

Cooperative Game Solution Concepts that Maximize Stability under Noise

Yuqian Li and Vincent Conitzer

Department of Computer Science

Duke University

{yuqian, conitzer}@cs.duke.edu

Abstract

In cooperative game theory, it is typically assumed that the value of each coalition is known. We depart from this, assuming that $v(S)$ is only a noisy estimate of the true value $V(S)$, which is not yet known. In this context, we investigate which solution concepts maximize the probability of *ex-post* stability (after the true values are revealed). We show how various conditions on the noise characterize the least core and the nucleolus as optimal. Modifying some aspects of these conditions to (arguably) make them more realistic, we obtain characterizations of new solution concepts as being optimal, including the partial nucleolus, the multiplicative least core, and the multiplicative nucleolus.

Introduction

Computational cooperative game theory (see the overview by Chalkiadakis, Elkind, and Wooldridge (2011)) is emerging as an important toolbox in the coordination of coalitions in multiagent systems. Much of cooperative game theory focuses on *stability*: no subset of agents should have an incentive to break off and work on its own. Solutions satisfying this criterion, straightforwardly interpreted, constitute the *core*, which is the standard stability-oriented solution concept. While it is a useful analytical tool, the core is likely to fall short in real multiagent settings, as it does not distinguish between solutions that are just barely stable and ones that are stable by a comfortable margin. This is particularly troublesome when the numbers used to compute the solution are just noisy estimates of the true values of coalitions. Existing solution concepts such as the *least core* and the *nucleolus* appear to provide some relief, as they in some sense give the solutions that are the “most” stable. But what is the right sense of “most” stable, particularly when we are worried about the fact that our input numbers are no more than noisy estimates? And does the right sense lead us to the least core, the nucleolus, or something else? These are the questions we set out to answer in this paper.

Our contribution is to lay out a framework to address these issues and obtain some first results. Specifically, we show that noise models can be designed that justify the least core, the nucleolus, and their multiplicative versions as the

most stable solution(s). We do not claim to thereby establish definitively that these concepts are the only ones that can reasonably be said to maximize stability in noisy environments. One may find aspects of the noise model unappealing and modify the noise model accordingly, which may lead either to an altogether new concept, or to a better justification of an existing concept. The first question to answer, though, is whether existing solution concepts can be justified in this manner at all, and what types of noise model are needed for this. This is the aim of our paper.¹

Definitions

In this section, we first review cooperative games without uncertainty, and then extend to games with uncertainty (as has been done in prior work). Finally, we discuss our framework for taking a cooperative game without uncertainty and considering “noisy” versions of it.

Cooperative Games

A *cooperative game* (or a *characteristic function game*) is specified by a pair (N, v) where N is a set of n agents and $v : 2^N \rightarrow \mathbb{R}$ is the characteristic function that assigns a value $v(S)$ to every subset $S \subseteq N$, representing the value that agents in S can obtain and distribute among themselves if they work (only) with each other.

A solution of such a cooperative game usually consists of a *payment vector* $x : N \rightarrow \mathbb{R}$, where $x(i)$ is the payment agent $i \in N$ receives. Let $x(S) = \sum_{i \in S} x(i)$ denote the total payment to a subset $S \subseteq N$. Generally it is required that $x(N) = v(N)$. Such a payment vector x is *stable* or *in the core* if $x(S) \geq v(S)$ for all $S \subseteq N$ (Gillies 1953). That is, the total payment $x(S)$ to a subset S should be no less than the value S can generate by itself. Otherwise, S would have an incentive to deviate and work on its own.

It is well known that the core can be empty; even when it is nonempty, we might like to select its “most” stable elements. The *least core* is an attempt to address all this.

Definition 1 (Least Core (Maschler, Peleg, and Shapley 1979)). *Let the ε -core be $\{x \mid x(N) = v(N) \text{ and } x(S) \geq v(S) - \varepsilon \text{ for all } \emptyset \subsetneq S \subsetneq N\}$. The least core is the nonempty*

¹The approach is analogous to that used to justify voting rules as maximum likelihood estimators of the “correct” outcome, which we will discuss in the related research section.

ε -core with the minimum ε (which can be positive or nonpositive, depending on whether the core is empty or not).

That is, solutions in the least core satisfy the core constraints by as large a margin as possible (possibly a negative margin if the core is empty). It is possible to take this reasoning a step further: sometimes, within the least core, it is possible to satisfy the constraints for some coalitions by a larger margin (without increasing ε). Taking this reasoning to its extreme leads to the definition of the *nucleolus*.

Definition 2 (Nucleolus (Schmeidler 1969)). *The excess vector of a payment vector x consists of the following $2^n - 2$ numbers in nonincreasing order: $(v(S) - x(S))_{\emptyset \subsetneq S \subsetneq N}$. The nucleolus is the (unique) payment vector x that lexicographically minimizes the excess vector.*

In an intuitive sense, when the core is nonempty, the least core can be seen as an attempt to make the solution more robust, and the nucleolus as an attempt to make it even more robust. But from the perspective of cooperative games without any uncertainty, it is hard to make this intuition precise, because it implies 0% chance of a deviation from any solution in the core, and 100% chance of a deviation from any solution outside of the core. Because our goal is to formally justify these and related notions in this paper, we will need to consider cooperative games with uncertainty.

Cooperative Games with Uncertainty

To model cooperative games with uncertainty, a natural approach is to specify that there is a random *state of the world* $\omega \in \Omega$ that affects the values of coalitions (see, e.g., Jeong and Shoham (2008)). Let \mathbb{P} be the probability measure that assigns a probability $\mathbb{P}(A) \in [0, 1]$ to each event $A \subseteq \Omega$.

Holding (Ω, \mathbb{P}) fixed, a cooperative game with uncertainty is specified by a pair (N, V) where N is the set of agents as before, and $V : 2^N \times \Omega \rightarrow \mathbb{R}$ is a *stochastic* characteristic function. That is, $V(S, \omega)$ is the value of subset S in realized world ω . (We use V instead of v now to indicate that each $V(S)$ is a random variable whose value depends on ω .) In most of our paper, we will not make the state space Ω explicit but rather work directly with distributions of the random variables $V(S)$.

Payment vectors $X : N \times \Omega \rightarrow \mathbb{R}$ can now in principle also become random. (Again, for $S \subseteq N$, let $X(S, \omega) = \sum_{i \in S} X(i, \omega)$ denote the total payment to S .) However, this paper mostly focuses on cases where both X and the grand coalition value $V(N)$ are deterministic: $X(i, \omega_1) = X(i, \omega_2) = x(i)$ and $V(N, \omega_1) = V(N, \omega_2) = v(N)$ for any $i \in N$ and $\omega_1, \omega_2 \in \Omega$. These assumptions are natural, for example, in the context where there is a well-established group of agents with a clear, cleanly worked out, and well-rehearsed plan of how to proceed (so that there is effectively no uncertainty about the grand coalition’s value), but it is much less clear how subsets of agents would do if they would break off and “go into the wild” on their own (see the examples in the following subsection).

Stability

In this paper, we focus on *ex-post* stability (Jeong and Shoham 2008), i.e., stability when the realized world ω

is revealed. Given a payment rule X and a game (N, V) , let $X \geq V$ denote the set of stable worlds $\{\omega \mid \forall S \subseteq N, X(S, \omega) \geq V(S, \omega)\}$. Our objective is then finding X (usually, a deterministic payment vector x) that maximizes $\mathbb{P}(X \geq V)$, the probability of *ex-post* stability. There are at least two interpretations of this objective.

The first one is motivated by games where the state of the world is *learned* over time after the contract (payment rule X) has been established. For example, consider an enterprise that hires a team of new employees for a new project under contract X (typically, this would be a deterministic contract x). At the point of establishing the contract, the team members may not know each other, and as a result do not know the value $V(S)$ that each subcoalition would generate in any single month. Over many months of cooperation, however, they will get to know each other well enough to accurately assess this value, and at that point they may decide to stay in the big enterprise (the grand coalition being stable) or a subteam may resign to open a new company.

The other interpretation concerns games where the different (subsets of) agents simply have different *perceptions* of their values, and the randomness derives from uncertainty about how agents perceive their values. In this case, $V(S)$ represents the value that coalition S perceives itself as being able to obtain on its own. Let us suppose that there is no uncertainty about $V(N) = v(N)$ and that we must have a deterministic payment vector x (as we do in most of this paper). From the perspective of the party tasked with determining x , $V(S)$ is a random variable, over which it has a subjective distribution. Then, if it turns out that $x(S) < V(S)$, S will deviate (regardless of whether S ’s perception of its own value was in fact correct). In this context, it is natural for the party tasked with determining the payment rule to try to maximize $\mathbb{P}(x \geq V)$, the probability that no coalition will deviate, with respect to its own subjective probability distribution. This appears to match well with the fact that in the real world, we will probably never be completely sure about what is going on in the minds of a particular coalition and, hence, about whether it will decide to deviate.

Noise

We assume that a point estimate (N, v) (a game without uncertainty) is known at the point where we determine payments. The true game differs from this point estimate by random noise D , where $D = v - V$, or equivalently, $V = v - D$. Hence, the true game is a game with uncertainty, which is obtained by subtracting the noise from the estimate.

Given the point estimate (N, v) , we can use traditional solution concepts for games without uncertainty (e.g., the least core or the nucleolus) to compute a solution x . In this paper, we study whether there is a distribution of the noise such that this solution x is optimal for the true game (N, V) in terms of *ex-post* stability. More precisely, what conditions on the distribution of D guarantee that x maximizes $\mathbb{P}(x \geq V)$ among all possible solutions x ?

We make the following assumptions on the distribution of D throughout this paper (except where otherwise specified).

Assumption 1. As mentioned and justified in the previous two subsections, we assume $D(N) = 0$, so that the grand

coalition's value $V(N) = v(N)$ is fixed. Hence, we can determine whether a deterministic payoff vector x is feasible.

Assumption 2. For any $\emptyset \subsetneq S_1, S_2 \subsetneq N$, the random variables $D(S_1)$ and $D(S_2)$ follow the same distribution F . This simplifies our specification to only one distribution F , rather than $2^n - 2$ distributions F_S . (Of course, the fact that we can justify the least core and the nucleolus even under this restriction makes the result stronger rather than weaker, because it immediately implies that they can be justified in the unrestricted case too, whereas the opposite is not necessarily the case.) One may well argue that this assumption is unrealistic, as one would expect a subset with higher estimated value $v(S)$ to also have higher expected absolute noise $\mathbb{E}[|D(S)|]$. We will revisit this in the section on multiplicative noise.

It bears emphasizing that the second assumption does not imply that $D(S_1)$ and $D(S_2)$ are identical—they are just identically *distributed*. Neither does this assumption say that $D(S_1)$ and $D(S_2)$ are independent. On the other hand, identical / i.i.d. random variables $D(S_1), D(S_2)$ constitute two special cases that satisfy this assumption. We will use them to generate the least core and the nucleolus, respectively.

Related Research

There are various previous axiomatizations of cooperative game-theoretic solution concepts (Peleg 1985; Potters 1991; Snijders 1995; Orshan 1993; Sudhölter 1997). The axioms (such as the reduced game property) are often quite complicated, arguably more complicated than the definition of the solution concept itself. In this paper, rather than characterizing the concepts axiomatically, we characterize them as solving an optimization problem. Similar approaches have been used in social choice to characterize voting rules, notably the distance rationalizability framework (Meskanen and Nurmi 2008; Elkind, Faliszewski, and Slinko 2009) and the maximum likelihood framework (Young 1988; 1995; Drissi-Bakhkhat and Truchon 2004; Conitzer and Sandholm 2005; Xia and Conitzer 2011; Pivato 2013). Our work particularly resembles the latter line of work due to its focus on noisy observations of an unobserved truth.

Stochastic cooperative games were first introduced in the 1970s (Charnes and Granot 1976; 1977), and received significantly more attention in recent work (Suijs et al. 1999; Suijs and Borm 1999; Chalkiadakis and Boutilier 2004; Chalkiadakis, Markakis, and Boutilier 2007). Much of their effort was spent on dealing with stochastic payment rules X . In contrast, our paper mostly focuses on deterministic payment rules x (so that we can analyze standard solution concepts in cooperative game theory), though we will briefly mention that the multiplicative least core (to be introduced later) can be characterized as the optimal stochastic payment rule X when only $V(N)$ is uncertain.

The multiplicative least core is similar to the weak least core (Shapley and Shubik 1966; Bejan and Gómez 2009; Meir, Rosenschein, and Malizia 2011) in the sense that larger coalitions are allowed to have larger excess. The difference is that the multiplicative least core compares the excess to $v(S)$, while the weak least core compares it to $|S|$. We argue that, at least in some cases, the former makes more

sense. For example, a single-agent coalition with a value of 1 billion may easily accept an excess of 100 and not deviate, while a single-agent (or even multiagent) coalition with a value of 100 and an excess of 100 would feel strongly inclined to deviate, thereby doubling its payment.

Finally, our notation is compatible with the partition and type models that have been predominantly used in previous work on cooperative games with uncertainty. Agent world partitions (see, e.g., Jeong and Shoham (2008)) and agent types (see, e.g., Myerson (2007) and Chalkiadakis and Boutilier (2007)) both represent the private information that agents possess (in addition to the common prior). However, we focus on stability *ex post* rather than *ex interim* (or *ex ante*, as studied by Bachrach et al. (2012)), so agents' private information does not play any significant role here.

Characterizing the Least Core

In this section, we show that under certain assumptions on the noise D , payments are optimal if and only if they are in the least core. The key assumption is that the noise is identical for all coalitions (e.g., if one coalition's realized value is 10 above its estimate, the same will be true for all other coalitions). The proof is quite straightforward, which arguably indicates that the solution concept is justified by our approach in a natural manner. On the other hand, one may well take issue with the key assumption—for example, we would expect coalitions with large values to also experience larger noise. We will revisit this in the section on multiplicative noise, leading to our definition of the multiplicative least core. This illustrates how our approach can do more than justify an existing solution concept; it can also naturally lead us to an improvement of it.

Theorem 1. *When the noise of any two subsets $\emptyset \subsetneq S_1, S_2 \subsetneq N$ is identical ($D(S_1) = D(S_2)$), and its distribution has full support (the cumulative density function (CDF) is strictly increasing everywhere—e.g., a normal distribution), then a payment x maximizes stability ($\mathbb{P}(x \geq V) = \mathbb{P}(x \geq v - D)$) if and only if it is in the least core. (Moreover, even when the distribution does not have full support, any x in the least core maximizes stability.)*

Proof. (“If” direction.) Let x^* be an element in the least core. If x^* were not optimal, then there must be another x such that $\mathbb{P}(x \geq v - D) > \mathbb{P}(x^* \geq v - D)$. That requires at least one world $\omega \in \Omega$ in which $x^*(S') < v(S') - D(S', \omega)$ for some $\emptyset \subsetneq S' \subsetneq N$ while $x(S) \geq v(S) - D(S, \omega)$ holds for all $\emptyset \subsetneq S \subsetneq N$. Because $D(S)$ is identical for all S , let $D(S, \omega) = \varepsilon$. Then x is in the ε -core while x^* is not in it. This violates the definition of the least core.

(“Only if” direction.) Let $\varepsilon(x)$ denote $\max_{\emptyset \subsetneq S \subsetneq N} v(S) - x(S)$. Then $\mathbb{P}(x \geq v - D) = \mathbb{P}(D \geq \varepsilon(x))$. Since D has full support, x must minimize ε to maximize the stability. \square

Characterizing the Nucleolus

In this section, we consider whether it is possible to characterize the nucleolus with assumptions on the noise D . It turns out that we can, if we move to a richer model where we consider what happens in the limit for a *sequence* of noise

distributions. Naturally, one may be somewhat dissatisfied with this, preferring instead to characterize the concept with a single distribution. Unfortunately, as we will show, this is impossible (under i.i.d. noise), necessitating the move to a richer model. In a sense, the result involving a sequence of distributions shows that we can come *arbitrarily close* to justifying the nucleolus with a single distribution, but due to its lexicographic nature we will always be slightly off.

The assumption that we made on the noise to characterize the least core is on the extreme end where the $D(S)$ are perfectly correlated (identical). In this section, we go to the other extreme where the $D(S)$ are independent. Hence, the $2^n - 2$ random variables $D(S)$ are i.i.d. (due to our Assumption 2). To make the notation more compact, let $F(\delta) = \mathbb{P}(D(S) \geq \delta)$. This results in the following lemma.

Lemma 1. *Let $\xi(S) = v(S) - x(S)$ be the excess of coalition $S \subset N$. Let $\vec{\delta} = (\delta_1, \delta_2, \dots, \delta_{2^n-2})$ denote the excess vector (sorted from large to small). If the $2^n - 2$ random variables $\{D(S) | \emptyset \subsetneq S \subsetneq N\}$ are i.i.d. with $\mathbb{P}(D(S) \geq \delta) = F(\delta)$, then the probability of stability is*

$$\mathbb{P}(x \geq v - D) = \prod_{\emptyset \subsetneq S \subsetneq N} \mathbb{P}(D(S) \geq \xi(S)) = \prod_{i=1}^{2^n-2} F(\delta_i)$$

Proof. Subset S 's stability $x(S) \geq V(S)$ is equivalent to $v(S) - V(S) \geq v(S) - x(S)$, and thus equivalent to $D(S) \geq \xi(S)$. Hence, the probability of S 's stability is $F(\xi(S))$. The lemma follows because the $\{D(S)\}$ are independent. \square

Does the nucleolus necessarily maximize the expression in Lemma 1? No. For example, consider the game where $N = \{a, b\}$, $v(N) = 2$, and $v(\{a\}) = v(\{b\}) = v(\emptyset) = 0$. The nucleolus gives $x(a) = x(b) = 1$, resulting in excess vector $\vec{\delta} = (-1, -1)$. Alternatively, we could set $x'(a) = 2, x'(b) = 0$ to obtain excess vector $\vec{\delta}' = (0, -2)$. Now suppose $D(S)$ follows a distribution with $F(-2) = 1, F(-1) = 0.7$ and $F(0) = 0.5$. Then, $\vec{\delta}'$ has a higher probability of stability, namely $F(0)F(-2) = 0.5 > F(-1)^2 = 0.49$. But this only proves that one particular distribution F will not work to characterize the nucleolus. What about others? Unfortunately, the following proposition proves that no reasonable distribution will work. (We omit a number of the proofs in this paper due to the space constraint; they are available in the full version.)

Proposition 1. *For any fixed distribution F , if there exists a point $q \in \mathbb{R}$ such that $F'(q) \neq 0$ (so that the probability density function (PDF) is defined and nonzero at q), then there exists a 3-player cooperative game whose nucleolus does not maximize stability. Hence, no fixed noise distribution with nonzero derivative at some point can guarantee the nucleolus' optimality.*

It follows that our characterization of the nucleolus will not be as clean as the one we obtained for the least core. The intuition for the way in which we characterize the nucleolus is as follows. The nucleolus corresponds to worrying more about large excesses than about small excesses, but to still worry about small excesses as a secondary objective.

To some extent, it turns out we can achieve this by choosing a noise distribution where decreasing a large excess has a much greater impact on the probability of stability than decreasing a small excess. But this will at best result in an approximation of the nucleolus, because the (exact) nucleolus would require the former impact to be *infinitely* greater, which is impossible if decreasing a small excess is still to have some impact (which is also necessary for the nucleolus). Nevertheless, we can create a *sequence* of noise distributions that correspond to ever better approximations, so that any other payment vector becomes less stable than the nucleolus after some point in the sequence. We can additionally require that the distributions in this sequence become ever more concentrated around our original estimate of the value function. Thus, in a sense the nucleolus corresponds to what we do in the limit as we become more and more confident in our estimate. The following definition makes this precise.

Definition 3. *Let $\{F_c\}$ be a sequence of distributions indexed by $c \in \{0, 1, 2, \dots\}$ (the confidence). Let D_c be drawn according to F_c . We require that the larger confidence c is, the less noisy F_c is: for any $c_1 < c_2$ and $\epsilon > 0$, we have $\mathbb{P}(D_{c_1}(S) > \epsilon) \geq \mathbb{P}(D_{c_2}(S) > \epsilon)$ and $\mathbb{P}(D_{c_1}(S) < -\epsilon) \geq \mathbb{P}(D_{c_2}(S) < -\epsilon)$. We say a payment x^* is strictly better than x under confidence (with respect to $\{F_c\}$) if there exists some confidence threshold C above which x^* is always strictly better than x : ($\forall c \geq C$), $\mathbb{P}(x^* \geq v - D_c) > \mathbb{P}(x \geq v - D_c)$. We say x^* is uniquely optimal under confidence (with respect to $\{F_c\}$) if for any other payment $x \neq x^*$, x^* is strictly better than x under confidence.*

In words, a payment x^* is uniquely optimal under confidence if for any other payment x , x^* is strictly more stable than x as long as we are confident enough about our estimate of the value function (but how much confidence is enough may depend on x). What conditions do we need on $\{F_c\}$ for the nucleolus to be uniquely optimal? The following lemma gives a sufficient condition.

Lemma 2. *If $\{F_c\}$ satisfies that for any $\delta \in \mathbb{R}$, $\epsilon > 0$, and $k \geq 1$, there exists a threshold C such that for all $c \geq C$, we have $F_c(\delta)^k > F_c(\delta + \epsilon)$, then the nucleolus is uniquely optimal under confidence with respect to $\{F_c\}$. The condition above is also denoted as $\lim_{c \rightarrow \infty} F_c(\delta)^k > F_c(\delta + \epsilon)$ for convenience.*

While the precondition in Lemma 2 is useful, it is difficult to interpret in terms of becoming ever more confident. The next two definitions do say something about how quickly we need to become more confident, and together they imply the precondition of Lemma 2.

Definition 4 (Overestimate Condition). *We say $\{F_c\}$ satisfies the overestimate condition if for any $\delta \geq 0$, $\epsilon > 0$, and $k \geq 1$, there exists a threshold C such that for all $c \geq C$, $F_c(\delta)^k > F_c(\delta + \epsilon)$.*

We call this the overestimate condition because for non-negative δ , $F_c(\delta)$ denotes the probability that we overestimate and the error $v(S) - V(S)$ is at least δ . The overestimate condition then indicates that when we are very confident ($\forall c \geq C$), this probability diminishes to 0 extremely

fast as δ increases: $F_c(\delta + \epsilon)$ is so much closer to 0 than $F_c(\delta)$ that even $F_c(\delta)^k$ is not as close to 0 as $F_c(\delta + \epsilon)$ is.

Definition 5 (Underestimate Condition). Let $G_c(\delta) = 1 - F_c(\delta)$. We say $\{F_c\}$ satisfies the underestimate condition if for any $\delta < 0$, $\epsilon > 0$, and $k \geq 1$, there exists a threshold C such that for all $c \geq C$, $kG_c(\delta - \epsilon) < G_c(\delta)$.

Here, for $\delta < 0$, $G_c(\delta) = 1 - F_c(\delta) = \mathbb{P}(D_c(S) = v(S) - V(S) \leq \delta)$ is the probability that we underestimate ($V(S) > v(S)$) and the error $V(S) - v(S)$ is greater than $|\delta|$, which is why we call it the underestimate condition. Like the overestimate condition, it indicates that when we are very confident, the probability $G_c(\delta)$ diminishes to 0 very fast as δ decreases: $G_c(\delta - \epsilon)$ is so much closer to 0 than $G_c(\delta)$ that even $kG_c(\delta - \epsilon)$ is still closer to 0 than $G_c(\delta)$ is.²

Lemma 3. If $\{F_c\}$ satisfies the overestimate condition and the underestimate condition, then it also satisfies the condition $\lim_{c \rightarrow \infty} F_c(\delta)^k > F_c(\delta + \epsilon)$ from Lemma 2.

Proof. For $\delta \geq 0$, the condition $\lim_{c \rightarrow \infty} F_c(\delta)^k > F_c(\delta + \epsilon)$ immediately follows from the overestimate condition. It remains to show that the underestimate condition implies that $\lim_{c \rightarrow \infty} F_c(\delta)^k > F_c(\delta + \epsilon)$ for $\delta < 0$. Since F_c is monotone, it is sufficient to prove that the underestimate condition implies $\lim_{c \rightarrow \infty} F_c(\delta)^k > F_n(\delta + \epsilon)$ for $\delta < 0$ and $\epsilon < -\delta$ (as this will immediately imply the same for larger ϵ). Let $\delta' = \delta + \epsilon < 0$. By the underestimate condition, there exists a C such that for all $c \geq C$, we have $kG_c(\delta' - \epsilon) < G_c(\delta')$. Then for all $c \geq C$, $F_c(\delta)^k = F_c(\delta' - \epsilon)^k = (1 - G_c(\delta' - \epsilon))^k > 1 - kG_c(\delta' - \epsilon) > 1 - G_c(\delta') = F(\delta + \epsilon)$, where the first inequality follows from Bernoulli's inequality. \square

Combining all of the preceding lemmas, we obtain the following theorem about the nucleolus.

Theorem 2. When the noise $\{D(S)\}$ is drawn i.i.d. across coalitions $\emptyset \subsetneq S \subsetneq N$, the nucleolus is uniquely optimal under confidence for any sequence of distributions $\{F_c\}$ that satisfies both the overestimate and underestimate conditions.

Such sequences indeed exist:

Proposition 2. For $\{F_c(\delta) = \exp(-e^{c\delta})\}$, both the overestimate and underestimate conditions are satisfied.

The CDF ($1 - F$) and PDFs ($-F'$) corresponding to distributions in Proposition 2 are plotted in Figure 1. Sequences where the probability of large noise decreases significantly more slowly, such as $\{\exp(-ce^\delta)\}$, will not suffice, even though in this sequence the probability of larger noise becomes arbitrarily smaller than that of smaller noise.

Partial Nucleolus

So far, we have used noise models to justify existing solution concepts. One benefit of doing so is that we can now investigate whether these noise models are reasonable (in all circumstances), adjust them when they are not, and perhaps obtain new solution concepts as a result.

²Note that the overestimate condition requires a faster diminishing speed than the underestimate condition: in the former, even exponentiation by k does not change the order, whereas in the latter, only multiplication by k does not change the order.

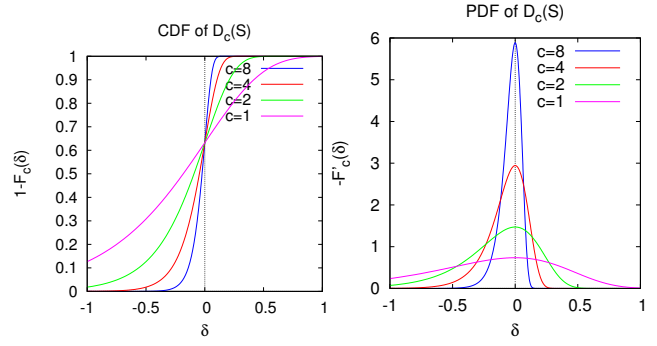


Figure 1: The CDF ($1 - F$) and PDF ($-F'$) of distributions in $\{F_c(\delta) = \mathbb{P}(D(S) \geq \delta) = \exp(-e^{c\delta})\}$ which satisfies both the overestimate condition and the underestimate condition.

For example: in some cases, we might expect that there is only overestimate error (the noise is nonnegative). For example, $D(S) = v(S) - V(S) \geq 0$ may correspond to an unknown cost that S has to pay if S deviates (so the value that S truly has after deviating is $V(S) = v(S) - D(S)$). We introduce the following solution concept:

Definition 6 (Partial Nucleolus). Given payment vector x , the nonnegative excess vector consists of the following $2^n - 2$ numbers in non-increasing order: $\max(0, v(S) - x(S))_{\emptyset \subsetneq S \subsetneq N}$. The partial nucleolus is the set of x that lexicographically minimize the nonnegative excess vector.

We can characterize the partial nucleolus as consisting exactly of the optimal payment vectors for certain sequences of distributions. The (omitted) proof is similar to that given for the nucleolus above.

Theorem 3. When the noise $\{D(S)\}$ is drawn i.i.d. across coalitions $\emptyset \subsetneq S \subsetneq N$, the partial nucleolus is optimal (with all payment vectors in the partial nucleolus performing equally well, and strictly better than any other payment vector) under confidence with respect to any sequence of distributions $\{F_c\}$ that satisfies the overestimate condition and $F_c(0) = 1$ (for nonnegativity of the noise).

Multiplicative Noise

In the noise models considered so far, $D(S)$ follows the same distribution F for all $\emptyset \subsetneq S \subsetneq N$. One may argue that this is an unreasonable aspect of these models. For example, it seems unreasonable to say that the probability that revenue turns out \$1B higher than expected is the same for Google and Yahoo (because Google's revenue is much bigger), but it seems reasonable to say that the probability that revenue turns out 2% higher than expected is the same for both. To address this, we now introduce *multiplicative* noise, which makes the noise distribution the same in *relative* terms.

Throughout this section, we assume that v and V are nonnegative. We then use a standard trick using logarithms (which will map these nonnegative numbers to the full set of real numbers) to adapt our results in the additive model to the multiplicative model.

Previously, our noise $D = v - V$ was defined in an additive manner. Now, we define the multiplicative noise D^\times to

be $\log(v/V) = \log v - \log V$. The multiplicative noise D^\times remains unchanged if v and V are both multiplied by the same amount ρ : $D^\times = \log(\rho v) - \log(\rho V) = \log v - \log V$. When $D^\times = 0$, $v = V$; as D^\times moves away from 0, v moves further away from V ; and when $D^\times(S)$ follows the same distribution F for all S , the expected absolute value of (additive) noise $\mathbb{E}(|v - V|) = v\mathbb{E}(|1 - e^{-D^\times}|)$ grows in proportion with v .

A minor problem for D^\times is that $\log(v/V)$ is not well defined when either v or V is 0. In this paper, we will be interested in *generative* models in which we draw a real (but finite) $D^\times(S)$ for each coalition S (so $0 < e^{D^\times(S)} < \infty$). Therefore, $v(S) = 0 \Leftrightarrow V(S) = 0$, resulting in $\log(v/V) = \log(0/0)$. This means that when $v(S) = V(S) = 0$, we will not be able to recover what exact amount of noise $D^\times(S)$ was drawn, but this will not matter for stability because $V(S) = 0$ for any $D^\times(S)$.

We now define the multiplicative excess. The multiplicative excess represents by how great a *factor* a coalition could make itself better (or worse) off by deviating. One could argue that this is intuitively a better measure of how likely a coalition is to deviate: a coalition that is currently receiving 1 billion will not feel greatly incentivized to deviate to obtain an additional 100, but one that is currently receiving 100 is likely to deviate if doing so means doubling its payoff. Moreover, it lines up nicely with the notion of multiplicative noise: a coalition will have no incentive to deviate iff its multiplicative noise is at least its multiplicative excess.

Definition 7 (Multiplicative Excess). *The multiplicative excess of coalition S is defined as $\xi^\times(S) = \log v(S) - \log x(S)$. When $v(S) = 0$, we define $\xi^\times(S) = -\infty$ (indicating complete stability). When $v(S) > 0$ and $x(S) = 0$, we define $\xi^\times = +\infty$ (indicating complete instability).*

Based on this, we can define the multiplicative least core.

Definition 8 (Multiplicative Least Core). *Let the multiplicative ε -core be $\{x | x(N) = v(N) \text{ and } \xi^\times(S) \leq \varepsilon \text{ for all } \emptyset \subsetneq S \subsetneq N\}$ ($+\infty$ and $-\infty$ may appear as we defined for the multiplicative excess). The multiplicative least core is the nonempty multiplicative ε -core with the minimum ε (which can be positive or nonpositive, depending on whether the core is empty or not).*

We then obtain the following analogue of Theorem 1.

Theorem 4. *When the multiplicative noise of any two subsets $\emptyset \subsetneq S_1, S_2 \subsetneq N$ is identical ($D^\times(S_1) = D^\times(S_2)$), and its distribution has full support (the CDF is strictly increasing everywhere—e.g., a normal distribution), then a payment x maximizes stability ($\mathbb{P}(x \geq V) = \mathbb{P}(x \geq v/e^{D^\times})$) if and only if it is in the multiplicative least core. (Moreover, even if the distribution does not have full support, any x in the multiplicative least core maximizes stability.)*

In fact, multiplying all subcoalition values $v(S)$ ($S \subsetneq N$) by a factor $\rho = 1/e^{D^\times}$ ($V(S) = \rho v(S)$) is, in some sense, equivalent to multiplying (only) the grand coalition's value $v(N)$ by the factor's inverse ($V(N) = v(N)/\rho$). In the latter scenario, (only) the grand coalition's value is not known at the outset; in such cases, it is natural to restrict attention to

payment rules that give each agent a fixed *share* of the eventually realized value $V(N)$. In this case, the multiplicative least core consists exactly of the optimal share rules.

We can also define the multiplicative nucleolus and obtain a result analogous to Theorem 2.

Definition 9 (Multiplicative Nucleolus). *The multiplicative excess vector of a payment x consists of the following $2^n - 2$ numbers in nonincreasing order: $(\xi^\times(S))_{\emptyset \subsetneq S \subsetneq N}$. The multiplicative nucleolus consists of the payment vectors x that lexicographically minimize the excess vector.*

Theorem 5. *When the multiplicative noise $\{D^\times(S)\}$ is drawn i.i.d. across coalitions $\emptyset \subsetneq S \subsetneq N$, the multiplicative nucleolus is optimal (with all payment vectors in the multiplicative nucleolus performing equally well, and strictly better than any other payment vector) under confidence with respect to any sequence of distributions $\{F_c\}$ for the multiplicative noise that satisfies both the overestimate condition and the underestimate condition.³*

Using similar techniques as for the regular nucleolus, it can be shown that the multiplicative nucleolus is unique as long as there exist n coalitions S with $v(S) > 0$ that are linearly independent.⁴ Very similarly, we can also obtain a multiplicative version of the partial nucleolus. The details are in the full version of the paper.

Conclusion

In this paper, we justified the least core and the nucleolus by proving that they maximize the probability of *ex-post* stability under particular noise models. In words, the least core is uniquely optimal when the noise of each coalition is perfectly correlated, and the nucleolus is uniquely optimal when the noise is independent and additionally, we are very confident that the noise is small—intuitively, the probability of larger noise is infinitesimal compared to the probability of smaller noise. This appears to nicely reflect the intuitive sense in which these solutions are “more” stable than the core. Moreover, by modifying the noise models in ways that make them arguably more realistic, we obtained several new solution concepts. For example, when there is only uncertainty about a *nonnegative* cost that a coalition experiences from deviating, the *partial* nucleolus is optimal; when the noise $D(S)$ of a coalition's value scales proportionately with its estimated value $v(S)$ (and the multiplicative noises $\{D^\times(S)\}$ are identically distributed), the *multiplicative* least core and nucleolus become optimal.

Acknowledgments

We thank ARO and NSF for support under grants W911NF-12-1-0550, W911NF-11-1-0332, IIS-0953756, CCF-1101659, and CCF-1337215.

³Even though the conditions are identical to those for the additive nucleolus, they are now applied to the distribution of the *multiplicative* noise, not the distribution of the additive noise. Thus, the set of distributions over true value V satisfying the conditions is different. This is why we obtain a different solution concept.

⁴ $v(S) = 0$ is equivalent to S being “forbidden” in the remark following the proof of Theorem 2 in Schmeidler (1969).

References

- Bachrach, Y.; Meir, R.; Feldman, M.; and Tennenholtz, M. 2012. Solving cooperative reliability games. *arXiv preprint arXiv:1202.3700*.
- Bejan, C., and Gómez, J. C. 2009. Core extensions for non-balanced tu-games. *International Journal of Game Theory* 38(1):3–16.
- Chalkiadakis, G., and Boutilier, C. 2004. Bayesian reinforcement learning for coalition formation under uncertainty. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems - Volume 3, AAMAS '04*, 1090–1097. Washington, DC, USA: IEEE Computer Society.
- Chalkiadakis, G., and Boutilier, C. 2007. Coalitional bargaining with agent type uncertainty. In *IJCAI*, 1227–1232.
- Chalkiadakis, G.; Elkind, E.; and Wooldridge, M. 2011. *Computational Aspects of Cooperative Game Theory*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Chalkiadakis, G.; Markakis, E.; and Boutilier, C. 2007. Coalition formation under uncertainty: Bargaining equilibria and the bayesian core stability concept. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, 64. ACM.
- Charnes, A., and Granot, D. 1976. Coalitional and chance-constrained solutions to n-person games. i: The prior satisficing nucleolus. *SIAM Journal on Applied Mathematics* 31(2):358–367.
- Charnes, A., and Granot, D. 1977. Coalitional and chance-constrained solutions to n-person games, ii: Two-stage solutions. *Operations research* 25(6):1013–1019.
- Conitzer, V., and Sandholm, T. 2005. Common voting rules as maximum likelihood estimators. In *Proceedings of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, 145–152.
- Drissi-Bakkkhat, M., and Truchon, M. 2004. Maximum likelihood approach to vote aggregation with variable probabilities. *Social Choice and Welfare* 23:161–185.
- Elkind, E.; Faliszewski, P.; and Slinko, A. 2009. On distance rationalizability of some voting rules. In *Proceedings of the 12th Conference on Theoretical Aspects of Rationality and Knowledge, TARK '09*, 108–117.
- Gillies, D. 1953. *Some theorems on n-person games*. Ph.D. Dissertation, Princeton University, Department of Mathematics.
- Jeong, S., and Shoham, Y. 2008. Bayesian coalitional games. In *AAAI*, 95–100.
- Maschler, M.; Peleg, B.; and Shapley, L. S. 1979. Geometric properties of the kernel, nucleolus, and related solution concepts. *Mathematics of Operations Research*.
- Meir, R.; Rosenschein, J. S.; and Malizia, E. 2011. Subsidies, stability, and restricted cooperation in coalitional games. In *IJCAI*, volume 11, 301–306.
- Meskanen, T., and Nurmi, H. 2008. Closeness counts in social choice. In *Power, Freedom, and Voting*. Springer. 289–306.
- Myerson, R. B. 2007. Virtual utility and the core for games with incomplete information. *Journal of Economic Theory* 136(1):260–285.
- Orshan, G. 1993. The prenucleolus and the reduced game property: equal treatment replaces anonymity. *International Journal of Game Theory* 22(3):241–248.
- Peleg, B. 1985. An axiomatization of the core of cooperative games without side payments. *Journal of Mathematical Economics* 14(2):203–214.
- Pivato, M. 2013. Voting rules as statistical estimators. *Social Choice and Welfare* 40(2):581–630.
- Potters, J. A. 1991. An axiomatization of the nucleolus. *International Journal of Game Theory* 19(4):365–373.
- Schmeidler, D. 1969. The nucleolus of a characteristic function game. *Society for Industrial and Applied Mathematics Journal of Applied Mathematics* 17:1163–1170.
- Shapley, L. S., and Shubik, M. 1966. Quasi-cores in a monetary economy with nonconvex preferences. *Econometrica: Journal of the Econometric Society* 805–827.
- Snijders, C. 1995. Axiomatization of the nucleolus. *Mathematics of Operations Research* 20(1):189–196.
- Sudhölter, P. 1997. The modified nucleolus: properties and axiomatizations. *International Journal of Game Theory* 26(2):147–182.
- Suijs, J., and Borm, P. 1999. Stochastic cooperative games: superadditivity, convexity, and certainty equivalents. *Games and Economic Behavior* 27(2):331–345.
- Suijs, J.; Borm, P.; De Waegenaere, A.; and Tijs, S. 1999. Cooperative games with stochastic payoffs. *European Journal of Operational Research* 113(1):193–205.
- Xia, L., and Conitzer, V. 2011. A maximum likelihood approach towards aggregating partial orders. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, 446–451.
- Young, H. P. 1988. Condorcet's theory of voting. *American Political Science Review* 82:1231–1244.
- Young, H. P. 1995. Optimal voting rules. *Journal of Economic Perspectives* 9(1):51–64.