^1)$ are the respective marginal distributions of x, y over X^1.

2. For every $i = n^1+1, \ldots, n^1+n^2$, it holds that $x \succeq_i y$ if and only if $x^2 \succeq_{i-n^1}^2 y^2$, where $x^2, y^2 \in \Delta(X^2)$ are the respective marginal distributions of x, y over X^2.

We inductively extend this definition to define the composition of more than two contexts.

Overall, so far, Axioms 1 through 3 facilitate the comparability of the social inefficiencies of different outcomes within the same context, and Axiom 4 facilitates the comparability of social inefficiencies across contexts that have the same number of individuals but larger or smaller sets of alternatives. The fifth axiom facilitates the comparability of social inefficiencies across contexts that have different numbers of individuals, by ensuring that the relative social inefficiency between two alternatives in a context remains the same, and does not arbitrarily change, upon self-composition of (multiple copies of) the context. This axiom pushes the social inefficiency to have a per-capita flavor, but is completely silent regarding whether this is achieved by

(weighted or unweighted, under many different possible norms) arithmetic averaging across individuals, geometric averaging, harmonic averaging, or in any other way.

Axiom 5 (Population-Size Stability)**.** Let $C = \left(X, (\succeq_i)_{i=1}^n\right)$ be a context. For every $x, x' \in X$ and $k \in \mathbb{N}$, it holds that $I\left(\otimes_{j=1}^k C, (x, \ldots, x)\right) - I\left(\otimes_{j=1}^k C, (x', \ldots, x')\right) = I(C, x) - I(C, x')$.

For the sixth and final axiom, if we had, in each context, an exogenously given optimal alternative, we would have required that this alternative have a social inefficiency of zero. In the absence of such an exogenous optimal alternative, this axiom merely ensures that in each context, there is *some* alternative that has a social inefficiency of zero—i.e., is considered most socially efficient.

Axiom 6 (Feasibility)**.** Let $C = \left(X, (\succeq_i)_{i=1}^n\right)$ be a context. There exists $x \in \Delta(X)$ such that $I(C, x) = 0$.

4 Characterization

In this section, we constructively define a social inefficiency function and prove that this social inefficiency function is uniquely characterized by our six axioms from Section 3, up to a global choice of unit of measure.

4.1 Explicit Construction

We start by explicitly constructing a social inefficiency function that satisfies our six axioms from Section 3 (see Figure 1 for an illustration of this construction when $n = 2$). Let $C = \left(X, (\succeq_i)_{i=1}^n\right)$ be a context, and for every $i = 1, \ldots, n$, let u_i be some vNM utility representation of $\succeq_i$ (see Figure 1a). Let $u_i^{max} \triangleq \max_{x \in \text{Pareto}(C)} u_i(x)$ and $u_i^{min} \triangleq \min_{x \in \text{Pareto}(C)} u_i(x)$ (see Figure 1b). Let $V(C, x) \triangleq \frac{1}{n} \sum_{i=1}^n \frac{u_i(x) - u_i^{min}}{u_i^{max} - u_i^{min}}$ for every $x \in \Delta(X)$ (see Figure 1c); the fractions in this expression are all well defined since no individual is indifferent along the entire Pareto frontier of C. We note that $V(C, x)$ does not depend on the choice of vNM utility representations of $(\succeq_i)_{i=1}^n$. One can view $V(C, \cdot)$ as a per-capita utilitarian social welfare function that normalizes the utility function of each individual so that its maximum and minimum over the Pareto frontier of C are 1 and 0, respectively. We define the social inefficiency at an

alternative $x \in \Delta(X)$ as the loss in the value of $V(C, \cdot)$ compared to its maximum possible value (see Figure 1d):

$$\hat{I}(C, x) \triangleq \max_{x' \in X} V(C, x') - V(C, x).$$

We note that $\hat{I}(C, x) \geq 0$ for every context C and $x \in \Delta(X)$. Hence (and since $\hat{I}(C, x)$ does not depend on the choice of vNM utility representations of $(\succeq_i)_{i=1}^n$), we have that $\hat{I}$ is a social inefficiency function.

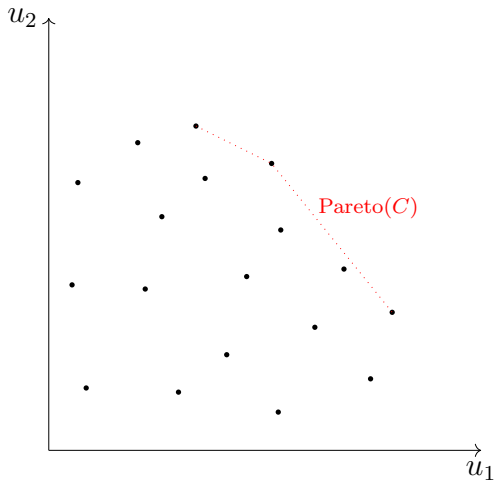
The social inefficiency function $\hat{I}$ has a natural interpretation: It is the per-capita additive utility loss (compared to the first best) when utility functions are endogenously normalized with respect to the magnitude of the distributional conflict, that is, normalized so that the unit of each utility function is the diameter of its range over the Pareto frontier. We therefore call $\hat{I}$ the *Frontier-normalized Loss in Average Welfare (FLAW)*.

It is straightforward to verify that FLAW satisfies our six axioms. An additional appealing property of FLAW is that for every $x \in \Delta(X)$, the individuals' preferences over the Pareto frontier of C and over x (but not over the rest of $\Delta(X)$) uniquely determine the FLAW of x—a property that we term *invariance to dominated alternatives*.

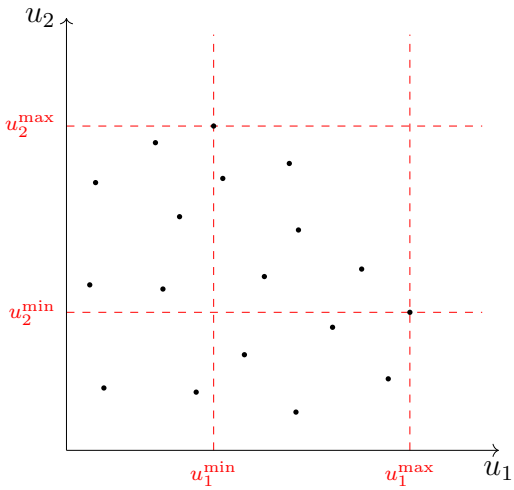
Definition 3 (Invariance to Dominated Alternatives)**.** A social inefficiency function I satisfies *invariance to dominated alternatives* if for every context $C = \left(X, (\succeq_i)_{i=1}^n\right)$ and every $X' \subset X$ such that $\Delta(X')$ weakly dominates $X \backslash X'$,[16] it holds that $I(C', x) = I(C, x)$ for every $x \in \Delta(X')$, where $C' \triangleq \left(X', (\succeq_i|_{\Delta(X')})_{i=1}^n\right)$.

Invariance to dominated alternatives—which we show in Lemma 2 in Section 4.4.1 below to be implied by Pareto monotonicity, IIA, and feasibility—is a desirable property, which relieves models that use social inefficiency functions that satisfy it (like FLAW) from the burden of, e.g., encompassing "universally worst" alternatives such as "all the members of the society die" (Kaneko and Nakamura, 1979) or having a set of alternatives that is "sufficiently encompassing as to include, besides the actual alternatives of interest, each person's best and worst alternative within the 'universal' set" (Dhillon and Mertens, 1999; for a discussion of these papers and other papers that make similar assumptions, see Section 6). Instead, invariance to dominated al-

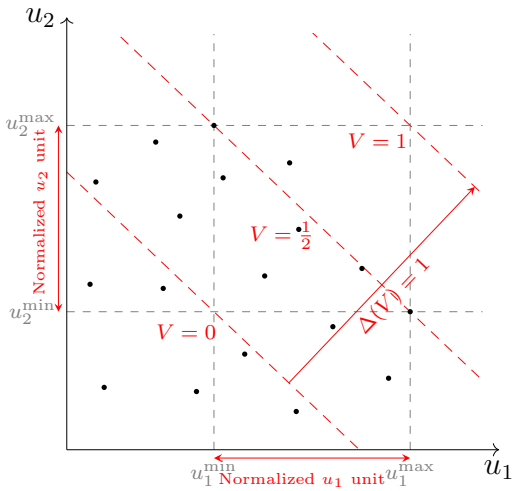
[16] That is, for every $x \in X \setminus X'$ there exists $x' \in \Delta(X')$ such that $x' \succeq_i x$ for every $i = 1, \ldots, n$.



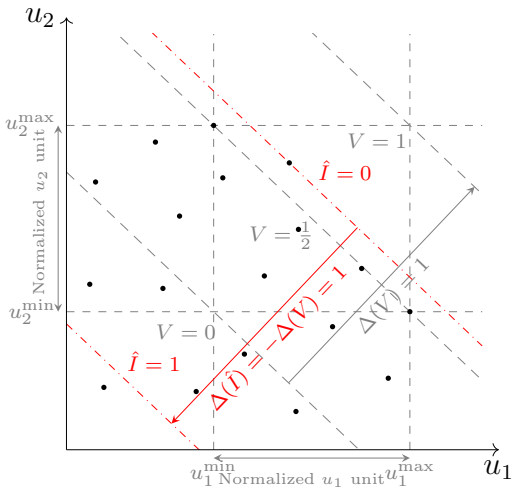
(a) A two-individual context. Scale and shift of vNM utility functions are chosen arbitrarily and do not affect calculation. Pareto frontier denoted by a dotted line.



(b) $u_i^{\max}$ and $u_i^{\min}$ are the respective maximum and minimum utilities along the Pareto frontier, for each individual i.



(c) Each utility function is normalized by respectively normalizing $u_i^{\max}$ and $u_i^{\min}$ to 1 and 0. At any alternative, V is the average normalized utility of the individuals.



(d) The FLAW of the alternative(s) with maximum V is $\hat{I} = 0$. For every other alternative, its FLAW $\hat{I}$ is the loss in V compared to the alternative(s) with $\hat{I} = 0$.

Figure 1: Illustration of FLAW calculation.

ternatives allows models to focus only on modeling undominated alternatives and other alternatives of interest.[17]

If we had at our disposal vNM utility functions that are exogenously normalized in a way that allows for interpersonal comparison, then a social inefficiency function that measures the loss in welfare compared to the welfare-maximizing alternative might naturally depend only on the (interpersonally comparable) utility values at the welfare-maximizing alternative and at the alternative whose inefficiency is evaluated.[18] Invariance to dominated alternatives is a natural generalization of this property for settings in which an exogenous way to compare utilities across individuals is not available. Indeed, invariance to dominated alternatives precisely requires that the social inefficiency of x is uniquely determined by the individuals' preferences over x as well as over all alternatives that are "potential welfare maximizers," since the Pareto frontier is precisely the set of all alternatives that are each a welfare-maximizer for *some* choice of vNM utility representations of the individuals' respective preferences (see, e.g., Mas-Colell et al., 1995, Prop. 16.E.2).

Another desirable property of FLAW (which is a direct consequence of Pareto monotonicity) is that zero FLAW is only possible on a (many times strict) subset of the Pareto frontier, with non-Pareto-frontier outcomes having positive FLAW. FLAW therefore refines the notion of Pareto efficiency.

Our main result is that our six axioms uniquely define FLAW, up to a single positive (finite) multiplicative constant that applies to all contexts and alternatives, i.e., up to a global choice of unit of measure.

Theorem 1. *A social inefficiency function I satisfies Pareto monotonicity, anonymity, expected inefficiency, IIA, population-size stability, and feasibility if and only if there exists a constant $c > 0$ such that for every context $C = \left(X, (\succeq_i)_{i=1}^n\right)$ and $x \in \Delta(X)$, it holds that $I(C, x) = c \cdot \hat{I}(C, x)$. Furthermore, dropping any of these six axioms invalidates this statement.*

[17] See Footnote 7 (on Page 6) for Arrow's (1951) criticism of a social choice function for not satisfying invariance to dominated alternatives.

[18] This is the case for the Price of Anarchy in computer science.

4.2 Cardinal Interpretation via Nonmonetary Compensating Differentials

Our social inefficiency function, FLAW ($\hat{I}$), conveys both ordinal and cardinal information. In this section, we show that the cardinal information has a natural interpretation in terms of the ordinal information via what we call *nonmonetary compensating differentials*.

In what follows, for every context $C = \big(X, (\succeq_i)_{i=1}^n\big)$, we denote by $\succeq_C$ the social preference order induced by FLAW over $\Delta(X)$ (i.e., $x \succeq_C y$ iff $\hat{I}(C,x) \leq \hat{I}(C,y)$), and we denote by opt_C an arbitrary $\succeq_C$-maximal alternative in X.

Recall from Theorem 1 that our axioms characterize FLAW up to a global choice of unit of measure. Therefore, we do not attempt to interpret the (unit-dependent) value of the FLAW of any specific alternative in isolation. Rather, we interpret the unit-independent part of FLAW, which consists of proportional comparisons between FLAW values; i.e., we interpret what it means that the FLAW of one alternative is some specific fraction of the FLAW of another alternative.

We start by considering two alternatives, $x^1, x^2 \in \Delta(X)$ in the same context $C = \big(X, (\succeq_i)_{i=1}^n\big)$. Assume w.l.o.g. that $x^1 \succeq_C x^2$. An immediate natural interpretation of the FLAW of x^1 being an α fraction of the FLAW of x^2 (where $\alpha \in [0,1]$) is as follows. $\hat{I}(C,x^1) = \alpha\hat{I}(C,x^2)$ if and only if

$$x^1 \quad \sim_C \quad \alpha \cdot x^2 + (1-\alpha) \cdot \text{opt}_C. \tag{1}$$

In other words, society is indifferent between x^1 and a lottery of an α probability of x^2 and otherwise opt_C. Mixing opt_C into each side of Eq. (1) in proportion $\alpha : 1$, we obtain that $\hat{I}(C,x^1) = \alpha\hat{I}(C,x^2)$ if and only if

$$\frac{1}{1+\alpha} \cdot x^1 + \frac{\alpha}{1+\alpha} \cdot \text{opt}_C \quad \sim_C \quad \frac{1}{1+\alpha} \cdot \text{opt}_C + \frac{\alpha}{1+\alpha} \cdot x^2. \tag{2}$$

Eq. (2) captures, using only $\succeq_C$ (and opt_C, which is itself defined using $\succeq_C$), the idea that going from opt_C (the first summand on the right-hand side) to x^1 (the first summand on the left-hand side) is "α times as socially bad" as going from opt_C (the second summand on the left-hand side) to x^2 (the second summand on the right-hand side). Eq. (2) captures this by stating that these two changes, with the latter scaled by α, socially cancel each other out. In other words, Eq. (2) says that the movement

from opt_C on the left-hand side to x^2 on the right-hand side, when weighted by α, forms a (nonmonetary) compensating differential for moving from x^1 on the left-hand side to opt_C on the right-hand side. This provides a natural interpretation of $\hat{I}(C, x^1) = \alpha\hat{I}(C, x^2)$, using only the order $\succeq_C$. We note that Eq. (2) holds for any $x^1, x^2 \in \Delta(X)$ and $\alpha \geq 0$ such that $\hat{I}(C, x^1) = \alpha\hat{I}(C, x^2)$, without the restriction of $\alpha \in [0, 1]$ that is required for Eq. (1) to hold.

The main reason we view the interpretation through Eq. (2) as a more salient interpretation of the cardinal information in FLAW than the interpretation through Eq. (1) is that the former extends also to comparison across contexts, i.e., to settings in which alternative x^1 is in some context C^1 and alternative x^2 is in a different context C^2. In this case, as we now show, there is a natural sense in which $\hat{I}(C^1, x^1) = \alpha\hat{I}(C^2, x^2)$ if and only if society is indifferent between $\frac{1}{1+\alpha} \cdot x^1 + \frac{\alpha}{1+\alpha} \cdot \text{opt}_{C^2}$ and $\frac{1}{1+\alpha} \cdot \text{opt}_{C^1} + \frac{\alpha}{1+\alpha} \cdot x^2$ (analogously to Eq. (2)). To make this precise, we must explain the meaning of lotteries that combine alternatives from both C^1 and C^2, and moreover, the sense in which the social order induced by FLAW is defined between such lotteries. We therefore start with a definition of when two contexts are naturally thought of as coexisting within a single large context.

We call two contexts with the same number of individuals $C^1 = \left(X^1, (\succeq_i^1)_{i=1}^n\right)$ and $C^2 = \left(X^2, (\succeq_i^2)_{i=1}^n\right)$ *compatible* if $X^1 \cap X^2 = \emptyset$ and there exists a context $U(C^1, C^2) = \left(X^1 \cup X^2, (\succeq_i)_{i=1}^n\right)$ such that the following two conditions hold for every $j = 1, 2$:

1. $\succeq_i |_{\Delta(X^j)} = \succeq_i^j$ for every $i = 1, \ldots, n$,

2. $U(C^1, C^2)$ and C^j have the same ideal point and the same point of minimal expectations (i.e., the two focal points of $\succeq_{U(C^1,C^2)}$ coincide with those of $\succeq_{C^j}$). That is, both $\max_{\succeq_i}(X^1 \cup X^2) \sim_i \max_{\succeq_i^j} X^j$ and $\min_{\succeq_i} \text{Pareto}\big(U(C^1, C^2)\big) \sim_i \min_{\succeq_i^j} \text{Pareto}(C^j)$ for every $i = 1, \ldots, n$.

Note that if such $U(C^1, C^2)$ exists, then it is uniquely defined.

The compensating-differential interpretation of the cardinal information in FLAW extends to alternatives in different compatible contexts C^1, C^2 through their embeddings in $U(C^1, C^2)$: Given two compatible contexts $C^1 = \left(X^1, (\succeq_i^1)_{i=1}^n\right)$ and $C^2 = \left(X^2, (\succeq_i^2)_{i=1}^n\right)$, for every $x^1 \in \Delta(X^1)$ and $x^2 \in \Delta(X^2)$, we have that $\hat{I}(C^1, x^1) = \alpha\hat{I}(C^2, x^2)$ (where $\alpha \geq 0$) if and only if in $U(C^1, C^2)$, going from opt_{C^1} to x^1 is "α

times as socially bad" as going from opt_{C^2} to x^2:

$$\frac{1}{1+\alpha}\cdot x^1+\frac{\alpha}{1+\alpha}\cdot \mathrm{opt}_{C^2} \quad \sim_{U(C^1,C^2)} \quad \frac{1}{1+\alpha}\cdot \mathrm{opt}_{C^1}+\frac{\alpha}{1+\alpha}\cdot x^2$$

(where each side is a lottery in $\Delta(X^1 \cup X^2)$), completely analogously to Eq. (2).

When C^1, C^2 are incompatible contexts with the same number of individuals, an interpretation via the transitive closure of the above relation trivially applies.[19] Importantly, this transitive closure stabilizes after only two steps, connecting any two alternatives, in any two contexts with the same number of individuals, through an intermediate context, without appealing to multiplication of other ratios of inefficiencies. That is, for every two contexts (compatible or incompatible, identical or distinct) $C^1 = \left(X^1, (\succeq_i^1)_{i=1}^n\right)$ and $C^2 = \left(X^2, (\succeq_i^2)_{i=1}^n\right)$ there exists a context $C^3 = \left(X^3, (\succeq_i^3)_{i=1}^n\right)$ that is compatible with C^1 and with C^2 such that the following holds for every $x^1 \in \Delta(X^1)$ and $x^2 \in \Delta(X^2)$. $\hat{I}(C^1, x^1) = \alpha \hat{I}(C^2, x^2)$ if and only if there exists an alternative $x^3 \in \Delta(X^3)$ such that both of the following hold:[20]

1. Going from opt_{C^1} to x^1 is "α times as socially bad" as going from opt_{C^3} to x^3:

$$\frac{1}{1+\alpha}\cdot x^1+\frac{\alpha}{1+\alpha}\cdot \mathrm{opt}_{C^3} \quad \sim_{U(C^1,C^3)} \quad \frac{1}{1+\alpha}\cdot \mathrm{opt}_{C^1}+\frac{\alpha}{1+\alpha}\cdot x^3.$$

2. Going from opt_{C^3} to x^3 is "as socially bad" (i.e., $\alpha = 1$) as going from opt_{C^2} to x^2:

$$\frac{1}{2}\cdot x^3+\frac{1}{2}\cdot \mathrm{opt}_{C^2} \quad \sim_{U(C^3,C^2)} \quad \frac{1}{2}\cdot \mathrm{opt}_{C^3}+\frac{1}{2}\cdot x^2.$$

[19] That is, for contexts C^1, C^2 and respective alternatives x^1, x^2, for every sequence of contexts $C^1 = D^1, D^2, \ldots, D^{k-1}, D^k = C^2$ in which every two consecutive contexts are compatible, and for every sequence of alternatives $x^1 = y^1, y^2, \ldots, y^{k-1}, y^k = x^2$ in these respective contexts, if for every $j = 1, \ldots, k-1$, going from opt_{D^j} to y^j is "α^j times as socially bad" as going from $\mathrm{opt}_{D^{j+1}}$ to y^{j+1} (i.e., $\hat{I}(D^j, y^j) = \alpha^j \hat{I}(D^{j+1}, y^{j+1})$), then $\hat{I}(C^1, x^1) = \alpha^1 \cdot \ldots \cdot \alpha^{k-1} \hat{I}(C^2, x^2)$.

[20] One way to construct such C^3 is as follows. Let $L \triangleq \max\left\{\hat{I}(C^j, x) \;\middle|\; j \in \{1,2\} \;\&\; x \in \Delta(X^j)\right\}$. Assume without loss of generality that $0, 1, \ldots, n \notin X^1 \cup X^2$, let $X^3 \triangleq \{0, 1, \ldots, n\}$, and define $C^3 \triangleq \left(X^3, (\succeq_i^3)_{i=1}^n\right)$ by defining $\succeq_i^3$ for every individual $i = 1, \ldots, n$ through the vNM utility representation u_i that for every $x \in X^3$ satisfies

$$u_i(x) \triangleq \begin{cases} -L & x = 0, \\ 1 & x = i, \\ 0 & \text{otherwise.} \end{cases}$$

The alternative x^3 can then be taken as the lottery over the alternatives $0 \in X^3$ and $1 \in X^3$ that satisfies $\hat{I}(C^3, x^3) = \hat{I}(C^2, x^2)$. Such a lottery exists since $\hat{I}(C^3, 1) = 0 \leq \hat{I}(C^2, x^2) \leq L < \hat{I}(C^3, 0)$. See the proof of Lemma 5 below for an in-depth analysis of a similar construction.

Finally, for two contexts with different numbers of individuals $C^1 = \left(X^1, (\succeq_i^1)_{i=1}^{n^1}\right)$ and $C^2 = \left(X^2, (\succeq_i^2)_{i=1}^{n^2}\right)$, the compensating-differential interpretation of the cardinal information in FLAW is again by embedding these two contexts in one large context. Because the two contexts have different numbers of individuals, such embedding requires equalizing the number of individuals in these contexts. This is done by duplicating each context an appropriate number of times so that the number of individuals in each of the two duplicated contexts is the same. That is, we let $\bar{C}^1 \triangleq \otimes_{j=1}^{n^2} C^1$ and $\bar{C}^2 \triangleq \otimes_{j=1}^{n^1} C^2$ (note that each of these contexts has $n^1 n^2$ individuals). For every $x^1 \in \Delta(X^1)$ and $x^2 \in \Delta(X^2)$ we have that $\hat{I}(C^1, x^1) = \alpha \hat{I}(C^2, x^2)$ if and only if $\hat{I}\left(\bar{C}^1, (x^1, \ldots, x^1)\right) = \alpha \hat{I}\left(\bar{C}^2, (x^2, \ldots, x^2)\right)$,[21] where the latter is interpreted as above via compensating differentials for two contexts with the same number of individuals (i.e., through $U(\bar{C}^1, \bar{C}^2)$ if $\bar{C}^1$ and $\bar{C}^2$ are compatible, and otherwise through $U(\bar{C}^1, C^3)$ and then $U(C^3, \bar{C}^2)$ for a context C^3 that is compatible with both $\bar{C}^1$ and $\bar{C}^2$).[22]

Overall, the social inefficiency function that emerges from our axioms, FLAW, is such that, based only on its ordinal information (its induced social order in every context), nonmonetary compensating differentials provide a complete interpretation of its unit-invariant cardinal information (proportional comparisons between FLAW values, even across contexts)—i.e., an interpretation of the cardinal comparison between alternatives that is uniquely defined by our axioms.

4.3 Logical Independence of Our Axioms

In this section, we show that our six axioms are logically independent. That is, if we drop any one of them, Theorem 1 ceases to hold (as claimed in the second part of that theorem). We do so by demonstrating, for each of the six axioms, a social inefficiency function that satisfies all of our axioms except that axiom.

Example 1 (Pareto Monotonicity is Logically Independent of the Other Axioms)**.**

[21] Here, for every $j = 1, 2$, we abuse notation by writing $(x^j, \ldots, x^j)$ for any lottery over alternatives in $\bar{C}^j$ whose marginals all equal x^j.

[22] The same interpretation is also obtained using any other pair of duplications of C^1 and C^2 that have the same number of individuals. For example, one might define each $\bar{C}^j$ as the duplication of C^j whose number of individuals is the least common multiple of n^1 and n^2.

The constant zero social inefficiency function, $I(C,x)=0$, satisfies all of our axioms except Pareto monotonicity.

Example 2 (Anonymity is Logically Independent of the Other Axioms)**.** Using the notation from the definition of $\hat{I}$, for every context C, first set $V'(C,x) \triangleq \frac{1}{1-2^{-n}} \sum_{i=1}^{n} 2^{-i} \cdot \frac{u_i(x)-u_i^{min}}{u_i^{max}-u_i^{min}}$ for all $x \in \Delta(X)$ and then $I(C,x) = \max_{x' \in X} V'(C,x') - V'(C,x)$. This social inefficiency function satisfies all of our axioms except anonymity.

Example 3 (Expected Inefficiency is Logically Independent of the Other Axioms)**.** Using the notation from the definition of $\hat{I}$, for every context C, first set $V'(C,x) \triangleq \frac{1}{n} \sum_{i=1}^{n} \left(\frac{u_i^{max}-u_i(x)}{u_i^{max}-u_i^{min}}\right)^2$ for all $x \in \Delta(X)$ and then $I(C,x) = V'(C,x) - \min_{x' \in \Delta(X)} V'(C,x')$. This social inefficiency function satisfies all of our axioms except expected inefficiency.

The remaining three (counter)examples of social inefficiency functions all induce, for each context, the same order over its alternatives as is induced by FLAW, and yet these examples differ cardinally from FLAW (and from each other). This once again showcases the richness of the design space of social inefficiency functions and that axiomatizing a cardinal social welfare function is a far stricter requirement than axiomatizing the order that it induces.

Example 4 (IIA is Logically Independent of the Other Axioms)**.** The social inefficiency function $I(C,x) = \mathrm{rad}\big(|X|\big) \cdot \hat{I}(C,x)$ (where X is the set of alternatives in C and $\mathrm{rad}\big(|X|\big)$, the radical of $|X|$, is the product of the distinct prime factors of $|X|$) satisfies all of our axioms except IIA.

Example 5 (Population-Size Stability is Logically Independent of the Other Axioms)**.** The social inefficiency function $I(C,x) = 2^n \cdot \hat{I}(C,x)$ (where n is the number of individuals in C) satisfies all of our axioms except population-size stability.[23]

Example 6 (Feasibility is Logically Independent of the Other Axioms)**.** The shifted social inefficiency function $I(C,x) = 1 + \hat{I}(C,x)$ satisfies all of our axioms except feasibility.

4.4 Proofs

In this section, we prove Theorem 1, as well as a number of intermediate results, some of which might be of independent interest. Several proofs are relegated to Section A.

[23] As part of our proof of Theorem 1, in Lemma 6 we characterize all social inefficiency functions that satisfy all of our axioms except population-size stability.

4.4.1 General Properties

In this section, we state and prove several notable properties of all social inefficiency functions that satisfy various subsets of our axioms. In addition to using these properties throughout the proof of Theorem 1 below, these properties are also of independent interest and might be viewed as additional desirable properties of social inefficiency functions that are implied by our axioms.

The first property, which immediately follows from Pareto monotonicity, is *Pareto indifference*: Two alternatives between which all individuals are indifferent must have the same social inefficiency.

Definition 4 (Pareto Indifference)**.** A social inefficiency function I satisfies *Pareto indifference* if for every context $C = \left(X, (\succeq_i)_{i=1}^n\right)$ and every $x, y \in \Delta(X)$ such that $x \sim_i y$ for every $i = 1, \ldots, n$, it holds that $I(C, x) = I(C, y)$.

Lemma 1. *If a social inefficiency function satisfies Pareto monotonicity, then it satisfies Pareto indifference.*

Recall from Section 4.1 the definition of invariance to dominated alternatives: for every $x \in \Delta(X)$, the individuals' preferences over the Pareto frontier of C and over x (but not over the rest of $\Delta(X)$) uniquely determine the social inefficiency of x. As noted in that section, this property is implied by Pareto monotonicity, IIA, and feasibility.

Lemma 2. *If a social inefficiency function satisfies Pareto monotonicity, IIA, and feasibility, then it satisfies invariance to dominated alternatives.*

We next establish that any social inefficiency function that satisfies Pareto indifference and invariance to dominated alternatives (or alternatively, by Lemmas 1 and 2, any social inefficiency function that satisfies Pareto monotonicity, IIA, and feasibility) also satisfies *neutrality*: it must be invariant to renaming of the alternatives in any context. To formally define neutrality, we first define what it means for two contexts to be *isomorphic*, i.e., identical up to renaming of the alternatives.

Definition 5 (Isomorphic Contexts)**.** Let $C^1 = \left(X^1, (\succeq_i^1)_{i=1}^n\right)$ and $C^2 = \left(X^2, (\succeq_i^2)_{i=1}^n\right)$ be two contexts with the same number of individuals. Let $\phi : X^1 \to X^2$ be a bijection between X^1 and X^2. We extend ϕ in the natural way to be defined over $\Delta(X^1)$, making it a bijection between $\Delta(X^1)$ and $\Delta(X^2)$. We say that ϕ is an *isomorphism*

between C^1 and C^2 if for every $x, y \in \Delta(X^1)$ and $i = 1, \ldots, n$, it holds that $x \succeq_i^1 y$ if and only if $\phi(x) \succeq_i^2 \phi(y)$. If there exists an isomorphism between C^1 and C^2, then we say that C^1 and C^2 are *isomorphic*.

Definition 6 (Neutrality)**.** A social inefficiency function I satisfies *neutrality* if for every pair $C^1 = \big(X^1, (\succeq_i^1)_{i=1}^n\big)$ and $C^2 = \big(X^2, (\succeq_i^2)_{i=1}^n\big)$ of isomorphic contexts with isomorphism ϕ between C^1 and C^2 and for every $x \in \Delta(X^1)$, it holds that $I(C^1, x) = I\big(C^2, \phi(x)\big)$.

Lemma 3. *If a social inefficiency function satisfies Pareto indifference and invariance to dominated alternatives, then it satisfies neutrality.*

4.4.2 Proof of Theorem 1

To prove Theorem 1, we go over our axioms, one at a time, at each point characterizing the social inefficiency functions that satisfy all axioms until that point, culminating in the characterization of the social inefficiency functions that satisfy all of our axioms, i.e., Theorem 1.

The first step of our proof of Theorem 1, which is formalized in the following lemma, has the following logic: expected inefficiency implies that a social inefficiency function is affine; adding Pareto monotonicity implies negative slopes; and adding feasibility pinpoints the affine shift so that the minimum inefficiency in the context is precisely zero. We note that while Pareto monotonicity immediately implies negative slopes when the image of $\Delta(X)$ in utility space has full dimension, the proof of this claim when this image is not of full dimension is more subtle, and employs a separating-hyperplane theorem.

Lemma 4. *For every context $C = \big(X, (\succeq_i)_{i=1}^n\big)$, fix some vNM utility representations $u_C^1, \ldots, u_C^n$ of $\succeq_1, \ldots, \succeq_n$. A social inefficiency function I satisfies Pareto monotonicity, expected inefficiency, and feasibility if and only if for every context $C = \big(X, (\succeq_i)_{i=1}^n\big)$ there exist constants $c_C^1, \ldots, c_C^n > 0$ such that*

$$I(C, x) = \max_{x' \in X} \left\{ \sum_{i=1}^n c_C^i u_C^i(x') \right\} - \sum_{i=1}^n c_C^i u_C^i(x)$$

for every $x \in \Delta(X)$.

Our next step is to apply IIA. Assume that for two given contexts C, C' with the same number of individuals and with disjoint sets of alternatives, a new context with the same number of individuals and the union of the sets of alternatives can be constructed, such that the new context and each of the original contexts have the same ideal point and the same point of minimal expectations. In this setting, for each individual i, if we define the slopes of i in these two contexts, $-c_C^i$ and $-c_{C'}^i$ (in the notation of Lemma 4), with respect to utility functions that coincide on the ideal point of each original context and the point of minimal expectations of each original context, then IIA implies that these two slopes are identical. While such a new context cannot be constructed for any arbitrary pair of contexts with the same number of individuals, it can be constructed for any pair of such contexts in which one of the contexts is a specific canonical context that we define for this number of individuals. Therefore, for each individual, the slopes of the social inefficiency function in any two contexts with the same number of individuals (if they are defined with respect to such utility functions) are the same. This logic is formalized in the following lemma, which characterizes all social inefficiency functions that satisfy the four axioms used so far: Pareto monotonicity, expected inefficiency, IIA, and feasibility.

Lemma 5. *A social inefficiency function I satisfies Pareto monotonicity, expected inefficiency, IIA, and feasibility if and only if for every $n \geq 2$ there exist constants $c_n^1, \ldots, c_n^n > 0$ such that*

$$I(C, x) = \max_{x' \in X} \left\{ \sum_{i=1}^{n} c_n^i \frac{u_i(x') - u_i^{min}}{u_i^{max} - u_i^{min}} \right\} - \sum_{i=1}^{n} c_n^i \frac{u_i(x) - u_i^{min}}{u_i^{max} - u_i^{min}}$$

for every context with n individuals $C = \left(X, (\succeq_i)_{i=1}^n\right)$, for every choice of vNM utility representations $u_1, \ldots, u_n$ of $\succeq_1, \ldots, \succeq_n$, and for every $x \in \Delta(X)$.

Our next step is to apply anonymity. This axiom implies that for symmetric contexts—ones in which every permutation of the individuals results in an isomorphic context—the slopes of the social inefficiency function (as defined in the discussion preceding Lemma 5) are all the same. Hence, by Lemma 5, this must be the case for all contexts. This step is formalized in the following lemma, which characterizes all social inefficiency functions that satisfy all of our axioms except population-size stability. In particular, fixing any number of individuals n, this lemma characterizes

FLAW as the social inefficiency function for contexts with n individuals that uniquely (up to the unit in which it is globally measured, across all such contexts) satisfies Pareto monotonicity, anonymity, expected inefficiency, IIA, and feasibility.

Lemma 6. *A social inefficiency function I satisfies Pareto monotonicity, anonymity, expected inefficiency, IIA, and feasibility if and only if for every $n \geq 2$ there exists a constant $c_n > 0$ such that $I(C, x) = c_n \cdot \hat{I}(C, x)$ for every context $C = \left(X, (\succeq_i)_{i=1}^n\right)$ with n individuals and for every $x \in \Delta(X)$.*

The final step is to leverage population-size stability (the final axiom that we have not used so far) to prove that the constants from Lemma 6 are the same for every number of individuals, characterizing $\hat{I}(C, x)$ as the unique social inefficiency function (up to the unit in which it is globally measured, across all contexts) that satisfies all six of our axioms, as claimed by Theorem 1. This step is formalized in the appendix.

5 Application to Object Allocation

In this section, we apply our social inefficiency function, FLAW, to a setting in which due to the absence of monetary transfers and outside options, and to the repugnance of measuring using money,[24] interpersonal comparison is challenging, and meaningful precise cardinal values (that is, not up to a monotone affine transformation) for utilities are challenging to justify: the object allocation setting. We explore this setting both from a theoretical perspective and from a computational one. Several proofs are relegated to Section B.

5.1 Model

In an object allocation context—to which, in line with the literature on matching theory, we henceforth refer as an *object allocation problem*—there are n individuals and n objects for some $n \geq 2$. The set of alternatives $X = X(n)$ is the set of all the perfect matchings of the n objects to the n individuals. The individuals' vNM preferences $\succeq = (\succeq_i)_{i=1}^n$ over (lotteries over) the matchings are such that each individual's preferences are induced by vNM preferences over (lotteries over) the object allocated to this individual, with no ties between any two objects. We denote

[24] See Roth (2007) for a survey of repugnance as a constraint in market design.

the set of all object allocation problems of size n (i.e., all contexts $C = \big(X, (\succeq_i)_{i=1}^n\big)$ such that $X = X(n)$ and the preferences are as above) by $\mathcal{M}_n$. We denote the set of all object allocation problems by $\mathcal{M} = \cup_{n=2}^{\infty} \mathcal{M}_n$.[25]

An *ordinal (matching) mechanism* is a mapping that takes a profile of arbitrary size n of (strict ordinal) rankings over n (pure) objects and outputs a distribution in $\Delta\big(X(n)\big)$. Given an object allocation problem $\big(X, (\succeq_i)_{i=1}^n\big) \in \mathcal{M}$, an ordinal mechanism μ induces a game between the n individuals, in which each individual i reports a ranking R_i of the n objects, and the outcome is the mechanism outcome $\mu\big((R_i)_{i=1}^n\big) \in \Delta(X)$. An ordinal mechanism is *truthful* if for every object allocation problem $\big(X, (\succeq_i)_{i=1}^n\big) \in \mathcal{M}$, for each individual i, *truthtelling* (i.e., playing the ranking induced by $\succeq_i$) is a dominant strategy in the induced game. To reduce clutter, we abuse notation by writing $\mu(\succeq)$ for the outcome of μ when every individual i reports the order induced by $\succeq_i$.

Random Serial Dictatorship (RSD) (Abdulkadiroğlu and Sönmez, 1998) is a truthful ordinal mechanism that has found wide use in allocating "objects" as varied as housing units (by both governments and colleges), internships, and school slots, just to name a few examples. This mechanism, in a uniformly random order over the individuals, matches each individual to the not-yet-matched object that is ranked highest on her reported ranking.

5.2 Robust FLAW Guarantees

The widely used Random Serial Dictatorship (RSD) mechanism is known to *not* be ex-ante efficient (Zhou, 1990; see also Bogomolnaia and Moulin, 2001).[26] In this section, we apply FLAW—which measures ex-ante inefficiency—to prove that RSD is nonetheless *approximately* ex-ante efficient.

We define the *robust FLAW* of a truthful ordinal mechanism to be the worst FLAW, over all object allocation problems, of the truthtelling equilibrium of that mechanism (ex-ante over the randomness of the mechanism). We prove that it is

[25]Since we require that an object allocation problem be a context, we exclude the possibility of individuals who are indifferent along the entire Pareto frontier. Given a setting with such individuals, unless there is no distributional conflict at all (i.e., each individual has a distinct most-preferred object), if we remove these individuals—and for each of them, the only object she receives in every allocation on the Pareto frontier—we obtain an essentially equivalent object allocation problem that is in $\mathcal{M}$.

[26]Manea (2009) further proves that RSD is frequently not ex-ante efficient. Dillenberger and Segal (2025, Proposition 1) prove that RSD is structurally far from all ex-ante efficient mechanisms.

impossible for any truthful ordinal mechanism to have a robust FLAW that is less than $\sim 72\%$ of that of RSD.

Theorem 2. *There does not exist any ordinal mechanism μ such that*

$$\sup_{(X,\succeq)\in\mathcal{M}} \hat{I}\big((X,\succeq),\mu(\succeq)\big) < \frac{1}{2\ln 2} \sup_{(X,\succeq)\in\mathcal{M}} \hat{I}\big((X,\succeq),RSD(\succeq)\big).$$

We note that while the statement of Theorem 2 does not require μ to be truthful, this theorem is most interesting when μ is truthful, as then it considers μ at equilibrium.

We recall that ratios between FLAW values, such as the ratio bounded by Theorem 2, were interpreted in Section 4.2. More generally, we emphasize that we would not have even been able to meaningfully interpret Theorem 2 had we not had a social inefficiency function that both allows for comparison across contexts (giving meaning to the supremum operation) and is cardinal (giving meaning to the notion of "$\frac{1}{2\ln 2}$ of the social inefficiency").

In the remainder of this section, we prove two bounds that together imply Theorem 2: A lower bound on the robust FLAW of truthful ordinal mechanisms in Theorem 3, and an upper bound on the robust FLAW of RSD in Theorem 4. In both bounds, we essentially adapt proofs from Filos-Ratsikas et al. (2014), who bound the "Price of Anarchy" of RSD and other ordinal mechanisms, that is, they bound the asymptotic order of magnitude of the ratio (rather than difference) between the sum of utilities at the mechanism output and the optimal sum of utilities. As we discuss below, these bounds of Filos-Ratsikas et al. (2014) depend on exogenously supplied cardinal utilities for specifying both an *exogenously given interpersonal comparison* to define welfare, and an *exogenously given "zero"* to give meaning to fractional comparison of welfare. These two features of their analysis—taking the ratio and having exogenous normalization for the utility functions—are common to many papers in computer science; by adapting the proofs of Filos-Ratsikas et al. (2014) with almost no conceptual change, we also demonstrate the potential of a large body of approximation theorems in computer science to be leveraged to give approximate-optimality results for FLAW, which in contrast to standard ones in computer science, depends only on the ordinal preferences and therefore assumes nothing about interpersonal comparison or an exogenously given, special zero utility value such as that of an outside option or a disagreement point.

5.2.1 Preliminaries

Before proving our bounds, we start with a few definitions and a lemma.

When n is clear from context, we denote by $\mathcal{R}$ the set of all (strict ordinal) rankings of the n (pure) objects. We call an element of $\mathcal{R}$ a *ranking* and an element of $\mathcal{R}^n$ a *profile of rankings*, denoting a general element of $\mathcal{R}^n$ by $R = (R_i)_{i=1}^n$. For an object allocation problem $(X, \succeq) \in \mathcal{M}_n$, for each individual $i = 1, \ldots, n$ denote by $R(\succeq_i) \in \mathcal{R}$ the ranking of the objects that is induced by $\succeq_i$, and denote $R(\succeq) = \big(R(\succeq_i)\big)_{i=1}^n \in \mathcal{R}^n$.

We say that a matching $x \in X = X(n)$ is *ex-post Pareto efficient* with respect to a profile of rankings $R \in \mathcal{R}^n$ if there does not exist $x' \in X$ that, according to R, is ranked weakly higher than x by all individuals and strictly higher than x by some individual. When the object allocation problem $(X, \succeq)$ discussed is clear from context, we say that a matching is *ex-post Pareto efficient* as a shorthand for saying that it is ex-post Pareto efficient with respect to $R(\succeq)$. An ordinal mechanism μ is *ex-post Pareto efficient* if for every profile of rankings R, it is the case that $\mu(R)$ is supported only on matchings that are ex-post Pareto efficient with respect to R. It is well known that RSD is ex-post Pareto efficient.

We emphasize that ex-post Pareto efficiency is defined with respect to Pareto dominance by pure matchings, whereas the Pareto frontier of $(X, \succeq)$ is defined with respect to Pareto dominance by randomized matchings. Nonetheless, for object allocation these two concepts are in fact equivalent.

Lemma 7. *Let $(X, \succeq)$ be an object allocation problem and let $x \in X$. The matching x is ex-post Pareto efficient if and only if $x \in \mathrm{Pareto}\big((X, \succeq)\big)$. Moreover, this equivalence holds even if some individuals are indifferent along the entire Pareto frontier.*

5.2.2 A Lower Bound on the Robust FLAW of Truthful Ordinal Mechanisms

We start by proving the following lower bound on the robust FLAW of every truthful ordinal mechanism.

Theorem 3. *For every $n \geq 2$ and every ordinal mechanism μ,*

$$\sup_{(X,\succeq)\in\mathcal{M}_n} \hat{I}\big((X, \succeq), \mu(\succeq)\big) \geq \frac{1}{2} - \frac{1}{n}.$$

Analogously to what we noted for Theorem 2, while the statement of Theorem 3 does not require μ to be truthful, this theorem is most interesting when μ is truthful, as then it considers μ at equilibrium.

The proof of Theorem 3, which is relegated to the appendix, is an adaptation of the proof of Lemma 5 in Filos-Ratsikas et al. (2014), which lower-bounds the order of magnitude of the *ratio* between the optimal sum of utilities and the sum of utilities at the mechanism output in ordinal mechanisms. The preference profile used in the proof of Theorem 3 in fact turns out to be slightly simpler than that in their Lemma 5.

5.2.3 An Upper Bound on the Robust FLAW of Random Serial Dictatorship

To complement Theorem 3, we prove the following *upper* bound on the robust FLAW of the Random Serial Dictatorship (RSD) mechanism.

Theorem 4. *For every $n \geq 2$,*

$$\sup_{(X,\succeq)\in\mathcal{M}_n} \hat{I}\big((X,\succeq), RSD(\succeq)\big) \leq \ln 2 \approx 0.693.$$

Our analysis is again heavily inspired by that of Filos-Ratsikas et al. (2014). As already noted, there are two main differences between FLAW and their measure of social inefficiency. First (as is common in computer science), the normalization of their utility functions is exogenous rather than endogenous: Theirs are normalized so that their minimum (over all matchings) is always 0 and their maximum is always 1, which exogenizes also interpersonal comparison and is affected by preferences over objects that would never be awarded by an ex-post Pareto efficient mechanism like RSD.[27] Second (as is also common in computer science), they analyze (and, for RSD, upper-bound) the ratio between the optimal and attained sum of utilities rather than the difference between the two.

We start by bridging this first difference. We do so by showing that the robust (worst-case) FLAW is smaller than the worst-case loss measured by normalizing each utility function to $[0, 1]$ over the entire set of alternatives.

[27] As discussed in Section 2, this is one of several social inefficiency functions that all correspond to relative utilitarianism, though Filos-Ratsikas et al. (2014) neither use the term relative utilitarianism nor appeal to any of the axiomatizations of relative utilitarianism.

For every individual $i = 1, \ldots, n$, let $\mathcal{S}_i$ be the set of utility functions over X that are each induced by a utility function over the object allocated to individual i that assigns distinct utilities to the objects. For every $i = 1, \ldots, n$ and every utility function $u_i \in \mathcal{S}_i$, we denote by $R(u_i) \in \mathcal{R}$ the ranking of the objects that is induced by u_i. Let $\mathcal{S} \triangleq \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$. For every profile $u = (u_i)_{i=1}^n \in \mathcal{S}$ of utility functions over X, we denote $R(u) = \big(R(u_i)\big)_{i=1}^n \in \mathcal{R}^n$; furthermore, for every ordinal mechanism μ, to reduce clutter we abuse notation by writing $\mu(u)$ instead of $\mu\big(R(u)\big)$. For every $i = 1, \ldots, n$, let $\mathcal{UR}_i \subset \mathcal{S}_i$ be the set of *unit range* utility functions in $\mathcal{S}_i$—that is, functions in $\mathcal{S}_i$ whose minimum is 0 and whose maximum is 1. Let $\mathcal{UR} \triangleq \mathcal{UR}_1 \times \cdots \times \mathcal{UR}_n$. Finally, for every profile $u = (u_i)_{i=1}^n$ of utility functions over X, we extend u over $\Delta(X)$ by taking expectation and define $U_u(x) \triangleq \sum_{i=1}^n u_i(x)$ and $I_u(x) \triangleq \max_{x' \in X} U_u(x') - U_u(x)$ for every $x \in \Delta(X)$.

Lemma 8. $\sup_{(X,\succeq) \in \mathcal{M}_n} \hat{I}\big((X,\succeq), \mu(\succeq)\big) \leq \frac{1}{n} \sup_{u \in \mathcal{UR}} I_u\big(\mu(u)\big)$ *for every ex-post Pareto efficient ordinal mechanism* μ *and* $n \geq 2$.

Proof. Given a profile $u = (u_i)_{i=1}^n \in \mathcal{S}$ and $i \in \{1, \ldots, n\}$, we say that u_i is *Pareto-frontier-normalized* (with respect to the profile u) if the maximum of u_i is 1 and the minimum of u_i over the Pareto frontier is 0. Denote by $\mathcal{N}$ the profiles $u = (u_i)_{i=1}^n \in \mathcal{S}$ such that u_i is Pareto-frontier-normalized for every $i = 1, \ldots, n$. By definition, we have that $\sup_{(X,\succeq) \in \mathcal{M}_n} \hat{I}\big((X,\succeq), \mu(\succeq)\big) \leq \frac{1}{n} \sup_{u \in \mathcal{N}} I_u\big(\mu(u)\big)$. It therefore suffices to show that $\sup_{u \in \mathcal{N}} I_u\big(\mu(u)\big) \leq \sup_{u \in \mathcal{UR}} I_u\big(\mu(u)\big)$. Since μ is ex-post Pareto efficient, this follows by induction from the following claim, which we now prove: For every profile $u = (u_i)_{i=1}^n \in \mathcal{S}$ such that each u_i is either Pareto-frontier-normalized with respect to u or in $\mathcal{UR}_i$, for every $i = 1, \ldots, n$, and for every $0 < \varepsilon < \frac{1}{2}$, there exists $u_i' \in \mathcal{UR}_i$ satisfying $R(u_i') = R(u_i)$, such that $I_u(x) \leq I_{(u_i', u_{-i})}(x) + \varepsilon$ for every matching $x \in X$ that is ex-post Pareto efficient with respect to $R(u)$. (Note that it follows from Lemma 7 that the extreme objects on the Pareto frontier for each individual do not change upon the replacement of u_i with u_i' since $R(u_i') = R(u_i)$, and therefore preferences that are Pareto-frontier-normalized before this replacement remain Pareto-frontier-normalized after it.)

If $u_i \in \mathcal{UR}_i$, then we can set $u_i' = u_i$ and the proof is complete; assume therefore that $u_i \notin \mathcal{UR}_i$; since u_i is Pareto-frontier-normalized, we therefore have that there exists $x \in X$ such that $u_i(x) < 0$. Let $L \triangleq \min_{x \in X} u_i(x) < 0$. Define u_i' as follows for

every $x \in X$:

$$u_i'(x) \triangleq \begin{cases} \varepsilon + (1-\varepsilon)u_i(x) & u_i(x) \geq 0, \\ \varepsilon\big(1 - u_i(x)/L\big) & u_i(x) < 0. \end{cases}$$

Note that by definition, $R(u_i') = R(u_i)$. Furthermore, since u_i is Pareto-frontier-normalized and there exists $x \in X$ such that $u_i(x) < 0$, we have that $u_i' \in \mathcal{UR}_i$. Note furthermore that for every matching $x \in X$ that is ex-post Pareto efficient with respect to $R(u)$, by Lemma 7 this matching is on the Pareto frontier (with respect to u), and hence $0 \leq u_i'(x) - u_i(x) \leq \varepsilon$ and therefore $0 \leq U_{(u_i', u_{-i})}(x) - U_u(x) \leq \varepsilon$. Fix some $x_{max} \in \arg\max_{x \in X} U_u(x)$, and note that x_{max} is ex-post Pareto efficient with respect to $R(u)$. For every matching $x \in X$ that is ex-post Pareto efficient with respect to $R(u)$, we have that

$$I_u(x) = U_u(x_{max}) - U_u(x) \leq U_{(u_i', u_{-i})}(x_{max}) - U_{(u_i', u_{-i})}(x) + \varepsilon \leq I_{(u_i', u_{-i})}(x) + \varepsilon. \quad \square$$

The second step of the proof is essentially unchanged from Filos-Ratsikas et al. (2014, Lemma 3): proving that the worst-case loss is obtained when all of the mass of the utility of each individual is near the minimum (0) or near the maximum (1) of the utility function. While Filos-Ratsikas et al. (2014) quantify the ratio, instead of difference, between the realized and optimal welfare, their proof of this step goes through, *mutatis mutandis*. For every $i = 1, \ldots, n$ and every $0 < \varepsilon < \frac{1}{2}$, let $\mathcal{UR}_i^\varepsilon$ be the set of utility functions $u_i \in \mathcal{UR}_i$ whose range is contained in $[0, \varepsilon) \cup (1-\varepsilon, 1]$. For every $0 < \varepsilon < \frac{1}{2}$, let $\mathcal{UR}^\varepsilon \triangleq \mathcal{UR}_1^\varepsilon \times \cdots \times \mathcal{UR}_n^\varepsilon$.

Lemma 9. $\sup_{u \in \mathcal{UR}} I_u\big(\mu(u)\big) = \sup_{u \in \mathcal{UR}^\varepsilon} I_u\big(\mu(u)\big)$ *for every ordinal mechanism* μ*,* $n \geq 2$*, and* $0 < \varepsilon < \frac{1}{2}$*.*

The last step of the proof is to quantify the right-hand side of the expression in Lemma 9. This is where we genuinely encounter the second difference, highlighted above, between the analysis of Filos-Ratsikas et al. (2014) and ours: that they quantify the ratio, instead of difference, between the realized and optimal welfare. While this difference invariably means that our calculation will differ from theirs, we manage to follow a significant part of their calculation, and only diverge after the majority of the heavy lifting. Indeed, we reuse all of the RSD-specific argumentation of Filos-Ratsikas et al. (2014, Lemma 4) verbatim, and only diverge from their analysis in the

algebraic simplification that follows.[28]

Lemma 10. $\lim_{\varepsilon\to 0}\sup_{u\in\mathcal{UR}^\varepsilon} I_u\big(RSD(u)\big) \leq (\ln 2)\cdot n$ *for every* $n \geq 2$.

Finally, Theorem 4 follows from the above three lemmas.

Proof of Theorem 4. Combining Lemmas 8 through 10, we have that

$$\sup_{(X,\succeq)\in\mathcal{M}_n} \hat{I}\big((X,\succeq), RSD(\succeq)\big) \leq \frac{1}{n}\sup_{u\in\mathcal{UR}} I_u\big(RSD(u)\big) = \frac{1}{n}\lim_{\varepsilon\to 0}\sup_{u\in\mathcal{UR}^\varepsilon} I_u\big(RSD(u)\big) \leq \ln 2.$$

□

5.3 FLAW Computation

Complementing the robust guarantee of the previous section, one might wish to compute the FLAW of some matching—such as a realized mechanism output—with respect to specific given preferences. This raises computational concerns, which we discuss and resolve in this section.

To compute the FLAW of a specific allocation in an object allocation problem, letting $u_1, \ldots, u_n$ be vNM utility representations of the individuals' preferences, we need to compute (recalling the notation of the definition of $\hat{I}$ from Section 4.1): (1) the per-individual maximum and minimum Pareto-frontier utilities u_i^{max} and u_i^{min}, and (2) the linear translation $\max_{x'\in X} V(C, x')$. Once the per-individual maximum and minimum Pareto-frontier utilities are known, the linear translation can be readily computed using any of the well-known computationally feasible algorithms for the optimal assignment problem (e.g., Kuhn, 1955; Munkres, 1957). For the per-individual maximum and minimum Pareto-frontier utilities, by Lemma 7 we need to compute, for each individual, her most preferred and least preferred objects among those matched to her in ex-post Pareto efficient matchings. While the former is trivially her favorite object, it is at first glance unclear how one might feasibly compute the latter. Indeed, Saban and Sethuraman (2015) show that it is NP-complete (i.e., computationally intractable) to answer, for a given individual i and object o, the question of whether there exists an ex-post Pareto efficient matching that matches o to i. We nonetheless prove that it is computationally feasible to compute each individual i's least preferred object among those matched to her in ex-post Pareto

[28]The proof of Lemma 10 in the appendix explicitly marks this point of divergence from the analysis of Filos-Ratsikas et al. (2014).

Algorithm 1 Finds the object least preferred by individual $\hat{i}$ among the objects matched to $\hat{i}$ in ex-post Pareto efficient matchings. Runs in time polynomial in n (i.e., computationally feasible).

function FindMinParetoMatch$(n, (\succeq_i)_{i=1}^n, \hat{i})$
  for all object o, from $\hat{i}$'s least preferred to $\hat{i}$'s most preferred, **do**
    if Test$(n, (\succeq_i)_{i=1}^n, \hat{i}, o)$ **then**
      return o
    end if
  end for
end function
function Test$(n, (\succeq_i)_{i=1}^n, \hat{i}, o)$
  $V \leftarrow$ all individuals except $\hat{i}$ and all objects except o
  $E \leftarrow$ all individual–object pairs $(i, o') \in V^2$ s.t. $o' \succ_i o$
  $G \leftarrow (V, E)$
  // The next line is calculable in polynomial time (Hopcroft and Karp, 1973).
  $\sigma \leftarrow$ max-cardinality matching in the bipartite graph G
  if $|\sigma| = n - 1$ **then**
    return True
  else
    return False
  end if
end function

efficient matchings, which in turn allows us to feasibly compute the FLAW of any allocation. Our algorithm is presented in Algorithm 1.

Proposition 1. *For every object allocation problem* $\left(X, (\succeq_i)_{i=1}^n\right) \in \mathcal{M}$ *and individual* $\hat{i} \in \{1, \ldots, n\}$*, the function call* FindMinParetoMatch$(n, (\succeq_i)_{i=1}^n, \hat{i})$ *returns the object least preferred by* $\hat{i}$ *among the objects matched to* $\hat{i}$ *in ex-post Pareto efficient matchings.*

Inspecting Algorithm 1, it might seem at first glance that for Proposition 1 to hold, the function call Test$(n, (\succeq_i)_{i=1}^n, \hat{i}, o)$ must check whether object o is matched to $\hat{i}$ in some ex-post Pareto efficient matching, i.e., this function call must be guaranteed to return True if and only if this is the case. Such a guarantee, however, would have contradicted the above-discussed computational intractability. Rather, the function Test has a somewhat weaker guarantee, which is nonetheless strong enough for Proposition 1 to still hold. First, a one-sided guarantee holds: if o is *not* matched to $\hat{i}$ in any ex-post Pareto efficient matching, then this function call indeed returns

False.

Lemma 11. *For every object allocation problem* $\left(X, (\succeq_i)_{i=1}^n\right) \in \mathcal{M}$*, individual* $\hat{i}$*, and object* o*, if* o *is not matched to* $\hat{i}$ *in any ex-post Pareto efficient matching, then the function call* $\textsc{Test}(n, (\succeq_i)_{i=1}^n, \hat{i}, o)$ *returns False.*

Second, while the other-sided guarantee does not hold—i.e., if o is matched to $\hat{i}$ in some ex-post Pareto efficient matching, then $\textsc{Test}(n, (\succeq_i)_{i=1}^n, \hat{i}, o)$ might return either True or False—this guarantee *does* hold when o is $\hat{i}$'s least preferred object among all those matched to her in ex-post Pareto efficient matchings.

Lemma 12. *For every object allocation problem* $\left(X, (\succeq_i)_{i=1}^n\right) \in \mathcal{M}$*, individual* $\hat{i}$*, and object* o*, if* o *is the object least preferred by* $\hat{i}$ *among the objects matched to* $\hat{i}$ *in ex-post Pareto efficient matchings, then the function call* $\textsc{Test}(n, (\succeq_i)_{i=1}^n, \hat{i}, o)$ *returns True.*

Since Algorithm 1 traverses the objects in $\hat{i}$'s reverse order of preference, these two guarantees suffice: First, as long as the algorithm traverses objects that are never matched to $\hat{i}$ in ex-post Pareto efficient matchings, Test returns False by Lemma 11. Second, once the algorithm reaches an object matched to $\hat{i}$ in some ex-post Pareto efficient matching, this object is $\hat{i}$'s least preferred such object, and then Test returns True by Lemma 12 and the algorithm returns this object and concludes.

6 Related Work

To our knowledge, our paper is the first that, for arbitrary sets of alternatives, axiomatically microfounds a social welfare function that depends *only* on the individuals' ordinal (vNM) preferences over the alternatives, whose *numeric value* (rather than only its induced order) is unique (whether or not up to a global choice of unit of measure) and cannot be arbitrarily additively shifted, and which can be meaningfully compared *across contexts*. We are furthermore not aware of any paper that uses or characterizes the *order* induced by our social welfare function, FLAW.

The question of whether and how utilities of different individuals could be aggregated in order to assess the welfare of society is long-standing in welfare economics and social choice theory, two momentous fields whose comprehensive summary is beyond the scope of this literature review. For many years, economists have been split on whether and to what extent interpersonal comparison of utility is justified. In his

1951 book, "Social Choice and Individual Values," Arrow takes the stand that "interpersonal comparison of utilities has no meaning." Later, Harsanyi (1955) argues to the contrary, for utilitarianism (taking the sum of the individuals' utilities) based on the idea that the individuals' utility functions can be meaningfully compared and aggregated. He expands social welfare functions to randomized alternatives and adopts the vNM axioms (von Neumann and Morgenstern, 1953) for social preferences. He then shows that the social welfare functions that are weighted sums of the individuals' utilities (with weights that may be arbitrarily chosen) are the only ones that are Pareto monotone and satisfy an additional axiom called independence. Following Harsanyi (1955), many other papers construct social welfare functions that take as input a profile of ordinal preferences of the individuals over the given set of alternatives and yield a social ranking of the alternatives. Among these papers, several influential ones assume given utility functions that are exogenously chosen to convey some additional information beyond only the ordinal preferences (e.g., DeMeyer and Plott, 1971; Maskin, 1978; Sen, 1970; Roberts, 1980; Hammond, 2005). Our paper, in contrast to all of the above, does not assume any exogenous information except for the ordinal (vNM) preferences, and yet pinpoints *unique* weights on the individuals' utility functions to be used for aggregation.

A different approach is taken by Kaneko and Nakamura (1979), who define the *Nash social welfare function*. They assume "the existence of a distinguished alternative [...] which represents one of the worst states for all individuals that we may imagine, e.g., we may imagine all the members of the society die," and evaluate social welfare by considering (the product of) the relative increases of the individuals' utilities from this distinguished alternative. The order induced by this function (which is what they axiomatize) is robust to monotone affine transformations of the utility functions that represent the individuals' preferences, while the numeric value depends on the precise choice of utility functions (and can be furthermore arbitrarily monotonically transformed without violating their axioms). Piacquadio (2017) assumes that alternatives are allocations of commodities and introduces an inequality-concerned "equal-preference transfer" axiom. He then shows that if the set of alternatives consists of a bundle that is considered worst by all individuals, the social welfare functions in a class that he terms "opportunity-equivalent utilitarianism" are the only ones that meet this axiom in conjunction with several other axioms. In contrast to both of these papers, our construction *does not* depend on a modeling choice of a

distinguished worst-case alternative or bundle for all members of society.

The closest social preference order to the order induced by FLAW is *relative utilitarianism*. This social preference order is obtained by normalizing each individual's (vNM) utility function to be between 0 and 1 (over all alternatives rather than over the Pareto frontier like FLAW does) and then summing the individuals' normalized utilities to define the social index according to which alternatives are compared. This social preference order received several axiomatic foundations (Karni, 1998; Dhillon and Mertens, 1999; Segal, 2000). Within this literature, the closest paper to ours is Dhillon and Mertens (1999).[29] As already noted by Arrow (1951) (see Footnote 7 on Page 6), the normalization to $[0, 1]$ over the entire set of alternatives makes relative utilitarianism (in contrast to FLAW) sensitive to the addition or deletion of dominated alternatives. Each of the above papers motivates this normalization differently: Karni (1998) uses an impartiality axiom to provide the basis for interpersonal utility comparisons. The other two papers do not assume interpersonally comparable utilities, but rather assume an exogenous meaningful minimum point: Dhillon and Mertens (1999) assume that "the set of alternatives is sufficiently encompassing as to include, besides the actual alternatives of interest, each person's best and worst alternative within the 'universal' set"; Segal (2000) restricts the analysis to allocation of an endowment, and allows for the entire endowment to remain unallocated. Our paper, like the latter two papers, does not assume interpersonally comparable utilities. However, since the normalization of FLAW depends only on the Pareto frontier, which is endogenously determined for each context, we do not need to assume such structure or such richness of the set of alternatives as in the latter two papers. Like all papers preceding ours that we are aware of, all three of the above papers axiomatize an order rather than a cardinally-meaningful real-valued index (i.e., any cardinal representation of the social preference order that they axiomatize can be arbitrarily monotonically transformed), and like most previous papers, this order does not allow for comparison across contexts.[30]

[29] Segal (2000) takes the alternatives to be different ways to distribute a divisible endowment between the individuals. In Karni (1998), the axiomatized social preference order is over allocations that may allocate a different (randomized) alternative to each individual. Dhillon and Mertens (1999) axiomatize the relative utilitarianism social preference order over a general fixed set of alternatives.

[30] Among our axioms, the only one that is violated by all social inefficiency functions that induce the relative utilitarianism order is IIA. Let IIA$'$ be the axiom obtained from IIA by replacing the requirement for having the same point of minimal expectations with that of having the same worst point—defined such that $C = (X, (\succeq_i)_{i=1}^n)$ and $C' = (X', (\succeq_i|_{\Delta(X')})_{i=1}^n)$ have *the same worst point*

Another paper that axiomatically motivates a specific normalization of utility functions is that of Fleurbaey and Zuber (2021). In this paper, an alternative is a specification of a positive (single-dimensional) consumption level for each individual "living" in that particular alternative; the set of alternatives is then the set of all such consumption profiles. Individuals' preferences are over their own consumption. Assuming differentiable vNM utility functions, this paper introduces axioms that justify the normalization of the utilities of all individuals so that their marginal utilities at a poverty line are the same. Our paper, in contrast, does not assume any structure on the set of alternatives (or on the vNM utility functions).

Fleurbaey and Tadenuma (2014) develop a social-choice theory that allows for comparison across societies, resembling our inter-context comparisons. In their paper, a *social situation* consists of some population, preferences for each individual in that population over her own consumption, and a consumption profile that specifies a consumption bundle for each individual in that population. Fleurbaey and Tadenuma (2014) seek to ordinally compare such social situations. (They are not interested in comparison of lotteries over social situations.) On a conceptual level, while their social situations technically resemble context–alternative pairs in our paper, they seek to answer the question of which of two social situations is preferred (e.g., which country is better off), while we seek to quantify, in every context, the social inefficiency of an alternative *as contextualized by the context in question*. On a technical level, this translates in their setting to each individual always having preferences over the same, structured set of consumption bundles for that individual; in contrast, our sets of alternatives are completely structureless (e.g., not even a metric or topological space) and arbitrarily change between contexts (indeed, in our setting in contrast to theirs, the social inefficiency of an alternative in some context might change if new possible alternatives are added to the context). In both Fleurbaey and Tadenuma (2014) and Fleurbaey and Zuber (2021), and in contrast to our paper, the set of alternatives

if $\min_{\succeq_i} X' \sim_i \min_{\succeq_i} X$ for every $i = 1, \ldots, n$. Our axiomatic system with IIA replaced by IIA$'$ gives rise, through an analysis similar to ours, to a social inefficiency function that induces the relative utilitarianism order. Beyond being a new characterization of relative utilitarianism (distinct from that of Dhillon and Mertens, 1999; e.g., working for any numbers of individuals and alternatives), the resulting social inefficiency function is cardinally unique (i.e., pinpointed among all the social inefficiency functions that induce the relative utilitarianism order, which are discussed in Section 2) up to a global choice of unit of measure, allowing for intra- and inter-context comparability. (The cardinal information in the arising social inefficiency function can be interpreted as in Section 4.2, replacing the point of minimal expectations with the worst point in the definition of compatible contexts.)

is unbounded, and therefore, since the preferences of individuals in these papers are monotone, neither an optimal or ideal point nor a Pareto frontier of alternatives (consumption profiles) exists. Furthermore, like all preceding papers of which we are aware, Fleurbaey and Tadenuma (2014) characterize orders that satisfy various axioms rather than characterize a cardinally-meaningful real-valued index.

A different stream of literature follows the Nash bargaining solution (Nash, 1953) and, in a setting where society consists of two individuals, searches for a socially preferred *alternative* rather than a social preference order. In this literature, it is usually assumed that such an alternative should be unique. Roth (1977) considers different "Independence of Irrelevant Alternatives (IIA)" axioms and the bargaining solutions they obtain. While Nash's bargaining solution is obtained for the axiom of "IIA other than the disagreement point," Roth (1977) shows that the bargaining solution of Kalai and Smorodinsky (1975) (see also Raiffa, 1953) is obtained for the axiom of "IIA other than both the disagreement point and the ideal point." Roth (1977) also introduces the (endogenous) point of minimal expectations, and characterizes the bargaining solutions that meet the axiom of "IIA other than the point of minimal expectations" or the axiom of "IIA other than the ideal point." However, he does not consider, as we do, the axiom of "IIA other than both the point of minimal expectations and the ideal point." Nalebuff (2021) looks for "perspective invariant" bargaining solutions in the sense of obeying a modified version of Nash's axioms with respect to the disagreement point as well as a flipped version of the same axioms with respect to the ideal point. He normalizes so that the (exogenous) disagreement point is $(0, 0)$ and the (endogenous) ideal point is $(1, 1)$ and exhibits such a modification of Nash's axioms that gives a bargaining solution that maximizes the sum of the two individuals' utilities.

Computer scientists have been interested in measuring the inefficiency of a given alternative compared to the optimal efficiency. Within games, Koutsoupias and Papadimitriou (1999) define the Price of Anarchy as the ratio between the worst-case performance at any equilibrium and the optimal performance at any strategy profile—both measured using some welfare (or cost) function (such as the cumulative driving time of all individuals, as studied by Roughgarden and Tardos, 2002). Following that paper, the Price of Anarchy (many times using the sum of exogenously normalized utilities as the welfare function) has been adopted not only in computer science but also in operations research, industrial engineering, and other areas as a measure of

the inefficiency of an interactive environment. For example, the Price of Anarchy has been applied to classical operations research areas such as supply chains (e.g., Perakis and Roels, 2007), queuing (see Ghosh and Hassin, 2021 for a survey), load balancing (e.g., Ayesta et al., 2010), and scheduling (e.g., Hoeksma and Uetz, 2019). In contrast to the Price of Anarchy, FLAW does not require exogenously given welfare or cost functions (or exogenously normalized utility functions) that are cardinally comparable within and between contexts.

7 Discussion

In this paper, we axiomatically define a *cardinal* social welfare function, or more specifically a cardinal social inefficiency function, which provides a unique cardinal comparison between alternatives, even across contexts. Like any axiomatization, our axioms can of course be debated. Changing one of them might lead to a very different social inefficiency function. We do not claim that our set of axioms is necessarily the only possible one or even the best possible one. Our aim is to demonstrate the applicability of the axiomatic approach for microfounding an economically meaningful quantitative comparison of the social inefficiency of different alternatives within and across contexts.

Beyond demonstrating the usefulness of our social inefficiency function, FLAW, and the types of theorems that it can facilitate, an additional purpose of our application section, Section 5, is to demonstrate the relative ease with which existing computer-science analyses can be adapted to yield FLAW results. For this reason, we make no attempt to develop further techniques to get a better fraction than $\frac{1}{2\ln 2}$ in Theorem 2. As is common in computer science, approximation theorems allow for progress to be made gradually, and we hence leave the question of tighter bounds in the object allocation application as an interesting open problem.[31]

While we demonstrate the usefulness of FLAW within noncooperative game theory, we emphasize that its applicability is potentially much broader. Indeed, it can be

[31] Of course, ours is not the first economics paper that derives an approximation result and leaves potentially tightening that result as an open problem. For example (for approximately optimizing a monopolist's revenue rather than an aggregate, social welfare function), Hart and Nisan (2017) prove that for a two-item monopolist, optimally pricing each item separately guarantees at least 50% of the optimal revenue, a guarantee that was later improved to $\sim 62\%$ by Hart and Reny (2019), with a tight bound on this guarantee still remaining an open problem.

applied to any setting that features a set of alternatives and preferences of individuals over them. For example, within a bargaining setting, FLAW essentially frames negotiations as, instead of bargaining with respect to an exogenous disagreement point *à la* Nash, bargaining with respect to an *endogenous* agreement *region* (the Pareto frontier). This might be seen as in line with many negotiation texts for executives that discuss the "Zone of Possible Agreement (ZOPA)" as a key object that frames negotiations.[32]

As noted in the introduction, our analysis in this paper is a demonstration of a potentially larger program of providing a microfoundation for what computer scientists call *approximation theorems*. These theorems, which ascertain that some quantity is at least a specific fraction of some reference quantity, are not inherently computational. Nonetheless, they are emblematic of theoretical computer science, and as such, tools and techniques for deriving them have been developed over many decades. While we focus on cardinally measuring (in)efficiency, one might imagine axiomatically defining cardinal measures of other objectives of interest such as (un)fairness, (in)equity, and so forth. Such cardinal measures, of (in)efficiency or otherwise, can also be used beyond bounding the loss in a certain objective in specific alternatives, e.g., for deriving *cardinal* comparative statics: Quantifying not only the direction, but also the magnitude, of change in a certain objective as certain model parameters are changed. Overall, our paper demonstrates a new direction for providing a sound economic foundation for approximation theorems and other quantitative theorems as we enter the second quarter-century of intense EconCS cross-pollination.

References

A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.

M. Akbarpour, S. D. Kominers, K. M. Li, S. Li, and P. Milgrom. Algorithmic mechanism design with investment. *Econometrica*, 91(6):1969–2003, 2023.

K. J. Arrow. *Social Choice and Individual Values*. John Wiley & Sons, 1951.

[32]See, e.g., https://online.hbs.edu/blog/post/understanding-zopa. We thank Yuval Gonczarowski for introducing us to this notion.

U. Ayesta, O. Brun, and B. J. Prabhu. Price of anarchy in non-cooperative load balancing. In *Proceedings of the 29th IEEE Conference on Computer Communications (INFOCOM)*, pages 1–5, 2010.

A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.

D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968.

F. DeMeyer and C. R. Plott. A welfare function using "relative intensity" of preference. *The Quarterly Journal of Economics*, 85(1):179–186, 1971.

A. Dhillon and J.-F. Mertens. Relative utilitarianism. *Econometrica*, 67(3):471–498, 1999.

P. A. Diamond. Cardinal welfare, individualistic ethics, and interpersonal comparison of utility: Comment. *Journal of Political Economy*, 75(5):765–766, 1967.

D. Dillenberger and U. Segal. Allocation mechanisms with mixture-averse preferences. *American Economic Review*, 115(12):4523–4547, 2025.

A. Filos-Ratsikas, S. K. S. Frederiksen, and J. Zhang. Social welfare in one-sided matchings: Random priority and beyond. In *Proceedings of the 7th International Symposium on Algorithmic Game Theory (SAGT)*, pages 1–12, 2014.

M. Fleurbaey and K. Tadenuma. Universal social orderings: An integrated theory of policy evaluation, inter-society comparisons, and interpersonal comparisons. *Review of Economic Studies*, 81(3):1071–1101, 2014.

M. Fleurbaey and S. Zuber. Fair utilitarianism. *American Economic Journal: Microeconomics*, 13(2):370–401, 2021.

S. Ghosh and R. Hassin. Inefficiency in stochastic queueing systems with strategic customers. *European Journal of Operational Research*, 295(1):1–11, 2021.

P. J. Hammond. Utility invariance in non-cooperative games. In U. Schmidt and S. Traub, editors, *Advances in Public Economics: Utility, Choice and Welfare: A Festschrift for Christian Seidl*, pages 31–50. Springer, 2005.

J. C. Harsanyi. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63(4):309–321, 1955.

S. Hart and N. Nisan. Approximate revenue maximization with multiple items. *Journal of Economic Theory*, 172:313–347, 2017.

S. Hart and P. J. Reny. The better half of selling separately. *ACM Transactions on Economics and Computation*, 7(4):18:1–18:18, 2019.

J. Hartline, D. Hoy, and S. Taggart. Price of anarchy for auction revenue. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pages 693–710, 2014.

R. Hoeksma and M. Uetz. The price of anarchy for utilitarian scheduling games on related machines. *Discrete Optimization*, 31:29–39, 2019.

J. E. Hopcroft and R. M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on Computing*, 2(4):225–231, 1973.

D. Hoy, S. Taggart, and Z. Wang. A tighter welfare guarantee for first-price auctions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 132–137, 2018.

Y. Jin and P. Lu. First price auction is $1-1/e^2$ efficient. *Journal of the ACM*, 70(5): 36:1–36:86, 2023.

E. Kalai and M. Smorodinsky. Other solutions to Nash's bargaining problem. *Econometrica*, 43(3):513–518, 1975.

M. Kaneko and K. Nakamura. The Nash social welfare function. *Econometrica*, 47 (2):423–435, 1979.

E. Karni. Impartiality: Definition and representation. *Econometrica*, 66(6):1405–1415, 1998.

E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.

H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1–2):83–97, 1955.

R. D. Luce and H. Raiffa. *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons, 1957.

M. Manea. Asymptotic ordinal inefficiency of random serial dictatorship. *Theoretical Economics*, 4(2):165–197, 2009.

A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

E. Maskin. A theorem on utilitarianism. *The Review of Economic Studies*, 45(1): 93–96, 1978.

R. P. McAfee. Coarse matching. *Econometrica*, 70(5):2025–2034, 2002.

J. Munkres. Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1):32–38, 1957.

B. Nalebuff. A perspective-invariant approach to Nash bargaining. *Management Science*, 67(1):577–593, 2021.

J. Nash. Two-person cooperative games. *Econometrica*, 21(1):128–140, 1953.

Z. Neeman. The effectiveness of English auctions. *Games and Economic Behavior*, 43(2):214–238, 2003.

G. Perakis and G. Roels. The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Management Science*, 53(8):1249–1268, 2007.

P. G. Piacquadio. A fairness justification of utilitarianism. *Econometrica*, 85(4): 1261–1276, 2017.

M. Rabin. Incorporating fairness into game theory and economics. *American Economic Review*, 83(5):1281–1302, 1993.

H. Raiffa. Arbitration schemes for generalized two-person games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games, Volume II*, pages 361–388. Princeton University Press, 1953.

K. W. S. Roberts. Interpersonal comparability and social choice theory. *The Review of Economic Studies*, 47(2):421–439, 1980.

W. P. Rogerson. Simple menus of contracts in cost-based procurement and regulation. *American Economic Review*, 93(3):919–926, 2003.

A. E. Roth. Independence of irrelevant alternatives, and solutions to Nash's bargaining problem. *Journal of Economic Theory*, 16(2):247–251, 1977.

A. E. Roth. Repugnance as a constraint on markets. *Journal of Economic Perspectives*, 21(3):37–58, 2007.

T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49 (2):236–259, 2002.

D. Saban and J. Sethuraman. The complexity of computing the random priority allocation matrix. *Mathematics of Operations Research*, 40(4):1005–1014, 2015.

U. Segal. Let's agree that all dictatorships are equally bad. *Journal of Political Economy*, 108(3):569–589, 2000.

A. Sen. Interpersonal aggregation and partial comparability. *Econometrica*, 38(3): 393–409, 1970.

V. Syrgkanis and E. Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC)*, pages 211–220, 2013.

J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, third edition, 1953.

L. Zhou. On a conjecture by Gale about one-sided matching problems. *Journal of Economic Theory*, 52:123–135, 1990.

A Proofs Omitted from Section 4

Proof of Lemma 1. Let $C = \big(X, (\succeq_i)_{i=1}^n\big)$ be a context and let $x, y \in \Delta(X)$ such that $x \sim_i y$ for every $i = 1, \ldots, n$. By Pareto monotonicity (applied twice) we have both that $I(C, x) \leq I(C, y)$ and that $I(C, x) \geq I(C, y)$, and hence $I(C, x) = I(C, y)$, and so I satisfies Pareto indifference. □

Proof of Lemma 2. Let $C = \big(X, (\succeq_i)_{i=1}^n\big)$ be a context and let $X' \subset X$ be a set of alternatives such that $\Delta(X')$ weakly dominates $X \setminus X'$. Let $C' \triangleq \big(X', (\succeq_i|_{\Delta(X')})_{i=1}^n\big)$.

We first claim that there exists $\hat{x} \in \Delta(X')$ such that $I(C, \hat{x}) = 0$. To see this, first note that by feasibility there exists $x \in \Delta(X)$ (not necessarily in $\Delta(X')$) such that $I(C, x) = 0$. Write $x = \sum_{\ell=1}^L \alpha_\ell \cdot x_\ell$ where $\alpha_1, \ldots, \alpha_L \in [0, 1]$ and $x_1, \ldots, x_L \in X$. By assumption, for every $\ell = 1, \ldots, L$ there exist $\hat{x}_\ell \in \Delta(X')$ such that $\hat{x}_\ell \succeq_i x_\ell$ for every $i = 1, \ldots, n$. Therefore, $\hat{x} \triangleq \sum_{\ell=1}^L \alpha_\ell \cdot \hat{x}_\ell \in \Delta(X')$ satisfies $\hat{x} \succeq_i x$ for every $i = 1, \ldots, n$. By Pareto monotonicity, $I(C, \hat{x}) = 0$ as claimed.

By feasibility, there exists $\hat{x}' \in \Delta(X')$ such that $I(C', \hat{x}') = 0$. By construction of C' and since $\Delta(X')$ weakly dominates $X \setminus X'$, the contexts C and C' have the same ideal point and the same point of minimal expectations. Hence, by IIA we have that

$$0 \leq I(C, \hat{x}') = I(C, \hat{x}') - I(C, \hat{x}) = I(C', \hat{x}') - I(C', \hat{x}) = -I(C', \hat{x}) \leq 0,$$

and therefore $I\big(C', \hat{x}\big) = 0$. Now, for every $x \in \Delta(X')$, by IIA we have that

$$I(C', x) = I(C', x) - I(C', \hat{x}) = I(C, x) - I(C, \hat{x}) = I(C, x). \qquad \square$$

Proof of Lemma 3. Let $C^1 = \big(X^1, (\succeq_i^1)_{i=1}^n\big)$ and $C^2 = \big(X^2, (\succeq_i^2)_{i=1}^n\big)$ be two isomorphic contexts and let ϕ be an isomorphism between C^1 and C^2.

We first prove that for the special case in which $X^1 \cap X^2 = \emptyset$, it indeed holds that $I(C^1, x) = I\big(C^2, \phi(x)\big)$ for every $x \in \Delta(X^1)$. Let $X \triangleq X^1 \cup X^2$ and define a context $C = \big(X, (\succeq_i)_{i=1}^n\big)$ by defining $\succeq_i$ for every i using the isomorphism ϕ as follows: Let $x, y \in \Delta(X)$. By definition of X, there exist $\alpha \in [0, 1]$, $x^1 \in \Delta(X^1)$, and $x^2 \in \Delta(X^2)$ such that $x = \alpha \cdot x^1 + (1 - \alpha) \cdot x^2$. Similarly, there exist $\beta \in [0, 1]$, $y^1 \in \Delta(X^1)$, and $y^2 \in \Delta(X^2)$ such that $y = \beta \cdot y^1 + (1 - \beta) \cdot y^2$. We define $\succeq_i$ so that $x \succeq_i y$ if and only if $\alpha \cdot \phi(x^1) + (1 - \alpha) \cdot x^2 \succeq_i^2 \beta \cdot \phi(y^1) + (1 - \beta) \cdot y^2$. (Equivalently, if and only if $\alpha \cdot x^1 + (1 - \alpha) \cdot \phi^{-1}(x^2) \succeq_i^1 \beta \cdot y^1 + (1 - \beta) \cdot \phi^{-1}(y^2)$.) Note that $C^1 = \big(X^1, (\succeq_i|_{\Delta(X^1)})_{i=1}^n\big)$ and that $C^2 = \big(X^2, (\succeq_i|_{\Delta(X^2)})_{i=1}^n\big)$.

Note that for every $x \in \Delta(X^1)$, we have that $x \sim_i \phi(x)$ for every $i = 1, \ldots, n$. By Pareto indifference, we therefore have that $I(C, x) = I\big(C, \phi(x)\big)$. Therefore, by invariance to dominated alternatives (applied twice), we have that

$$I(C^1, x) = I(C, x) = I\big(C, \phi(x)\big) = I\big(C^2, \phi(x)\big),$$

completing the proof for the special case in which $X^1 \cap X^2 = \emptyset$.

For the general case in which possibly $X^1 \cap X^2 \neq \emptyset$, let X^3 be a set such that $|X^3| = |X^1|$ (and hence $|X^3| = |X^2|$) and $(X^1 \cup X^2) \cap X^3 = \emptyset$. Since $|X^3| = |X^1|$, there exists a bijection ψ between X^3 and X^1. We extend ψ in the natural way to be defined over $\Delta(X^3)$, making it a bijection between $\Delta(X^3)$ and $\Delta(X^1)$. For every $i = 1, \ldots, n$, we define $\succeq_i^3$ over $\Delta(X^3)$ using the bijection ψ as follows: Let $x, y \in \Delta(X^3)$. We define $\succeq_i^3$ so that $x \succeq_i^3 y$ if and only if $\psi(x) \succeq_i^1 \psi(y)$. Note that ψ is an isomorphism between $C^3 \triangleq \big(X^3, (\succeq_i^3)_{i=1}^n\big)$ and C^1, and therefore $\phi \circ \psi$ is an isomorphism between C^3 and C^2. For every $x \in \Delta(X^1)$, we therefore have by the above special case (applied twice) that $I(C^1, x) = I\big(C^3, \psi^{-1}(x)\big) = I\big(C^2, \phi(x)\big)$, and so I satisfies neutrality. □

Proof of Lemma 4. The "if" direction is straightforward to check. We prove the "only if" direction.

Let $C = \big(X, (\succeq_i)_{i=1}^n\big)$ be a context and let $u_C^1, \ldots, u_C^n$ be vNM utility representations of $\succeq_1, \ldots, \succeq_n$. Denote $u_C = (u_C^1, \ldots, u_C^n)$ and let $U \triangleq u_C\big(\Delta(X)\big)$. By Pareto indifference (Lemma 1), $I(C, x)$ depends on x only through the values $u_C^1(x), \ldots, u_C^n(x)$, i.e., there exists a function $f : U \to \mathbb{R}_{\geq 0}$ such that $I(C, x) = f\big(u_C^1(x), \ldots, u_C^n(x)\big)$. By expected inefficiency and since all the functions u_C^i are affine, f preserves convex combinations on U; hence, f is an affine function on U. Therefore, f extends to an affine function on the affine hull of U, which in turn extends to an affine function on $\mathbb{R}^n$. Therefore, f can be written in the form

$$f(u) = a - c \cdot u \tag{3}$$

for some constants $a \in \mathbb{R}$ and $c = (c^1, \ldots, c^n) \in \mathbb{R}^n$. We emphasize that since U is not necessarily of full dimension, the choice of a, c need not be unique, and we fix one such choice arbitrarily. By Pareto monotonicity, f is strictly decreasing, in the sense that $f(u) > f(u')$ for every $u, u' \in U$ such that $u \leq u'$ and $u \neq u'$.

Let $D \triangleq \mathrm{Span}(U - U) \subseteq \mathbb{R}^n$ be the subspace spanned by all directions between utilities in U. Since the set $U - U$ is convex (by convexity of U) and symmetric, contains 0, and spans D, it is absorbing in D. We claim that there exists $\varepsilon > 0$ such that for every $d \in D \setminus \{0\}$ such that $d \geq 0$, it holds that $c \cdot d \geq \varepsilon \mathbf{1} \cdot d$ (where $\mathbf{1} \triangleq (1, \ldots, 1) \in \mathbb{R}^n$). Let $V \triangleq \{d \in D \mid d \geq 0 \ \& \ \mathbf{1} \cdot d = 1\}$. If $V = \emptyset$, then one may choose any $\varepsilon > 0$ and the claim vacuously holds. Assume, therefore, that $V \neq \emptyset$. For every $d \in V$, because $d \in D$ and since $U - U$ is absorbing in D, there exist $u, u' \in U$ such that $u - u' = \delta d$ for some $\delta > 0$. Since $d \neq 0$, we have that $u \neq u'$. Since $d \geq 0$, we furthermore have that $u \geq u'$. Therefore, since f is strictly decreasing, $\delta c \cdot d = c \cdot (u - u') = f(u') - f(u) > 0$ and hence $c \cdot d > 0$. Since $c \cdot d > 0$ for every $d \in V$ and since $d \mapsto c \cdot d$ is a continuous function, by compactness of V we have that there exists $\varepsilon > 0$ such that $c \cdot d \geq \varepsilon$ for every $d \in V$. Since $\mathbf{1} \cdot d = 1$ for every $d \in V$, we therefore have that $c \cdot d \geq \varepsilon \mathbf{1} \cdot d$ for every $d \in V$, and since this inequality is homogeneous, we have that it holds for every $d \in D \setminus \{0\}$ such that $d \geq 0$ (regardless of whether or not $\mathbf{1} \cdot d = 1$), as claimed.

Let A be a matrix whose columns span $D^\perp$. We claim that there exists a vector λ such that $A\lambda \geq \varepsilon \mathbf{1} - c$. Assume for contradiction that this is not the case. Then, by Farkas's lemma, there exists $d \geq 0$ such that both $A^\top d = 0$ and $(\varepsilon \mathbf{1} - c) \cdot d > 0$. By the latter, $d \neq 0$. Since $A^\top d = 0$, by definition of A we have that $d \in (D^\perp)^\perp = D$. Therefore, since $d \geq 0$, by the previous argument we have that $c \cdot d \geq \varepsilon \mathbf{1} \cdot d$, contradicting the conclusion of Farkas's lemma that $(\varepsilon \mathbf{1} - c) \cdot d > 0$. Therefore, there indeed exists a vector λ such that $A\lambda \geq \varepsilon \mathbf{1} - c$, as claimed.

Let $c_C \triangleq c + A\lambda$. By definition of λ, we have that $c_C \geq \varepsilon \mathbf{1} > 0$. By definition of A, we have that $A\lambda \in D^\perp$. Therefore, $(A\lambda) \cdot u$ is constant for all $u \in U$. Set $a_C \triangleq a + (A\lambda) \cdot u$ for some (equivalently, every) $u \in U$. Therefore, for every $u \in U$, we have that:

$$f(u) = a - c \cdot u = a + (A\lambda) \cdot u - (c + A\lambda) \cdot u = a_C - c_C \cdot u.$$

Therefore, $I(C, x) = a_C - c_C \cdot u_C(x)$ for every $x \in \Delta(X)$.

Since by definition of a social inefficiency function, $I(C, x) \geq 0$ for every $x \in X$, we have that $a_C \geq \max_{x' \in X} \{c_C \cdot u_C(x')\}$. By feasibility, there exists $x \in \Delta(X)$ such that $a_C - c_C \cdot u_C(x) = 0$; hence, there exists a pure alternative x' in the support of x such that $a_C - c_C \cdot u_C(x') \leq 0$. Therefore, $a_C \leq \max_{x' \in X} \{c_C \cdot u_C(x')\}$. By both of

these, $a_C = \max_{x' \in X} \{c_C \cdot u_C(x')\}$. $\square$

Proof of Lemma 5. The "if" direction is straightforward to check. We prove the "only if" direction.

Let $n \geq 2$. Let $\hat{X}_n \triangleq \{0, 1, \ldots, n\}$, and define a context $\hat{C}_n \triangleq \left(\hat{X}_n, (\hat{\succeq}_i)_{i=1}^n\right)$ by defining $\hat{\succeq}_i$ for every individual $i = 1, \ldots, n$ through the vNM utility representation $\hat{u}_i$ that for every $x \in \hat{X}_n$ satisfies

$$\hat{u}_i(x) \triangleq \begin{cases} 1 & x = i, \\ 0 & x \neq i. \end{cases}$$

(In the discussion preceding Lemma 5, this context is referred to as the "canonical context with n individuals.") By Lemma 4, there exist constants $c^1_{\hat{C}_n}, \ldots, c^n_{\hat{C}_n} > 0$ such that $I(\hat{C}_n, x) = \max_{x' \in \hat{X}_n} \left\{\sum_{i=1}^n c^i_{\hat{C}_n} \hat{u}_i(x')\right\} - \sum_{i=1}^n c^i_{\hat{C}_n} \hat{u}_i(x)$ for every $x \in \Delta(\hat{X}_n)$. Denote $c^i_n \triangleq c^i_{\hat{C}_n}$ for every $i = 1, \ldots, n$.

Let $C = \left(X, (\succeq_i)_{i=1}^n\right)$ be a context with n individuals. By neutrality (Lemma 3, where Pareto indifference is by Lemma 1 and invariance to dominated alternatives is by Lemma 2), assume without loss of generality that $X \cap \hat{X}_n = \emptyset$.

Let $X' \triangleq X \cup \hat{X}_n$ and define a context $C' \triangleq \left(X', (\succeq'_i)_{i=1}^n\right)$ where for every $i = 1, \ldots, n$, the preference $\succeq'_i$ is obtained from $\succeq_i$ by stipulating for every $x \in \hat{X}_n$ that $x \sim'_i \max_{\succeq_i} X$ if $\hat{u}_i(x) = 1$ and $x \sim'_i \min_{\succeq_i} \text{Pareto}(C)$ if $\hat{u}_i(x) = 0$. Note that $\hat{C}_n = \left(\hat{X}_n, (\succeq'_i|_{\Delta(\hat{X}_n)})_{i=1}^n\right)$ as well as that $\hat{C}_n$ and C' have the same ideal point and the same point of minimal expectations.

For every $i = 1, \ldots, n$, choose a vNM utility representation u'_i of $\succeq'_i$ that coincides with $\hat{u}_i$ over $\hat{X}_n$. By Lemma 4, there exist constants $c^1_{C'}, \ldots, c^n_{C'} > 0$ such that $I(C', x) = \max_{x' \in X'} \left\{\sum_{i=1}^n c^i_{C'} u'_i(x')\right\} - \sum_{i=1}^n c^i_{C'} u'_i(x)$ for every $x \in \Delta(X')$. For every $i = 1, \ldots, n$, by IIA we have that

$$c^i_{C'} = I(C', 0) - I(C', i) = I(\hat{C}_n, 0) - I(\hat{C}_n, i) = c^i_{\hat{C}_n} = c^i_n.$$

For every $i = 1, \ldots, n$, let u_i be some vNM utility representation of $\succeq_i$. Observing that $\Delta(X)$ weakly dominates $\hat{X}_n$ in C', by invariance to dominated alternatives (Lemma 2), for every $x \in \Delta(X)$ we have that

$$I(C,x) = I(C',x) = \max_{x' \in X} \left\{ \sum_{i=1}^{n} c^i_{C'} u'_i(x') \right\} - \sum_{i=1}^{n} c^i_{C'} u'_i(x) =$$
$$\max_{x' \in X} \left\{ \sum_{i=1}^{n} c^i_n u'_i(x') \right\} - \sum_{i=1}^{n} c^i_n u'_i(x) =$$
$$\max_{x' \in X} \left\{ \sum_{i=1}^{n} c^i_n \cdot \frac{u_i(x') - u^{min}_i}{u^{max}_i - u^{min}_i} \right\} - \sum_{i=1}^{n} c^i_n \cdot \frac{u_i(x) - u^{min}_i}{u^{max}_i - u^{min}_i}. \qquad \square$$

Proof of Lemma 6. The "if" direction follows by construction of $\hat{I}$. We prove the "only if" direction.

Let $n \geq 2$. For every $i = 1, \ldots, n$, let $c^i_n > 0$ be the constant guaranteed to exist by Lemma 5. Let $c_n \triangleq n \cdot c^1_n$. Let $\hat{C}_n = \big(\hat{X}_n, (\hat{\succeq}_i)_{i=1}^n\big)$ be as in the proof of Lemma 5.

Let $i \in \{2, \ldots, n\}$. Let π be a permutation on $\{1, \ldots, n\}$ such that $\pi(1) = i$ and let $C' \triangleq \big(\hat{X}_n, (\hat{\succeq}_{\pi(i)})_{i=1}^n\big)$. By Lemma 5 (applied twice) and by anonymity,

$$c^i_n = I(\hat{C}_n, 0) - I(\hat{C}_n, i) = I(C', 0) - I(C', i) = c^1_n = \frac{c_n}{n}.$$

Let $C = \big(X, (\succeq_i)_{i=1}^n\big)$ be a context with n individuals. For every $i = 1, \ldots, n$, let u_i be some vNM utility representation of $\succeq_i$. Since we have shown that $c^i_n = \frac{c_n}{n}$ for all $i = 1, \ldots, n$, by Lemma 5 we have for every $x \in \Delta(X)$ that

$$I(C,x) = \max_{x' \in X} \left\{ \sum_{i=1}^{n} \frac{c_n}{n} \frac{u_i(x') - u^{min}_i}{u^{max}_i - u^{min}_i} \right\} - \sum_{i=1}^{n} \frac{c_n}{n} \frac{u_i(x) - u^{min}_i}{u^{max}_i - u^{min}_i} = c_n \cdot \hat{I}(C,x). \qquad \square$$

Proof of Theorem 1. The "if" direction follows by construction of $\hat{I}$. We prove the "only if" direction. (For the "furthermore" part, see Section 4.3.) For every $n \geq 2$, let $c_n > 0$ be the constant guaranteed to exist by Lemma 6. Let $c \triangleq c_2$. By Lemma 6, it suffices to prove for all integers $n \geq 3$ that $c_n = c$.

Let $n \geq 3$, and let $\hat{C}_n$ and $\hat{C}_2$ be as in the proof of Lemma 5. By definition of $\hat{I}$, by Lemma 6 (applied four times), and by population-size stability (applied twice),

$$c_n = c_n \cdot n\big(\hat{I}(\hat{C}_n, 0) - \hat{I}(\hat{C}_n, 1)\big) = n\big(I(\hat{C}_n, 0) - I(\hat{C}_n, 1)\big) =$$
$$n\left(I\big(\hat{C}_n \otimes \hat{C}_n, (0,0)\big) - I\big(\hat{C}_n \otimes \hat{C}_n, (1,1)\big)\right) =$$
$$c_{2n} \cdot n\left(\hat{I}\big(\hat{C}_n \otimes \hat{C}_n, (0,0)\big) - \hat{I}\big(\hat{C}_n \otimes \hat{C}_n, (1,1)\big)\right) = c_{2n} =$$

$$c_{2n} \cdot 2\left(\hat{I}\big(\otimes_{j=1}^{n}\hat{C}_2,(0,\ldots,0)\big) - \hat{I}\big(\otimes_{j=1}^{n}\hat{C}_2,(1,\ldots,1)\big)\right) =$$
$$2\left(I\big(\otimes_{j=1}^{n}\hat{C}_2,(0,\ldots,0)\big) - I\big(\otimes_{j=1}^{n}\hat{C}_2,(1,\ldots,1)\big)\right) =$$
$$2\big(I(\hat{C}_2,0) - I(\hat{C}_2,1)\big) = c_2 \cdot 2\big(\hat{I}(\hat{C}_2,0) - \hat{I}(\hat{C}_2,1)\big) = c_2 = c. \qquad \square$$

B Proofs Omitted from Section 5

Proof of Lemma 7. The "if" direction is immediate: An alternative on the Pareto frontier is not strictly dominated by any alternative in $\Delta(X)$, and therefore also not strictly dominated by any alternative in X, and hence it is ex-post Pareto efficient.

For the "only if" direction, we need to prove that for every ex-post Pareto efficient matching $x \in X$ and for every $y \in \Delta(X)$ such that $y \succeq_i x$ for all $i = 1,\ldots,n$, it holds that $y = x$. We prove this by induction over the number of individuals n. The case of $n = 1$ is immediate. Assume that $n \geq 2$ and that the claim holds for $n-1$.

Since x is ex-post Pareto efficient, there exists $\hat{i} \in \{1,\ldots,n\}$ that is matched to her favorite object in x (otherwise, there would exist a Pareto-improving trading cycle, contradicting the ex-post Pareto efficiency of x). Since $y \succeq_{\hat{i}} x$ and since the individuals' preferences over the objects are strict, $\hat{i}$ is matched with probability 1 to her favorite object in y as well. Let $N' \triangleq \{1,\ldots,n\} \setminus \{\hat{i}\}$ and let O' be the set of all objects except for the one matched to $\hat{i}$ by x. Let X' be the set of all perfect matchings between N' and O', and denote by $x' \in X'$ and $y' \in \Delta(X')$ the respective restrictions of x and y to N'. Note that x' is ex-post Pareto efficient with respect to $R\big((\succeq_i |_{\Delta(X')})_{i\in N'}\big)$ and that $y' \succeq_i x'$ for every $i \in N'$. By the induction hypothesis, $y' = x'$. Therefore, $y = x$. $\square$

Proof of Theorem 3. We say that μ is *anonymous* if for every profile of rankings $R = (R_i)_{i=1}^n \in \mathcal{R}^n$ and for every permutation π on $\{1,\ldots,n\}$, it holds that $\pi^{-1}\big(\mu(\pi(R))\big) = \mu(R)$, where $\pi(R) \triangleq (R_{\pi(1)},\ldots,R_{\pi(n)})$ and where $\pi^{-1}\big(\mu(\pi(R))\big)$ is defined so that if $\mu\big(\pi(R)\big) = \sum_{\ell=1}^{L} p_\ell \cdot x_\ell$ where $x_\ell \in X$ for all $1 \leq \ell \leq L$, then $\pi^{-1}\big(\mu(\pi(R))\big) \triangleq \sum_{\ell=1}^{L} p_\ell \cdot \pi^{-1}(x_\ell)$, where $\pi^{-1}(x_\ell)$ matches i to whichever object $\pi^{-1}(i)$ is matched to in x_ℓ, for every $i = 1,\ldots,n$.

We start with the case in which μ is anonymous. Let $0 < \varepsilon < 1/n$. For each $i = 1,\ldots,n$, define a utility function u_i over objects as follows for every object

$j = 1, \ldots, n$:

$$u_i(j) \triangleq \begin{cases} 1-(j-1)\varepsilon & j \leq i, \\ \frac{n-j}{n}\varepsilon & j > i, \end{cases}$$

and let $\succeq_i$ be the preference over $\Delta(X)$ whose vNM utility representation is u_i. Note that $R(\succeq_i) = R(\succeq_j)$ for all $i, j \in \{1, \ldots, n\}$. Therefore, all matchings are ex-post Pareto efficient, and by Lemma 7 are on the Pareto frontier, and so $u_i^{\max} = 1$ and $u_i^{\min} = 0$ for every $i = 1, \ldots, n-1$. Additionally, $u_n^{\max} = 1$ and $u_n^{\min} = 1-(n-1)\varepsilon$. Let $\hat{u}_i(\cdot) \triangleq \frac{u_i(\cdot)-u_i^{\min}}{u_i^{\max}-u_i^{\min}}$ be the Pareto-frontier-normalized utility function for every $i = 1, \ldots, n$. We have that $\hat{u}_i = u_i$ for every $i = 1, \ldots, n-1$, and that

$$\hat{u}_n(j) = \frac{1-(j-1)\varepsilon-\big(1-(n-1)\varepsilon\big)}{1-\big(1-(n-1)\varepsilon\big)} = \frac{(n-j)\varepsilon}{(n-1)\varepsilon} = \frac{n-j}{n-1}$$

for every $j = 1, \ldots, n$.

Denote $\succeq \triangleq (\succeq_1, \ldots, \succeq_n)$. Since μ is ordinal and anonymous and since $R(\succeq_i) = R(\succeq_j)$ for all $i, j \in \{1, \ldots, n\}$, when the reported profile of rankings is $R(\succeq)$, every matching in X is output by μ with the same probability. Therefore, the sum of Pareto-frontier-normalized utilities at $\mu(\succeq)$ is

$$\frac{1}{n}\sum_{i=1}^{n}\sum_{j=1}^{n}\hat{u}_i(j) = \frac{1}{n}\sum_{i=1}^{n-1}\sum_{j=1}^{n}\hat{u}_i(j)+\frac{1}{n}\sum_{j=1}^{n}\hat{u}_n(j) \leq \frac{1}{n}\sum_{i=1}^{n-1}\big(i+(n-i)\varepsilon\big)+\frac{1}{2} = \frac{n}{2}+\frac{n-1}{2}\varepsilon.$$

In contrast, the sum of Pareto-frontier-normalized utilities at the matching that matches each individual $i = 1, \ldots, n$ to the same-indexed object $j = i$ is

$$\sum_{i=1}^{n}\hat{u}_i(i) = \sum_{i=1}^{n-1}\big(1-(i-1)\varepsilon\big)+0 \geq n-1-\frac{(n-1)(n-2)}{2}\varepsilon.$$

Therefore, we have that[33]

$$\hat{I}\big((X,\succeq),\mu(\succeq)\big) \geq \frac{1}{n}\left(n-1-\frac{(n-1)(n-2)}{2}\varepsilon-\left(\frac{n}{2}+\frac{n-1}{2}\varepsilon\right)\right) = \frac{1}{2}-\frac{1}{n}-\frac{(n-1)^2}{2n}\varepsilon \xrightarrow[\varepsilon\to 0]{} \frac{1}{2}-\frac{1}{n}.$$

[33]Recall that $\hat{I}$ is a per-capita measure, i.e., it averages over the individuals' utility losses.

We move on to the case in which μ is not anonymous. Let $0 < \varepsilon < 1/n$. Let $\hat{\mu}$ be such that for every $R \in \mathcal{R}^n$, it holds that $\hat{\mu}(R) = \sum_\pi \frac{1}{n!} \cdot \pi^{-1}\big(\mu(\pi(R))\big)$, where the sum is over all permutations on $\{1,\ldots,n\}$. By the anonymous case, there exists an object allocation problem $\big(X,\succeq\big) \in \mathcal{M}_n$ such that $\hat{I}\big((X,\succeq),\hat{\mu}(\succeq)\big) \geq \frac{1}{2} - \frac{1}{n} - \frac{(n-1)^2}{2n}\varepsilon$. Therefore, there exists a permutation π on $\{1,\ldots,n\}$ such that $\hat{I}\big((X,\succeq),\pi^{-1}(\mu(\pi(\succeq)))\big) \geq \frac{1}{2} - \frac{1}{n} - \frac{(n-1)^2}{2n}\varepsilon$, where $\pi(\succeq)$ denotes the preference profile in which each individual i's preferences over objects are those of individual $\pi(i)$ in $\succeq$. We conclude the proof by focusing on the preference profile $\pi(\succeq)$ and noting that

$$\hat{I}\big((X,\pi(\succeq)),\mu(\pi(\succeq))\big) = \hat{I}\big((X,\succeq),\pi^{-1}(\mu(\pi(\succeq)))\big) \geq \frac{1}{2} - \frac{1}{n} - \frac{(n-1)^2}{2n}\varepsilon \xrightarrow[\varepsilon\to 0]{} \frac{1}{2} - \frac{1}{n}.$$

□

Proof of Lemma 9. The result follows by induction from the following claim, which we now prove: For every profile $u = (u_i)_{i=1}^n \in \mathcal{UR}$ and $i = 1,\ldots,n$ such that[34] $u_i(X) \cap [\varepsilon, 1-\varepsilon] \neq \emptyset$, there exists $u_i' \in \mathcal{UR}_i$ satisfying $R(u_i') = R(u_i)$ and $\big|u_i'(X) \cap [\varepsilon,1-\varepsilon]\big| < \big|u_i(X) \cap [\varepsilon,1-\varepsilon]\big|$, such that $I_u\big(\mu(u)\big) \leq I_{(u_i',u_{-i})}\big(\mu(u)\big)$.

Let $[0,\bar{l}] \subseteq [0,\varepsilon)$, $[l,r] \subseteq [\varepsilon,1-\varepsilon]$, and $[\bar{r},1] \subseteq (1-\varepsilon,1]$ be the smallest such segments such that $u_i(X) \subseteq [0,\bar{l}] \cup [l,r] \cup [\bar{r},1]$. These are well defined because the three intersections $u_i(X) \cap [0,\varepsilon)$, $u_i(X) \cap [\varepsilon,1-\varepsilon]$, and $u_i(X) \cap (1-\varepsilon,1]$ are all nonempty by assumption. Let $\hat{l} = \frac{\bar{l}+\varepsilon}{2}$ and $\hat{r} = \frac{\bar{r}+1-\varepsilon}{2}$. For every $\delta \in \{\hat{l}-l, 0, \hat{r}-r\}$, define the utility function u_i^δ as follows for every $x \in X$:

$$u_i^\delta(x) \triangleq \begin{cases} u_i(x) & u_i(x) \notin [l,r], \\ u_i(x)+\delta & u_i(x) \in [l,r]. \end{cases}$$

Note that $u_i^0 = u_i$ and that $R(u_i^\delta) = R(u_i)$ for every δ. Note also that for every $\delta \in \{\hat{l}-l, \hat{r}-r\}$, we have that $\big|u_i^\delta(X) \cap [\varepsilon,1-\varepsilon]\big| < \big|u_i(X) \cap [\varepsilon,1-\varepsilon]\big|$. It remains to show that $I_{(u_i^0,u_{-i})}\big(\mu(u)\big) \leq \max_{\delta\in\{\hat{l}-l,\hat{r}-r\}}\big\{I_{(u_i^\delta,u_{-i})}\big(\mu(u)\big)\big\}$.

Let $x_{max} \in \arg\max_{x\in X} U_u(x)$. Note that the function $f(\delta) \triangleq U_{(u_i^\delta,u_{-i})}(x_{max}) - U_{(u_i^\delta,u_{-i})}\big(\mu(u)\big)$ is a linear function. Therefore, at one of the extreme values $\hat{\delta} \in \{\hat{l}-l, \hat{r}-r\}$, the function f attains a weakly higher value than at the intermediate value $\delta = 0$. Let $u_i' \triangleq u_i^{\hat{\delta}}$. We conclude by observing that $I_u\big(\mu(u)\big) = f(0) \leq f(\hat{\delta}) \leq$

[34] We denote $u_i(X) \triangleq \big\{u_i(x) \bigm| x \in X\big\}$.

$I_{(u_i',u_{-i})}\big(\mu(u)\big)$. □

Proof of Lemma 10. Let $0 < \varepsilon < 1/n^3$ and let $u = (u_i)_{i=1}^n \in \mathcal{UR}^\varepsilon$. It is enough to prove that $I_u\big(RSD(u)\big) - 2n\varepsilon < (\ln 2) \cdot n$. By definition of $\mathcal{UR}^\varepsilon$, there exists $k \in \{1, \ldots, n\}$ such that $\big|k - \max_{x \in X} U_u(x)\big| < \varepsilon n < 1/n^2$. Filos-Ratsikas et al. (2014) prove that under these conditions, we have (in our notation) that

$$U_u\big(RSD(u)\big) \geq \sum_{l=0}^{\lfloor k/2 \rfloor} \frac{k-2l}{n-l} - n\varepsilon.$$

Since $\max_{x \in X} U_u(x) - n\varepsilon < k$, we have that[35]

$$I_u\big(RSD(u)\big) - 2n\varepsilon < k - \sum_{l=0}^{\lfloor k/2 \rfloor} \frac{k-2l}{n-l}.$$

Note that $d\big(\frac{k-2x}{n-x}\big)/dx = \frac{k-2n}{(n-x)^2} < 0$. Therefore,

$$\sum_{l=0}^{\lfloor k/2 \rfloor} \frac{k-2l}{n-l} > \int_0^{k/2} \frac{k-2x}{n-x} dx.$$

Combining these two, we have that

$$I_u\big(RSD(u)\big) - 2n\varepsilon < k - \int_0^{k/2} \frac{k-2x}{n-x} dx = \int_0^{k/2} \left(2 - \frac{k-2x}{n-x}\right) dx =$$
$$\int_0^{k/2} \frac{2n-2x-k+2x}{n-x} dx = \int_0^{k/2} \frac{2n-k}{n-x} dx = (2n-k) \int_0^{k/2} \frac{1}{n-x} dx =$$
$$(2n-k)\Big(\ln n - \ln\big(n - \tfrac{k}{2}\big)\Big) = (2n-k)\left(\ln \frac{1}{1 - \frac{k}{2n}}\right).$$

Writing $k' = \frac{k}{n} \leq 1$, we then have that

$$I_u\big(RSD(u)\big) - 2n\varepsilon < n \cdot (2-k') \left(\ln \frac{1}{1 - \frac{k'}{2}}\right) \leq n \cdot \ln 2.$$ □

Proof of Lemma 11. Assume for contradiction that True is returned. Then, the max-

[35] At this point, our analysis diverges from that of Filos-Ratsikas et al. (2014), who go on to bound the order of magnitude of the ratio $\frac{U_u(RSD(u))}{\max_{x \in X} U_u(x)}$.

cardinality matching σ in G is of size $n-1$. Therefore, $x \triangleq \sigma \cup (\hat{i}, o) \in X$ is a perfect matching of the n individuals to the n objects such that each individual except $\hat{i}$ is matched to an object that she strictly prefers to o. Let $x' \in X$ be an ex-post Pareto efficient matching such that $x' \succeq_i x$ for every $i = 1, \ldots, n$. Then, in x', each individual except $\hat{i}$ is matched to an object that she strictly prefers to o. Therefore, o is matched to $\hat{i}$ in x', in contradiction to o not being matched to $\hat{i}$ in any ex-post Pareto efficient matching. □

Proof of Lemma 12. Let $x \in X$ be an ex-post Pareto efficient matching that matches o to $\hat{i}$. Assume for contradiction that in x, some individual $j \neq \hat{i}$ is matched to an object that she strictly prefers less than o. Let $y \in X$ be the matching obtained from x by exchanging the objects of $\hat{i}$ and j. Then, j strictly prefers y to x. Since x is ex-post Pareto efficient, we therefore have that $\hat{i}$ is matched in y to an object she strictly prefers less than o. Let $y' \in X$ be an ex-post Pareto efficient matching such that $y' \succeq_i y$ for every $i = 1, \ldots, n$. We claim that $\hat{i}$ is matched in y' to an object she strictly prefers less than o; indeed, if this were not the case, then it would hold that $y' \succeq_i x$ for every $i = 1, \ldots, n$ and $y' \succ_j x$, in contradiction to x being ex-post Pareto efficient. We conclude that y' is an ex-post Pareto efficient matching that matches to $\hat{i}$ an object that she strictly prefers less than o, contradicting the definition of o. Therefore, in x each individual $j \neq \hat{i}$ is matched to an object that she strictly prefers to o. Hence, during the computation of the function call $\textsc{Test}(n, (\succeq_i)_{i=1}^n, \hat{i}, o)$, the max-cardinality matching σ in G is of size $n-1$. Thus, True is returned. □