

Recovering Games from Perturbed Equilibrium Observations Using Convex Optimization

Juba ZIANI* Venkat CHANDRASEKARAN† Katrina LIGETT‡

March 3, 2016

Abstract

We study the problem of reconstructing a game that is consistent with observed equilibrium play, a fundamental problem in econometrics. Our contribution is to develop and analyze a new methodology based on convex optimization to address this problem for many classes of games and observation models of interest. Our approach provides the flexibility to solve a number of variants and specializations of this problem, such as an evaluation of the power of games from a particular class (e.g., zero-sum, potential, linearly parameterized) to explain player behavior or the extent to which a particular set of observations tightly constrains the space of consistent explanations; it can also simply provide a compact summary of observed behavior. The framework underlying the development in this paper also differs from much of the literature on econometrics, as we do not make strong distributional assumptions on the observations of player actions.

We illustrate our approach with numerical simulations.

1 Introduction

In this paper, we study the problem of recovering properties and—when possible—payoffs of a game, based on observations of the actions of the players. There are several reasons why one might want to extrapolate beyond observations of player behavior. Reconstructing a game payoff matrix that explains observed behavior well (say, assuming observed behavior was generated by equilibrium play under perturbations of an underlying game) provides a *compact*, interpretable summary of said behavior. The reconstruction process may also yield insight into how tightly the observed behavior constrains the space of possible explanatory games—*are there multiple, wildly differing possible explanations for the observed behavior?* When the observations tightly constrain reconstruction of the underlying game, they may also yield *predictive power*; an observer who understands the payoff matrix of a game may be able to predict how player behavior will change under modifications to the underlying game, and may also be better able to manipulate game outcomes. Even when observations do not tightly constrain the space of explanatory games, one may wish to verify whether the observed behavior is *consistent with certain assumptions*—could the observed behavior have been generated by a zero-sum game? a potential game? How well is it explained by other models?

These questions are solidly within the domain of econometrics, an area that largely focuses on the identification (i.e., parameter-fitting) of simple models given observational data (data not generated by controlled experiments). Our approach is to cast the task of reconstructing and understanding games consistent with player behavior as an *optimization problem*, where observations of equilibrium play act as constraints on the space of possible explanatory games. This approach allows us to sidestep issues of model selection (we need not decide which aspects of the data to include in a model). Our approach may be viewed as complementary to a model-driven approach, in that the tools we provide here may be used to objectively

*California Institute of Technology, jziani@caltech.edu

†California Institute of Technology, venkac@caltech.edu

‡Hebrew University of Jerusalem and California Institute of Technology, katrina.ligett@mail.huji.ac.il

evaluate the quality of fit one achieves under certain modeling assumptions. Our approach also allows us to explore a variety of assumptions about the information that might be available to an observer, and the effect that it has on constraining the space of consistent games.

1.1 Summary of results

We study a setting in which, at each timestep, an observer observes the actions selected by the players in a finite, two-player game¹. We assume that players play according to some correlated equilibrium (a more permissive concept than Nash equilibria). In order to justify that the observer sees more than one fixed equilibrium of the underlying game, we assume that on each timestep, players experience (and play some equilibrium of) a slightly perturbed version of the underlying game. As discussed in the related work section, this is a standard approach.² We assume that the perturbations are “small,” but, unlike in previous literature, we do not make distributional assumptions on the perturbations. We also make no assumptions on the selection rule used by the players to decide which equilibrium to play, when multiple equilibria are present.

In this setting, we give a computationally efficient algorithm (Section 3.1), using convex optimization techniques, to identify the game(s) that best explain the observations. That is, we return a game that minimizes a measure of the extent of the perturbations necessary for the observed behaviors to correspond to equilibria of perturbed versions of the game. We can extend our framework (Section 3.2) to find the best explanatory game when we are restricted to recover games with certain linear properties, e.g., zero-sum games, potential games, and games whose utilities can be parametrized by linear functions; this allows us to determine to what extent the observed behavior is consistent with such assumptions on the underlying game.

Some sequences of observations may not sufficiently constrain the set of games that are consistent with them; in such a case, one may not be able to draw sharp conclusions about the game that generated them. To this end, we give an efficient algorithm (Section 3.3) that, given a set of observations, determines whether accurate recovery of the underlying game is impossible, i.e., whether there exist several, quite different, games that are all good explanations for the observations. Additionally (Section 5), we explore an enriched observation model where the observer also learns the expected payoff of each observed equilibrium strategy, and we give structural conditions on the sets of observations in this model that allow for accurate recovery. We also exhibit examples in which said conditions do not hold, and accurate recovery is not possible. Another source of information one might assume the observer has access to, inspired by approaches taken in econometrics, is the “payoff shifters” that affect the game payoffs from round to round; we see that these also fit nicely into our approach (Section 4).

If the payoffs and shifters of the played equilibria are not observed and no additional properties of the underlying game are assumed, the all-constant game always provides an optimal explanation for the observed behavior. However, one presumably wishes to recover a *nontrivial* consistent game. We provide a framework (Section 6) that eliminates such games by controlling the level of *degeneracy* of the explanations, and provide bounds on the trade-off between recovering non-trivial games and recovering games for which the perturbations that need be added in order to explain the observations are small.

Finally, in Section 7, we illustrate our approach with some simple simulations.

1.2 Related work

An important thread of theoretical economics takes an empirical perspective, with the goal of understanding what properties of agents are consistent with given, observed data on their economic behavior. While part of this literature focuses on discrete choice in single-agent problems, another significant line of research aims to rationalize the behavior of several agents in game theoretic settings, where their decisions impact each other.

In both the single-agent discrete choice and the multi-agent game theory settings, one important modeling issue is whether and why one would ever observe multiple, differing agent behaviors. A common approach

¹Our framework can be extended to multi-player games with succinct representations; for clarity, we focus here on the two-player case.

²This approach can also be used to model multiple observations coming from play of a game in slightly different settings, e.g., in different cities.

is to assume that the agents’ behaviors are observed in several perturbed versions of the same game. A natural, well-established approach models different observations found in the data as stemming from random perturbations to the agents’ utilities, as in [1, 4, 9, 10, 15, 16]. We adopt a similar approach here.

It is common (see [1–3, 5, 7] for example) to assume that the payoff perturbations and covariates have an observable part that is seen by the observer—usually observable economic parameters like costs or taxes—and a non-observable part, often assuming the probability distribution is known to him; the data he observes can be used to estimate the probability of different outcomes conditioned on the observed *payoff shifters* (perturbations). We demonstrate how a version of such payoff shifter information can be incorporated into our approach.

Much of the econometric literature on making inferences from observed equilibrium play restricts players’ utility functions to be parametrized, often as a linear function of covariates, and aims to find the best value of the parameter. Additionally, most of the literature does not deal with the possible presence of multiple equilibria. In an interesting departure from common approaches, [7] use random set theory and convex optimization to give a tractable representation of the *sharp identification region*—the collection of parameter values that could generate the same distribution of observations as found in the data—for a wide class of incomplete information problems with convex predictions; their results encompass finite games with multiple equilibria, and rely on information on the payoff shifters, and also on the parametric form of the utility function. Although their setting is different from ours, we use the setting of their simulations as the jumping off point for our own experimental section.

Recently, [13] study a dynamic sponsored search auction game, and provide a characterization of the *rationalizable set* consisting in the set of private parameters of a parametrized utility function that are consistent with the observations, under the relaxed assumption that players need not follow equilibrium play, but rather use some form of no-regret learning. They do so under no assumptions on the distribution of the non-observable parameters, nor on the selection rule used to choose among several possible no-regret plays.

[8, 14] study a model in which equilibrium behavior can be observed, but in a very concrete network setting, and study the query complexity of devising a variant of the game to induce desired target flows as equilibria. [6] adopt a different model in which the observer observes what joint strategies are played when restricting the actions of the players in a complete information game with no perturbations, and show that data with certain properties can be rationalized by games with low complexity.

2 Model and setting

Consider a finite two-player game G ; we will refer to it as the *true* or *underlying* game. Let $\mathcal{A}_1, \mathcal{A}_2$ be the finite sets of actions available to players 1 and 2, respectively, and let $m_1 = |\mathcal{A}_1|$ and $m_2 = |\mathcal{A}_2|$ be the number of actions available to them. For every $(i, j) \in \mathcal{A}_1 \times \mathcal{A}_2$, we denote by $G_p(i, j)$ the payoff of player p when player 1 chooses action i and player 2 chooses action j . $G_p \in \mathbb{R}^{m_1 \times m_2}$ is the vector representation of the utility of player p , and we often abuse notation and write $G = (G_1, G_2)$. The *strategies* available to player p are simply the distributions over \mathcal{A}_p . A *strategy profile* is a pair of strategies (distributions over actions), one for each player. A *joint strategy profile* is a distribution over pairs of actions (one for each player); it is not required to be a product distribution. We refer to strategies as *pure* when they place their entire probability mass on a single action, and *mixed* otherwise.

We consider l perturbed versions of the game G , indexed by $k \in [l]$ so that the k^{th} perturbed game is denoted G^k ; one can for instance imagine each G^k as a version of the game G played in a different location k . The same notation as for G applies to the G^k ’s.

Throughout the paper, we assume that for each k , the players’ strategies are given by a correlated equilibrium of the complete information game G^k . In the presence of several such equilibria, no assumption is made on the selection rule the players use to pick which equilibrium to play (though we assume they both play according to the same equilibrium). Correlated equilibria are defined as follows:

Definition 1. A probability distribution e is a correlated equilibrium of game $G = (G_1, G_2)$ if and only if

$$\begin{aligned} \sum_{j=1}^d G_1(i, j) e_{ij} &\geq \sum_{j=1}^d G_1(i', j) e_{ij} \quad \forall i, i' \in \mathcal{A}_1 \\ \sum_{i=1}^d G_2(i, j) e_{ij} &\geq \sum_{i=1}^d G_2(i, j') e_{ij} \quad \forall j, j' \in \mathcal{A}_2 \end{aligned}$$

One can see a correlated equilibrium as being a joint strategy profile recommended by some coordination device—think, for example, of a traffic light at an intersection; the traffic light coordinates the players by making sure that they never cross the intersection at the same time. The definition implies that a player does not gain anything (in expectation) by deviating to a different strategy than that recommended to her, given the conditional distribution this induces on the other player’s actions. In the traffic light example, deviating from the recommended strategy could result in a collision, severely reducing the utilities of the players. Note that the definition protects against unilateral deviations but does not consider the case in which the two players form a coalition and both deviate from their recommended strategies. The notion of correlated equilibrium extends the classical notion of Nash equilibrium by allowing players to act jointly; as every Nash equilibrium of a game is a correlated equilibrium of the same game, our results also hold when player behavior is restricted to Nash plays.

We model an observer as observing the entire correlated equilibrium distribution $e^k \in \mathbb{R}^{m_1 \times m_2}$ describing the strategies followed by the players in each perturbed game G^k , where $e^k(i, j)$ denotes the joint probability of player 1 playing action $i \in \mathcal{A}_1$ while player 2 plays action $j \in \mathcal{A}_2$. Note that as e^k represents a probability distribution, we require $e^k(i, j) \geq 0 \forall (i, j)$ and $\sum_{i,j} e_{ij}^k = 1$. The reader may interpret this assumption as describing a situation in which each perturbed game is repeatedly played over time with the two players repeatedly using the same strategies, allowing the observer to infer the probability distribution over actions that is followed by the players.

We always assume that the observer does not have access to the payoffs of the underlying game G nor of the perturbed games G^k for all k in $[l]$. We do assume that the observer has a prior on the order of magnitude of said perturbations, encoded by a parameter δ , and make no assumption on the distribution of the . In the paper, we consider three variants of the model of observations:

- In the *partial payoff information* setting, the observer has access to equilibrium observations e^1, \dots, e^l , and to the expected payoff of equilibrium e^k on perturbed games G^k , for all players p and all $k \in [l]$; we denote said payoff v_p^k and note that $v_p^k = e^k \cdot G_p^k$. See Section 5 for a discussion of the “partial payoff information” setting.
- In the *no payoff information* setting, the observer only sees e^1, \dots, e^l . See Section 6 for a discussion of the “no payoff information” setting.
- In the *payoff shifter information* setting, at each step k , a *payoff shifter* $S^k = (S_1^k, S_2^k) \in \mathcal{R}^{m_1 \times m_2} \times \mathcal{R}^{m_1 \times m_2}$ is added to game $G = (G_1, G_2)$, and the perturbed games G^k result from the further addition of small perturbations to the $G + S^k$ ’s. The observer knows S_1, \dots, S_l and observes e^1, \dots, e^l of perturbed games G^1, \dots, G^l . This is derived from a common model in the economics literature: it represents a situation in which changes in the behavior of agents are observed as a function of changes in observable economic parameters (taxes, etc.) See Section 4 for a discussion of the “payoff shifter information” setting.

Our paper aims to characterize the games that explain observations under the partial payoff, no payoff, and the payoff shifter information settings when the perturbations are known to be “small” and the perturbed games are thus “close” to the underlying game. The next two definitions formalize our notion of closeness:

Definition 2. A vector G is δ -close to vectors G^1, \dots, G^l with respect to metric d for $\delta > 0$ if and only if

$$d(G|G^1, \dots, G^l) \leq \delta.$$

For the above definition to make sense in the context of this paper, we need a notion of metric whose value on a set of payoff vectors G_p, G_p^1, \dots, G_p^l is small when G_p, G_p^1, \dots, G_p^l are close in terms of payoffs. With a model where the G^k 's are obtained from G through small and possibly bounded perturbations in mind, we consider the following two metrics:

Definition 3. *The sum-of-squares distance between vectors G and G^1, \dots, G^l is given by*

$$d_2(G^1, \dots, G^l | G) = \sum_{k=1}^l (G - G^k)'(G - G^k).$$

The maximum distance between vectors G and G^1, \dots, G^l is defined as

$$d_\infty(G^1, \dots, G^l | G) = \max_{k \in [l]} \|G - G^k\|_\infty,$$

where $\|\cdot\|_\infty$ denotes the usual infinity norm.

Both distances are useful, in different situations. The sum-of-squares distance is small when the average perturbation to G is small, but allows for worst-case perturbations to be large. An example is when the G^k 's are randomly sampled from a distribution with mean G , unbounded support, and small covariance matrix, in which case some of the perturbations may deviate from the mean with low probability, while the average squared perturbation remains small. If the distribution of perturbations is known to be i.i.d Gaussian, then the sum-of-squares norm is a natural choice, as it replicates the log-likelihood of the estimations and follows a known, Chi-square distribution. The maximum distance is small when it is known that all perturbations are small or bounded; one example is when the perturbations are known to be uniform in a small interval $[-\delta, \delta]$. While no distributional assumptions on the perturbations will be made anywhere in the paper, the observer may have a sense of which of the two distance notions best fits the game whose realized plays he is observing, or he may use our framework to understand how well the observations fit each of the two metrics.

It may be the case that, given a set of equilibrium observations with no additional assumption on the distribution of perturbations nor on the rule used to select among multiple equilibria, one cannot recover the game that generated these observations, as highlighted in the following example:

Example 1. *Take any set of observations e^1, \dots, e^l under the “no payoff information” observation model, and let \hat{G} be the all-constant game, i.e. $\hat{G}_1(i, j) = \hat{G}_2(i, j) = c$ for some $c \in \mathbb{R}$ and for all $(i, j) \in \mathcal{A}_1 \times \mathcal{A}_2$. Let $\hat{G}^1 = \dots = \hat{G}^l = \hat{G}$. Then for all $k \in [l]$, e^k is an equilibrium of \hat{G}^k , and $d_2(G^1, \dots, G^l | G) = d_\infty(G^1, \dots, G^l | G) = 0$. That is, \hat{G} is a trivial game, and it perfectly explains any arbitrary observations. Even when e^1, \dots, e^l are generated by a non-trivial G , without any additional observations, an observer cannot determine whether G or \hat{G} is the actual underlying game.*

Clearly, while our optimization framework can recover a game that explains observed equilibrium strategies, the recovered game is not necessarily unique. We thus adopt an observation-driven view that describes a class of games that are consistent with the observed behavior and bounds on the perturbations; from there, it is natural to ask what such a class of games looks like, and more specifically when the class of games that explain the observations is small and can lead to approximate point identification. We propose the following definition to formalize whether approximate point identification is possible:

Definition 4. *We say a set of observations is (ϵ, δ) -non-identifiable under a fixed observation model with respect to player p and distance metric d if there exist at least two sets of games $(\tilde{G}, \tilde{G}^1, \dots, \tilde{G}^l)$ and $(\hat{G}, \hat{G}^1, \dots, \hat{G}^l)$ such that*

- *For all k , e^k is an equilibrium of both \tilde{G}^k and \hat{G}^k . In the “partial payoff information” model of observations, we additionally require $\tilde{G}_p^k e^k = v_p^k$ and $\hat{G}_p^k e^k = v_p^k$ for all players p .*
- *In the “partial payoff information” or “no payoff information” models, we require that \tilde{G}_p is δ -close to $\tilde{G}_p^1, \dots, \tilde{G}_p^l$ and \hat{G}_p is δ -close to $\hat{G}_p^1, \dots, \hat{G}_p^l$ under metric d . In the “payoff shifter information” model, we require that \tilde{G}_p is δ -close to $\tilde{G}_p^1 - S_p^1, \dots, \tilde{G}_p^l - S_p^l$ and \hat{G}_p is δ -close to $\hat{G}_p^1 - S_p^1, \dots, \hat{G}_p^l - S_p^l$ under metric d .*

- $\|\tilde{G}_p - \hat{G}_p\|_\infty \geq \epsilon$

Informally, a set of observations is non-identifiable if there exist several games that are δ -good explanations of the observations, but are ϵ -far from each other in terms of payoff. It is then impossible to distinguish which (if either) of the candidate games is the underlying one, as they are both good explanations of the observations.

3 A convex optimization framework

In this section, we will see how techniques from convex optimization can be used to recover the best explanation for a set of observations, determine the extent to which observations are consistent with certain assumptions on the underlying game, and determine whether a set of observations tightly constrains the set of games that could explain it well. The results in this section are not tied to a specific observation model: they apply to all models considered in this paper, up to addition of linear constraints for the “partial payoff information” and “no payoff information” settings (see Programs (4) and (5)), and linear shift of the variables in the objective function for the “payoff shifter information” setting (see Program (3)).

3.1 Optimization program for recovering a compatible underlying game

One of our primary objectives is to recover a game that is a good explanation—according to the chosen distance metric—for a given set of observations from perturbed games, subject to the constraint that, for all k , each observation e^k is a correlated equilibrium of a perturbed game G^k . Note that the set of games compatible with equilibrium observation e^k is the set of games G^k that satisfy

$$\begin{aligned} \sum_{j=1}^d G_1^k(i, j) e_{ij}^k &\geq \sum_{j=1}^d G_1^k(i', j) e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1 \\ \sum_{i=1}^d G_2^k(i, j) e_{ij}^k &\geq \sum_{i=1}^d G_2^k(i, j') e_{ij}^k \quad \forall j, j' \in \mathcal{A}_2, \end{aligned}$$

This set is a convex cone defined by a polynomial (in m_1 and m_2) number of inequalities that each are linear in G^k . It is also easy to see that the distance metrics we consider are convex in (G, G^1, \dots, G^l) , thus our problem can be seen as optimizing a convex function over a convex set³:

$$\begin{aligned} \min_{G^k, G} \quad & d(G_1^1, \dots, G_1^l | G_1) + d(G_2^1, \dots, G_2^l | G_2) \\ \text{s.t.} \quad & e^k \text{ is an equilibrium of } G^k \quad \forall k \in [l] \end{aligned}$$

This optimization program simply finds the game \hat{G} that minimizes the distance d between \hat{G}_1, \hat{G}_2 and the cones of candidate perturbed games defined by equilibrium observations e^1, \dots, e^l . If the objective value of the above optimization program is small, then we have found a good explanation in the sense that we have found perturbed games $\hat{G}^1, \dots, \hat{G}^l$ that are compatible with the equilibrium observations, along with a game \hat{G} such that the distance between \hat{G} and the \hat{G}^k 's is small.

We dedicate the rest of this subsection to writing the above convex program in an efficiently solvable form for both metrics.

When we choose the metric to be the sum-of-squares distance, we obtain the following optimization problem:

³This can be extended to multi-player games with succinct representations: if there are n players with m actions each, then $n \cdot m$ equilibrium constraints are needed. Thus, a similar optimization problem with a tractable number of variables and constraints can be formulated for multi-player games so long as its utility function can be represented using a tractable number of variables.

$$\begin{aligned}
\min_{G^k, G} \quad & \sum_{k=1}^l (G_1^k - G_1)'(G_1^k - G_1) + \sum_{k=1}^l (G_2^k - G_2)'(G_2^k - G_2) \\
\text{s.t.} \quad & \sum_{j=1}^d G_1^k(i, j)e_{ij}^k \geq \sum_{j=1}^d G_1^k(i', j)e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\
& \sum_{i=1}^d G_2^k(i, j)e_{ij}^k \geq \sum_{i=1}^d G_2^k(i, j')e_{ij}^k \quad \forall i \in \mathcal{A}_1, \forall j, j' \in \mathcal{A}_2, \forall k \in [l]
\end{aligned} \tag{1}$$

The objective is separable and no constraint depends on both the G_1 's and the G_2 's, so the above problem can be decoupled and written as two problems of the same form:

$$\begin{aligned}
\min_{G^k, G} \quad & \sum_{k=1}^l (G_1^k - G_1)'(G_1^k - G_1) \\
\text{s.t.} \quad & \sum_{j=1}^d G_1^k(i, j)e_{ij}^k \geq \sum_{j=1}^d G_1^k(i', j)e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l]
\end{aligned}$$

and

$$\begin{aligned}
\min_{G^k, G} \quad & \sum_{k=1}^l (G_2^k - G_2)'(G_2^k - G_2) \\
\text{s.t.} \quad & \sum_{i=1}^d G_2^k(i, j)e_{ij}^k \geq \sum_{i=1}^d G_2^k(i, j')e_{ij}^k \quad \forall i \in \mathcal{A}_1, \forall j, j' \in \mathcal{A}_2, \forall k \in [l]
\end{aligned}$$

Both problems are quadratic optimization problems with a tractable number of constraints, therefore efficient algorithms to solve them are known.

When we choose our metric to be the maximum norm, the convex optimization program becomes

$$\begin{aligned}
\min_{G^k, G} \quad & \max_{k \in [l]} \|G_1 - G_1^k\|_\infty + \max_{k \in [l]} \|G_2 - G_2^k\|_\infty \\
\text{s.t.} \quad & \sum_{j=1}^d G_1^k(i, j)e_{ij}^k \geq \sum_{j=1}^d G_1^k(i', j)e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\
& \sum_{i=1}^d G_2^k(i, j)e_{ij}^k \geq \sum_{i=1}^d G_2^k(i, j')e_{ij}^k \quad \forall i \in \mathcal{A}_1, \forall j, j' \in \mathcal{A}_2, \forall k \in [l]
\end{aligned} \tag{2}$$

This program is again separable and we need only solve two programs of the form

$$\begin{aligned}
\min_{G_1^k, G_1} \quad & \max_{k \in [l]} \|G_1 - G_1^k\|_\infty \\
\text{s.t.} \quad & \sum_{j=1}^d G_1^k(i, j)e_{ij}^k \geq \sum_{j=1}^d G_1^k(i', j)e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l]
\end{aligned}$$

for each player. Note that this program can be written as a LP and therefore solved efficiently:

$$\begin{aligned}
\min_{G_1^k, G_1} \quad & \delta \\
\text{s.t.} \quad & \sum_{i=1}^d G_1^k(i, j)e_{ij}^k \geq \sum_{i=1}^d G_1^k(i, j')e_{ij}^k \quad \forall j, j' \in [m], \forall k \in [l] \\
& G_1^k(i, j) - G_1(i, j) \leq \delta \quad \forall (i, j) \in \mathcal{A}_1 \times \mathcal{A}_2, \forall k \in [l] \\
& G_1(i, j) - G_1^k(i, j) \leq \delta \quad \forall (i, j) \in \mathcal{A}_1 \times \mathcal{A}_2, \forall k \in [l]
\end{aligned}$$

3.2 Can observations be explained by linear properties?

This convex optimization-based approach can be used to determine whether there exists a game that is compatible with the observations and also has certain properties, as long as these properties can be written as a tractable number of linear equalities and inequalities. This section contains a non-exhaustive list of interesting properties that fit this framework:

3.2.1 Zero-sum games

Definition 5. A 2-player game $G = (G_1, G_2)$ is a zero-sum game if and only if $G_1 = -G_2$.

In other words, a zero-sum game is a game in which for each pure strategy (i, j) , the sum of the payoff of player 1 and the payoff of player 2 for (i, j) is 0. One can restrict the set of games we look for to be zero-sum games, at the cost of separability of Program (1), by solving:

$$\begin{aligned} \min_{G^k, G} \quad & d(G_1^1, \dots, G_1^l | G_1) + d(G_2^1, \dots, G_2^l | G_2) \\ \text{s.t.} \quad & \sum_{j=1}^d G_1^k(i, j) e_{ij}^k \geq \sum_{j=1}^d G_1^k(i', j) e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\ & \sum_{i=1}^d G_2^k(i, j) e_{ij}^k \geq \sum_{i=1}^d G_2^k(i, j') e_{ij}^k \quad \forall i \in \mathcal{A}_1, \forall j, j' \in \mathcal{A}_2, \forall k \in [l] \\ & G_1(i, j) = -G_2(i, j) \quad \forall (i, j) \in \mathcal{A}_1 \times \mathcal{A}_2 \end{aligned}$$

Note that depending on the choice of distance function, this is either a linear or a quadratic program with a tractable number of linear constraints, and can therefore be solved efficiently.

3.2.2 Exact potential games

Definition 6. A game G is an exact potential game if and only if there exists a function Φ such that for all players p , all pure strategy profiles (a_p, a_{-p}) and all deviations $a'_p \in \mathcal{A}_p$ for player p ,

$$\Phi(a_i, a_{-i}) - \Phi(a'_i, a_{-i}) = G_p(a_p, a_{-p}) - G_p(a'_p, a_{-p}).$$

Φ is called an exact potential function for game G .

In particular, for 2-player games, the definition translates to the existence of a function Φ such that

$$\begin{aligned} \Phi(i, j) - \Phi(i', j) &= G_1(i, j) - G_1(i', j) \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2 \\ \Phi(i, j) - \Phi(i, j') &= G_2(i, j) - G_2(i, j') \quad \forall i \in \mathcal{A}_1, \forall j, j' \in \mathcal{A}_2 \end{aligned}$$

In order to restrict the set of games we are searching over to the set of potential games, one can introduce $m_1 m_2$ variables $\Phi(i, j)$ in Program (1) and solve

$$\begin{aligned} \min_{G^k, G, \Phi} \quad & d(G_1^1, \dots, G_1^l | G_1) + d(G_2^1, \dots, G_2^l | G_2) \\ \text{s.t.} \quad & \sum_{j=1}^d G_1^k(i, j) e_{ij}^k \geq \sum_{j=1}^d G_1^k(i', j) e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\ & \sum_{i=1}^d G_2^k(i, j) e_{ij}^k \geq \sum_{i=1}^d G_2^k(i, j') e_{ij}^k \quad \forall i \in \mathcal{A}_1, \forall j, j' \in \mathcal{A}_2, \forall k \in [l] \\ & \Phi(i, j) - \Phi(i', j) = G_1(i, j) - G_1(i', j) \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2 \\ & \Phi(i, j) - \Phi(i, j') = G_2(i, j) - G_2(i, j') \quad \forall i \in \mathcal{A}_1, \forall j, j' \in \mathcal{A}_2 \end{aligned}$$

Once again, this is—depending on the choice of objective function—either a linear or a quadratic program with a tractable number of variables and constraints, thus it can be solved efficiently.

3.2.3 Games generated through linear parameter fitting

It is common in the economics literature to recover a game with the help of a parametrized function whose parameters are calibrated using the observations. In many applications, linear functions in the actions of the players (and in the state variables, when they exist) are considered—entry games come to mind. Our framework allows one to determine whether there exist parameters for such a linear function that provide good explanation for the observations. When such parameters exist, one can use it to find a set of parameters that describe a game which is consistent with the observations.

ALGORITHM 1: Measuring (ε, δ) non-identifiability for player 1

Input: Equilibrium observations e^1, \dots, e^l , parameter ε , metric d

Output: δ such that e^1, \dots, e^l is (ε, δ) -non-identifiable

$\delta = +\infty$;

for $(i, j) \in \mathcal{A}_1 \times \mathcal{A}_2$ **do**

$$\begin{aligned} P_\varepsilon(i, j) &= \min_{G_1^k, \tilde{G}_1^k, G, \hat{G}} d(\tilde{G}_1^1, \dots, \tilde{G}_1^l | \tilde{G}) + d(\hat{G}_1^1, \dots, \hat{G}_1^l | \hat{G}_1) \\ \text{s.t.} \quad &\sum_{j=1}^d \tilde{G}_1^k(i, j) e_{ij}^k \geq \sum_{j=1}^d \tilde{G}_1^k(i', j) e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\ &\sum_{j=1}^d \hat{G}_1^k(i, j) e_{ij}^k \geq \sum_{j=1}^d \hat{G}_1^k(i', j) e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\ &\tilde{G}_1(i, j) - \hat{G}_1(i, j) \geq \varepsilon \end{aligned}$$

if $P_\varepsilon(i, j) \leq \delta$ **then**

$\delta = P_\varepsilon(i, j)$

end

end

Take a linear function f_α that takes pure strategies (i, j) as input and parameters given by vector α , we can efficiently solve for player 1 (a similar program holds for player 2) as

$$\begin{aligned} \min_{G_1^k, G_1, \alpha} \quad &d(G_1^1, \dots, G_1^l | G_1) \\ \text{s.t.} \quad &\sum_{j=1}^d G_1^k(i, j) e_{ij}^k \geq \sum_{j=1}^d G_1^k(i', j) e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\ &G_1(i, j) = f_\alpha(i, j) \quad \forall i \in \mathcal{A}_1, \forall j \in \mathcal{A}_2 \end{aligned}$$

for $d = d_2$ or $d = d_\infty$. Note that unlike in previous work in economics, such a framework allows us to account for multiple equilibria at the same time, and does not rely on assumptions on the distribution of perturbations.

3.3 Deciding whether a set of observations is identifiable

In this section, we provide an algorithm for characterizing the level of identifiability of a set of observations for a player, Algorithm 1. Algorithm 1 is clearly computationally efficient, as it solves $m_1 m_2$ linear or quadratic—depending on the chosen metric—optimization programs with a tractable number of constraints, and has the following properties:

Lemma 1. *Let δ be the output of Algorithm 1 run with inputs $\varepsilon, e^1, \dots, e^l$. Then e^1, \dots, e^l is (ε, δ) -non-identifiable. Furthermore, suppose that e^1, \dots, e^l is (ε, γ) -non-identifiable, then $\gamma \geq \frac{\delta}{2}$.*

It immediately follows that Algorithm 1 returns a 2-approximation of the minimum value of δ such that the observations are (ε, δ) -non-identifiable.

Proof. See Appendix A. □

Note that while this optimization program is written with respect to player 1, a similar optimization program can be solved to characterize the level of non-identifiability of the observations for player 2. This algorithm can also be applied when additional constraints are included, as long as they are linear in the variables of the optimization program, or when the objective value is modified, as long as it remains linear or quadratic in the variables; we consider such modifications in Sections 3.2, 4, 5 and 6.

Remark 1. *It is important to note that (ε, δ) -non-identifiability is a property of the observations, not of the underlying game nor of our framework. On the one hand, if a set of observations is non-identifiable, then*

by definition there is no way of distinguishing between different explanations for the same observations and accurate point recovery is impossible for any algorithm that does not make additional assumptions on how the observations are generated; the best one can do in such a situation is to give a characterization of the games compatible with the observations. On the other hand, if a set of observations is identifiable, this paper provides an algorithm to obtain an approximation of the underlying game that generated the observations.

4 Consistent games with payoff shifter information

In this section, we briefly show how to adapt our framework to the “payoff shifter information” observation model. Here, our primary objective is to recover a game that is a good explanation—according to the chosen distance metric—for a given set of observations from perturbed games, subject to the constraint that, for all k , each observation e^k is a correlated equilibrium of a perturbed game G^k . As before, the set of games compatible with equilibrium observation e^k is the set of games G^k that satisfy

$$\begin{aligned} \sum_{j=1}^d G_1^k(i, j)e_{ij}^k &\geq \sum_{j=1}^d G_1^k(i', j)e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1 \\ \sum_{i=1}^d G_2^k(i, j)e_{ij}^k &\geq \sum_{i=1}^d G_2^k(i, j')e_{ij}^k \quad \forall j, j' \in \mathcal{A}_2, \end{aligned}$$

However, in the “payoff shifter information” model, we aim to minimize the distance between the game G and the “unshifted” G^1, \dots, G^l , as the perturbed games are obtained from G by adding known payoff shifters S^1, \dots, S^l and then also small, unknown perturbations. To do so, we simply consider objective function $d(G_1^1 - S_1^1, \dots, G_1^l - S_1^l | G_1) + d(G_2^1 - S_2^1, \dots, G_2^l - S_2^l | G_2)$ and solve

$$\begin{aligned} \min_{G^k, G} \quad & d(G_1^1 - S_1^1, \dots, G_1^l - S_1^l | G_1) + d(G_2^1 - S_2^1, \dots, G_2^l - S_2^l | G_2) \\ \text{s.t.} \quad & \sum_{j=1}^d G_1^k(i, j)e_{ij}^k \geq \sum_{j=1}^d G_1^k(i', j)e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\ & \sum_{i=1}^d G_2^k(i, j)e_{ij}^k \geq \sum_{i=1}^d G_2^k(i, j')e_{ij}^k \quad \forall i \in \mathcal{A}_1, \forall j, j' \in \mathcal{A}_2, \forall k \in [l] \end{aligned} \quad (3)$$

5 Consistent games with partial payoff information

This section considers the *partial payoff information* variant of the observation model described in Section 2. Recall that in this setting, for an equilibrium e^k observed from perturbed game G^k , the observer also learns the expected payoff v_p^k of player p in said equilibrium strategy on game G^k , on top of observing e^k . Similar to the previous sections, we are interested in computing a game \hat{G} that is close to some perturbed games $\hat{G}^1, \dots, \hat{G}^l$ that (respectively) have equilibria e^1, \dots, e^l with payoffs v^1, \dots, v^l ; to do so, one can solve the following convex optimization problem for player 1 (a similar optimization problem can be solved for player 2):

$$\begin{aligned} P(\epsilon) = \quad & \min_{G_1^k, G_1} d(G_1^1, \dots, G_1^l | G_1) \\ \text{s.t.} \quad & \sum_{j=1}^d G_1^k(i, j)e_{ij}^k \geq \sum_{j=1}^d G_1^k(i', j)e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\ & e^k \cdot G_1^k = v_1^k \quad \forall k \in [l] \end{aligned} \quad (4)$$

The first constraint means that e^k must be an equilibrium of the recovered \hat{G}^k , while the second constraint imposes that the expected payoff of strategy e^k on game \hat{G}^k is v^k . Note that unlike in Example 1, if G is far from a trivial game, information on the difference in the payoffs of different strategies is given by the v^k s, and the game we recover will not have a trivial structure. All results from Section 3 can be extended to

the partial payoff information setting by adding the linear expected payoff constraint to each optimization program.

5.1 Quantifying identifiability of many linearly independent equilibrium observations

We take $l \geq m_1 m_2$ and make the following assumption for the remainder of this subsection, unless otherwise specified:

Assumption 1. *There exists a subset $E \subset \{e^1, \dots, e^l\}$ of size $m_1 m_2$ such that the vectors in E are linearly independent.*

We abuse notation and denote by E the $m_1 m_2 \times m_1 m_2$ matrix in which row i is given by the i^{th} element of set E , for all $i \in [m_1 m_2]$. For every $p \in \mathbb{N} \cup \{+\infty\}$, let $\|\cdot\|_p$ be the p-norm. We can define the corresponding induced matrix norm $\|\cdot\|_p$ that satisfies $\|M\|_p = \sup_{x \neq 0} \frac{\|Mx\|_p}{\|x\|_p}$ for any matrix $M \in \mathbb{R}^{m_1 m_2 \times m_1 m_2}$.

The following statement highlights that if one has $m_1 m_2$ linearly independent observations (among the l equilibrium observations) such that the induced matrix of observations E is well-conditioned, and the perturbed games are obtained from the underlying game through small perturbations, any optimal solution of Program (4) necessarily recovers a game whose payoffs are close to the payoffs of the underlying game. The statements are given for both metrics introduced in Section 2.

Lemma 2. *Let G be the underlying game, and G^1, \dots, G^l be the games generating observations e^1, \dots, e^l , where $l = m_1 m_2$. Suppose that for player p , $d_2(G_p^1, \dots, G_p^l | G_p) \leq \delta$. Let $(\hat{G}_p, \hat{G}_p^1, \dots, \hat{G}_p^l)$ be an optimal solution of Program (4) for player p with distance function d_2 . Then*

$$\|G_p - \hat{G}_p\|_2 \leq \sqrt{2\|E^{-1}\|_2 \cdot \delta}.$$

Proof. For simplicity of notation, we drop the p indices. We first remark that (G, G^1, \dots, G^l) is feasible for Program (4); as $(\hat{G}, \hat{G}^1, \dots, \hat{G}^l)$ is optimal, it is necessarily the case that

$$\sum_{k=1}^l \|\hat{G} - \hat{G}^k\|_2^2 \leq \sum_{k=1}^l \|G - G^k\|_2^2 \leq \delta.$$

Let us write $\Delta G = G - \hat{G}$. We know that for all k , $e^k{}'G^k = e^k{}'\hat{G}^k = v^k$, and thus $e^k{}'(G^k - \hat{G}^k) = 0$. We can write

$$\begin{aligned} E\Delta G &= (e_1'(G - \hat{G}) \dots e_l'(G - \hat{G}))' \\ &= (e_1'(G - G^1 + G^1 - \hat{G}^1 + \hat{G}^1 - \hat{G}) \dots e_l'(G - G^l + G^l - \hat{G}^l + \hat{G}^l - \hat{G}))' \\ &= (e_1'(G - G^1 + \hat{G}^1 - \hat{G}) \dots e_l'(G - G^l + \hat{G}^l - \hat{G}))'. \end{aligned}$$

Let $x_k = G - G^k + \hat{G}^k - \hat{G}$. We then have $\|E\Delta G\|_2^2 = \sum_{k=1}^l x_k' e_k e_k' x_k$, as $e_k e_k'$ is a symmetric, positive semi-definite, stochastic matrix, all its eigenvalues are between 0 and 1 and

$$\|E\Delta G\|_2^2 \leq \sum_{k=1}^l x_k' x_k = \sum_{k=1}^l \|x_k\|_2^2 \leq 2\delta.$$

It immediately follows that $\|\Delta G\|_2 \leq \sqrt{2\|E^{-1}\|_2 \cdot \delta}$. □

Lemma 3. *Let G be the underlying game, and G^1, \dots, G^l be the games generating observations e^1, \dots, e^l , where $l = m_1 m_2$. Suppose that for player p , $d_\infty(G_p^1, \dots, G_p^l | G_p) \leq \delta$. Let $(\hat{G}_p, \hat{G}_p^1, \dots, \hat{G}_p^l)$ be an optimal solution of Program (4) for player p with distance function d_∞ . Then*

$$\|G_p - \hat{G}_p\|_\infty \leq 2\|E^{-1}\|_\infty \cdot \delta.$$

Proof. See Appendix B. □

When E is far from being singular, as long as the perturbations are small, we can accurately recover the payoff matrix of each player. An extreme example arises when we take E to be the identity matrix, in which case we observe every single pure strategy of the game and an approximation of the payoff of each of these strategies, allowing us to approximately reconstruct the game. It is also the case that there are examples in which $\|E^{-1}\|_\infty$ is large and there exist two games that are far from one another, yet both explain the observations, making our bound essentially tight:

Example 2. Consider the square matrix $E \in \mathbb{R}^{4 \times 4}$ with probability $0.25 + \epsilon$ on the diagonal and $\frac{0.75 + \epsilon}{3}$ off the diagonal, i.e., we get four equilibrium observations with a different action profile that has probability slightly higher than 0.25 for each equilibrium; the first equilibrium has a higher probability on action profile (1,1), the second on (1,2), the third on (2,1) and the last one on (2,2). Suppose the vector of observed payoffs is $v = (\delta, -\delta, \delta, -\delta)'$, where $v(i)$ is the payoff for the i^{th} equilibrium. Note that there exists a constant C such that for all $\epsilon > 0$ small enough, $\|E^{-1}\|_{+\infty} \leq \frac{C}{\epsilon}$.

In the rest of the example, we fix the payoff matrix of player 2 for all considered games to be all zero so that it is consistent with every equilibrium observation, and describe a game through the payoff matrix of player 1. Let G be the all-zero game, $G^1 = G^3$ be the game with payoff $\frac{\delta}{0.5 + 2\epsilon/3}$ on actions (1,1) and (1,2) and 0 everywhere else, and $G^2 = G^4$ be the game with payoff $-\frac{\delta}{0.5 + 2\epsilon/3}$ on actions (2,1) and (2,2) and 0 everywhere else. It is clear that the G^i 's are consistent with the payoff observations as the payoffs are constant across rows on the same column, making no deviation profitable, and that the payoff of each equilibrium is indeed δ . We have

$$d_\infty(G^1, G^2, G^3, G^4 | G) = \frac{\delta}{0.5 + 2\epsilon/3} \leq 2\delta$$

and

$$\lim_{\epsilon \rightarrow 0} d_\infty(G^1, G^2, G^3, G^4 | G) = 2\delta.$$

Now, take \hat{G} to be the game that has payoff δ/ϵ for action profiles (1,1) and (1,2), and $-\delta/\epsilon$ for (2,1) and (2,2). Take $\hat{G}^1 = \hat{G}^3$ to be the game with payoffs $\frac{\delta}{\epsilon}$ in the first column, and $-\frac{\delta}{\epsilon} \frac{3-2\epsilon}{3-4\epsilon}$ in the second column; similarly, take $\hat{G}^2 = \hat{G}^4$ to be the game with payoffs $\frac{\delta}{\epsilon} \frac{3-2\epsilon}{3-4\epsilon}$ in the first column and $-\frac{\delta}{\epsilon}$ in the second column. It is clear that the observations are equilibria of the \hat{G}^i 's and yield payoff δ . Now, note that for $\epsilon < 3/4$,

$$d_\infty(G^1, G^2, G^3, G^4 | \hat{G}) = \frac{\delta}{\epsilon} \left| 1 - \frac{3-2\epsilon}{3-4\epsilon} \right| = \frac{2}{3-4\epsilon} \delta$$

Therefore, both G and \hat{G} are good explanations of the equilibrium observations, in the sense that for $\epsilon \leq 1/4$, G is δ -close to G^1, \dots, G^l and \hat{G} is δ -close to $\hat{G}^1, \dots, \hat{G}^l$ that have e^1, \dots, e^l as equilibria, respectively. However,

$$\|G - \hat{G}\|_\infty = \frac{\delta}{\epsilon} - \frac{\delta}{0.5 + 2\epsilon/3} \geq \delta \left(\frac{1}{\epsilon} - 2 \right),$$

which immediately implies

$$\|G - \hat{G}\|_\infty = \Omega_{\epsilon \rightarrow 0} \left(\frac{\delta}{\epsilon} \right) = \Omega_{\epsilon \rightarrow 0} (\|E^{-1}\|_\infty \delta).$$

Remark 2. In the case of sparse games, in which some action profiles are never profitable to the players, and are therefore never played, one can reduce the number of linearly independent, well-conditioned observations needed for accurate recovery. Under the assumption that the action profiles that are never played with positive probability have payoffs strictly worse than the lowest payoff of any action profile played with non-zero probability, one can solve the optimization problem on the restricted set of action profiles that are observed in at least one equilibrium, and set the payoffs of the remaining action profiles to be lower than the lowest

payoff of the recovered subgame, without affecting the equilibrium structure of the game. While the recovered game may not be the unique good explanation of the observations when looking at the full payoff matrix, it is unique with respect to the subgame of non-trivial actions when one has access to sufficiently many linearly independent, well-conditioned equilibrium observations.

6 Finding consistent games without additional information

This section focuses on the *no payoff information* variant of the observation model given in Section 2. Recall that in this setting, the observer only observes what equilibrium e^k is played for each perturbed game G^k .

6.1 Finding non-degenerate games

In this section, we focus on the optimization problem that recovers the payoffs of player 1 (by symmetry, all results can be applied to the optimization program that recovers the payoffs player 2), and drop the player indices for notational simplicity. Since no payoff information is given, throughout this section, we assume w.l.o.g that the games are normalized to have all payoffs between 0 and 1. As mentioned in Example 1, the all-constant game $G = G^1 = \dots = G^l$ gives an optimal solution to our optimization problem, as such a game is compatible with all equilibrium observations and has an objective value $d(G^1, \dots, G^l | G) = 0$. It is therefore the case that solving our optimization problem might output a degenerate game, so in this section, we provide a framework that allows us to control the degree of degeneracy of the game we recover and to avoid trivial, all-constant games. To do so, we require some of the equilibria of the games to be “strict,” in the sense that

$$\sum_{j=1}^d G(i, j)x_{ij} \geq \sum_{j=1}^d G(i', j)x_{ij} + \varepsilon_{ii'} \quad \forall i, i'$$

with $\varepsilon_{ii'} \geq 0$ and with the condition that at least one of the $\varepsilon_{ii'}$ is non-zero. All-constant games do not have strict equilibria, thus this avoids such games. Note that such a technique only affect the payoffs of pure strategies that are played with positive probability, and does not accord any importance to strategies that are never played. Let us now consider the new problem:

$$\begin{aligned} \min_{G^k, G} \quad & d(G^1, \dots, G^l | G) \\ \text{s.t.} \quad & e^k \text{ is a “strict” equilibrium of } G^k, \forall k \\ & 0 \leq G(i, j) \leq 1, \forall (i, j) \end{aligned}$$

which can be rewritten as

$$\begin{aligned} \min_{G^k, G} \quad & d(G^1, \dots, G^l | G) \\ \text{s.t.} \quad & \sum_{j=1}^d G^k(i, j)e_{ij}^k = \sum_{j=1}^d G^k(i', j)e_{ij}^k + \varepsilon_{ii'}^k \quad \forall (i, i'), \forall k \\ & 0 \leq G(i, j) \leq 1, \forall (i, j) \end{aligned}$$

We introduce a positive parameter ε that controls the level of non-degeneracy of the game and let the optimization program decide how to split ε among the $\varepsilon_{ii'}^k$'s in a way that minimizes the objective. The optimization program can now be written as

$$\begin{aligned} P(\varepsilon) = \min_{G^k, G} \quad & d(G^1, \dots, G^l | G) \\ \text{s.t.} \quad & \sum_{j=1}^d G^k(i, j)e_{ij}^k = \sum_{j=1}^d G^k(i', j)e_{ij}^k + \varepsilon_{ii'}^k \quad \forall (i, i'), \forall k \\ & \sum_{k=1}^l \sum_{i, i'} \varepsilon_{ii'}^k = \varepsilon \\ & 0 \leq G \leq 1 \\ & \varepsilon_{ii'}^k \geq 0 \quad \forall (i, i'), \forall k \end{aligned}$$

For all $i, i' \in \mathcal{A}_1$ such that $i \neq i'$ and $k \in [l]$, we introduce vectors $\tilde{e}_{ii'}^k$ whose entries are defined as follows:

$$\tilde{e}_{ii'}^k(h, j) = \begin{cases} -e^k(i, j) & \text{if } h = i \\ e^k(i, j) & \text{if } h = i' \\ 0 & \text{if } h \neq i, i' \end{cases}$$

This allows us to rewrite the optimization program under the following form:

$$\begin{aligned} P(\varepsilon) = & \min_{G^k, G} d(G^1, \dots, G^l | G) \\ \text{s.t.} & \sum_{k, i, i'} \tilde{e}_{ii'}^k G^k = -\varepsilon \\ & \tilde{e}_{ii'}^k G^k \leq 0 \quad \forall (i, i'), \forall k \\ & 0 \leq G \leq 1 \end{aligned} \tag{5}$$

This optimization problem is, depending on the chosen metric, either a linear or quadratic optimization program with a tractable number of constraints, and can therefore be efficiently solved.

6.2 A duality framework

In this section, we give a duality framework under distance d_2 that offers insight into the solutions to the optimization program. Throughout the section, we let $D(\varepsilon)$ be the dual of Program (5).

6.2.1 Sufficient conditions for strong duality

Claim 1. *If there exist G^1, \dots, G^l such that*

$$\tilde{e}_{ii'}^k G^k < 0 \quad \forall (i, i') \in cA_1, \forall k \in [l] \text{ s.t. } \tilde{e}_{ii'}^k \neq 0,$$

then strong duality holds and $P(\varepsilon) = D(\varepsilon)$.

Proof. Slater's condition holds iff there exists a solution G, G^1, \dots, G^l such that

$$\begin{aligned} \sum_{k, i, i'} \tilde{e}_{ii'}^k G^k &= -\varepsilon \\ \tilde{e}_{ii'}^k G^k &< 0 \quad \forall (i, i'), \forall k \text{ s.t. } \tilde{e}_{ii'}^k \neq 0. \end{aligned}$$

It is enough to find G^1, \dots, G^l such that

$$\tilde{e}_{ii'}^k G^k < 0 \quad \forall (i, i'), \forall k \text{ s.t. } \tilde{e}_{ii'}^k \neq 0$$

as we can then renormalize the G^k 's such that $\sum_{k=1}^l \sum_{i, i'} \varepsilon_{ii'}^k = \varepsilon$. □

Note that the previous sufficient condition is not necessarily tractable to check. We give a stronger sufficient condition such that for any fixed k , $\tilde{e}_{ii'}^k G^k < 0 \quad \forall (i, i')$ has a solution:

Lemma 4. *Let $k \in [l]$. Let $e^k(i, \cdot) = (e^k(i, 1), \dots, e^k(i, m_2)) \quad \forall i \in \mathcal{A}_1$. If the non-null $e^k(1, \cdot), \dots, e^k(m_1, \cdot)$ are linearly independent, then the non-null $\tilde{e}_{ii'}^k$'s are linearly independent. In particular, there exists G^k such that*

$$\tilde{e}_{ii'}^k G^k < 0, \quad \forall i, i' \in \mathcal{A}_1.$$

If this holds for all $k \in [l]$, then $P(\varepsilon) = D(\varepsilon)$.

Proof. Let $\alpha(h, h')$'s be such that $\sum_{h, h'} \alpha(h, h') \tilde{e}_{hh'}^k = 0$, and so $\sum_{h, h'} \alpha(h, h') \tilde{e}_{hh'}^k(i, j) = 0 \forall (i, j)$. Recall that for a fixed (i, j) , $\tilde{e}_{h, h'}^k(i, j) \neq 0$ only if $h = i$ or $h' = i$, but not both at the same time. Therefore,

$$\sum_{h, h'} \alpha(h, h') \tilde{e}_{hh'}^k(i, j) = \sum_{h' \neq i} \alpha(i, h') \tilde{e}_{i, h'}^k(i, j) + \sum_{h \neq i} \alpha(h, i) \tilde{e}_{h, i}^k(i, j).$$

As $\tilde{e}_{i, h'}^k(i, j) = -e^k(i, j)$ and $\tilde{e}_{h, i}^k(i, j) = e^k(h, j)$, we have for all (i, j) that

$$-e^k(i, j) \sum_{h' \neq i} \alpha(i, h') + \sum_{h \neq i} \alpha(h, i) e^k(h, j) = \sum_{h, h'} \alpha(h, h') \tilde{e}_{hh'}^k(i, j) = 0.$$

Since this holds for all values of j , it immediately follows that for all i ,

$$-e^k(i, :) \sum_{h' \neq i} \alpha(i, h') + \sum_{h \neq i} \alpha(h, i) e^k(h, :) = 0.$$

Take any i, i' such that $e_{ii'}^k \neq 0$. Then $e^k(i, :) \neq 0$ and $e^k(i', :) \neq 0$. By the previous equation, we have

$$\begin{aligned} -e^k(i', :) \sum_{h' \neq i'} \alpha(i', h') + \alpha(i, i') e^k(i, :) + \sum_{h \neq i, i'} \alpha(h, i') e^k(h, :) \\ = -e^k(i', :) \sum_{h' \neq i'} \alpha(i', h') + \sum_{h \neq i'} \alpha(h, i') e^k(h, :) \\ = 0 \end{aligned}$$

and by the linear independence assumption, we necessarily have $\alpha(i, i') = 0$. Therefore, the $\tilde{e}_{ii'}^k \neq 0$'s are linearly independent, completing the proof. \square

Note that in the worst case, we want $m_1 \leq m_2$, as there can be up to m_1 non-null $e^k(i, :)$ of size m_2 , and by symmetry, we want $m_2 \leq m_1$ for the program that recovers the payoffs of player 2. From now on, we require $m_1 = m_2$, which can be obtained by adding dummy actions to \mathcal{A}_p of player p with the least available actions. When the condition does not hold in such a setting, it can be obtained through small perturbations of the equilibrium observations.

6.2.2 Dual program

The dual of program (5) is given by:

Theorem 1. *The dual of optimization problem 5 is given by:*

$$\begin{aligned} D(\varepsilon) = \max_{\mu_{ii'}^k, \lambda_0, \lambda_1} & -\frac{1}{4} \sum_{k=1}^l (\sum_{i, i'} \mu_{ii'}^k \tilde{e}_{ii'}^k)' (\sum_{i, i'} \mu_{ii'}^k \tilde{e}_{ii'}^k) - \mathbb{1}' \lambda_1 - \mu \varepsilon \\ \text{s.t.} & \lambda_1 - \lambda_0 + \sum_{k, i, i'} \mu_{ii'}^k \tilde{e}_{ii'}^k = 0 \\ & \mu + \mu_{ii'}^k \geq 0 \\ & \lambda_0, \lambda_1 \geq 0 \end{aligned} \tag{6}$$

The KKT conditions imply that if $(G^{1*}, \dots, G^{l*}, G^*)$ is a primal optimal solution and $(\lambda_0^*, \lambda_1^*, \mu^*, \mu_{ii'}^{k*})$ is a dual optimal solution, then

$$\begin{aligned} \forall k, G^{k*} &= A - \frac{1}{2} (\frac{\lambda_1^* - \lambda_0^*}{l} + \sum_{i, i'} \mu_{ii'}^{k*} \tilde{e}_{ii'}^k) \\ G^* &= A - \frac{1}{2l} (\lambda_1^* - \lambda_0^*) \end{aligned} \tag{7}$$

for some matrix $A \in \mathbb{R}^{l \times l}$

Proof. See Appendix C. \square

This duality framework will allow us to obtain bounds on the trade-off between degeneracy and accuracy in the next subsection.

6.3 Properties of the optimization program

6.3.1 Equilibria of the recovered game

Lemma 5. *Let $\hat{G}, \hat{G}^1, \dots, \hat{G}^l$ be a solution of Program (5). Then for all k , e^k is an $2\epsilon_k$ -approximate equilibrium of \hat{G} , where $\sum_{k=1}^l \epsilon_k^2 = P(\varepsilon)$.*

Proof. Take $\epsilon_k^2 = (\hat{G} - \hat{G}^k)'(\hat{G} - \hat{G}^k)$. It immediately follows that $\|\hat{G} - \hat{G}^k\|_\infty \leq \epsilon_k$. Since e^k is an equilibrium of \hat{G}^k , it must be a $2\epsilon_k$ -approximate equilibrium of \hat{G} . \square

Therefore, as long as $P(\varepsilon)$ is small, the ϵ_k 's are small, the observed equilibria are equilibria of the recovered game G up to small perturbations in average, and G has all of the observed equilibria of the underlying game and its close-by games. However, the statement does not make sense unless it is possible to obtain a small value of $P(\varepsilon)$ for a non-trivial degree of degeneracy. Otherwise, it is only natural that G has all of the observed equilibria of the underlying game, as it also has all of the possible equilibria, being a trivial game. Understanding the trade-off between degeneracy and accuracy lets us understand whether it is possible to obtain a non-trivial game that has most of the observations as approximate equilibria.

6.3.2 Trade-off between degeneracy and objective value

Definition 7. *We define the degeneracy threshold ε^* of a set of observations as*

$$\varepsilon^* = \sup\{\varepsilon \text{ s.t. } P(\varepsilon) = 0\}$$

Claim 2. *The degeneracy threshold is given by*

$$\begin{aligned} \varepsilon^* = & - \min_{G, \varepsilon_{ii'}^k} \sum_{k, i, i'} \tilde{e}_{ii'}^k / G \\ \text{s.t.} & \quad \tilde{e}_{ii'}^k / G \leq 0 \quad \forall (i, i'), \forall k \\ & \quad 0 \leq G \leq 1 \end{aligned} \tag{8}$$

The claim gives a tractable linear program to solve for the degeneracy threshold.

Proof. Remark that ε^* solves

$$\begin{aligned} \varepsilon^* = & \max_{G^k, G, \varepsilon_{ii'}^k} \varepsilon \\ \text{s.t.} & \quad \tilde{e}_{ii'}^k / G^k + \varepsilon_{ii'}^k = 0 \quad \forall (i, i'), \forall k \\ & \quad \sum_{k=1}^l (G^k - G)'(G^k - G) = 0 \\ & \quad \sum_{k, i, i'} \varepsilon_{ii'}^k = \varepsilon \\ & \quad \varepsilon_{ii'}^k \geq 0 \quad \forall (i, i'), \forall k \\ & \quad 0 \leq G \leq 1 \end{aligned}$$

From the fact that $\sum_{k=1}^l (G^k - G)'(G^k - G) = 0$ implies $G^1 = \dots = G^l = G$, we have

$$\begin{aligned} \varepsilon^* = & \max_{G, \varepsilon_{ii'}^k} \sum_{k, i, i'} \varepsilon_{ii'}^k \\ \text{s.t.} & \quad \tilde{e}_{ii'}^k / G + \varepsilon_{ii'}^k = 0 \quad \forall (i, i'), \forall k \\ & \quad \varepsilon_{ii'}^k \geq 0 \quad \forall (i, i'), \forall k \\ & \quad 0 \leq G \leq 1 \end{aligned}$$

The result follows immediately. \square

Claim 3. ε^* is finite, $\forall \varepsilon \leq \varepsilon^*$, $P(\varepsilon) = 0$, and $\forall \varepsilon > \varepsilon^*$, $P(\varepsilon) > 0$.

Proof. The proof follows immediately from claim 8. $P(\varepsilon^*) = 0$ comes from the fact that the feasible set of Program (8) is bounded: indeed, for any point in its feasible set, $0 \leq G \leq 1$ and $\varepsilon_{ii'}^k = -\tilde{\varepsilon}_{ii'}^k G$, forcing the $\varepsilon_{ii'}^k$ to also be bounded. Thus, ε^* is a solution of a linear program on a bounded polytope and is therefore finite, and attained at an extreme point of this polytope. \square

Note that if we solve optimization Program (8) and find that ε^* is large, then it is possible to find a large value of ε such that $P(\varepsilon) = 0$, and we can therefore recover a non-degenerate game that has all of the observed equilibria, i.e., we recover a game that has equilibrium properties similar in some sense to those of the true, underlying game. For smaller values of ε^* , we refer to the following statement:

Theorem 2. For every $\varepsilon_0 > \varepsilon^*$, and for all $\varepsilon \geq \varepsilon_0$, we have $f(\varepsilon) \leq P(\varepsilon) \leq g(\varepsilon)$ where f and g are given by

$$f(\varepsilon) = P(\varepsilon_0) \frac{\varepsilon^2}{\varepsilon_0^2} \quad (9)$$

$$g(\varepsilon) = \left((\sqrt{P(\varepsilon_0)} + \frac{\sqrt{lm}}{2}) \frac{\varepsilon}{\varepsilon_0} - \frac{\sqrt{lm}}{2} \right)^2 \quad (10)$$

Proof. See Appendix D. \square


7 Simulations for entry games

In this section, we run simulations for a concrete setting, to illustrate the power of our approach. We consider an *entry game*, in which each of two players have to decide whether to enter a market. Each player p has two actions: $\mathcal{A}_p = \{0, 1\}$; $a_p = 0$ if player p does not enter the market, $a_p = 1$ if he does. The utility of a player is given by $G_p(a_p, a_{-p}) = a_p((1 - a_{-p})\gamma_p + a_{-p}\theta_p)$ for some parameters $\gamma_p \geq 0$ and $\theta_p \leq \gamma_p$, similarly to [16]: if player p does not enter the market, his utility is zero; if he enters the game but the other player does not, p has a monopoly on the market and gets non-negative utility; finally, if both players enter the game, they compete against each other and get less utility than if they had a monopoly. In our simulations, we fix values for the parameters (γ_p, θ_p) and generate the perturbed games as follows:

- In the “no payoff information” and the “partial payoff information” settings, we add Gaussian noise with mean 0 and standard deviation σ (we vary the value of σ) to $G_p(a_p = 1, a_{-p})$ to obtain the perturbed games G^1, \dots, G^l .
- In the “payoff shifter information” case, we sample the payoff shifters S^1, \dots, S^l such that for all $k \in [l]$, for all players p , $S_p^k(a_p = 1, a_{-p})$ follows a normal distribution of mean 0 and standard deviation σ_s . We then add Gaussian noise with mean 0 and standard deviation σ to $G_p(a_p = 1, a_{-p})$ to obtain the perturbed games G^1, \dots, G^l .
- In all observation models, no observed payoff shifter nor unknown noise is added to the payoff of action $a_p = 0$ for player p ; action $a_p = 0$ is always assumed to yield payoff 0 for player p , independently of a_{-p} .

Once the perturbed games are generated, we find the set of equilibria S^k of each of the G^k , and sample a point e^k uniformly at random in S_k . In the payoff information case, we also compute $v^k = e^k G^k$.

We assume the observer knows the form of the utility function, i.e., that $G_p(0, 0) = 0$ and $G_p(a_p = 0, a_{-p} = 1) = 0$, and that he aims to recover the values of γ_p and θ_p . Thus, we add linear constraints $G_p(0, 0) = 0$ and $G_p(a_p = 0, a_{-p} = 1) = 0$ in each of the optimization programs (3), (4) and (5) that we solve for player p in the ‘payoff shifters information’, the “partial payoff information” and the “no payoff information” settings respectively. Furthermore, we assume that the observer knows that perturbations are only added to γ and θ , and therefore we add linear constraints $G_p^k(0, 0) = 0$ and $G_p^k(a_p = 0, a_{-p} = 1) = 0$ for all $k \in [l]$ to the optimization problems for player p in each of the observation models. All optimization problems are solved in Matlab, using CVX (see [12]).



(a) Compatible region in the partial payoff information setting with $\sigma = 0.5$ (b) Compatible region in the payoff shifter information setting $\sigma = 0.5, \sigma_s = 10$

Figure 1: Plots of the compatible region for the “payoff shifter information” and “partial payoff information settings”

Our model for entry-games is similar to the ones presented in [16] and used in simulations in [7], so as to facilitate informal comparisons of the simulation results of both papers; in particular, the parametrization of the utility functions of the players in our simulations is inspired from [7], and noise is generated and added in a similar fashion. However, formal comparison is difficult because of the differences in the observation models and distributional information available in both papers, as well as in the parameters we recover.

7.1 Simulation results

We fix $l = 150$, $\gamma = 5$, $\theta = -10$ in all simulations, and vary the values of σ and σ_s . We solve all optimization problems with respect to player 1, and drop the player indices for simplicity of notation. In all plots, the colored region in the plots is the set of parameters (γ, θ) , i.e., of the payoffs of the recovered game for which there exist G^1, \dots, G^l such that the metric $d_2(G^1, \dots, G^l | G)$ is at most C times larger than the optimal value it can reach subject to the constraints on G, G^1, \dots, G^l presented in the section dealing with the corresponding observation model; we call this region the “compatible” region. The darker the region, the smaller the objective value of the best explanation for the corresponding values of γ and θ . The light-colored center of the region represents the value of (γ, θ) that minimizes $d_2(G^1, \dots, G^l | G)$. We fix $C = 5$ in all simulations.

First, we take $\sigma = 0.5$, $\sigma_s = 10$. In the “no payoff information” setting, we remark that when using the framework from Section 6 and looking for a game G whose payoffs are rescaled between -20 and 20 and $\varepsilon = 600$, we recover $\gamma = 4.8820$, $\theta = -9.9043$; however, accurate recovery is highly dependent on the chosen ratio of ε to scaling constants, and, in general, it may be difficult for an observer to determine a good value of ε . Figure 1 plots the compatible region in both the “partial payoff information” and “payoff shifter information” settings. Note that in both cases, the compatible region is centered on the true value of the parameters, and that the diameter of the black region, which contains the best games given the observations, is small.

Figure 2 shows the evolution of the compatible region when varying σ and σ_s in the “payoff shifter information” setting. The smaller the standard deviation σ of the unknown noise, the tighter the compatible region. On the other hand, reasonably increasing the value of σ_s can be beneficial, at least when it come to centering the compatible region on the true values of the parameters: this comes from the fact that when the game is sufficiently perturbed, new equilibria arise and new behavior is observed, while not adding significant additional uncertainty to the payoffs of the game

Figure 3 shows the evolution of the compatible region when varying σ . The larger the value of σ , the larger




Figure 2: Plots of the compatible region for different values of σ, σ_c in the “payoff shifter information” observation model

the compatible region, and the further away its center is from the underlying, true value of the parameters.




Figure 3: Plots of the compatible region for different values of σ in the “partial payoff information” observation model

Acknowledgements

We thank Federico Echenique and Matt Shum for extremely helpful comments and suggestions.

Bibliography

- [1] A. Aradillas-Lopez. Semiparametric estimation of a simultaneous game with incomplete information. *Journal of Econometrics*, 157(2):409 – 431, 2010.

- [2] A. Aradillas-Lopez. Nonparametric probability bounds for nash equilibrium actions in a simultaneous discrete game. *Quantitative Economics*, 2(2):135–171, 2011.
- [3] A. Aradillas-Lopez. Pairwise-difference estimation of incomplete information games. *Journal of Econometrics*, 168(1):120 – 140, 2012. The Econometrics of Auctions and Games.
- [4] P. Bajari, J. Hahn, H. Hong, and G. Ridder. A Note On Semiparametric Estimation Of Finite Mixtures Of Discrete Choice Models With Application To Game Theoretic Models. *International Economic Review*, 52(3):807–824, 08 2011.
- [5] P. Bajari, H. Hong, J. Krainer, and D. Nekipelov. Estimating Static Models of Strategic Interactions. *Journal of Business & Economic Statistics*, 28(4):469–482, 2010.
- [6] S. Barman, U. Bhaskar, F. Echenique, and A. Wierman. The empirical implications of rank in bimatrix games. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, EC '13, pages 55–72, New York, NY, USA, 2013. ACM.
- [7] A. Beresteanu, I. Molchanov, and F. Molinari. Sharp identification regions in models with convex moment predictions. *Econometrica*, 79(6):1785–1821, 2011.
- [8] U. Bhaskar, K. Ligett, L. J. Schulman, and C. Swamy. Achieving target equilibria in network routing games without knowing the latency functions. *CoRR*, abs/1408.1429, 2014.
- [9] P. A. Bjorn and Q. H. Vuong. Simultaneous equations models for dummy endogenous variables: A game theoretic formulation with an application to labor force participation. Working Papers 537, California Institute of Technology, Division of the Humanities and Social Sciences, 1984.
- [10] T. Bresnahan and P. C. Reiss. Empirical models of discrete games. *Journal of Econometrics*, 48(1-2):57–81, 1991.
- [11] A. V. Fiacco and J. Kyparisis. Convexity and concavity properties of the optimal value function in parametric nonlinear programming. *Journal of Optimization Theory and Applications*, 48(1):95–126, 1986.
- [12] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, Mar. 2014.
- [13] D. Nekipelov, V. Syrgkanis, and É. Tardos. Econometrics for learning agents. *CoRR*, abs/1505.00720, 2015.
- [14] R. Rogers, A. Roth, J. Ullman, and Z. S. Wu. Inducing approximately optimal flow using truthful mediators. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, EC '15, pages 471–488, New York, NY, USA, 2015. ACM.
- [15] K. Seim. An empirical model of firm entry with endogenous producttype choices. *RAND Journal of Economics*, 37(3):619–640, 2006.
- [16] E. Tamer. Incomplete simultaneous discrete response model with multiple equilibria. *The Review of Economic Studies*, 70(1):147–165, 2003.

APPENDIX

A Proof of performance of the algorithm

Consider the following optimization program:

$$\begin{aligned}
P_\epsilon &= \min_{\tilde{G}^k, \hat{G}^k, \tilde{G}, \hat{G}} d(\tilde{G}_1^1, \dots, \tilde{G}_1^l | \tilde{G}_1) + d(\hat{G}_1^1, \dots, \hat{G}_1^l | \hat{G}_1) \\
\text{s.t.} \quad & \sum_{j=1}^d \tilde{G}_1^k(i, j) e_{ij}^k \geq \sum_{j=1}^d \tilde{G}_1^k(i', j) e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\
& \sum_{j=1}^d \hat{G}_1^k(i, j) e_{ij}^k \geq \sum_{j=1}^d \hat{G}_1^k(i', j) e_{ij}^k \quad \forall i, i' \in \mathcal{A}_1, \forall j \in \mathcal{A}_2, \forall k \in [l] \\
& \|\tilde{G}_1 - \hat{G}_1\|_\infty \geq \epsilon
\end{aligned}$$

If $P_\epsilon = \delta$, then clearly the set of observations is (ϵ, δ) -non-identifiable by definition. Now suppose the set of observations is (ϵ, γ) -non-identifiable, then there exist $(\tilde{G}_1, \hat{G}_1, \tilde{G}_1^k, \hat{G}_1^k)$ such that \tilde{G}_1 is γ -close to \hat{G}_1 , $\tilde{G}_1^1, \dots, \tilde{G}_1^l, \hat{G}_1$ is γ -close to $\hat{G}_1^1, \dots, \hat{G}_1^l$, $\|\tilde{G}_1 - \hat{G}_1\|_\infty \geq \epsilon$, and $\tilde{G}_1^k, \hat{G}_1^k$ are compatible with $e^k \forall k \in [l]$. Such a $(\tilde{G}_1, \hat{G}_1, \tilde{G}_1^k, \hat{G}_1^k)$ is feasible for P_ϵ and satisfies $d(\tilde{G}_1^1, \dots, \tilde{G}_1^l | \tilde{G}_1) + d(\hat{G}_1^1, \dots, \hat{G}_1^l | \hat{G}_1) \leq 2\gamma$, hence $P_\epsilon = \delta \leq 2\gamma$. To conclude, we just make the easy remark that $P_\epsilon = \min_{(i,j) \in m_1 m_2} P_\epsilon(i, j)$.

B Proof of recovery lemma under infinite norm and payoff information

Proof. For simplicity of notation, we drop the indices p . We first remark that (G, G^1, \dots, G^l) is feasible for Program (4); as $(\hat{G}, \hat{G}^1, \dots, \hat{G}^l)$ is optimal, it is necessarily the case that

$$\max_k \|\hat{G} - \hat{G}^k\|_\infty \leq \max_k \|G - G^k\|_\infty \leq \delta.$$

Let us write $\Delta G = G - \hat{G}$. We know that for all k , $e^k \prime G^k = e^k \prime \hat{G}^k = v^k$, and thus $e^k \prime (G^k - \hat{G}^k) = 0$. We can write

$$\begin{aligned}
E\Delta G &= (e_1'(G - \hat{G}) \dots e_l'(G - \hat{G}))' \\
&= (e_1'(G - G^1 + G^1 - \hat{G}^1 + \hat{G}^1 - \hat{G}) \dots e_l'(G - G^l + G^l - \hat{G}^l + \hat{G}^l - \hat{G}))' \\
&= (e_1'(G - G^1 + \hat{G}^1 - \hat{G}) \dots e_l'(G - G^l + \hat{G}^l - \hat{G}))'.
\end{aligned}$$

Let $x_k = G - G^k + \hat{G}^k - \hat{G}$. We then have $\|E\Delta G\|_\infty \leq \max_k \|x_k\|_\infty$ as e^k has only elements between 0 and 1. Therefore, by the triangle inequality,

$$\|E\Delta G\|_\infty \leq 2\delta.$$

It immediately follows that $\|\Delta G\|_\infty \leq 2\|E^{-1}\|_\infty \cdot \delta$. □

C Obtaining the dual program

We have

$$\begin{aligned}
& L(G^k, G, \lambda_{ii'}^k, \lambda_0, \lambda_1, \mu) \\
&= d_2(G^1, \dots, G^k | G) + \sum_{k,i,i'} \lambda_{ii'}^k \tilde{e}_{ii'}^k G^k + \mu \left(- \sum_{k,i,i'} \tilde{e}_{ii'}^k G^k - \varepsilon \right) + \lambda_1' (G - \mathbb{1}) - \lambda_0' G \\
&= d_2(G^1, \dots, G^k | G) + \sum_{k,i,i'} (\lambda_{ii'}^k - \mu) \tilde{e}_{ii'}^k G^k + (\lambda_1 - \lambda_0)' G - \mu \varepsilon - \mathbb{1}' \lambda_1 \\
&= d_2(G^1, \dots, G^k | G) + \sum_{k,i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k G^k + (\lambda_1 - \lambda_0)' G - \mu \varepsilon - \mathbb{1}' \lambda_1
\end{aligned}$$

with

$$\mu_{ii'}^k + \mu = \lambda_{ii'}^k \quad (11)$$

Our goal is to find $h(\lambda_{ii'}^k, \lambda_0, \lambda_1, \mu, \mu_{ii'}^k) = \inf_{G, G^k} L(G^k, G, \lambda_0, \lambda_1, \mu, \mu_{ii'}^k)$ in order to write the dual. Since L is a convex function of G^1, \dots, G^l, G , the first order condition needs to hold at a minimum in G^1, \dots, G^l, G , unless this minimum is $-\infty$. Remark that for all k ,

$$\frac{\partial L}{\partial G^k}(G^k, G, \lambda_0, \lambda_1, \mu, \mu_{ii'}^k) = 2(G^k - G) + \sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k$$

and

$$\frac{\partial L}{\partial G}(G^k, G, \lambda_0, \lambda_1, \mu, \mu_{ii'}^k) = 2 \sum_{k=1}^l (G - G^k) + (\lambda_1 - \lambda_0)$$

Therefore, the first order condition is given by

$$\begin{aligned}
2(G^k - G) + \sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k &= 0 \quad \forall k \\
2 \sum_{j=1}^l (G - G^j) + (\lambda_1 - \lambda_0) &= 0
\end{aligned}$$

that can be rewritten

$$G^k = G - \frac{1}{2} \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) \quad \forall k \quad (12)$$

$$G = \frac{1}{l} \sum_{j=1}^l G^j - \frac{1}{2l} (\lambda_1 - \lambda_0) \quad (13)$$

and implies the following system of equalities that must hold whenever the first order condition is satisfied:

$$G^k = \frac{1}{l-1} \sum_{j \neq k}^l G^j - \frac{l}{2(l-1)} \left(\frac{\lambda_1 - \lambda_0}{l} + \sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) \quad \forall k \quad (14)$$

$$G = \frac{1}{l} \sum_{j=1}^l G^j - \frac{1}{2l} (\lambda_1 - \lambda_0) \quad (15)$$

The system has a solution if and only if the system of equations in (14) has a solution. Let us write $x(i, j) = (G^1(i, j), \dots, G^l(i, j))'$, $b^k = -\frac{l}{2(l-1)}(\frac{\lambda_1 - \lambda_0}{l} + \sum_{i, i'} \mu_{ii'}^k \tilde{e}_{ii'}^k)$ for all k , $b(i, j) = (b^1(i, j), \dots, b^l(i, j))$, and $A \in \mathbb{R}^{l \times l}$ the matrix that has 1's on the diagonal and $-\frac{1}{l-1}$ for every other coefficient. Furthermore, let $\mathcal{R}(A)$ denote the range of A , and $\mathcal{N}(A)$ its nullspace. Then there exists a solution to (14) iff there exists a solution to $Ax(i, j) = b(i, j)$ for all (i, j) , i.e., if and only if $b(i, j) \in \mathcal{R}(A)$ for all (i, j) . The following statements characterize $\mathcal{R}(A)$ and $\mathcal{N}(A)$.

Claim 4. $\text{rank}(A) = l - 1$, $\dim \mathcal{N}(A) = 1$

Proof. Let us write $A = (a_1, a_2, \dots, a_l)$ where $a_k \in \mathbb{R}^l$ has 1 as a k^{th} coordinate and has $-\frac{1}{l-1}$ for all other coordinates. Therefore, for all i ,

$$\sum_{k=1}^l a_k(i) = 1 - \sum_{k \neq i} \frac{1}{l-1} = 0,$$

so $\sum_{k=1}^l a_k = 0$ and, necessarily, $\text{rank}(A) \leq l - 1$. Now $\text{rank}(A) \geq l - 1$ because $(-1, 0, \dots, 0, 1)'$, $(-1, 0, \dots, 0, 1, 0)'$, $(-1, 0, \dots, 0, 1, 0, 0)'$, \dots , $(-1, 1, 0, \dots, 0)'$ are $l - 1$ linearly independent vectors that are in the range of A , as they are eigenvectors for eigenvalue $\frac{k}{k-1}$. $\dim \mathcal{N}(A) = 1$ follows from the rank-nullity theorem. \square

Claim 5. $\mathcal{R}(A) = \{x \in \mathbb{R}^l / \sum_{k=1}^l x_k = 0\}$

Proof. Let $x \in \mathcal{R}(A)$, $x = Ay$. Write $A = (a_1, \dots, a_l)'$, then $x = (a_1'y, a_2'y, \dots, a_l'y)'$, so $\sum_{k=1}^l x_k = (\sum_{k=1}^l a_k)'\ y = 0'y = 0$. Therefore, $\mathcal{R}(A) \subseteq \{x \in \mathbb{R}^l / \sum_{k=1}^l x_k = 0\}$. The rest follows from $\{x \in \mathbb{R}^l / \sum_{k=1}^l x_k = 0\}$ being a linear subspace of $\mathcal{R}(A)$ that has dimension $l - 1$. \square

Corollary 1. *There exists a solution to the first order conditions if and only if*

$$\sum_{k=1}^l b^k = -\frac{l}{2(l-1)} \sum_{k=1}^l \left(\frac{\lambda_1 - \lambda_0}{l} + \sum_{i, i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) = 0 \quad (16)$$

Proof. Follows immediately from claim 5. \square

Claim 6. $\mathcal{N}(A) = \text{span}(1, \dots, 1)'$

Proof. $A(1, \dots, 1)' = 0$ so $\text{span}(1, \dots, 1)' \subseteq \mathcal{N}(A)$ and $\dim \text{span}(1, \dots, 1)' = \dim \mathcal{N}(A) = 1$. \square

Corollary 2. *If equation (16) holds, the set of solutions $S(i, j)$ of $Ax(i, j) = b(i, j)$ is given by*

$$S(i, j) = \{(\alpha_{ij} + \tilde{G}^1(i, j), \dots, \alpha_{ij} + \tilde{G}^l(i, j))' / \alpha_{ij} \in \mathbb{R}\}$$

for any $(\tilde{G}^1, \dots, \tilde{G}^l)$ that satisfies the first order conditions. In particular, the set of solutions S to the first order conditions is given by

$$S = \{(M + \tilde{G}^1, \dots, M + \tilde{G}^l) / M \in \mathbb{R}^{l \times l}\}$$

for any $(\tilde{G}^1, \dots, \tilde{G}^l)$ that satisfies the first order conditions.

Claim 7. $\forall k$, let $\tilde{G}^k = \frac{l-1}{l} b^k$. Then $(\tilde{G}^1, \dots, \tilde{G}^l)$ satisfy the first order conditions.

Proof. Take any k , $\frac{1}{l-1} \sum_{j \neq k} G^j + b^k = \frac{1}{l-1} \cdot \frac{l-1}{l} \sum_{j \neq k} b^j + b^k = -\frac{1}{l} b^k + b^k = \frac{l-1}{l} b^k = G^l$ as $\sum_{j \neq k} b^j = -b^k$ from equation 16. \square

Putting it all together, we obtain the following lemma:

Lemma 6. *The first order conditions are satisfied if and only if*

$$-\frac{l}{2(l-1)} \sum_{k=1}^l \left(\frac{\lambda_1 - \lambda_0}{l} + \sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) = 0$$

in which case the set S of (G^1, \dots, G^l) satisfying the first order conditions is given by

$$S = \left\{ \left(M - \frac{1}{2} \left(\frac{\lambda_1 - \lambda_0}{l} + \sum_{i,i'} \mu_{ii'}^1 \tilde{e}_{ii'}^1 \right), \dots, M + \frac{1}{2} \left(\frac{\lambda_1 - \lambda_0}{l} + \sum_{i,i'} \mu_{ii'}^l \tilde{e}_{ii'}^l \right) \right) / M \in \mathbb{R}^{l \times l} \right\}$$

We now have, when constraints (11) and (16) are satisfied, and recalling that equation (12) must hold, that

$$\begin{aligned} h(\mu_{ii'}^k, \lambda_{ii'}^k, \lambda_0, \lambda_1) &= \frac{1}{4} \sum_{k=1}^l \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right)' \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) + \sum_{k,i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \left(G - \frac{1}{2} \sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) + (\lambda_1 - \lambda_0)' G - \mathbb{1}' \lambda_1 - \mu \varepsilon \\ &= \frac{1}{4} \sum_{k=1}^l \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right)' \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) + \sum_{k,i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \left(-\frac{1}{2} \sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) - \mathbb{1}' \lambda_1 - \mu \varepsilon \\ &= -\frac{1}{4} \sum_{k=1}^l \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right)' \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) - \mathbb{1}' \lambda_1 - \mu \varepsilon \end{aligned}$$

and otherwise, $h(\mu_{ii'}^k, \lambda_{ii'}^k, \lambda_0, \lambda_1) = -\infty$. Recall $\lambda_{ii'}^k, \lambda_0, \lambda_1 \geq 0 \forall k, i, i'$, and get the following dual:

$$\begin{aligned} (D) &= \max_{\mu_{ii'}^k, \lambda_0, \lambda_1} && -\frac{1}{4} \sum_{k=1}^l \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right)' \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) - \mathbb{1}' \lambda_1 - \mu \varepsilon \\ \text{s.t.} &&& \lambda_1 - \lambda_0 + \sum_{k,i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k = 0 \\ &&& \mu + \mu_{ii'}^k = \lambda_{ii'}^k \\ &&& \lambda_{ii'}^k, \lambda_0, \lambda_1 \geq 0 \end{aligned}$$

This can further be rewritten as:

$$\begin{aligned} (D) &= \max_{\mu_{ii'}^k, \lambda_0, \lambda_1} && -\frac{1}{4} \sum_{k=1}^l \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right)' \left(\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \right) - \mathbb{1}' \lambda_1 - \mu \varepsilon \\ \text{s.t.} &&& \lambda_1 - \lambda_0 + \sum_{k,i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k = 0 \\ &&& \mu + \mu_{ii'}^k \geq 0 \\ &&& \lambda_0, \lambda_1 \geq 0 \end{aligned}$$

D Proof of the degeneracy-accuracy trade-off

Claim 8. *$P(\varepsilon)$ is a non-decreasing function of ε . In particular, if $\varepsilon_2 > \varepsilon_1 \geq 0$, then $\frac{\varepsilon_2}{\varepsilon_1} P(\varepsilon_2) \geq P(\varepsilon_1)$.*

Proof. Since $\frac{\varepsilon_1}{\varepsilon_2} \leq 1$, we have $0 \leq \frac{\varepsilon_1}{\varepsilon_2} G \leq 1$. Therefore, one can take an optimal solution of $P(\varepsilon_2)$ and multiply all variables by $\frac{\varepsilon_1}{\varepsilon_2}$, to get a solution that is feasible for $P(\varepsilon_1)$; this solution clearly has objective $(\frac{\varepsilon_1}{\varepsilon_2})^2 P(\varepsilon_2)$. \square

This immediately gives the first part of the theorem.

Lemma 7. Let $\varepsilon_2 \geq \varepsilon_1 > 0$, and suppose $P(\varepsilon_2) > 0$. Then:

$$P(\varepsilon_1) \geq (1 - 2\frac{\varepsilon_2 - \varepsilon_1}{\varepsilon_2})P(\varepsilon_2) - \sqrt{lm}\frac{\varepsilon_2 - \varepsilon_1}{\varepsilon_2}\sqrt{P(\varepsilon_2)}$$

Proof. Recall that

$$\begin{aligned} D(\varepsilon) = & \max_{\mu_{ii'}, \lambda_0, \lambda_1} -\frac{1}{4} \sum_{k=1}^l (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k)' (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k) - \mathbb{1}' \lambda_1 - \mu \varepsilon \\ \text{s.t.} & \lambda_1 - \lambda_0 + \sum_{k,i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k = 0 \\ & \mu + \mu_{ii'}^k \geq 0 \\ & \lambda_0, \lambda_1 \geq 0 \end{aligned} \tag{17}$$

Take any optimal solution $(\mu_{ii'}^k, \mu, \lambda_1)$ of $D(\varepsilon_2)$, it is clearly feasible for $D(\varepsilon_1)$ as the constraints in the dual do not depend on the value of ε . Therefore,

$$-\frac{1}{4} \sum_{k=1}^l (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k)' (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k) - \mathbb{1}' \lambda_1 - \mu \varepsilon_1 \leq D(\varepsilon_1)$$

Note that since strong duality holds, by the KKT conditions,

$$\frac{1}{4} \sum_{k=1}^l (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k)' (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k) = P(\varepsilon_2) = D(\varepsilon_2)$$

and therefore

$$-D(\varepsilon_2) - \mathbb{1}' \lambda_1 - \mu \varepsilon_1 \leq D(\varepsilon_1)$$

Since

$$D(\varepsilon_2) = -\frac{1}{4} \sum_{k=1}^l (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k)' (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k) - \mathbb{1}' \lambda_1 - \mu \varepsilon_2 = -D(\varepsilon_2) - \mathbb{1}' \lambda_1 - \mu \varepsilon_2$$

we have

$$2D(\varepsilon_2) + \mu \varepsilon_2 = -\mathbb{1}' \lambda_1$$

and therefore,

$$D(\varepsilon_2) + \mu(\varepsilon_2 - \varepsilon_1) = -D(\varepsilon_2) - \mathbb{1}' \lambda_1 - \mu \varepsilon_1 \leq D(\varepsilon_1)$$

Now, let us try to lower bound μ . We first remark that necessarily, $\mu \leq 0$. If not,

$$D(\varepsilon_2) = -\frac{1}{4} \sum_{k=1}^l (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k)' (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k) - \mathbb{1}' \lambda_1 - \mu \varepsilon < 0$$

and strong duality cannot hold as $P(\varepsilon_2) \geq 0$. Since

$$\mu = \frac{1}{\varepsilon_2} (-2D(\varepsilon_2) - \mathbb{1}' \lambda_1)$$

it is enough to upper-bound $\mathbb{1}' \lambda_1$. Note that since λ_1 is always chosen to be as small as possible as a function of the $\mu_{ii'}^k$ in order to minimize the objective, we have the following coordinate by coordinate inequality:

$$\lambda_1 = \max(0, \sum_{k,i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k) \leq \sum_{k,i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k \leq \sum_k \sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k$$

by the triangle inequality. For simplicity, let us denote $X_k = |\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k|$.

$$\sum_{k=1}^l |\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k|' |\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k| = \sum_{k=1}^l (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k)' (\sum_{i,i'} \mu_{ii'}^k \tilde{e}_{ii'}^k) = \sum_{k=1}^l X_k' X_k = 4D(\varepsilon_2)$$

An upper bound on $\mathbb{1}'\lambda_1 \leq \sum_k \mathbb{1}'X_k$ is therefore given by

$$\begin{aligned} \max_{X_k} \quad & \sum_k \mathbb{1}'X_k \\ \text{s.t.} \quad & \sum_{k=1}^l X_k' X_k \leq 4D(\varepsilon_2) \end{aligned}$$

We can find an exact solution to this convex optimization problem by looking at its dual (Slater and therefore strong duality hold); the Lagrangian is given by $L(X_k, \lambda) = \sum_k \mathbb{1}'X_k - \lambda \sum_{k=1}^l X_k' X_k + 4\lambda D(\varepsilon_2)$ with $\lambda \geq 0$ and the first order condition is $X_k = \frac{1}{2\lambda} \mathbb{1}$. Therefore,

$$h(\lambda) = \inf_{X_k} L(X_k, \lambda) = \frac{1}{4\lambda} \sum_{k=1}^l \mathbb{1}'\mathbb{1} + 4\lambda D(\varepsilon_2)$$

and the dual is given by

$$\begin{aligned} \min_{\lambda} \quad & \sum_k \frac{1}{4\lambda} \sum_{k=1}^l \mathbb{1}'\mathbb{1} + 4\lambda D(\varepsilon_2) \\ \text{s.t.} \quad & \lambda \geq 0 \end{aligned}$$

The solution to the dual is $\lambda^* = \sqrt{\frac{\sum_{k=1}^l \mathbb{1}'\mathbb{1}}{16D(\varepsilon_2)}} \geq 0$ by the first order condition, as $P(\varepsilon_2) = D(\varepsilon_2) > 0$ and we get

$$h(\lambda^*) = \sqrt{\sum_{k=1}^l \mathbb{1}'\mathbb{1} \sqrt{D(\varepsilon_2)}} \leq \sqrt{lm^2} \sqrt{D(\varepsilon_2)} = \sqrt{lm} \sqrt{D(\varepsilon_2)}.$$

So, $0 \leq \mathbb{1}'\lambda_1 \leq \sqrt{lm} \sqrt{D(\varepsilon_2)}$ leading to

$$\mu = \frac{1}{\varepsilon_2} (-2D(\varepsilon_2) - \mathbb{1}'\lambda_1) \geq \frac{1}{\varepsilon_2} (-2D(\varepsilon_2) - \sqrt{lm} \sqrt{D(\varepsilon_2)})$$

and therefore, as $\varepsilon_2 - \varepsilon_1 \geq 0$,

$$(1 - 2\frac{\varepsilon_2 - \varepsilon_1}{\varepsilon_2})D(\varepsilon_2) - \sqrt{lm} \frac{\varepsilon_2 - \varepsilon_1}{\varepsilon_2} \sqrt{D(\varepsilon_2)} \leq D(\varepsilon_2) + \mu(\varepsilon_2 - \varepsilon_1) \leq D(\varepsilon_1)$$

□

Claim 9. $\forall \varepsilon \geq \varepsilon^*$,

$$\lim_{h \rightarrow 0^+} P(\varepsilon + h) = P(\varepsilon)$$

i.e. $P(\cdot)$ is right-continuous on $[\varepsilon^*, +\infty[$.

Proof. Take $h > 0$. $P(\varepsilon + h) \geq (\frac{\varepsilon+h}{\varepsilon})^2 P(\varepsilon)$ from claim 8, and $(\frac{\varepsilon+h}{\varepsilon})^2 P(\varepsilon) \rightarrow P(\varepsilon)$ when h tends to 0. From Lemma 7, since $P(\varepsilon + h) > 0$ as $\varepsilon + h > \varepsilon^*$, we have

$$P(\varepsilon) \geq (1 - 2\frac{h}{\varepsilon+h})P(\varepsilon + h) - \sqrt{lm} \frac{h}{\varepsilon+h} \sqrt{P(\varepsilon + h)}$$

and so

$$\frac{1}{1 - 2\frac{h}{\varepsilon+h}}(P(\varepsilon) + \sqrt{lm}\frac{h}{\varepsilon+h}\sqrt{P(\varepsilon+h)}) \geq P(\varepsilon+h)$$

Note that $P(\varepsilon+h) \leq P(\varepsilon+\alpha)$ for any constant α and small enough h ; fix an α , we have for h small that

$$\frac{1}{1 - 2\frac{h}{\varepsilon+h}}(P(\varepsilon) + \sqrt{lm}\frac{h}{\varepsilon+h}\sqrt{P(\varepsilon+\alpha)}) \geq P(\varepsilon+h)$$

As $P(\varepsilon+\alpha)$ is finite (this is clear from the linear program and previous calculations), we have $\frac{1}{1 - 2\frac{h}{\varepsilon+h}}(P(\varepsilon) + \sqrt{lm}\frac{h}{\varepsilon+h}\sqrt{P(\varepsilon+\alpha)}) \rightarrow P(\varepsilon)$ when h tends to 0. \square

Claim 10. $\forall \varepsilon > \varepsilon^*$,

$$\lim_{h \rightarrow 0^-} P(\varepsilon+h) = P(\varepsilon)$$

i.e. $P(\cdot)$ is left-continuous on $]\varepsilon^*, +\infty[$.

Proof. The proof is similar to the right-continuity one. The main difference comes from the fact that we now require $P(\varepsilon) > 0$ to satisfy the condition of lemma 7 (as $\varepsilon > \varepsilon+h$ or $h < 0$), so we cannot include the $\varepsilon = \varepsilon^*$ case. \square

Claim 11. $P(\cdot)$ is continuous on $[0, +\infty[$.

Proof. It is clear that $P(\cdot)$ is continuous on $[0, \varepsilon^*[$ and $]\varepsilon^*, +\infty[$ by claims 9 and 10. We just need to check that $P(\cdot)$ is continuous at ε^* ; it is right-continuous at ε^* by claim 9, and the left-continuity follows from the fact that $P(\varepsilon) = 0 \forall \varepsilon \leq \varepsilon^*$. \square

Claim 12. $P(\cdot)$ is convex on $[0, +\infty[$.

Proof. Let $S(\varepsilon)$ be the feasible region of optimization program (5). It is easy to see that the objective function of (5) is convex, and that the mapping $\varepsilon \rightarrow P(\varepsilon)$ is convex according to the definition of [11]. Therefore, as seen in [11], the optimal value function $P(\cdot)$ of Program (5) is convex. \square

Claim 13. Let $\varepsilon \geq \varepsilon^*$, we have for all $h > 0$

$$\frac{h+2\varepsilon}{\varepsilon^2}P(\varepsilon) \leq \frac{P(\varepsilon+h) - P(\varepsilon)}{h} \leq \frac{2}{\varepsilon+h}P(\varepsilon+h) + \sqrt{lm}\frac{1}{\varepsilon+h}\sqrt{P(\varepsilon+h)}$$

Proof. From claim 8, for $h > 0$,

$$\frac{P(\varepsilon+h) - P(\varepsilon)}{h} \geq \frac{P(\varepsilon)}{h} \left(\left(\frac{\varepsilon+h}{\varepsilon} \right)^2 - 1 \right) = \frac{h+2\varepsilon}{\varepsilon^2}P(\varepsilon) \quad (18)$$

From Lemma 7,

$$\begin{aligned} \frac{P(\varepsilon+h) - P(\varepsilon)}{h} &\leq \frac{1}{h} \left(2\frac{h}{\varepsilon+h}P(\varepsilon+h) + \sqrt{lm}\frac{h}{\varepsilon+h}\sqrt{P(\varepsilon+h)} \right) \\ &= \frac{2}{\varepsilon+h}P(\varepsilon+h) + \sqrt{lm}\frac{1}{\varepsilon+h}\sqrt{P(\varepsilon+h)}. \end{aligned}$$

\square

Since $P(\cdot)$ is a convex and continuous on $[0, +\infty[$, its right derivative $\frac{dP(\varepsilon)}{d\varepsilon}$ exists at every point $\varepsilon \geq 0$. By Claim 13, we have

$$l(\varepsilon) \leq \frac{dP(\varepsilon)}{d\varepsilon} \leq L(\varepsilon)$$

where

$$l(\varepsilon) = \lim_{h=0^+} \frac{h+2\varepsilon}{\varepsilon^2} P(\varepsilon)$$

$$L(\varepsilon) = \lim_{h=0^+} \frac{2}{\varepsilon+h} P(\varepsilon+h) + \sqrt{lm} \frac{1}{\varepsilon+h} \sqrt{P(\varepsilon+h)}$$

if they exist. Clearly, $l(\varepsilon) = \frac{2}{\varepsilon} P(\varepsilon)$ exists; $L(\varepsilon) = \frac{2}{\varepsilon} P(\varepsilon) + \frac{\sqrt{lm}}{\varepsilon} \sqrt{P(\varepsilon)}$ exists because by claim 9, $P(\cdot)$ is right continuous and $\lim_{h=0^+} P(\varepsilon+h) = P(\varepsilon)$. This implies that given an initial condition $P(\varepsilon_0)$ for some $\varepsilon_0 > \varepsilon^*$, $P(\cdot)$ lies between the function f with $f(\varepsilon_0) = P(\varepsilon_0)$ and $\frac{df(\varepsilon)}{d\varepsilon} = l(\varepsilon)$ and the function g with $g(\varepsilon_0) = P(\varepsilon_0)$ and $\frac{dg(\varepsilon)}{d\varepsilon} = L(\varepsilon)$ for all $\varepsilon \geq \varepsilon_0$. We can find f and g by solving differential equations

$$\frac{df(\varepsilon)}{d\varepsilon} = \frac{2}{\varepsilon} f(\varepsilon) \tag{19}$$

$$\frac{dg(\varepsilon)}{d\varepsilon} = \frac{2}{\varepsilon} g(\varepsilon) + \frac{\sqrt{lm}}{\varepsilon} \sqrt{g(\varepsilon)} \tag{20}$$

for all $\varepsilon \geq \varepsilon_0$. It is easy to see that with initial condition $f(\varepsilon_0) = P(\varepsilon_0)$, ODE (19) has as a unique solution

$$f(\varepsilon) = P(\varepsilon_0) \frac{\varepsilon^2}{\varepsilon_0^2}$$

To solve ODE (20), let us write $g(\varepsilon) = \sqrt{f(\varepsilon)}$, and note that the differential equation can be rewritten:

$$2g(\varepsilon) \frac{dg(\varepsilon)}{d\varepsilon} = \frac{2}{\varepsilon} g(\varepsilon)^2 + \frac{\sqrt{lm}}{\varepsilon} g(\varepsilon)$$

Noting that by the choice of initial condition, $g(\varepsilon_0) = P(\varepsilon_0) > 0$ and that the solution of the differential equation is necessarily increasing as the derivative is always non-negative, we have $g(\varepsilon) > 0$ for all $\varepsilon \geq \varepsilon_0$. Therefore, on $[\varepsilon_0, +\infty[$,

$$\frac{dg(\varepsilon)}{d\varepsilon} = \frac{1}{\varepsilon} g(\varepsilon) + \frac{\sqrt{lm}}{\varepsilon}$$

The initial condition being fixed, this differential equation has a unique solution. Note that solutions to the homogeneous ODE $\frac{dg(\varepsilon)}{d\varepsilon} = \frac{1}{\varepsilon} g(\varepsilon)$ are of the form $g(\varepsilon) = C\varepsilon$, and that $g_0(\varepsilon) = -\frac{\sqrt{lm}}{2}$ is a particular solution of the ODE. Therefore, given the initial condition $g(\varepsilon_0) = \sqrt{P(\varepsilon_0)}$, we have

$$g(\varepsilon) = \left(\sqrt{P(\varepsilon_0)} + \frac{\sqrt{lm}}{2} \right) \frac{\varepsilon}{\varepsilon_0} - \frac{\sqrt{lm}}{2}$$

and thus

$$f(\varepsilon) = \left(\left(\sqrt{P(\varepsilon_0)} + \frac{\sqrt{lm}}{2} \right) \frac{\varepsilon}{\varepsilon_0} - \frac{\sqrt{lm}}{2} \right)^2.$$