Comparison and Continuity Properties of Equilibrium Values in Information Structures for Stochastic Games

by

Ian Jay Hogeboom-Burr

A thesis submitted to the
Department of Mathematics and Statistics
in conformity with the requirements for
the degree of Master of Applied Science

Queen's University
Kingston, Ontario, Canada
June 2022

Abstract

In stochastic games where players measure a cost-relevant exogenous state variable through measurement channels, an information structure is the joint probability measure induced on the state space and player measurement spaces.

For single-player decision problems in finite spaces, a theorem due to Blackwell leads to a complete characterization of when one information structure is "better" than another. For zero-sum games with finite state, measurement, and action spaces, Pęski produced necessary and sufficient conditions for ordering information structures. In this thesis, we obtain an infinite dimensional (standard Borel) generalization of Pęski's result. A corollary is that more information cannot hurt a decision maker taking part in a zero-sum game. We also establish two supporting results which are essential and explicit, though modest contributions to the literature: (i) a partial converse to Blackwell's ordering in the standard Borel setup and (ii) an existence result for equilibria in zero-sum games with incomplete information.

Then we study continuity properties of stochastic game problems with respect to various notions of convergence of information structures. For zero-sum games, team problems, and general games, we will establish continuity properties of the value function under total variation, setwise, and weak convergence of information structures.

Acknowledgments

I would like to thank my mother for pushing me to take Grade 12 Functions and Calculus, even when I was convinced I would never use math outside of high school. She always knows best.

I would also like to thank my partner, Mikaela, for her unwavering support and boundless enthusiasm.

I am incredibly grateful to my supervisor, Serdar Yüksel, for his mentorship over the last four years. He has always provided thoughtful guidance on both mathematics and life, and has helped me find the right path to balance the two.

Finally, I would like to extend my gratitude to Yanglei Song, Tamás Linder, and Gregory Smith for their time and effort on my thesis committee.

Contents

List of Figures

Chapter 1

Introduction

1.1 Motivation

In every decision problem, information arises as a fundamental attribute. In decision and control theory, information always has a positive value, meaning that more information can never hurt. Nonetheless, the dependence on information of the optimal cost in a decision problem can be rather technical.

In game theory, the analysis is even more complicated: the value of information may not be positive and even the strongest notions of informational convergence may not entail continuity or regularity.

In this thesis, we will study information structures in the context of stochastic zero-sum games, team problems, and general games. Information structures in stochastic games capture the information available to decision-makers regarding a cost-relevant exogenous state variable. Each decision-maker's policy, which defines how the decision-maker selects actions, is a measurable function of the decision-maker's personal information.

In the context of stochastic games, a natural question that arises is how to characterize when one information structure is "better" than another over a large class of games. This topic was first addressed for single-player decision problems by David Blackwell in [10], who defined an information structure as being better than another if a decision-maker is guaranteed to not perform worse under the former than under the latter in any valid game. This definition can be extended to classes of game problems in which there exists a unique value that defines the outcome of a game, such as zero-sum games. Finding necessary and sufficient conditions to compare information structures according to this definition allows for a partial ordering on the space of information structures.

Another natural question to pose is: Under what notions of convergence is the equilibrium value of a game continuous under perturbations of the information structure? Such a question has consequences for robustness of models to incorrect information.

We will address the first question in the context of zero-sum games in Chapter 3, while we will present results on the second question for zero-sum games, team problems, and general non-zero-sum games in Chapter 4.

1.2 Stochastic Games and Information Structures

In a general static stochastic game, there are $N \in \mathbb{N}$ decision-makers (DMs), sometimes also known as *agents* or *players*. The set of all DMs in a game will be denoted by $\mathcal{N} := \{1, \dots, N\}$. The outcome of the game depends on actions that the DMs select, as well as an exogenous random variable known as the *state of nature*. We use $\mathbb{X}$ to denote the *state space* from which the state of nature x is drawn, and we will assume $\mathbb{X}$ is a standard Borel space (that is, a Borel subset of a complete, separable, metric space). x is drawn according to a *prior* distribution ζ, which is a probability distribution defined on $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$, where $\mathcal{B}(\mathbb{X})$ denotes the Borel sigma field on $\mathbb{X}$.

In our setup, the distribution ζ is common knowledge to all DMs, but the realized value of x is not. Instead, each DM makes a private *measurement* of x, which is available only to the DM that made it. We will denote this measurement by y^i, where the superscript indicates the measurement belongs to DM i from $\mathcal{N}$. y^i takes value in DM i's standard Borel *measurement space* $\mathbb{Y}^i$ and is defined by:

$$y^i = g^i(x, \omega^i),$$

for some measurable function g^i and noise variable ω^i which is independent of x (and which, without any loss, can be taken to be $[0, 1]$-valued).

Using stochastic realization results (see Lemma 1.2 in [27], or Lemma 3.1 of [12]), it follows that the functional representation in $y^i = g^i(x, \omega^i)$ is equivalent to a stochastic kernel description of an information structure. Thus, in the above, we can view g^i as inducing a measurement channel Q^i, which is a stochastic kernel or a regular conditional probability measure from $\mathbb{X}$ to $\mathbb{Y}$ in the sense that $Q^i(\,\cdot\,|x)$ is a probability measure on the (Borel) σ-algebra $\mathcal{B}(\mathbb{Y}^i)$ on $\mathbb{Y}^i$ for every $x \in \mathbb{X}$, and $Q^i(A|\,\cdot\,) : \mathbb{X} \to [0, 1]$ is a Borel measurable function for every $A \in \mathcal{B}(\mathbb{Y}^i)$.

For notation, we let $\mathcal{P}(\mathbb{X})$ denote the set of all probability measures on (the Borel sigma field over) $\mathbb{X}$. For $\zeta \in \mathcal{P}(\mathbb{X})$ and kernel Q, we let ζQ denote the joint distribution induced on $(\mathbb{X} \times \mathbb{Y}, \mathcal{B}(\mathbb{X} \times \mathbb{Y}))$ by channel Q with input distribution ζ:

$$\zeta Q(A) = \int_A Q(dy|x)\zeta(dx), \quad A \in \mathcal{B}(\mathbb{X} \times \mathbb{Y}).$$

We denote the joint probability measure on $\mathbb{X} \times \mathbb{Y}^1 \times \cdots \times \mathbb{Y}^N$ induced by the functions g^i by $\mu(dx, dy^1, \ldots, dy^N)$, and define this measure as the *information structure*. The information structure is common knowledge to all DMs.

Note that it is not necessarily true that the DMs' measurements are conditionally independent, as no assumption was made that the measurement noise variables ω^i were independent of each other, only of x. Thus, we will typically work with the general information structure representation of information, rather than the measurement channel representation. We will use the channel representation under certain circumstances, such as when a conditional independence assumption is made on the measurements.

Each DM also has a standard Borel action space $\mathbb{U}^i$, and a personal cost function $c^i(x, u^1, \ldots, u^N) : \mathbb{X} \times \mathbb{U}^1 \times \cdots \times \mathbb{U}^N \to \mathbb{R}$.

For fixed prior, state space, and measurement spaces ζ, $\mathbb{X}$, $\mathbb{Y}^1, \ldots, \mathbb{Y}^N$, a *game* is a $2N$-tuple consisting of a measurable and bounded cost function and an action space for each DM, $g = (c^1, \ldots, c^N, \mathbb{U}^1, \ldots, \mathbb{U}^N)$.

The DM's goal is to minimize their expected cost functional for a given game g:

$$J^i(g, \mu, \gamma^1, \ldots, \gamma^N) := E^{\mu,\bar{\gamma}}[c^i(x, \gamma^1(y^1), \ldots, \gamma^N(y^N))]$$

where the DMs select their policies from the set of all admissible *policies* $\Gamma^i := \{\gamma : \mathbb{Y}^i \to \mathbb{U}^i\}$, which are measurable functions from a DM's measurement space to their action space. We refer to $u^i = \gamma^i(y^i)$ as the *action* of the DM, and γ^i as their *policy*. We use $\bar{\gamma} = \{\gamma^1, \ldots, \gamma^N\}$ to denote a collection of policies for all N DMs.

For our analysis, we will allow policies to be randomized with independent randomness. Which is to say, the set of all admissible measurable policies Γ^i will be the set of all measurable functions γ^i, where $u^i = \gamma^i(y^i, \nu^i)$ for some independent $[0,1]$-valued noise variable ν^i (which is independent not only of x, but also of any other DM's noise variable ν^j, $j \in \{\mathcal{N} \setminus i\}$). Admissible randomized policies can be viewed as stochastic kernels from $\mathbb{Y}^i$ to $\mathbb{U}^i$.

Note that in this setup, each DM's measurement is conditionally independent of other DMs' measurements given the state of nature.

We now present a brief example to serve as an illustration of this setup.

Example 1.1 (A General Game)**.** *Let $\mathbb{X} = [0,1]$, with prior distribution ζ defined by the continuous uniform distribution. Let $\mathbb{Y}^1 = [0,1]$ and $\mathbb{Y}^2 = \{0,1,2\}$, where the measurements for the respective DMs are defined by:*

$$y^1 = \omega^1$$

$$y^2 = \begin{cases} 0, & x \in [0, 1/3), \\ 1, & x \in [1/3, 2/3], \\ 2, & x \in (2/3, 1] \end{cases}$$

where ω^1 is a uniformly distributed random variable on $[0,1]$. In this setup, DM 1 receives no information regarding x, as y^1 is generated solely by random noise. DM 2's channel is a quantizer consisting of 3 bins.

We define the action spaces of the respective DMs as $\mathbb{U}^1 = \mathbb{U}^2 = [0,1]$, and the cost functions as:

$$c^1(x, u^1, u^2) = |u^1 - u^2|$$

$$c^2(x, u^1, u^2) = (x - u^2)^2$$

Here, DM 1 has no effect on DM 2's cost, so DM 2 is faced with a single-DM decision problem and will attempt to minimize $E^{\mu}[(x-u^2)^2]$. *For each observation* y^2, *DM 2 knows which third of the interval [0,1] that* x *falls into. For instance, if* $y^2 = 0$, *DM 2 knows that* $x \in [0, 1/3)$. *Given this information, DM 2's problem then becomes selecting* u^2 *to minimize* $E[(z-u^2)^2]$ *where* z *is uniformly distributed on* $[0, 1/3)$; *this is minimized when* $u^2 = 1/6$ *is selected. It follows that DM 2's optimal strategy is:*

$$\gamma^{2,*}(y^2) = \begin{cases} 1/6, & y^2 = 0 \\ 1/2, & y^2 = 1 \\ 5/6, & y^2 = 2. \end{cases}$$

This leads to the expected cost for DM 2 being equal to the variance of z, *which is* $1/108$.

For DM 1, the only way to reduce the cost is to play an action as close in absolute value to DM 2's action as possible. Under DM 2's optimal strategy, DM 2 will play $1/6$, $1/2$, *and* $5/6$ *each with probability* $1/3$, *respectively. Thus, DM 1 wants to minimize the term* $1/3\,(|1/6 - u^1| + |1/2 - u^1| + |5/6 - u^1|)$.

This is minimized by DM 1 playing $u^1 = \gamma(y^1) = 1/2$, *for all* $y^1 \in \mathbb{Y}^1$, *resulting in an expected cost of* $2/9$ *for DM 1.*

So far, we have discussed the setup for a general static stochastic game. We will now look at two special classes of games within this setup.

1.2.1 Zero-Sum Games

In the zero-sum game case there are two DMs who share a common cost function. DM 1's goal is to minimize the expected cost, while DM 2's goal is to maximize it. The DMs are commonly referred to as the minimizer and the maximizer, respectively.

Given fixed $\mathbb{X}$, $\mathbb{Y}^1$, $\mathbb{Y}^2$, and ζ such that $x \sim \zeta$, a zero-sum game $g = (c, \mathbb{U}^1, \mathbb{U}^2)$ is a triple of a measurable cost function $c : \mathbb{X} \times \mathbb{U}^1 \times \mathbb{U}^2 \to \mathbb{R}$ and standard Borel action spaces for each DM $\mathbb{U}^1, \mathbb{U}^2$.

The induced cost under policies γ^1, γ^2 for the players is given by

$$J(g, \mu, \gamma^1, \gamma^2) \coloneqq \int c(x, \gamma^1(y^1), \gamma^2(y^2))\mu(dx, dy^1, \dots, dy^N).$$

Since the problem has been reformulated such that DM 1 wants to minimize J and DM 2 wants to maximize J, the single value $J(g, \mu, \gamma^1, \gamma^2)$ fully describes the expected outcome of the game for both players.

We will assess the outcome or *value* of a zero-sum game by the expected value of the cost function of the game at a saddle-point equilibrium. This equilibrium is defined as follows:

Definition 1.1. *Given an information structure μ, we say that $\gamma^{1,*}, \gamma^{2,*}$ is an equilibrium for a zero-sum game g if*

$$\inf_{\gamma^1 \in \Gamma^1} J(g, \mu, \gamma^1, \gamma^{2,*}) = J(g, \mu, \gamma^{1,*}, \gamma^{2,*}) = \sup_{\gamma^2 \in \Gamma^2} J(g, \mu, \gamma^{1,*}, \gamma^2).$$

The value function at equilibrium is denoted by $J^*(g, \mu)$. Existence results for such saddle-point equilibria will be discussed in Chapter 2.

1.2.2 Team Problems

In a team problem, there are N DMs working together to minimize a common cost function. A game for a team-problem is an $N+1$-tuple consisting of a common cost function $c : \mathbb{X} \times \mathbb{U}^1 \times \cdots \times \mathbb{U}^N \to \mathbb{R}$ and standard Borel action spaces $\mathbb{U}^i$ for each DM $i \in \{1, \ldots, n\}$.

The induced cost under a collection of policies $\bar{\gamma} := \{\gamma^1, \cdots, \gamma^N\}$ for a team problem is given by

$$J(g, \mu, \bar{\gamma}) := \int c(x, \gamma^1(y^1), \ldots, \gamma^N(y^N))\mu(dx, dy^1, \ldots, dy^N).$$

For a team problem, an equilibrium solutions is a *team-optimal policy* $\bar{\gamma}^* := (\gamma^{1,*}, \ldots, \gamma^{N,*})$ which minimizes $J(g, \mu, \cdot)$. As in zero-sum games, the value function at equilibrium is denoted by $J^*(g, \mu)$. General existence results for team-optimal policies will be discussed further in Chapter 4.

1.3 Organization of Thesis

In Chapter 2 we study the existence of saddle-point equilibria in zero-sum games. We provide an overview of existing results, and present a theorem which demonstrates existence of equilibria under mild conditions, as well as a further relaxation of this theorem. These results will be used to ensure the existence of an equilibrium value when discussing zero-sum games in Chapters 3 and 4.

In Chapter 3 we examine the problem of comparing information structures in zero-sum games. First, we will extend a partial converse to Blackwell's original ordering of information structures for single-DM decision problems to problems with standard

Borel spaces, and we will demonstrate the relation of this result to one developed by Strassen. We then present necessary and sufficient conditions for the comparison of information structures in zero-sum games with standard Borel spaces, extending results developed by Peski for finite-space problems.

In Chapter 4 we study regularity properties of the value function under perturbations of the information structure. In particular, we demonstrate continuity and semicontinuity properties of the equilibrium value for zero-sum games and team problems, using weak, setwise, and total variation convergence of information structures. We then present a counterexample to show continuity does not hold in general for general non-zero-sum games even under strong conditions.

In this thesis, we will see that zero-sum games share many of the positive attributes of stochastic teams. Zero-sum games have unique values and share the property that 'additional information can not hurt a decision-maker' (though with a much more tedious argument as compared to the team theoretic setup, where a 'choose to ignore the additional information' argument applies).

1.4 Contributions

The contributions of this thesis are as follows:

(i) As a minor technical contribution, in Chapter 2 we present sufficient conditions for the existence of saddle points in Bayesian zero-sum games with incomplete information in standard Borel spaces (Theorem 2.1). This will build on placing an appropriate topology on the space of policies adopted by the decision-makers. Our analysis is nearly equivalent to others in the literature, such as [6], but is presented for completeness[1].

(ii) In Chapter 3, as a supporting theorem, we will present a partial converse to Blackwell's ordering theorem for standard Borel spaces in Chapter 3, using a separating hyperplane argument and properties of locally convex spaces (Theorem 3.3). This presents an explicit, self-sufficient derivation for a converse theorem to be utilized in the comparison theorem, though related comprehensive results have been reported in the literature, as we note in the chapter.

(iii) As the main result of Chapter 3, we will derive a theorem characterizing an ordering of information structures for zero-sum games in standard Borel spaces (Theorem 3.5). This extends results shown in [44] for finite-space problems.

(iv) In Chapter 4, we show that the value function for a zero-sum game is continuous in total variation convergence of the information structure if the game has a bounded measurable cost function (Theorem 4.3). In addition, when the prior is fixed, the value function is either upper or lower semicontinuous in setwise convergence of an information structure, if the sequence of information structures is a minimizer-garbling or maximizer-garbling sequence (Theorem 4.5). The same results hold for weak convergence of the information structures, when the assumptions on the cost function are such that it is continuous and bounded, and the DM action spaces are convex (Theorem 4.4). If the channels are fixed and continuous in total variation, continuity under weak convergence of the prior holds (Theorem 4.6).

[1] We note that our results were derived independently and were presented in [31] with an accompanying literature review prior to our awareness of the nearly equivalent results in [6].

(v) In Chapter 4, we also show that the value function for team problems is continuous in total variation convergence of the information structure for bounded cost functions, upper semicontinuous in setwise convergence for measurable and bounded cost functions, and upper semicontinuous in weak convergence for games with continuous and bounded cost functions and convex action spaces (Theorems 4.7, 4.10, 4.9). If the channels are fixed and continuous in total variation, continuity under weak convergence of the prior holds (Theorem 4.8).

(vi) A counterexample in Chapter 4 reveals that DMs in general non-zero-sum games may not have value functions that are continuous in total variation convergence of information structures, even for bounded cost functions (Example 4.2).

Chapter 2

Existence of Equilibria in Zero-Sum Games

2.1 Introduction

Prior to focusing on the ordering and regularity properties of information structures, we present a supporting result regarding when equilibrium solutions to zero-sum games exist. In the finite case, equilibrium solutions always exist [60] (through e.g. [5, Theorem 4.4]), but this does not hold true in general. Theorems 2.1 and 2.2 below gives sufficient conditions for equilibrium solutions to exist for games with incomplete information.

The existence of a value for games with incomplete information has been studied rather extensively. For readers' convenience, and as a direct proof, we present the results below. Our results essentially replicate those of [6, Theorem 3.1], which presents an existence result under an absolute continuity condition on information structures with respect to a reference measure. We also refer the reader to [40] which presents complementary conditions where the absolute continuity condition is relaxed. Finally, [6, Theorem 3.4] presents a generalization where action sets are information dependent. The essence of these results is the same as those below, although the results below were arrived at independently.

Additionally, Milgrom and Weber present an existence result for more general games in [43, Theorem 1], which presents conditions whose generality is difficult to interpret: a careful look at condition R1 in [43, p. 625] leads to the conclusion that the authors have nearly (but not exactly) the same condition (ii) we note below; that is continuity of the cost function in the actions for every fixed hidden state variable x is sufficient, though the statements given in [43] imposes conditions that are not conclusive on this; we attribute this to the fact that the authors utilize [43, Prop. 1(c)] without establishing its relation to item (ii) below (due to the measurability requirement in the statement of [43, Prop. 1(c)]). Our analysis affords the simplicity and generality in the condition, since we build on the w-s topology, rather than weak topology and directly Lusin's theorem [20] as followed in [43] (we also note that the relation between weak and w-s topologies on probabilities defined on product spaces with a fixed marginal can in fact be established using Lusin's theorem). Hence, in a strict sense, our conditions are more direct and general as stated.

The comprehensive book [41, Proposition III.4.2.] imposes continuity in all the variables (unlike what is presented below). Furthermore, [41, Proposition III.4.2.] builds on a topology construction on policies which is different from what we present here; regarding the construction in [41] we would like to caution that in the absence of absolute continuity conditions on the information structure, this construction may lead to a lack of closedness on the sets of admissible policies (or *strategic measures*) as the counterexample [67, Theorem 2.7] reveals: in this counterexample, which would reduce to the setup studied here with $y^1 = y^2 = y$, a sequence of policies is constructed so that for each element of the sequence, the action variables of the two decision-makers are conditionally independent given their measurements, but the setwise (and hence, weak) limit of the sequence is not conditionally, or otherwise, independent; and thus the limit measure does not belong to the original information structure. For a more detailed discussion, we refer the reader to [52, Section 7.2].

While the results presented below are not novel given the wealth of previous results on existence of equilibria for standard zero-sum games, the general proof framework is quite simple and robust, and the proofs demonstrate the utility of both the independent measurements reduction and the w-s topology in this setting.

2.2 On Existence of Saddle-Points and Equilibria

We now present two conditions for the information structure μ in a zero-sum game. We note that Assumption 2.2 implies Assumption 2.1, but this assumption often allows for a simpler interpretation. That this implication holds is a consequence of the independent measurements reduction formulation to be explained in greater detail later in Theorem 2.1. Theorem 2.1 will be presented under the more general Assumption 2.1. However, the conditional independence assumption in Assumption 2.2 will allow for a useful characterization of the results in Chapter 3.

Assumption 2.1. *The information structure is absolutely continuous with respect to a product measure:*

$$P(dy^1, dy^2, dx) \ll \bar{Q}^1(dy^1)\bar{Q}^2(dy^2)\zeta(dx),$$

for some reference probability measures $\bar{Q}^i$, $i = 1, 2$. That is, there exists an integrable *f which satisfies for every Borel A, B, C*

$$P(y^1 \in B, y^2 \in C, x \in A) = \int_{A,B,C} f(x, y^1, y^2)\zeta(dx)\bar{Q}^1(dy^1)\bar{Q}^2(dy^2).$$

Assumption 2.2. *The following conditional independence (or Markov) condition holds:*

$$P(dy^1, dy^2, dx) = Q^1(dy^1|x)Q^2(dy^2|x)\zeta(dx).$$

where the measurements of agents are absolutely continuous so that for $i = 1, 2$, there exists a non-negative function f^i and a reference probability measure $\bar{Q}^i$ such that for all Borel S:

$$Q^i(y^i \in S|x) = \int_S f^i(y^i, x)\bar{Q}^i(dy^i).$$

We now present the results on the existence of saddle-point equilibria in stochastic zero-sum games. The first result utilizes Assumption 2.1, while the second result relaxes this condition. The first result remains included for comparison to the literature, namely [6, Theorem 3.1], as well as to demonstrate the applicability of the independent measurements reduction, which is utilized frequently in stochastic team problems.

Theorem 2.1 (Existence of Equilibria)**.** *For a fixed information structure μ and a given game g, assume that Assumption 2.1 holds. Further, let the following hold:*

(i) The action spaces of DMs, $\mathbb{U}^1, \mathbb{U}^2$, are compact.

(ii) The cost function c is bounded and continuous in DMs' actions, for every state of nature x.

Then an equilibrium exists under possibly randomized policies, and so there exists a value of the zero-sum game.

Proof. ***Step (1)***: By Assumption 2.1, we can reformulate the problem in a new probability space in which the measurements are independent from the unknown variable x. This reformulation, called an *independent-measurements reduction*, is essentially due to Witsenhausen [62], with a detailed discussion in [64, Section 2.2.]. Figure 2.1 provides an illustration of the concept.

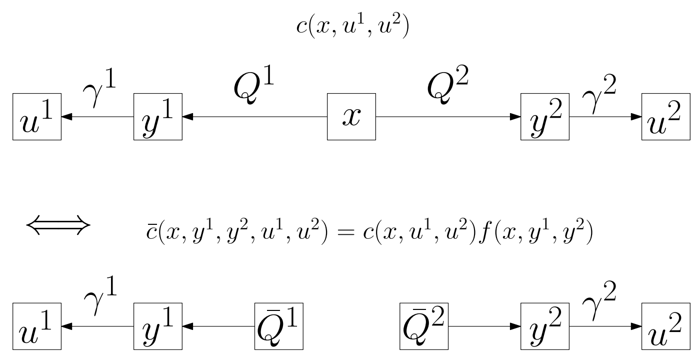

Figure 2.1: Reformulation of two information structures (with respect to an independent measurements reduction)

The main benefit of this approach is to define a compact/convex policy space for the DMs (e.g. see [67, Section 2.2]). To complete this reformulation, we note the following holds for some function f and reference probability measures $\bar{Q}^i$:

$$P(dx, dy^1, dy^2, du^1, du^2) = \zeta(dx) f(x, y^1, y^2) \bar{Q}^1(dy^1) 1_{\{\gamma^1(y^1) \in du^1\}} \bar{Q}^2(dy^2) 1_{\{\gamma^2(y^2) \in du^2)\}},$$

where $1_{\{.\}}$ is the indicator function. Thus, the value function for the game g can be written as:

$$J(g,\mu,\gamma^1,\gamma^2) = \int f(x,y^1,y^2)c(x,u^1,u^2)\bar{Q}^1(dy^1)\bar{Q}^2(dy^2)\zeta(dx).$$

We then create a new cost function $c(x,u^1,u^2,y^1,y^2) := c(x,u^1,u^2)f(x,y^1,y^2)$.

Step (2): Let $x \in \mathbb{X}$ be the random state of nature. Let γ^1, γ^2 be the policies for the DMs, and u^1, u^2 be the resulting actions chosen by the DMs. We allow for policies γ^i where u^i is chosen in a random way, i.e. $u^i = \gamma^i(y^i, \omega^i)$, where ω^i is some $[0,1]$-valued independent random variable (we note that any randomized policy, defined as a stochastic kernel from $\mathbb{Y}^i$ to $\mathbb{U}^i$, admits such a stochastic realization; see [27, Lemma 1.2], or [12, Lemma 3.1]).

Step (3): Let c be the reformulated cost function of this game, under the new product probability measure, we have:

$$J(g,\mu,\gamma^1,\gamma^2) = \int c(x,u^1,u^2,y^1,y^2)(\bar{Q}^1\gamma^1)(dy^1,du^1)(\bar{Q}^2\gamma^2)(dy^2,du^2)\zeta(dx).$$

Here, $(\bar{Q}^1\gamma^1)(dy^1,du^1)$ and $(\bar{Q}^2\gamma^2)(dy^2,du^2)$ are the probability measures induced on the measurement and the action variables. By independence due to the reduction, we can consider the expected cost as a function of the reduced-form policies: $J(g,\mu,\gamma^1,\gamma^2) = J(g,\mu,\bar{Q}^1\gamma^1,\bar{Q}^2\gamma^2)$. Now, without loss of generality, we fix $\bar{Q}^1\gamma^1$, allowing us to express the above equation in the following form:

$$J(g,\mu,\bar{Q}^1\gamma^1,\bar{Q}^2\gamma^2) = \int (\bar{Q}^2\gamma^2)(dy^2,du^2) \int c(x,u^1,u^2,y^1,y^2)(\bar{Q}^1\gamma^1)(dy^1,du^1)\zeta(dx).$$

Let $\int c(x,u^1,u^2,y^1,y^2)(\bar{Q}^1\gamma^1)(dy^1,du^1)\zeta(dx)$ be defined as $\bar{c}(u^2,y^2)$.

Now that we have an independent-measurements reduction, we will (similar to the analysis from [43, 12, 67]), identify, almost surely, every admissible policy with a probability measure on the product space: we adopt the view that, given game (g, μ), $\bar{Q}^i\gamma^i$ is a probability measure on $\mathbb{Y}^i \times \mathbb{U}^i$ with fixed marginal $\bar{Q}^i(dy^i)$ on $\mathbb{Y}^i$. Let Γ^i denote the space of all such measures since every $\bar{Q}^i\gamma^i$ can be identified with an element in Γ^i almost surely. The pairing of an information structure and a policy induces a probability measure P on the five-tuple: $(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2, \mathbb{U}^1, \mathbb{U}^2)$, with

$$P(dx, dy^1, dy^2, du^1, du^2) = \gamma^1(du^1|y^1)\gamma^2(du^2|y^2)Q^1(dy^1|x)Q^2(dy^2|x)\zeta(dx).$$

This construction allows us to obtain a proper topology to work with for spaces of policies with desirable convexity and compactness properties.

We now recall the *w-s* topology [54] on the set of probability measures $\mathcal{P}(\mathbb{X} \times \mathbb{U})$; this is the coarsest topology under which $\int f(x, u)\nu(dx, du) : \mathcal{P}(\mathbb{X} \times \mathbb{U}) \to \mathbb{R}$ is continuous for every measurable and bounded f which is continuous in u for every x (but unlike weak topology, f does not need to be continuous in x). We note that functions which are continuous in one argument and measurable in the other are sometimes referred to as Carathéodory functions. Now, since the exogenous variables are fixed, weak convergence in this setting is equivalent to *w-s* convergence (see [64]), and continuity in the exogenous variable is not needed here. Consider a sequence of actions $(\bar{Q}^2\gamma^2)_m(dy^2, du^2)$ which converges to $(\bar{Q}^2\gamma^2)(dy^2, du^2)$ weakly. We have that $\bar{c}(u^2, y^2)$ is continuous in u^2. Since μ is fixed, the marginals on $\mathbb{Y}^2$ are fixed. Therefore, by [54, Theorem 3.10] (or [7, Theorem 2.5]), we can use the *w-s* topology on the set of probability measures $\mathcal{P}(\mathbb{Y}^2 \times \mathbb{U}^2)$. And so we have continuity of $J(g, \mu, \bar{Q}^1\gamma^1, \bar{Q}^2\gamma^2)$ in $\bar{Q}^2\gamma^2$ in the *w-s* topology and, by the equivalence in this setting, the weak topology.

This also holds for continuity in $\bar{Q}^1\gamma^1$ in the reverse case where we fix $\bar{Q}^2\gamma^2$. Therefore, in general, we have that $J(g,\mu,\cdot,\cdot)$ is continuous in $(\bar{Q}^i\gamma^i)$ when $(\bar{Q}^{-i}\gamma^{-i})$ is fixed.

Step (4): Let $\Gamma = \{\lambda \in \mathcal{P}(\mathbb{Y}\times\mathbb{U}) : \lambda_{\mathbb{Y}} = \bar{Q}\}$ be our reduced policy space, where $\bar{Q}$ is the fixed marginal of the measure λ on $\mathbb{Y}$. Following from [67, Section 2.1], the space of all $\bar{Q}^i\gamma^i$ (which we denote by Γ^i) is compact under weak convergence.

Step (5): We observe that for fixed g and μ, $J(g,\mu,\bar{Q}^1\gamma^1,\bar{Q}^2\gamma^2)$ is linear and hence is both concave and convex in the third and fourth entries. For completeness, we establish this linearity result. Take $\theta \in (0,1)$. Then, without loss of generality, we fix $\bar{Q}^1\gamma^1$ and obtain the following:

$$
\begin{aligned}
&J(g,\mu,\bar{Q}^1\gamma^1,\theta\bar{Q}^2\gamma^2+(1-\theta)\tilde{Q}^2\tilde{\gamma}^2)\\
=&\int(\theta\bar{Q}^2\gamma^2+(1-\theta)\tilde{Q}^2\tilde{\gamma}^2)(dy^2,du^2)\int_{y^2,u^2}\bar{c}(x,u^2,y^2)\\
=&\int(\theta\bar{Q}^2\gamma^2)(dy^2,du^2)\int_{y^2,u^2}\bar{c}(x,u^2,y^2)+\int(1-\theta)(\tilde{Q}^2\tilde{\gamma}^2)(dy^2,du^2)\int_{y^2,u^2}\bar{c}(x,u^2,y^2)\\
=&\theta J(g,\mu,\bar{Q}^1\gamma^1,\bar{Q}^2\gamma^2)+(1-\theta)J(g,\mu,\bar{Q}^1\gamma^1,\tilde{Q}^2\tilde{\gamma}^2).
\end{aligned}
$$

Lastly, we recall that, under the weak topology, the space of probability measures is a metric space, and thus our spaces Γ^i are Hausdorff spaces.

Since $J(g,\mu,\bar{Q}^1\gamma^1,\bar{Q}^2\gamma^2)$ is continuous, and convex/concave in the compact Hausdorff spaces Γ^i, we have the following equality [24, Theorem 1]:

$$\min_{Q^1\gamma^1}\max_{Q^2\gamma^2} J(g,\mu,Q^1\gamma^1,Q^2\gamma^2) = \max_{Q^2\gamma^2}\min_{Q^1\gamma^1} J(g,\mu,Q^1\gamma^1,Q^2\gamma^2).$$

This establishes a (saddle-point) equilibrium for the game. □

Thus, we have obtained an existence result for the value of the games considered, and also provided an approach to topologize and convexify/compactify the policy spaces.

We also present the following theorem, which is a mild relaxation of the theorem above. The proof below follows similarly to Theorem 2.1.

Theorem 2.2 (Existence of Equilibria with a Further Relaxation)**.** *For fixed information structure μ and a given game g, assume the following hold.*

(i) The action spaces of DMs, $\mathbb{U}^1, \mathbb{U}^2$, are compact.

(ii) The cost function c is bounded and continuous in DMs' actions, for every state of nature x.

Then an equilibrium exists under possibly randomized policies, and so there exists a value of the zero-sum game.

Proof. (Sketch.) Here, we let $\mu_{\mathbb{Y}^i}$ denote the marginal of the information structure μ on $\mathbb{Y}^i$. We will combine our policies with these marginals to form the product measures $\mu_{\mathbb{Y}^1}(dy^1)\gamma^1(du^1|y^1)$ and $\mu_{\mathbb{Y}^2}(dy^2)\gamma^2(du^2|y^2)$ on the DMs' measurement and action spaces. We will denote these measures by $(\mu_{\mathbb{Y}^1}\gamma^1)$ and $(\mu_{\mathbb{Y}^2}\gamma^2)$

Similar to the previous proof, without loss of generality, we fix DM 2's strategy γ^2. Then we have:

$$
\begin{aligned}
&J(g, \mu, \gamma^1, \gamma^2) \\
&\quad = \int \mu_{\mathbb{Y}^1}(dy^1)\gamma^1(du^1|y^1) \left(\int_{\mathbb{X}\times\mathbb{Y}^2\times\mathbb{U}^2} \mu(dx, dy^2|dy^1)\gamma^2(du^2|y^2)c(x, u^1, u^2) \right)
\end{aligned}
$$

Let $\int_{\mathbb{X}\times\mathbb{Y}^2\times\mathbb{U}^2} \mu(dx, dy^2|dy^1)\gamma^2(du^2|y^2)c(x, u^1, u^2)$ be defined as $\bar{c}(u^1, y^1)$.

We can observe that, by assumption, $\bar{c}(u^1, y^1)$ is bounded and is continuous in u^1. Furthermore, it is also evident that $\bar{c}(u^1, y^1)$ is measurable in y^1. By the same arguments of **Step (2)** of the preceding theorem, via the machinery of the w-s topology [54], we can show that $\bar{c}(u^1, y^1)$ is continuous in $(\mu_{\mathbb{Y}^1}\gamma^1)$ under w-s convergence, and thus also under weak convergence. This also holds for continuity in $(\mu_{\mathbb{Y}^2}\gamma^2)$ in the reverse case where we fix γ^1.

Following from [67, Section 2.1], the space of all $\mu^i_{\mathbb{Y}^i}\gamma^i$ is compact under weak convergence, and we can observe that $J(g, \mu, \mu_{\mathbb{Y}^1}\gamma^1, \mu_{\mathbb{Y}^2}\gamma^2)$ is linear and hence is both concave and convex in the third and fourth entries.

The existence of a (saddle-point) equilibrium for the game then follows by [24, Theorem 1]. □

Chapter 3

Comparison of Information Structures

3.1 Introduction

Characterizing the value of information structures is a problem in many disciplines involving decision making under uncertainty. In stochastic control theory, it is well-known that more information cannot hurt a given decision-maker since the decision-maker can always choose to ignore this information. In statistical decision theory involving a single decision-maker, one says that an information structure is better than another one if for any given measurable and bounded cost function involving a hidden state variable and an action variable which is restricted to be only a function of some measurement, the solution value obtained under optimal policies under the former is not worse than the value obtained under the latter. For finite probability spaces, Blackwell's celebrated theorem [10] on the ordering of information structures obtains a precise characterization of when an information structure is better. This finding has inspired much further research as reviewed in e.g. [15, 47].

Since Blackwell's seminal 1953 paper [10], significant work has been done to extend Blackwell's results to team problems and games. Stochastic team problems (known also as identical interest games) were studied in a finite-space setting by Lehrer, Rosenberg, and Shmaya [21]; see also [65, Chapter 4]. The value of information in various types of repeated games has also been explored in [34], [35], and [37].

In general games, information can have both positive and negative value to a DM since additional information can lead to a perturbation which is not necessarily monotone due to the presence of competitive equilibrium, unlike in a team setup. Some of the earlier accounts on such phenomena are [30] and [4], where the latter studied the comparison of information structures for team-like (LQG) and zero-sum like (quadratic duopoly) games.

As noted above, for general non-zero sum game problems, informational aspects are very challenging to address and more information can hurt some or even all of the DMs in a system, see [30, 28, 33, 3]. To make this discussion more concrete, we provide the following example from the literature, Example 3.1, found in some form in e.g. [9, 33]. While Example 3.1 simply illustrates the desired point, it utilizes a dynamic game rather than a static one; for completeness, we have also included a longer original example demonstrating the same phenomenon in a static game.

Example 3.1. *Consider a card drawn at random from a deck, where its colour can be either red or black, each with probability* 1/2*. DM 1 first declares his guess of the colour, and then, after hearing what DM 1 guessed, DM 2 submits her guess for the colour. If both DMs guess the same colour, the payout is $2 each, whereas if one DM guesses correctly, that DM receives a payout of $6 and the other DM receives $0.*

In the case where both DMs are uninformed about the colour of the card, the expected payout is $3 each, as DM 1's optimal strategy is arbitrary, and DM 2's optimal strategy is to guess the opposite colour of what DM 1 guessed.

In the case where both DMs are informed of the colour of the card prior to declaring their guess, the equilibrium for the game occurs when both DMs guess the true colour of the card. In this case, the expected payout becomes \$2 for each DM. ◇

In the case where only the first DM is informed of the colour of the card prior to declaring his guess, the pure-strategy equilibrium also occurs when both DMs guess the true colour of the card with an expected payout of \$2 each, and so DM 1 gaining information results in a loss for both players.

Example 3.2. *Let* $\mathbb{X} = [0, 1]$*, with prior distribution* ζ *defined by the continuous uniform distribution. Let* $\mathbb{Y}^1 = \mathbb{Y}^2 = [0, 1]$*.*

We define the action spaces of the respective DMs as $\mathbb{U}^1 = \mathbb{U}^2 = [0, 1]$*, and the cost functions as:*

$$c^1(x, u^1, u^2) = \begin{cases} (x - u^1)^2 + 2, & u^2 = x \\ 2(x - u^1)^2 + (u^1 - u^2)^2, & u^2 \neq x \end{cases}$$

$$c^2(x, u^1, u^2) = \begin{cases} (x - u^2)^2 + 2, & u^1 = x \\ 2(x - u^2)^2 + (u^1 - u^2)^2, & u^1 \neq x \end{cases}$$

Under measurement channels where neither DM has any information regarding x*, (which could be achieved for instance by measurement channels returning* $y^1 = 0 = y^2, \forall x \in \mathbb{X}$*), there exists a unique Nash equilibrium solution given with* $u^1 = \frac{1}{2} = u^2$*. Uniqueness can be seen as follows: when the DMs have no information regarding* x*, the event that* $u^1 = x$ *or* $u^2 = x$ *has zero measure and therefore the complementary conditions are active for defining the cost in the above. Note though that in this case the cost functions essentially turn the problem into a team problem since minimization* *over u^1 of*

$$2(x-u^1)^2+(u^1-u^2)^2$$

will lead to the same solution as the minimization of

$$2(x-u^1)^2+(u^1-u^2)^2+2(x-u^2)^2.$$

The same applies for u^2, and so we can view the DMs as solving essentially the same problem. Therefore, we have a standard static quadratic team problem which admits a unique optimal solution [65, Theorem 2.6.3] with $u^1=\frac{1}{2}=u^2$. This leads to an expected cost of $\frac{1}{6}$ for each DM.

However, when the DM measurement channels are such that both DMs have perfect information regarding x (e.g. through $y^1=x=y^2$), we first see that another equilibrium solution arises: with $u^1=x=u^2$. One can observe this is a Nash equilibrium by noting that if DM 2 holds their strategy constant as $u^2=x$, then DM 1 will play the game with cost function $(x-u^1)^2+2$, which is minimized by playing $u^1=x$. The same holds true for DM 2 in the reverse case where DM 1's strategy is fixed. We now make the point that this is the unique deterministic Nash equilibrium:

The equilibrium under no information of $u^1=\frac{1}{2}=u^2$ is no longer a Nash equilibrium under perfect information. We can observe this by noting that if DM 2 holds their strategy constant as $u^2=\frac{1}{2}$, then DM 1 will almost surely be tasked with minimizing $2(x-u^1)^2+(u^1-\frac{1}{2})^2$. However, now that DM 1 knows x, and because the first term is weighted more heavily, the expected value of the cost function is minimized by DM 1 playing $u^1=\frac{2}{3}x+\frac{1}{6}$.

In fact, if the condition $u^1=x=u^2$ is not active, we would again reduce to a *team problem and in this case, we would have*

$$\gamma^1(x) = \frac{2}{3}x + \frac{1}{3}\gamma^2(x), \qquad \gamma^2(x) = \frac{2}{3}x + \frac{1}{3}\gamma^1(x)$$

leading to, by uniqueness [65, Theorem 2.6.3], $\gamma^1(x) = x = \gamma^2(x)$*, which however would take us back to the inefficient equilibrium given above. Thus, the unique pure strategy Nash equilibrium for the full-information problem is at* $u^1 = x = u^2$*. This results in an expected cost of* 2 *for each DM.*

This example demonstrates a static game in which more information hurts both DMs.

Bassan et al. provided sufficient conditions for games to have the 'positive value of information property', where providing additional information to some or all DMs results in greater or equal payoffs for all DMs [9]. Gossner and Mertens highlighted zero-sum games as a particularly interesting class to study in the context of ordering information structures in games and did preliminary work on this ordering [28]; zero-sum games provide a worthwhile class of games to study due to the fact that, under mild conditions, every game has a value (achieved at a saddle point), as discussed in Chapter 2.

For comparison of information structures in zero-sum games with finite measurement and action spaces, Pęski provided necessary and sufficient conditions, and thus a complete characterization [46]. Prior to Pęski's results, De Meyer, Lehrer, and Rosenberg had shown that the value of information is positive in zero-sum games, albeit with a slightly different setup than Pęski, where their payoff depended on an individual 'type' for each DM rather than a common state of nature; their results were applicable for infinite action spaces and finite type spaces [19]. Furthermore, Lehrer and Shmaya studied a 'malevolent nature' zero-sum game played between nature and a DM in a finite setting, and characterized a partial ordering of information structures for these games [42]. A recent comprehensive study on the value and topological properties of information structures in zero-sum games, which also generalizes [46] to the countably infinite probability space setup, is [26].

In this chapter, we generalize Pęski's results to a broad class of zero-sum games with standard Borel measurement and action spaces.

Toward this goal, we also present sufficient conditions for a partial converse to Blackwell's ordering when the DM has standard Borel measurement and action spaces and the unknown variable also takes values from a standard Borel space.

3.2 Preliminaries and Literature Review

3.2.1 Comparison of information structures in single-agent problems

Consider a stochastic game following the setup of Section 1.2 with $N = 1$. That is, where there is only one decision-maker. We refer to this as a single-DM decision problem.

The comparison question for a single DM is the following: when can one compare two measurement channels Q_1, Q_2 (which induce information structures μ_1, μ_2, respectively) such that

$$\inf_{\gamma \in \mathbf{\Gamma}} J(g, \mu_1, \gamma) \leq \inf_{\gamma \in \mathbf{\Gamma}} J(g, \mu_2, \gamma),$$

for all games g in a large class of single-DM decision problems?

We now recall the notion of garbling. We note that garbling is sometimes defined to be equivalent to physical degradedness of communication channels (as opposed to stochastic degradedness) [17], however we will take stochastic degradedness and garbling to be equivalent.

Definition 3.1. *An information structure induced by some channel Q_2 is garbled (or stochastically degraded) with respect to another one, Q_1, if there exists a channel Q' on $\mathbb{Y} \times \mathbb{Y}$ such that*

$$Q_2(B|x) = \int_{\mathbb{Y}} Q'(B|y) Q_1(dy|x),\ B \in \mathcal{B}(\mathbb{Y}),\ \zeta\ a.s.\ x \in \mathbb{X}.$$

We also define the notion of *more informative than* and introduce a useful result:

Definition 3.2. *An information structure μ is more informative than another information structure ν for some class of single DM decision problems $\mathcal{G}$ if*

$$\inf_{\gamma \in \Gamma} E_{\zeta}^{\nu,\gamma}[c(x,u)] \geq \inf_{\gamma \in \Gamma} E_{\zeta}^{\mu,\gamma}[c(x,u)],$$

for all single DM decision problems $(c(x,u), \mathbb{U})$ in $\mathcal{G}$.

Proposition 3.1. *The function*

$$V(\zeta) := \inf_{u \in \mathbb{U}} \int c(x,u) \zeta(dx),$$

is concave in ζ, under the assumption that c is measurable and bounded.

For a proof of this proposition see [65, Theorem 4.3.1].

We emphasize that in Definition 3.2, $\mathbb{U}$ is also a design variable for the decision problem. With this in mind, and in view of Proposition 3.1, we state Blackwell's classical result in the following.

Theorem 3.1. *[Blackwell [10]] Let $\mathbb{X}, \mathbb{Y}$ be finite spaces. The following are equivalent:*

(i) Q_2 is stochastically degraded with respect to Q_1 (that is, a garbling of Q_1).

(ii) The information structure induced by channel Q_1 is more informative than the one induced by channel Q_2 for all single DM decision problems with finite $\mathbb{U}$.

That (i) implies (ii) for general spaces follows from Proposition 3.1, which is an immediate finding in statistical decision theory, and Jensen's inequality [65, Theorem 4.3.2]. We also note that this result will hold, and the proof will follow in an identical manner, if the DM is allowed to use randomized policies, i.e. $u = \gamma(y, \omega)$, where ω is an independent noise variable.

The converse, ii) implies i), is significantly more challenging. For the case with general spaces, related results are attributed to [11], and [16], [57], which relate an ordering of information structures in terms of dilatations and their relation with comparisons under concave functions defined on conditional probability measures. A very concise yet informative review is in [15, p. 130-131] and a more comprehensive review is in [58]. We will present a direct proof that will be utilized in our main result of the chapter and present a comparative discussion.

In single DM setups, a related result to that of Blackwell is due to Le Cam [14], to be discussed further in Chapter 4 in some detail in the context of regularity problems. In addition, there is a comparison criterion for channels (and the joint map from a source and channel output after coding and decoding) developed by

Shannon in [56] and expanded upon in [47]. This leads to a comparison criterion for information structures in single-DM decision problems, based on communication-theoretic relaxations of such problems.

3.2.2 Comparison of information structures in zero-sum game problems

Now, consider a zero-sum game generalization of the problem above, with two DMs, following the setup of Section 1.2.1.

Recall that the joint measure $P(dy^1, dy^2, dx)$ defines the information structure for the game and let us denote this with μ. For a zero-sum game with the conditional independence assumption in Assumption 2.2, an information structure μ consists of private information structures μ^1 and μ^2 defined with $Q^i, i = 1, 2$. Define μ^i as the joint probability measure induced on $\mathcal{P}(\mathbb{X}\times\mathbb{Y}^i)$ by measurement channel Q^i with input distribution ζ. We will present the main results of this chapter with an additional characterization using Q^1 and Q^2, when the conditional independence assumption holds.

Recall that DM 1 (the minimizer) wishes to minimize the cost and DM 2 (the maximizer) wishes to maximize the cost. In the spirit of Blackwell, we now define the notion of an information structure being "better" than another in the zero-sum game setup.

Definition 3.3. *For fixed $\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2$ and ζ such that $x \sim \zeta$, we say that an information structure μ is better for the maximizer than information structure ν (written as $\nu \lesssim \mu$) over all games in a class of games $\mathbb{G}$ if and only if for all games g in $\mathbb{G}$:*

$$J^*(g, \mu) \geq J^*(g, \nu).$$

Definition 3.4. *We denote by $\kappa^i\mu$ the information structure in which DM i's information from μ is garbled by a stochastic kernel κ^i. We let -i denote the other DM in the game. Explicitly, this means the information structure becomes:*

$$(\kappa^i\mu)(B, dy^{-i}, dx) = \int_{\mathbb{Y}^i} \kappa^i(B|y^i)\mu(dy^i, dy^{-i}, dx),\ B \in \mathcal{B}(\mathbb{Y}^i).$$

We use K^i to denote the space of all such stochastic kernels κ^i for DM i.

Theorem 3.2 (Pęski [46])**.** *Let $\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2$ be finite. For any two information structures μ and ν, μ is better for the maximizer than ν over all games with finite action spaces $\mathbb{U}^1, \mathbb{U}^2$ if and only if there exist kernels $\kappa^i \in K^i, i = 1, 2,$ such that*

$$\kappa^1\nu = \kappa^2\mu,$$

In particular, under Assumption 2.2, we have the more explicit characterization with

$$\kappa^1 Q^1_\nu = Q^1_\mu \quad \text{and} \quad Q^2_\nu = \kappa^2 Q^2_\mu.$$

Where Q^i_μ and Q^i_ν are the measurement channels for DM i under information structures μ and ν, respectively.

In this chapter we will obtain a standard Borel generalization of this result. We also note that [26] extends the above result to a setup with countably-infinite spaces; this result will be discussed further in Chapter 4 in the context of continuity of the value function.

3.2.3 Team Theoretic Setup

For completeness, we also discuss the team theoretic setup in our review.

Lehrer, Rosenberg and Shmaya extended Blackwell's ordering of information structures to team problems in finite-space settings for various solution concepts, including Nash equilibrium and several forms of *correlated* equilibrium, in [39]. For these results to hold, various degrees of correlation between the DMs' private signals is allowed. These solution concepts for correlated equilibrium are adopted from [25], which builds on ideas first introduced in [2]. These provide an ordering of information structures for static stochastic team problems. Related results are discussed in [65, Chapter 4]. Recently, advances have been made in understanding the topological properties of strategic measures in team problems in [67].

3.3 On a Partial Converse to Blackwell Ordering in the Standard Borel Setup

We need to address the extension of Blackwell's ordering of information structures to the infinite case, as this will form a key aspect of the proof of the main result of this chapter, Theorem 3.5.

Here, we present a partial converse to Blackwell's theorem.

The forward direction to Blackwell's theorem holds in the infinite case (see [65, Theorem 4.3.2]), i.e. when $\mathbb{X}$, $\mathbb{Y}$ are standard Borel spaces for a single-DM setup, ν being a stochastically degraded version of μ implies that μ is more informative than ν over all single-DM decision problems with standard Borel action spaces and bounded cost functions that are continuous in the DM's action for every state of nature.

As noted earlier, related results were presented by C. Boll in 1955 in an unpublished thesis paper [11]. Le Cam presents a summary of these results in [15], with a detailed review reported in [58]. The approach in the literature often builds on the construction of *dilatations* of conditional probability measures, which is related to Blackwell's comparison of experiments theorem through what is known as the Blackwell-Sherman-Stein theorem. A detailed comparative analysis is provided further below. Our main contribution here is an explicit converse compatible with the conditions on existence results presented in the previous section and a comparison to be presented in the next section. This result serves as a supporting step with a direct proof; the results reported in the literature are often very technical and the explicit implication for our setup is not evident *a priori* as we discuss in the next subsection.

We note that our setup differs slightly from that of Blackwell in [10], contributing to the fact that this is a *partial* converse to Blackwell's result. In Blackwell's original setup with finite $\mathbb{X}$, information structures could be compared over different priors on $\mathbb{X}$ as the comparison would apply uniformly to all such prior measures that satisfy a positivity condition on each of the finitely many outcomes. In our setup, since the space is possibly uncountable, we consider a fixed prior measure on $\mathbb{X}$.

Theorem 3.3. *Let us consider a single DM whose goal is to minimize the value of the cost function c for a set of single-DM decision problems. If $\mathbb{Y}$ is compact and an information structure μ is more informative than another information structure ν over all single-DM decision problems with compact standard Borel action spaces and bounded cost functions $c : \mathbb{X} \times \mathbb{U} \to \mathbb{R}$ that are continuous in u for every x, then ν must be a garbling of μ in the sense of Definition 3.1.*

Proof. We note that under the conditions of the theorem, an optimal policy (which is also deterministic) exists for every information structure (see Theorem 3.1 in [66]).

Step (1): Let ζ be the fixed probability distribution on $\mathbb{X}$ for any given decision problem in our set. Take information structures $\mu, \nu \in \{\mathcal{P}(\mathbb{X} \times \mathbb{Y}) : P_{\mathbb{X}} = \zeta\}$, where μ is more informative than ν in Blackwell's sense (i.e. $J^*(g, \mu) \geq J^*(g, \nu)$ over all games with a compact standard Borel action space and a bounded cost function c that is continuous in u).

Take the space $K\mu$, a subset of $\mathcal{P}(\mathbb{X} \times \mathbb{Y})$, to be the space of all possible garblings of μ, where the garblings are from $\mathbb{Y}$ to $\mathbb{Y}$.

Step (2): We now establish the weak compactness of the space of all garbled information structures.

First, observe that the set of all induced garblings on the product space (involving all of K) inducing probability measures of the form

$$P_\kappa(dx, dy, d\tilde{y}) := \mu(dx, dy)\kappa(d\tilde{y}|y),$$

leads to a weakly pre-compact space in the space of probability measures on $\mathbb{X}\times\mathbb{Y}\times\mathbb{Y}$. If closedness can also be established, this would lead to a weakly compact space. This follows from the proof of [64, Theorem 5.6] or [52, Theorem 5.2]: since the marginals on $\mathbb{X} \times \mathbb{Y}$ are fixed, any limit of a weakly converging sequence will also satisfy the property that the limit is a garbling of the original information structure. For completeness, we present the following: With $P_\kappa(dx, dy, d\tilde{y}) = \kappa(d\tilde{y}|y)\mu(dx, dy)$, consider a weakly converging sequence $P_{\kappa_n}(dx, dy, d\tilde{y})$. We will show that the weak limit also admits such a garbled structure. Let $P_{\kappa_n}(dx, dy, d\tilde{y})$ converge weakly to

$P(dx, dy, d\tilde{y})$. Then, for every continuous and bounded h

$$\int h(x, y, \tilde{y}) P_{\kappa_n}(dx, dy, d\tilde{y}) = \int \left(\int h(x, y, \tilde{y}) \mu(dx|y) \right) P_{\kappa_n}(dy, d\tilde{y}).$$

Since the marginal on y is fixed, even though the function $\int h(x, y, \tilde{y}) \mu(dx|y)$ is only measurable and bounded in y and is continuous in $\tilde{y}$, w-s convergence is equivalent to the weak convergence of $P_{\kappa_n}(dy, d\tilde{y})$ and as a result we have that

$$\int \left(\int h(x, y, \tilde{y}) \mu(dx|y) \right) P_{\kappa_n}(dy, d\tilde{y}) \to \int \left(\int h(x, y, \tilde{y}) \mu(dx|y) \right) P(dy, d\tilde{y}).$$

As a result, P decomposes as $P(dx, dy, d\tilde{y}) = \mu(dx, dy)\tilde{\kappa}(d\tilde{y}|y)$ for some $\tilde{\kappa}$. This establishes the weak compactness of the garbled information structure in the product space $\mathbb{X} \times \mathbb{Y} \times \mathbb{Y}$.

Now, take the projection of this space onto the measures on the first and the third coordinate; as a continuous image of a weakly compact set, this map will also be compact and gives us our space $K\mu$.

Finally, $K\mu$ is convex, since the space of stochastic kernels is convex. As a result, the space $K\mu$ of all possible garblings of μ is a convex and compact subset of $\mathcal{P}(\mathbb{X} \times \mathbb{Y})$ under the weak convergence topology.

Now, assume there does not exist a stochastic kernel $\kappa \in K$ such that $\nu = \kappa\mu$. Which is to say, we assume ν is not a garbling of μ and proceed with a proof by contradiction. Then, $K\mu \cap \nu = \phi$. That is, $\nu \notin K\mu$.

Step (3): We now use the Hahn-Banach Separation Theorem for Locally Convex Spaces by treating the space of probability measures $\mathcal{P}(\mathbb{X} \times \mathbb{Y})$ as a locally convex space of measures (see [51, Theorem 3.4]). As such, since our spaces $K\mu$ and $\{\nu\}$

are subsets of this space and are convex, closed and compact, in addition to being disjoint, we can separate them using a continuous linear map from $\mathcal{P}(\mathbb{X} \times \mathbb{Y})$ to $\mathbb{R}$.

To apply [51, Theorem 3.4], we require local convexity of $\mathcal{P}(\mathbb{X} \times \mathbb{Y})$, and so we define the locally convex space of probability measures with the following notion of convergence: We say that $\nu_n \to \nu$ if $\int f(x,y)\nu_n(dx,dy) \to \int f(x,y)\nu(dx,dy)$ for every measurable and bounded function which is continuous in y for every x. We note that our measures must still have fixed marginal ζ on $\mathbb{X}$.

Since continuous and bounded functions *separate* probability measures (in the sense that, if the integrations of two measures with respect to continuous functions are equal, the measures must be equal), it follows from [51, Theorem 3.10] that we can represent every continuous linear map on $\mathcal{P}(\mathbb{X} \times \mathbb{Y})$ using the form $\int f(x,y)\nu(dx,dy)$ for some measurable and bounded function $f(x,y)$ continuous in y for every x. It also follows from [51, Theorem 3.10] that, given this notion of convergence, $\mathcal{P}(\mathbb{X} \times \mathbb{Y})$ is a locally convex space.

Therefore, we have the following statement from combining [51, Theorem 3.4] and [51, Theorem 3.10]: there exists a measurable and bounded function (continuous in y) $f : \mathbb{X} \times \mathbb{Y} \to \mathbb{R}$ and constants $D_1, D_2 \in \mathbb{R}$ where $D_1 < D_2$ such that:

$$\langle \nu, f \rangle \leq D_1, \ \langle \kappa\mu, f \rangle \geq D_2, \ \ \forall \, \kappa \in K.$$

Where we use the following notation:

$$\langle \nu, f \rangle = \int_{\mathbb{X} \times \mathbb{Y}} f(x,y)\nu(dx,dy).$$

This gives us the following inequality: $\langle \nu, f \rangle < \langle \kappa\mu, f \rangle$, $\forall \kappa \in K$.

Step (4): Now consider the class of decision problems with bounded cost functions continuous in the actions, with compact $\mathbb{Y}$, $\mathbb{U}$, where $\mathbb{U} = \mathbb{Y}$. This is clearly a subset of all decision problems considered so far in the proof. Now let $f(x, y)$ be the separating function found above. Consider a game in this particular subclass where $f(x, y)$ is the cost function (which is valid since $\mathbb{U} = \mathbb{Y}$ and $f(x, y)$ is bounded continuous in y). We note that $\langle \nu, f \rangle$ gives the expected value of the game with cost function $f(x, y)$ under information structure ν when the DM plays the identity policy $\gamma^{id}(y) = y$. We can observe the following:

$$\int_{\mathbb{X}\times\mathbb{Y}} f(x,y)\nu(dx,dy) < \int_{\mathbb{X}\times\mathbb{Y}} f(x,y)\kappa\mu(dx,dy), \ \ \forall \kappa \in K,$$

and hence,

$$\begin{aligned}\int_{\mathbb{X}\times\mathbb{Y}} f(x,y)\nu(dx,dy) &< \inf_{\kappa\in K} \int_{\mathbb{X}\times\mathbb{Y}} f(x,y)\kappa\mu(dx,dy) \\ &= \inf_{\kappa\in K} \int_{\mathbb{X}\times\mathbb{Y}} f(x,y') \int_{\mathbb{Y}} \kappa(dy'|y)\mu(dx,y) \\ &= \inf_{\kappa\in K} \int_{\mathbb{X}\times\mathbb{Y}} f(x,\kappa(\cdot|y))\mu(dx,dy).\end{aligned}$$

Where we define:

$$f(x,\kappa(\cdot|y)) := \int_{\mathbb{Y}} f(x,y')\kappa(dy'|y)$$

Recalling that $\kappa(\cdot|y)$ has a functional representation $\gamma(y) = g(y, \omega)$ for some independent noise variable ω, and since K is the space of all stochastic kernels from $\mathbb{Y}$ to $\mathbb{Y}$, we can observe that this gives us:

$$\inf_{\kappa\in K} \int_{\mathbb{X}\times\mathbb{Y}} f(x,\kappa(\cdot|y))\mu(dx,dy) = \inf_{\gamma\in\Gamma} \int_{\mathbb{X}\times\mathbb{Y}} f(x,\gamma(y))\mu(dx,dy).$$

Since we allow for randomized policies, this minimization is equivalent to finding the optimal policy $\gamma^* \in \Gamma$ for the cost function $f(x, y)$ under information structure μ. We refer to this single-DM game with cost function f and action space $\mathbb{Y}$ as $\tilde{g}$. And so we have:

$$\int_{\mathbb{X}\times\mathbb{Y}} f(x,y)\nu(dx,dy) = J(\tilde{g},\nu,\gamma^{id}) < J(\tilde{g},\mu,\gamma^*) = \inf_{\kappa\in K}\int_{\mathbb{X}\times\mathbb{Y}} f(x,\kappa(y))\mu(dx,dy).$$

Since we have found a game where, when playing its optimal policy, μ performs worse than ν does under some policy, we have contradicted the fact that μ is better than ν. Therefore, there must exist a $\kappa \in K$ such that $\nu = \kappa\mu$, and so ν is a garbling of μ. □

This result will allow us to use both directions of Blackwell's ordering of information structures in the standard Borel-type setup we are considering for DMs in zero-sum games.

Dilatations as measures for comparisons of experiments and Strassen's theorem. Strassen, in [57, Theorem 2], presents a related result that is often invoked when comparison of experiments is studied in infinite dimensional probability spaces, although the direct implication on Blackwell's ordering (in the sense needed in our main result to be presented in the next section) is not explicit as we note in the following. Likewise, Cartier, Fell, and Meyer relate an ordering of information structures in terms of dilatations (where the hidden variable x does not appear explicitly in the analysis) in [16]. A very concise yet informative review is in [15, p. 130-131]. A detailed discussion on comparisons of information structures along the same approach is present in the comprehensive book [58]. Both for completeness as well as to compare the findings, we present a discussion in the following.

Let Ω be a convex compact metrizable subset of a locally convex topological vector space. For Borel probability measures μ and ν write $\mu \prec \nu$ if and only if for all $y \in \mathcal{S} = \{$all continuous concave functionals on $\Omega\}$

$$\int y \, d\mu \geq \int y \, d\nu.$$

Theorem 3.4. *[57, Theorem 2] $\mu \prec \nu$ if and only if there is a dilatation P such that $\nu = P\mu$, where a dilatation P is a Markov kernel from Ω to Ω such that for all continuous affine functions z on Ω, $zP = z$, where $zP = z$ means that for any continuous affine function z on Ω:*

$$\int_\Omega z(r)P(dr,\omega) = z(\omega), \quad \forall \omega \in \Omega.$$

Theorem 3.4 does not lead to a converse to Blackwell's theorem in the generality presented in Theorem 3.3: Let Ω be the space of probability measures on $\mathbb{X}$. Let μ be an information structure that is more informative than another information structure ν in Blackwell's sense. Let us restrict ourselves to decision problems where $\mathbb{U}$ is compact. Let Q_μ and Q_ν be the measurement channels for the DM under information structures μ and ν, respectively. By definition, we have for all measurable and bounded cost functions c continuous in the actions:

$$\inf_{\gamma \in \Gamma} \int \zeta(dx) Q_\mu(dy|x) c(x,\gamma(y)) \leq \inf_{\eta \in \Gamma} \int \zeta(dx) Q_\nu(dy|x) c(x,\eta(y)).$$

Let $P^{\mu}(dy)Q(dx|y)$ be the alternative disintegration of the information structure μ following Bayes' rule. Likewise, perform the same disintegration for ν. Then we can rewrite the above equation as (due to the measurable selection conditions as in the proof of Theorem 3.1 in [66]):

$$\int P^{\mu}(dy)(\inf_{u\in\mathbb{U}}\int Q(dx|y)c(x,u)) \leq \int P^{\nu}(dy)(\inf_{u\in\mathbb{U}}\int Q(dx|y)c(x,u)), \tag{3.1}$$

Now, we define:

$$\Pi^{\mu}(A) := \int_{\mathbb{Y}} P^{\mu}(dy)1_{Q^{\mu}(\cdot|A)}$$

We note that Π^{μ} is a probability measure on Ω. Define Π^{ν} similarly. Then (3.1) becomes

$$\int \Pi^{\mu}(d\pi)(\inf_{u\in\mathbb{U}}\int \pi(dx)c(x,u)) \leq \int \Pi^{\nu}(d\pi)(\inf_{u\in\mathbb{U}}\int \pi(dx)c(x,u)),$$

with the interpretation that $\pi(dx) = Q(dx|y)$. Let $W^*(\pi) = \inf_{u\in\mathbb{U}}\int \pi(dx)c(x,u)$. Then we can rewrite this once again as:

$$\int \Pi^{\mu}(d\pi)W^*(\pi) \leq \int \Pi^{\nu}(d\pi)W^*(\pi).$$

Since Π^{μ} and Π^{ν} give probability distributions on Ω, and W^* is a function over Ω, we will have $\nu \prec \mu$ in Strassen's sense if the above inequality holds for all continuous and concave functions over Ω.

We can show that W^* is continuous and concave in π provided that additionally c is continuous *both* in x and u. Let $\pi_n \to \pi$ weakly, and let u_n^* be optimal for π_n.

Then:

$$
\begin{aligned}
&|\int c(x,u_n^*)\pi_n(dx) - \int c(x,u^*)\pi(dx)| \\
&\le \max(\int c(x,u_n^*)(\pi_n(dx)-\pi(dx)), \int c(x,u^*)\pi_n(dx)-\pi(dx)).
\end{aligned}
$$

We note that $\int c(x,u_n^*)(\pi_n(dx)-\pi(dx))$ goes to 0 following [55, Theorem 3.5] or [38, Theorem 3.5] (since the action space is compact, there always is a converging subsequence $u_{n_m}^* \to \bar{u}$ for some $\bar{u}$, and since for $x_n \to x$ we have that $c(x_{n_m}, u_{n_m}^*) \to c(x,\bar{u})$ the result follows from a generalized convergence theorem under weak convergence). The second term converges to zero by the weak convergence of π_n to π. We emphasize the requirement that c is continuous in both x and u, in Theorem 3.3 only continuity in u was required (one can construct a simple counterexample, even when $\mathbb{U}$ is a singleton to show that continuity in x is necessary for this argument to hold). Concavity of W^* in the conditional measure $\pi(dx)$ follows from Proposition 3.1.

Now, if one can show that by using all bounded continuous cost functions c and compact action spaces $\mathbb{U}$, the space of all continuous and concave functions on Ω is spanned by the space of all W^* functions, then a converse can be attained through Strassen's result. We note here that every concave and upper semi-continuous W can be written as an infimum of a family of affine functions (Fenchel-Moreau Theorem [48]) and an analysis can be pursued towards this direction at least for the case where c can be assumed to be continuous in both variables and the condition on W is to be relaxed in Strassen's theorem. However, due to the conditions of upper semi-continuity of W^* and the joint continuity of c noted earlier in both the state and actions, the applicability of Strassen's theorem to our setup does not hold in the generality reported.

In summary, we have presented a general condition and a direct proof, while we recognize that Strassen's theorem and accordingly its proof could be further modified to allow for additional relaxations for arriving at a similar result.

3.4 Comparison of Information Structures for Zero-Sum standard Borel Bayesian Games

We are now prepared to order information structures in the spirit of Theorem 3.2 for this standard Borel setup. We note that the following lemmas, theorem, and corollary also hold in the general finite case studied by Pęski, as they rely solely on the existence of equilibria (which are guaranteed to exist in the finite setup by von Neumann's min-max theorem, see [60]) and Blackwell's ordering of information structures. Therefore, these results also serve as a strict generalization of Theorem 3.2 to standard Borel Bayesian Games.

Definition 3.5. *For fixed $\mathbb{X}$ with $x \sim \zeta$, and fixed $\mathbb{Y}^1, \mathbb{Y}^2$, we define a class of games $\tilde{\mathbb{G}}_\zeta(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$ to be all games for which the DMs have compact action spaces and the cost function is bounded and continuous in DMs' actions for every state x.*

Lemma 3.1. *Given fixed $\mathbb{X}$, ζ, $\mathbb{Y}^1$, and $\mathbb{Y}^2$, for any information structure μ and any kernels $\kappa^i \in K^i$:*

$$\kappa^2\mu \lesssim \mu \quad \textit{and} \quad \mu \lesssim \kappa^1\mu,$$

over all games in $\tilde{\mathbb{G}}_\zeta(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$.

Proof. Let us consider the first relation. This states essentially that garbling the maximizer's information results in lower or equal cost over all games.

Take an arbitrary zero-sum game $g \in \tilde{\mathbb{G}}_\zeta(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$ with cost function c and action spaces $\mathbb{U}^1$ and $\mathbb{U}^2$. Let (γ^1, γ^2) be the Nash equilibrium policies for the DMs under information structure $\kappa^2\mu$ and (η^1, η^2) be the Nash equilibrium policies under information structure μ. By our assumption on $\tilde{\mathbb{G}}_\zeta(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$, these policies exist [Theorem 2.2]. Let Q^i be the measurement channel for DM i under information structure μ.

The expected value of the cost for the maximizer under the first information structure is:

$$J_g^{\kappa^2\mu}(\gamma^1, \gamma^2) \coloneqq \int_{\mathbb{X}\times\mathbb{Y}^1\times\mathbb{Y}^2} c(x, \gamma^1(y^1), \gamma^2(y^2))\kappa^2\mu(dx, dy^1, dy^2).$$

Since, the game g and information structure μ are fixed, we use $J_g^{\mu}(\gamma^1, \gamma^2) \coloneqq J(g, \mu, \gamma^1, \gamma^2)$ to simplify notation and highlight dependence of expected cost on the DMs' policies.

By definition, the equilibrium solution (γ^1, γ^2) for g under $\kappa^2\mu$ is given by the solution to the min-max problem:

$$\min_{\theta^1\in\Gamma^1} \max_{\theta^2\in\Gamma^2} J_g^{\kappa^2\mu}(\theta^1, \theta^2).$$

Therefore, since γ^1 is the minimizing policy under $\kappa^2\mu$, by perturbing the minimizer's policy γ^1 to be the policy $\eta^1 \in \Gamma^1$ we have the following inequality (i.e. we make the minimizer no longer play her optimal policy):

$$J_g^{\kappa^2\mu}(\gamma^1, \gamma^2) \leq J_g^{\kappa^2\mu}(\eta^1, \gamma^2).$$

We now wish to compare the two quantities $J_g^{\kappa^2\mu}(\eta^1, \gamma^2)$ and $J_g^{\mu}(\eta^1, \eta^2)$. To do so, fix η^1 across both terms and consider a cost function $\tilde{c}(x, \theta^2(y^2)) : \mathbb{X} \times \mathbb{U}^2 \to \mathbb{R}$ such that $\tilde{c}(x, \theta^2(y^2)) = c(x, \eta^1(y^1)), \theta^2(y^2)) \,\forall \theta^2 \in \Gamma^2$. I.e., by holding the minimizer's strategy constant as η^1, we reduce c to $\tilde{c}$ such that we now have a cost function that only reflects dependence on the maximizer's policy when the minimizer's policy is held at η^1. Such a function $\tilde{c}$ clearly exists, as the value of $\eta^1(y^1)$ is only dependent on x (potentially in some stochastic way, in that it depends on $Q^1(y|x)$), when η^1 (and μ^1) are constant, and so can be absorbed into the dependency of $\tilde{c}$ on x.

We can now compare the single-DM decision problem for the maximizer given by cost function $\tilde{c}$ and information structures $(\kappa^2\mu)^2$ and μ^2 (which we use to denote the maximizer's private information structures present in $\kappa^2\mu$ and μ, respectively, i.e. the marginals on $(\mathbb{X} \times \mathbb{Y}^2)$). This is a single-DM decision problem and as such can be treated using the forward direction to Blackwell's ordering of information structures [10, Theorem 2], which holds in this infinite-dimensional case [65]. Since $\tilde{c}$ and c are equal over all strategies in Γ^2, we know that γ^2 and η^2 are still optimal policies for the maximizer to play under the respective information structures for this game. Thus, since $(\kappa^2\mu)^2$ is a garbling of μ^2 by channel κ^2, and since $\tilde{c}(x, \eta^2(y^2)) = c(x, \eta^1(\mu^1(x)), \eta^2(y^2))$, we have that:

$$J_g^{\kappa^2\mu}(\eta^1, \gamma^2) = \int_{\mathbb{X}\times\mathbb{Y}^2} \tilde{c}(x, \gamma^2(y^2))\kappa^2\mu(dx, dy^2)$$

$$\leq \int_{\mathbb{X}\times\mathbb{Y}^2} \tilde{c}(x, \eta^2(y^2))\mu(dx, dy^2) = J_g^{\mu}(\eta^1, \eta^2).$$

Putting this all together, we have $J_g^{\kappa^2\mu}(\gamma^1, \gamma^2) \leq J_g^{\kappa^2\mu}(\eta^1, \gamma^2) \leq J_g^{\mu}(\eta^1, \eta^2)$. Since this is true for any arbitrary game $g \in \tilde{\mathbb{G}}_\zeta(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$, we have that $\kappa^2\mu \lesssim \mu$.

A nearly identical argument can be applied to show that $\mu \lesssim \kappa^1\mu$. □

Using a similar reasoning, we also develop the following converse result:

Lemma 3.2. *Take fixed $\mathbb{X}, \zeta$, fixed and compact $\mathbb{Y}^1, \mathbb{Y}^2$, and information structures ν and μ. If $\nu \lesssim \mu$ over all games in $\tilde{\mathbb{G}}_\zeta(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$, then there exist kernels $\kappa^i \in K^i$ such that:*

$$\kappa^1\nu = \kappa^2\mu.$$

In particular, under Assumption 2.2, we have the more explicit characterization with

$$\kappa^1 Q_\nu^1 = Q_\mu^1 \quad \text{and} \quad Q_\nu^2 = \kappa^2 Q_\mu^2.$$

Where Q_μ^i and Q_ν^i are the measurement channels for DM i under information structures μ and ν, respectively.

Proof. Let (γ^1, γ^2) be the equilibrium solution under ν and let (η^1, η^2) be the equilibrium solution under μ. As in Lemma 3.1, these equilibria exist and are the solutions of the standard min-max problem.

Therefore, we have the following inequality:

$$J_g^{\nu}(\gamma^1, \eta^2) \leq \max_{\theta^2 \in \Gamma^2} J_g^{\nu}(\gamma^1, \theta^2) = J_g^{\nu}(\gamma^1, \gamma^2).$$

Likewise, we can determine the following:

$$J_g^{\mu}(\eta^1, \eta^2) = \min_{\alpha^1 \in \Gamma^1} J_g^{\nu}(\alpha^1, \eta^2) \leq J_g^{\mu}(\gamma^1, \eta^2).$$

In addition, by assumption that $\nu \lesssim \mu$, we have that for all $g \in \tilde{\mathbb{G}}_{\zeta}(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$:

$$J_g^{\nu}(\gamma^1, \gamma^2) \leq J_g^{\mu}(\eta^1, \eta^2).$$

Putting this all together, one observes that $J_g^{\nu}(\gamma^1, \eta^2) \leq J_g^{\mu}(\eta^1, \eta^2)$. In the same manner as in Lemma 3.1, we hold η^2 constant across both terms and develop a reduced single-DM cost function $\tilde{c}$ for DM 1, with the action space remaining $\mathbb{U}^1$; we denote this new single-DM decision problem by $\tilde{g}$. Once again, we use ν^i and μ^i to denote the private (i.e. marginal) information structure for DM i under ν and μ, respectively. We then have a single-DM decision problem where we observe that γ^1 and η^1 are still the optimal policies for the minimizer for each respective information structure:

$$\begin{aligned} J^*(\tilde{g}, \nu^1) = J(\tilde{g}, \nu^1, \gamma^1) &= \int_{\mathbb{X} \times \mathbb{Y}^1} \tilde{c}(x, \gamma^1(y^1))\nu(dx, dy^1) \\ &\leq \int_{\mathbb{X} \times \mathbb{Y}^1} \tilde{c}(x, \eta^1(y^1))\mu(dx, dy^1) = J(\tilde{g}, \mu^1, \eta^1) = J^*(\tilde{g}, \mu^1). \end{aligned} \tag{3.2}$$

Since the inequality $J_g^{\nu}(\gamma^1, \eta^2) \leq J_g^{\mu}(\eta^1, \eta^2)$ holds true for every arbitrary zero-sum game $g \in \tilde{\mathbb{G}}_{\zeta}(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$, it holds for every game in the subclass $\hat{\mathbb{G}}$, defined here to be all games in $\tilde{\mathbb{G}}_{\zeta}(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$ where the action space of the maximizer is fixed as $\mathbb{U}^2 = \{0\}$. Moreover, we observe that for any action space $\mathbb{U}^1$ and any arbitrary bounded single-DM cost function that is continuous in the DM's action $\bar{c}(x, u^1) : \mathbb{X} \times \mathbb{U}^1 \to \mathbb{R}$, there exists a two-DM cost function $\hat{c}(x, u^1, u^2)$ corresponding to some game in $\hat{\mathbb{G}}$ with action space $\mathbb{U}^1$ such that $\hat{c}(x, u^1, u^2) = \bar{c}(x, u^1)\ \forall\, u^1 \in \mathbb{U}^1$ (following naturally from the fact that $u^2 = 0$ for these games). One such construction of $\hat{c}$ would be $\hat{c}(x, u^1, u^2) = \bar{c}(x, u^1) + u^2$; when played in $\hat{\mathbb{G}}$, $\hat{c}(x, u^1, u^2) = \bar{c}(x, u^1)\ \forall\, u^1 \in \mathbb{U}^1$. We note that since $\bar{c}(x, u^1)$ is continuous in u^1 for all x and is bounded, it is a valid for a game in $\hat{\mathbb{G}} \subset \tilde{\mathbb{G}}_\varsigma(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$ to use $\hat{c}(x, u^1, u^2)$.

Therefore, any single-DM cost function $\bar{c}$ is a valid single-DM reduction of some zero-sum cost function $\hat{c}$. We can thus observe that, when the maximizer's policy is held constant, the reduction of all games in $\tilde{\mathbb{G}}_\varsigma(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$ to single-DM problems for the minimizer is surjective on the entire space of single-DM problems with compact action spaces and bounded cost functions that are continuous in the DM's action. Therefore, applying the reasoning for (3.2), it follows that $J^*(\tilde{g}, \nu^1) \leq J^*(\tilde{g}, \mu^1)$ for all single-DM problems $\tilde{g}$ with compact action spaces and cost functions that are bounded and continuous in the action.

Lastly, we observe that since c is framed such that a higher quantity is better for the maximizer, the minimizer wants to minimize the value of $\tilde{c}$. Therefore, from the minimizer's perspective, the inequality $J^*(\tilde{g}, \nu^1) \leq J^*(\tilde{g}, \mu^1)$ indicates that she can never perform worse under ν than under μ. Thus, by the converse direction to Blackwell's ordering of information structures [10, Theorem 6], which by the lemma assumptions and the restrictions on the class $\tilde{\mathbb{G}}_\varsigma(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$ (namely compactness of $\mathbb{Y}^i$ and $\mathbb{U}^i$) holds in this infinite setup due to Theorem 3.3, we have that μ^1 must be a garbling of ν^1.

In a similar manner, by observing that $J_g^\nu(\gamma^1, \gamma^2) \leq J_g^\mu(\gamma^1, \eta^2)$ one discovers that ν^2 must be a garbling of μ^2.

Therefore, we have that the minimizer's information from ν is garbled in μ and the maximizer's information from μ is garbled in ν. Combining these two conditions yields the desired equality for some $\kappa^i \in K^i$:

$$\kappa^1\nu = \kappa^2\mu.$$

□

The following is our main result.

Theorem 3.5. *Take fixed* $\mathbb{X}, \zeta$, *fixed and compact* $\mathbb{Y}^1, \mathbb{Y}^2$, *and information structures* ν *and* μ. *Then* μ *is better for the maximizer than* ν ($\nu \lesssim \mu$) *over all games in* $\tilde{\mathbb{G}}_\zeta(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$ *if and only if there exist kernels* $\kappa^i \in K^i$ *such that:*

$$\kappa^1\nu = \kappa^2\mu.$$

Proof. The *if* direction follows directly from Lemma 3.1:

$$\nu \lesssim \kappa^1\nu = \kappa^2\mu \lesssim \mu$$

The *only if* direction is given in Lemma 3.2. □

Corollary 3.1. *Take fixed* $\mathbb{X}, \zeta$, *and fixed and compact* $\mathbb{Y}^1, \mathbb{Y}^2$. *The value of additional information to a decision-maker is never negative for that decision-maker in any zero-sum game in* $\tilde{\mathbb{G}}_\zeta(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$.

Proof. If μ is an information structure which has more information for the maximizer than another information structure ν, then there exists a kernel κ^2 such that $\nu = \kappa^2\mu$, and so we know that $J^*(g,\mu) \geq J^*(g,\nu)$, and κ^2 is a well-defined map since we can map any additional information to a fixed number. $\square$

We note that the value of information is not always strictly positive to a DM, since many situations (such as where the action set is a singleton) will result in no change in performance despite additional information.

Corollary 3.1 is consistent with the work of De Meyer, Lehrer and Rosenberg [19, Theorem 3.1], who found this result when studying the value of information in zero-sum games with incomplete information with a slightly different setup, where the 'state of nature' was replaced by an individual 'type' for each DM drawn from a finite space, and where the cost function depended on both DMs' types.

We note that for Theorem 3.5, the proof will follow for any class of zero-sum games for which every game has an equilibrium solution and Blackwell's ordering of information structures holds for each DM when holding the other DM's policy constant. Therefore, the ordering result can be generalized to be applicable for more general classes of zero-sum games than $\tilde{\mathbb{G}}_\zeta(\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2)$.

3.4.1 Discussion

This main result comes with the following intuitive interpretation:

An information structure μ is better for the maximizer than ν if and only if one of the following holds:

1. ν^2 is a non-identity garbling of the maximizer's channel from μ, and the minimizer's channel is identical.

2. μ^1 is a non-identity garbling of the minimizer's channel from ν, and the maximizer's channel is identical.

3. ν^2 is a non-identity garbling of the maximizer's channel from μ, and μ^1 is a non-identity garbling of the minimizer's channel from ν.

4. The information structures are identical.

In plain terms, this has the following interpretation: In zero-sum games, improving or hurting both DMs' information structures will never give a general benefit to either DM over all games. The only time a DM will not do worse under a new information structure is if it only makes his channel better, only makes his opponent's channel worse, makes his channel better and his opponent's channel worse, or is identical to the previous information structure (and the DM is guaranteed to not do worse if any of these conditions holds).

In the following, we present an example showing that we cannot *view garbling from decision-maker to decision-maker in isolation from the entire information structure*. Consider a finite probability space game with $\mathbb{X} = \mathbb{U} = \mathbb{Y}^1 = \mathbb{Y}^2 = \{1, 2, 3, 4\}$, with x distributed according to the uniform distribution, and cost function:

$$c(x, u^1, u^2) = \begin{cases} -12, \quad u^1 = x \quad \text{and} \quad u^1 \neq u^2 \\ -5, \quad u^1 = x \quad \text{and} \quad u^1 = u^2 \\ 0, \quad \text{otherwise} \end{cases}$$

DM 1 (the minimizer) gets rewarded for guessing x correctly, and DM 2 (the maximizer) can only limit his losses by playing the same action as DM 1. We can observe that DM 1's optimal strategy will always be to attempt to guess x correctly, since she is only penalized for guessing incorrectly, while DM 2's optimal strategy will always be to attempt to copy DM 1's action since that is the only way he can positively affect the outcome for himself.

Now consider the following two information structures:

μ_1: Under this information structure, both DMs receive *the same random measurement* $y^1 = y^2 = y$, where $y = x$ with probability 0.9 and y is any of the other three incorrect values with probability 0.1/3. Under this information structure, the best strategy for DM 1 (and thus also for DM 2) is to guess her observation, so $u^1 = u^2 = y$ and the expected payoff is $-5(0.9) = -4.5$.

μ_2: Under this information structure, both DMs receive *conditionally (given x) independent measurements*. For DM 1, $y^1 = x$ with probability 0.85, and is any of the three incorrect values of x with probability 0.05 each. DM 2 has the same structure as under μ_1, with a 0.9 chance of success, albeit now uncoupled with DM 1's chance of success. The optimal strategies remain the same under this information structure, but the expected payoff is now $-5(0.9)(0.85) + (-12)(0.85)(0.1) = -4.845$.

Therefore, μ_2 is better for the minimizer than μ_1. But, we can observe that the minimizer's channel in μ_2 is garbled from μ_1, in the sense that the distribution on $\mathbb{Y}^1$ for DM 1 can be run through a stochastic kernel to get the distribution under μ_2. The maximizer's channel is identical in both games in the sense that the distribution on $\mathbb{Y}^2$ is unchanged. Yet, the ordering of information structures rule from Theorem 3.2 appears to have been violated, since the minimizer performs better under the garbled information structure. This demonstrates that we cannot consider garbling in isolation and the comparison should be in view of the entire information structure.

While μ_2 appears to be a garbling of μ_1 for the minimizer, it is not a garbling in the sense of this chapter, since μ_1 features dependent measurements between the DMs, while μ_2 has independent measurements between the DMs. Definition 3.4

specifies that garblings are done in view of the entire information structure, and so a garbling could not decouple dependence when going from μ_1 to μ_2. If the garbling had been done in accordance with the results of this chapter so that $y^1 = y^2$ but y^1 is then garbled to arrive at some $\tilde{y}^1$ whose probability measure is as specified under μ_2, then the DMs' measurements would still contain dependence after the garbling. Under this construction, naturally DM 1 would perform worse in the equilibrium under the garbled information structure, since DM 2 has maintained a good ability to copy DM 1's actions when DM 1 is correct due to the dependence being maintained, while DM 1 has received a disadvantage in being able to accurately guess x. If the stochastic kernel garbling DM 1's information is as given by $\tilde{\kappa}$ below, where the (i, j) entry is the probability of DM 1 measuring $\tilde{y}^1 = i$ given that the DMs originally measured $y^1 = y^2 = j$, then the expected equilibrium payoff in this situation would be $-5(0.9)(0.9423) + (-12)(0.1)(0.0192) = -4.263$, which is worse for DM 1, as expected. Under this garbling, DM 1 has a probability of 0.85 of observing the correct measurement $\tilde{y}^1 = x$ and a 0.05 probability of observing any of the three incorrect measurements, matching the distribution specified in the definition of μ_2.

$$\tilde{\kappa} = \begin{bmatrix} 0.9423 & 0.0192 & 0.0192 & 0.0192 \\ 0.0192 & 0.9423 & 0.0192 & 0.0192 \\ 0.0192 & 0.0192 & 0.9423 & 0.0192 \\ 0.0192 & 0.0192 & 0.0192 & 0.9423 \end{bmatrix}$$

Chapter 4

Continuity of Equilibrium Value in Information

4.1 Introduction

A related problem to that of comparison of information structures, which was studied in Chapter 3, involves continuity properties of optimal solutions/equilibrium solutions in information structures under various topologies, which will be studied in this chapter.

4.2 Preliminaries and Literature Review

4.2.1 Information Structure Models and Assumptions

We follow the setup for N DMs playing in single-stage stochastic game, as in Section 1.2.

Throughout this chapter, we allow for the entire information structure μ to vary as it converges, meaning the prior on the state space ζ is not fixed. The sole exceptions are Theorem 4.4 and Theorem 4.5, where the prior must be fixed for the proofs to hold.

We now define the following assumptions, which will be used within the results of this chapter:

Assumptions.

A1: The cost functions for DMs are measurable and bounded.

A2: The cost functions for DMs are continuous and bounded.

A3: The information structure is absolutely continuous with respect to a product measure:

$$P(dy^1, dy^2, \cdots, dy^n, dx) \ll \bar{Q}^1(dy^1)\bar{Q}^2(dy^2)\cdots\bar{Q}^n(dy^n)\zeta(dx),$$

for reference probability measures $\bar{Q}^i$, $i = 1, 2, \cdots$. That is, e.g., for $n = 2$, there exists an integrable f which satisfies for every Borel A, B, C

$$P(y^1 \in B, y^2 \in C, x \in A) = \int_{A,B,C} f(x, y^1, y^2)\zeta(dx)\bar{Q}^1(dy^1)\bar{Q}^2(dy^2).$$

A4: The action space of each agent is compact.

A5: The bounded measurable cost function is continuous in DMs' actions for every state of nature x.

A6: The individual DM channels Q^i are fixed and independent (i.e., given x, all the measurement variables y^i are conditionally independent). Furthermore, each Q^i is continuous in total variation, i.e., as $x_m \to x$ then $\|Q^i(\cdot|x) - Q^i(\cdot|x_m)\|_{TV} \to 0$.

A7: The action spaces for the DMs are convex subsets of $\mathbb{R}^k$, for some k.

Note that A3 is simply Assumption 2.1 extended to allow for N DMs.

4.2.2 Convergence of information structures

For a standard Borel space $\mathbb{X}$, recall that we let $\mathcal{P}(\mathbb{X})$ denote the family of all probability measures on $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$, where $\mathcal{B}(\mathbb{X})$ denotes the Borel sigma-field over $\mathbb{X}$. Let $\{\mu_n,\, n \in \mathbb{N}\}$ be a sequence in $\mathcal{P}(\mathbb{X})$. Recall that $\{\mu_n\}$ is said to converge to $\mu \in \mathcal{P}(\mathbb{X})$ *weakly* if

$$\int_{\mathbb{X}} c(x)\mu_n(dx) \to \int_{\mathbb{X}} c(x)\mu(dx) \tag{4.1}$$

for every continuous and bounded $c : \mathbb{X} \to \mathbb{R}$. On the other hand, $\{\mu_n\}$ is said to converge to $\mu \in \mathcal{P}(\mathbb{X})$ *setwise* if (4.1) holds for every measurable and bounded $c : \mathbb{X} \to \mathbb{R}$. Setwise convergence can also be defined through pointwise convergence on Borel subsets of $\mathbb{X}$, that is $\mu_n(A) \to \mu(A)$, for all $A \in \mathcal{B}(\mathbb{X})$. For two probability measures $\mu, \nu \in \mathcal{P}(\mathbb{X})$, the *total variation* metric is given by

$$\begin{aligned} \|\mu - \nu\|_{TV} &:= 2 \sup_{B \in \mathcal{B}(\mathbb{X})} |\mu(B) - \nu(B)| \\ &= \sup_{f:\, \|f\|_\infty \leq 1} \left| \int f(x)\mu(dx) - \int f(x)\nu(dx) \right|. \end{aligned} \tag{4.2}$$

A sequence $\{\mu_n\}$ is said to converge to $\mu \in \mathcal{P}(\mathbb{X})$ in total variation if $\|\mu_n - \mu\|_{TV} \to 0$. Setwise convergence is equivalent to pointwise convergence on Borel sets whereas total variation requires uniform convergence on Borel sets. Thus these three convergence notions are in increasing order of strength.

We also recall the *w*-*s* topology introduced by Schäl [54], and used earlier in Chapter 2: the *w*-*s* topology on the set of probability measures $\mathcal{P}(\mathbb{X} \times \mathbb{Y})$ is the coarsest topology under which, for every measurable and bounded $f(x, y)$ which is continuous in y for every x, the map $\int f(x,y)\mu(dx,dy) : \mathcal{P}(\mathbb{X}\times\mathbb{Y}) \to \mathbb{R}$ is continuous (but unlike the weak convergence topology, f does not need to be continuous in x). An important relevant result [54, Theorem 3.10] or [7, Theorem 2.5] is that if $\mu_n \to \mu$ in weakly but if the marginals $\mu_n(dx\times\mathbb{Y}) \to \mu(dx\times\mathbb{Y})$ setwise, then the convergence is also in the w-s sense.

Definition 4.1. *A sequence of information structures μ_n in $\mathcal{P}(\mathbb{X}\times\mathbb{Y}^1\times\cdots\times\mathbb{Y}^n)$ converges to μ weakly/setwise/in total variation if $\mu_n \to \mu$ in the corresponding sense.*

4.2.3 Literature review

Single decision-maker/Player Problems

For a single decision-maker setup, [66, Theorem 3.4], for a fixed prior, showed that the value function is continuous under total variation convergence of information structures. Counterexamples also revealed that continuity is not necessarily preserved under weak or setwise convergence, but the value function is upper semicontinuous under both under mild conditions.

The case where priors also change has been studied in [36], where conditions for optimality under both total variation and weak convergence were established: total variation convergence of the priors always leads to convergence [36, Theorem 2.9], but continuity under the weak convergence of priors requires a total variation continuity condition on the information channels [36, Theorem 2.5].

A related result due to Le Cam [14] will be discussed further below in the context of zero-sum games. We also note that a related result due to Wu and Verdú [63] establishes continuity of the quadratic (minimum mean-square estimation) error under weak convergence of the priors when the channel is additive, the additive noise has a finite variance and it admits a continuous and bounded density function (such as a Gaussian). [1], [29] study the effects of uncertainties in the input distribution in the quantizer design, which may be viewed as a decision problem. [45, 59] study the effects of variations in system models, and thus also information structures, in a general relative entropy perturbation framework in the context of minimax LQG control; we note that relative entropy is a more stringent distance measure via Pinsker's inequality [18].

Zero-Sum Games

The specific form of an information structure has been shown to have subtle impact on (different types of) equilibria in games, as well as on their existence, uniqueness, and characterization (see for example [61, 5, 8]). In spite of the existence of several studies on the impact of information structure on equilibria of games, there does not exist a complete theory of such impact. The theory of and regularity properties on information structures in stochastic games is significantly more challenging when compared with that for stochastic team theory, since the value of information is not necessarily positive and informational changes lead to variations in the equilibrium behaviour in intricate ways [9, 33, 3]).

Zero-sum games, however, form a class of problems where strong regularity properties can be established. For example, under mild conditions as reviewed in Chapter 2, every zero-sum game, and accordingly every information structure, has a value. Furthermore, as discussed earlier in Chapter 3, a theory of ordering of information structures can be obtained and, in particular, no additional information can hurt a DM in a zero-sum game.

In view of the above, we can interpret zero-sum games as an intermediate category between team problems and non-zero sum game problems as far as their informational properties are concerned. The findings of this chapter will further contribute towards such an interpretation.

Among contributions, [22] shows uniform continuity of the value function for zero-sum games under convergence of DM information fields using a notion of convergence of sigma fields developed by Boylan in [13], under Lipschitz continuity assumptions on the cost function; while [23] establishes continuity of the value function in the prior distribution on the state space for zero-sum games using the total variation metric. In [26], for countable $\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2$, a *value-distance* is introduced, which is a method of comparison of information structures that can be characterized using the total variation distance in zero-sum games.

Definition 4.2 ([26])**.** *Take countable* $\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2$*. Let* $\tilde{\mathbb{G}}$ *be all games with cost functions satisfying A1 (and where the cost is bounded by* 1*) and countable action spaces* $\mathbb{U}^1$ *and* $\mathbb{U}^2$*. The value-distance between two information structures for zero-sum games is:* $d_2(\mu, \nu) =: \sup_{g \in \tilde{\mathbb{G}}} |J^*(g, \mu) - J^*(g, \nu)|$.

We recall the definition and notation of garbling introduced in Definition 3.4.

Theorem 4.1 (Theorem 1, [26])**.** *Take countable* $\mathbb{X}, \mathbb{Y}^1, \mathbb{Y}^2$*. The following equality holds:*

$$d_2(\mu, \nu) = \max\{ \min_{\kappa^1 \in K^1, \kappa^2 \in K^2} \|\kappa^1 \mu - \nu \kappa^2\|_{TV}, \min_{\kappa^1 \in K^1, \kappa^2 \in K^2} \|\mu \kappa^2 - \kappa^1 \nu\|_{TV} \}.$$

Via this distance, it follows that convergence of $\{\mu_n\}$ to μ in total variation results in convergence of $J^*(g, \mu_n)$ to $J^*(g, \mu)$ when the state, measurement and action spaces are countable. In this chapter we will consider general standard Borel spaces and demonstrate that if $\{\mu_n\}$ converges to μ in total variation, then $J^*(g, \mu_n)$ converges to $J^*(g, \mu)$ for all games g for which saddle-point equilbria exist under $\{\mu_n\}$ and μ. Furthermore, we will establish upper semi-continuity properties that $J^*(g, \cdot)$ exhibits under weak and setwise convergence of measurement channels for each DM.

We also note that Theorem 4.1 is closely related in the single-DM case to a result from Le Cam in [14]. We recall the definition of the Le Cam distance for a single-DM decision problem with information structures μ and ν.

Definition 4.3. *The Le Cam distance between two information structures μ and ν is*

$$\begin{aligned}\Delta(\mu,\nu) &:= \max\{\delta(\mu,\nu),\delta(\nu,\mu)\} \\ &:= \max\{\inf_{\kappa\in K}\sup_{x\in\mathbb{X}}\|\kappa\mu(\cdot|x)-\nu(\cdot|x)\|_{TV}, \inf_{\kappa\in K}\sup_{x\in\mathbb{X}}\|\mu(\cdot|x)-\kappa\nu(\cdot|x)\|_{TV}\}.\end{aligned}$$

$\delta(\mu,\nu)$ is referred to as the *Le Cam deficiency of μ with respect to ν*. We now recall the following theorem adapted from Le Cam [14], which holds in the standard Borel setup.

Theorem 4.2 (Theorem 3, [14])**.** *For a given $\epsilon > 0$, if $\delta(\mu,\nu) \le \epsilon$ then for any policy γ_ν under ν and any single-DM game with a bounded cost function $\|c\|_\infty \le 1$, there exists a policy $\bar{\gamma}_\mu$ under μ such that*

$$E^{\mu}_{\zeta}[c(x,\bar{\gamma}_\mu(y))] \le E^{\nu}_{\zeta}[c(x,\gamma_\nu(y))] + \epsilon.$$

Remark 4.1. *We observe that the "value-distance" in Definition 4.2 is similar to the Le Cam distance, with the difference being that the Le Cam distance looks at maximizing the total variation distance over individual states, rather than incorporating* *the prior. Theorem 4.2 is also closely related to Theorem 4.1 in the single-DM case. Theorem 4.2 implies that if $\delta(\mu, \nu) = \epsilon$ then the maximum possible decrease in value over all valid games when changing from μ to ν is upper-bounded by ϵ. Similarly, if $\delta(\nu, \mu) = \alpha$ then the maximum possible increase in value going from μ to ν is upper-bounded by α. If we define d_1 similar to Definition 4.2, as the supremum of the absolute difference in value over all valid single-DM games, as a natural restriction of d_2 to the single-DM case, we would have that $d_1(\mu, \nu) \leq \max\{\delta(\mu, \mu), \delta(\nu, \mu)\}$, and so $d_1(\mu, \nu) \leq \Delta(\mu, \nu)$. Thus, the Le Cam distance gives an upper-bound on the value-distance.*

Stochastic Team Problems

We recall the team problem setup introduced in Section 1.2.2.

In team problems, equilibrium solutions are team policies $\bar{\gamma}^* := (\gamma^{1,*}, \ldots, \gamma^{n,*})$ which minimize the value function, which is the expected value of the common cost function. The value function at equilibrium is denoted by $J^*(c, \mu)$. General existence results for team-optimal policies can be found in [64, Section 5]; in particular, an equilibrium solution exists under assumptions A1, A3, A4, and A5.

General Stochastic Games

The value of information may not be well-posed for general games, since unlike team and zero-sum game problems, one cannot talk about the unique value of a general non-zero-sum game. Each DM has a personal value function in a general game. Furthermore, more information may have a positive or negative value to a DM who receives it, unlike in zero-sum games and team problems. An example of such an occurrence was discussed in Example 3.1. The electronic mail game from [50] is a further relevant classic example where equilibria do not converge (although the information does not converge in total variation in this example). We note that a somewhat different notion of ϵ-equilibria (a uniform-ϵ equilibrium concept where ϵ-proximity, with ϵ being uniform for each conditional expected cost/reward given the private information realization, regardless of how unlikely the realizations are) has been studied in [32] and [49] for finite games, where it has been shown that total variation continuity does not hold under this concept.

4.3 Continuity Results for Zero-Sum Games

We have the following theorem.

Theorem 4.3. *(i) For a fixed zero-sum game which satisfies Assumptions A1, A4 and A5, the value function is continuous in information structures under total variation convergence of information structures.*

(ii) The value function is not necessarily continuous under weak or setwise convergence of information structures even under Assumptions A2 and A4 and with a fixed prior.

Proof. (i) By Theorem 2.2, A1, A4, and A5 guarantee that an equilibrium exists under μ and ν. Assume that μ and ν are such that $\|\mu - \nu\|_{TV} \leq \epsilon$ for some $\epsilon > 0$. Without loss of generality, assume that $J^*(g, \mu) - J^*(g, \nu) \geq 0$ for some game g (a symmetric argument can be applied in the case where $J^*(g, \mu) - J^*(g, \nu) \leq 0$). Then we have the following:

$$J^*(g, \mu) - J^*(g, \nu)$$

$$= \int c(x,\gamma_\mu^{1,*}(y^1),\gamma_\mu^{2,*}(y^2))\mu(dx,dy^1,dy^2) - \int c(x,\gamma_\nu^{1,*}(y^1),\gamma_\nu^{2,*}(y^2))\nu(dx,dy^1,dy^2)$$
$$\leq \int c(x,\gamma_\nu^{1,*}(y^1),\gamma_\mu^{2,*}(y^2))\mu(dx,dy^1,dy^2) - \int c(x,\gamma_\nu^{1,*}(y^1),\gamma_\mu^{2,*}(y^2))\nu(dx,dy^1,dy^2).$$

The inequality comes from perturbing the minimizer's equilibrium strategy in the first term, making the expected cost larger, and perturbing the maximizer's equilibrium strategy in the second term, making the expected cost smaller. Then, viewing $\gamma_\nu^{1,*}$ and $\gamma_\mu^{2,*}$ as fixed, we can view $c(\cdot,\gamma_\nu^{1,*}(\cdot),\gamma_\mu^{2,*}(\cdot))$ as a measurable and bounded function from $\mathbb{X}\times\mathbb{Y}^1\times\mathbb{Y}^2\to\mathbb{R}$. Let $M\in\mathbb{R}_{\geq 0}$ be such that $\|c\|_\infty\leq M$. We have:

$$\int c(x,\gamma_\nu^{1,*}(y^1),\gamma_\mu^{2,*}(y^2))\mu(dx,dy^1,dy^2) - \int c(x,\gamma_\nu^{1,*}(y^1),\gamma_\mu^{2,*}(y^2))\nu(dx,dy^1,dy^2)$$
$$\leq M\|\mu-\nu\|_{TV}\leq M\epsilon.$$

Thus, if $\|\mu_n-\mu\|_{TV}\to 0$, then $|J^*(g,\mu)-J^*(g,\nu)|\to 0$.

(ii) For stochastic control problems (which can be viewed as single-DM games), [66, Sections 3.1.1. and 3.2.1] establish that value functions are not necessarily continuous under weak or setwise convergence. These counterexamples directly extend to the zero-sum game case: one can extend such a single-DM game to a two-DM zero-sum game where the action of one DM has no impact on the cost: e.g. $|\mathbb{U}^2|=1$ so that DM 2 has no freedom to select a control action and the information structure of DM 1 changes. □

We now study weak and setwise convergences of information structures. First, we make the following definitions.

Definition 4.4. *For a fixed prior ζ, a sequence of information structures $\{\mu\}_{k=1}^\infty$ is a maximizer-garbling (or a minimizer-degarbling) sequence if for every $j\in\{1,\dots,\infty\}$*

there exist kernels κ_j^1 *and* κ_j^2 *such that* $\kappa_j^1\mu_{j+1} = \kappa_j^2\mu_j$.

Definition 4.5. *For a fixed prior* ζ*, a sequence of information structures* $\{\mu\}_{k=1}^{\infty}$ *is a minimizer-garbling (or a maximizer-degarbling) sequence if for every* $j \in \{1, \ldots, \infty\}$ *there exist kernels* κ_j^1 *and* κ_j^2 *such that* $\kappa_j^2\mu_{j+1} = \kappa_j^1\mu_j$.

We note that, under appropriate assumptions, following the results of Lemma 3.1, the value function for any zero-sum game will be monotonically decreasing for a maximizer-garbling sequence and monotonically increasing for a minimizer-garbling sequence. For the purpose of illustration, we present examples of maximizer-garbling sequences.

Example 4.1 (Maximizer-Garbling Sequences)**.** *Consider two DMs in a zero-sum game with state space* $\mathbb{X} = [0,1]$ *endowed with the uniform distribution as* ζ*. Let* $\mathbb{Y}^1 = \mathbb{Y}^2 = \mathbb{R}$.

(i) For $m \in \mathbb{Z}_{\geq 0}$*, define* μ_m *by the following independent measurement channels for the DMs: DM 1 receives measurement* y^1*, which is drawn randomly from a Gaussian distribution with mean* x *and variance* $1 + \frac{1}{m}$*; DM 2 receives measurement* y^2 *which is drawn randomly from a Gaussian distribution with mean* x *and variance* $1 - \frac{1}{m+1}$*. Since Gaussian distributions with higher variances are garblings of Gaussian distributions with lower variances, this sequence is a maximizer-garbling sequence, since for any* $i \in \mathbb{Z}_{\geq 0}$*, the maximizer's channel under* μ_{i+1} *is a garbling of his channel under* μ_i*, while the minimizer's channel under* μ_i *is a garbling of her channel under* μ_{i+1}*. Furthermore, this sequence converges weakly to the information structure in which both DMs' channels have variance* 1*. One can also construct such an example where we have* $y^1 = x + v_m^1$,

and if the variance of v_m^1 decreases to zero, we will have weak convergence to a point distribution.

(ii) *Another example involves noiseless but quantized measurement channels. Let $y^2 = 1$ be a constant measurement for DM 2, and $y^1 = Q_m^1(x)$ for DM 1, where Q_m^1 is a uniform quantization of $[0,1]$ into 2^m bins. In this case, the sequence of information structures is a maximizer-garbling sequence (as the channels for the minimizer are successive refinements as m increases, and the maximizer's information is constant) and the weak limit consists of the fully informative channel for DM 2: $Q^1(dy^1|x) = \delta_x(dy^1)$.*

We emphasize that not all information structures can be related to each other as either minimizer or maximizer-garbling sequences.

Theorem 4.4. *Let Assumptions A1, A4, and A7 hold and the prior be fixed. Let μ_m be a sequence of information structures converging weakly to information structure μ. If the sequence is a maximizer-garbling, then the value function is lower semi-continuous. If the sequence is a minimizer-garbling, then the value function is upper semicontinuous.*

Proof. Consider a minimizer-garbling sequence. By A2, our cost function is bounded; let $M \in \mathbb{R}_{\geq 0}$ such that $\|c\|_\infty = M$. Let $\mu_{\mathbb{Y}^j}$ be the marginal of μ on its $(j+1)$th component. Let γ^j be an arbitrary policy for DM j. We note that every Polish space is second countable (i.e. has a countable basis), since it is separable and metrizable. Thus, every standard Borel space is second countable. Then, by Lusin's theorem, using the fact that $\mathbb{U}^j$ is convex by assumption, we have that for any $\epsilon > 0$, there exists a continuous function $f^j : \mathbb{Y}^j \to \mathbb{U}^j$ such that $\mu_{\mathbb{Y}^j}(\{y^j : f^j(y^j) \neq \gamma^j(y^j)\}) < \epsilon$.

Letting $B^2 := \{y^2 : f^2(y^2) \neq \gamma^2(y^2)\}$, we have:

$$\begin{aligned}
&\int_{\mathbb{X}\times\mathbb{Y}^1\times\mathbb{Y}^2} c(x,\gamma^1(y^1),\gamma^2(y^2))\mu(dx,dy^1,dy^2) - M\epsilon \\
&< \int_{\mathbb{X}\times\mathbb{Y}^1\times\mathbb{Y}^2} c(x,\gamma^1(y^1),f^2(y^2))\mu(dx,dy^1,dy^2) \\
&< \int_{\mathbb{X}\times\mathbb{Y}^1\times\mathbb{Y}^2} c(x,\gamma^1(y^1),\gamma^2(y^2))\mu(dx,dy^1,dy^2) + M\epsilon \qquad (4.3)
\end{aligned}$$

This holds since the cost function is bounded by $[-M, M]$ and so the expected cost on $\mathbb{X} \times B^i \times \mathbb{Y}^{-i}$ is bounded by $[-M\epsilon, M\epsilon]$. Denote by Γ^j_C the policy space for DM j in which the DM's policy is continuous from $\mathbb{Y}^j$ to $\mathbb{U}^j$. We note that if DM 1's strategy is fixed, DM 2 receives the same cost if supremizing over policies in Γ^2 or Γ^2_C: that the value must be greater when Γ^2 is used is a consequence of $\Gamma^2_C \subset \Gamma^2$. That the value must be greater than or equal to that under Γ^2 when Γ^2_C is used is a consequence of (4.3), since the policies, as well as ϵ, were arbitrary. The same reasoning applies in the reverse case.

We now observe that, for any $m \in \mathbb{Z}_{>0}$:

$$\begin{aligned}
&\inf_{\gamma^1\in\Gamma^1_C}\sup_{\gamma^2\in\Gamma^2}\int c(x,\gamma^1(y^1),\gamma^2(y^2))\mu_m(dx,dy^1,dy^2) \\
&= \inf_{\gamma^1\in\Gamma^1}\sup_{\gamma^2\in\Gamma^2}\int c(x,\gamma^1(y^1),\gamma^2(y^2))\mu_m(dx,dy^1,dy^2). \qquad (4.4)
\end{aligned}$$

To prove this, let $\gamma^{1,*}$ and $\gamma^{2,*}$ be the equilibrium policies under μ_m. Then, following the discussion on Lusin's theorem, for any $\epsilon > 0$, there exists a continuous policy $\hat{\gamma}^1 \in \Gamma^1_C$, which is equal to $\gamma^{1,*}$ except on a set B^1 such that $\mu_{m,\mathbb{Y}^1}(B^1) = \epsilon$. If DM 1 plays this strategy, then for any strategy that DM 2 selects in Γ^2, the value for the game will increase by at most $2M\epsilon$. This is because, if DM 2's arbitrary new strategy $\hat{\gamma}^2$ results in a change in performance on $A := \{\mathbb{X} \times B^1 \times \mathbb{Y}^2\}$, the maximum possible difference in expected cost on this set is $2M\epsilon$. If DM 2's new strategy results in better performance on $C := \{\mathbb{X} \times (\mathbb{Y}^1 \setminus B^1) \times \mathbb{Y}^2\}$, then the difference between playing $\hat{\gamma}^2$ and $\gamma^{2,*}$ is at most $2M\epsilon$, because $\hat{\gamma}^1$ is identical to $\gamma^{1,*}$ on this set, so a larger difference would necessarily contradict the fact that $\gamma^{2,*}$ is DM 2's best response to $\gamma^{1,*}$ under μ (since DM 2 could lose at most $2M\epsilon$ on set A while employing this strategy against $\gamma^{1,*}$ instead of playing $\gamma^{2,*}$, and so by gaining more than $2M\epsilon$ on C, DM 2 would guarantee better expected performance playing $\hat{\gamma}^2$ against $\gamma^{1,*}$ over playing $\gamma^{2,*}$, violating the Nash-equilibrium condition). Thus, since DM 1 has a continuous strategy that guarantees a worst possible outcome of a $2M\epsilon$ increase in expected cost, we know that:

$$\begin{aligned}&\inf_{\gamma^1 \in \Gamma^1_C} \sup_{\gamma^2 \in \Gamma^2} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu_m(dx, dy^1, dy^2) \\ &\leq \inf_{\gamma^1 \in \Gamma^1} \sup_{\gamma^2 \in \Gamma^2} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu_m(dx, dy^1, dy^2) - 2M\epsilon\end{aligned}$$

We also know that since $\Gamma^1_C \subset \Gamma^1$, we have:

$$\begin{aligned}&\inf_{\gamma^1 \in \Gamma^1_C} \sup_{\gamma^2 \in \Gamma^2} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu_m(dx, dy^1, dy^2) \\ &\geq \inf_{\gamma^1 \in \Gamma^1} \sup_{\gamma^2 \in \Gamma^2} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu_m(dx, dy^1, dy^2).\end{aligned}$$

Since ϵ was arbitrary, (4.4) follows. Applying a similar argument again, it follows that:

$$\inf_{\gamma^1 \in \Gamma^1_C} \sup_{\gamma^2 \in \Gamma^2_C} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu_m(dx, dy^1, dy^2)$$

$$= \inf_{\gamma^1 \in \Gamma^1} \sup_{\gamma^2 \in \Gamma^2} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu_m(dx, dy^1, dy^2). \tag{4.5}$$

Now take arbitrary large $M \in \mathbb{Z}_{\geq 0}$. Then, we have:

$$\begin{aligned}
&\inf_{\gamma^1 \in \Gamma^1} \sup_{\gamma^2 \in \Gamma^2} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu_M(dx, dy^1, dy^2) \\
&= \inf_{\gamma^1 \in \Gamma^1_C} \sup_{\gamma^2 \in \Gamma^2_C} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu_M(dx, dy^1, dy^2) \\
&\leq \inf_{\gamma^1 \in \Gamma^1_C} \sup_{\gamma^2 \in \Gamma^2_C} \limsup_{M \to \infty} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu_M(dx, dy^1, dy^2) \qquad (4.6) \\
&= \inf_{\gamma^1 \in \Gamma^1_C} \sup_{\gamma^2 \in \Gamma^2_C} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu(dx, dy^1, dy^2) \\
&= \inf_{\gamma^1 \in \Gamma^1} \sup_{\gamma^2 \in \Gamma^2} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu(dx, dy^1, dy^2)
\end{aligned}$$

The first and final equalities follow from (4.5). The inequality (4.6) follows from the fact that the value function will be monotonically increasing in M: this is because the sequence of information structures is a minimizer-garbling sequence[1]. The second equality follows from the weak convergence of $\{\mu\}_m$ to μ; since the prior is fixed, weak convergence is equivalent to w-s convergence in this setting, so continuity of c in x is not required, and thus A1 and A5 are sufficient for this equality to hold. Since the above holds for all $M \in \mathbb{Z}_{\geq 0}$, it will also hold in the limit of M, and thus:

$$\begin{aligned}
&\limsup_{m \to \infty} \inf_{\gamma^1 \in \Gamma^1} \sup_{\gamma^2 \in \Gamma^2} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu_m(dx, dy^1, dy^2) \\
&\leq \inf_{\gamma^1 \in \Gamma^1} \sup_{\gamma^2 \in \Gamma^2} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu(dx, dy^1, dy^2).
\end{aligned}$$

[1]It is achieving this inequality which requires the sequence of information structures to be a minimizer-garbling sequence, rather than an arbitrary weakly converging sequence of information structures.

Thus, the value function is upper semicontinuous. A similar proof applies for maximizer-garbling sequences. □

We now consider setwise convergence of information structures; the proof follows closely that of Theorem 4.4.

Theorem 4.5. *Consider a fixed prior and let Assumptions A1, A4, and A5 hold. Let μ_m be a sequence of information structures converging setwise to information structure μ. If the sequence is a maximizer-garbling, then the value function is lower semicontinuous. If the sequence is a minimizer-garbling, then the value function is upper semicontinuous.*

Proof. Consider a maximizer-garbling sequence. We have that:

$$\begin{aligned}&\liminf_{m\to\infty} \inf_{\gamma^1\in\Gamma^1} \sup_{\gamma^2\in\Gamma^2} \int c(x,\gamma^1(y^1),\gamma^2(y^2))\mu_m(dx,dy^1,dy^2)\\ &\geq \sup_{\gamma^2\in\Gamma^2} \liminf_{m\to\infty} \inf_{\gamma^1\in\Gamma^1} \int c(x,\gamma^1(y^1),\gamma^2(y^2))\mu_m(dx,dy^1,dy^2).\end{aligned}$$

The inequality above follows from interchanging the limit inferior and the supremum. Following a similar reasoning as the proof of Theorem 4.4, we have the following for any fixed policy $\gamma^2 \in \Gamma^2$:

$$\begin{aligned}&\liminf_{m\to\infty} \inf_{\gamma^1\in\Gamma^1} \int c(x,\gamma^1(y^1),\gamma^2(y^2))\mu_m(dx,dy^1,dy^2)\\ &\geq \inf_{\gamma^1\in\Gamma^1} \int c(x,\gamma^1(y^1),\gamma^2(y^2))\mu(dx,dy^1,dy^2).\end{aligned}$$

This gives us

$$\liminf_{m\to\infty} \inf_{\gamma^1\in\Gamma^1} \sup_{\gamma^2\in\Gamma^2} \int c(x,\gamma^1(y^1),\gamma^2(y^2))\mu_m(dx,dy^1,dy^2)$$

$$\geq \sup_{\gamma^2 \in \Gamma^2} \inf_{\gamma^1 \in \Gamma^1} \int c(x, \gamma^1(y^1), \gamma^2(y^2)) \mu(dx, dy^1, dy^2).$$

Thus, the value function is lower semicontinuous. A similar proof can show the result for minimizer-garbling sequences. □

We can also show that when the DM channels are fixed, independent and total-variation continuous, the value function is continuous under weak convergence of the priors. This generalizes a single-DM result from [36].

Theorem 4.6. *Under Assumptions A2, A4, and A6, the value function is continuous under weak convergence of the priors.*

Proof. Let ζ_m be a sequence of priors on $\mathbb{X}$ converging weakly to ζ. Following the same reasoning as in Theorem 4.3, if for some j, $J^*(g, Q^1\zeta Q^2) - J^*(g, Q^1\zeta_j Q^2) \geq 0$, then we have:

$$\begin{aligned}
&J^*(g, Q^1\zeta Q^2) - J^*(g, Q^1\zeta_j Q^2) \\
&\leq \int c(x, \gamma_j^{1,*}(y^1), \gamma^{2,*}(y^2)) Q^1(dy^1|x) Q^2(dy^2|x) \zeta(dx) \\
&\qquad - \int c(x, \gamma_j^{1,*}(y^1), \gamma^{2,*}(y^2)) Q^1(dy^1|x) Q^2(dy^2|x) \zeta_j(dx) \\
&= \int (\zeta - \zeta_j)(dx) \int c(x, \gamma_j^{1,*}(y^1), \gamma^{2,*}(y^2)) Q^1(dy^1|x) Q^2(dy^2|x).
\end{aligned}$$

Considering the case $J^*(g, Q^1\zeta Q^2) - J^*(g, Q^1\zeta_j Q^2) \leq 0$ also as above we get

$$\begin{aligned}
&|J^*(g, Q^1\zeta Q^2) - J^*(g, Q^1\zeta_m Q^2)| \\
&\leq \max\{|\int (\zeta - \zeta_m)(dx) \int c(x, \gamma_m^{1,*}(y^1), \gamma^{2,*}(y^2)) Q^1(dy^1|x) Q^2(dy^2|x)|, \\
&\qquad |\int (\zeta - \zeta_m)(dx) \int c(x, \gamma^{1,*}(y^1), \gamma_m^{2,*}(y^2)) Q^1(dy^1|x) Q^2(dy^2|x)|\}. \quad (4.7)
\end{aligned}$$

Then, following the proof of [36, Theorem 2.5], since the channels are continuous in total variation,

$$\int c(x, \gamma_m^{1,*}(y^1), \gamma^{2,*}(y^2)) Q^1(dy^1|x) Q^2(dy^2|x),$$

and

$$\int c(x, \gamma^{1,*}(y^1), \gamma_m^{2,*}(y^2)) Q^1(dy^1|x) Q^2(dy^2|x),$$

are equicontinuous families of functions. Using this fact and that ζ_m converges weakly to ζ, by [20, Corollary 11.3.4] it follows that both terms in (4.7) converge to zero, and so $|J^*(g, Q^1\zeta Q^2) - J^*(g, Q^1\zeta Q^2)| \to 0$. □

4.4 Continuity Results for Stochastic Teams

Now we evaluate the regularity properties of the value function for team problems in information structures under total variation, weak, and setwise convergence.

Theorem 4.7. *Consider an n decision-maker team and let Assumptions A1, A3, A4 and A5 hold for a fixed team problem. Then the value function is continuous under total variation convergence of information structures.*

Proof. An optimal solution exists for any given information structure under the assumptions [64, Section 5]. Let μ and ν be two information structures, and let $\gamma_\mu^{i,*}$ and $\gamma_\nu^{i,*}$ denote DM i's individual strategy as part of the team-optimal strategies $\bar{\gamma}_\mu^*$, $\bar{\gamma}_\nu^*$ under each respective information structure. Let the cost function be bounded in absolute value by M.

Without loss of generality, assume $J^*(c, \mu) - J^*(c, \nu) \geq 0$. Then

$$J^*(c, \mu) - J^*(c, \nu) = \int c(x, \gamma_\mu^{1,*}(y^1), \ldots, \gamma_\mu^{n,*}(y^n)) \mu(dx, dy^1, \ldots, dy^n)$$

$$
\begin{aligned}
&\qquad - \int c(x, \gamma_\nu^{1,*}(y^1), \ldots, \gamma_\nu^{n,*}(y^n))\nu(dx, dy^1, \ldots, dy^n) \\
&\leq \int c(x, \gamma_\nu^{1,*}(y^1), \ldots, \gamma_\nu^{n,*}(y^n))\mu(dx, dy^1, \ldots, dy^n) \\
&\qquad - \int c(x, \gamma_\nu^{1,*}(y^1), \ldots, \gamma_\nu^{n,*}(y^n))\nu(dx, dy^1, \ldots, dy^n) \leq M\|\mu - \nu\|_{TV}.
\end{aligned}
$$

Where the first inequality holds by perturbing the strategy under μ to be $\bar{\gamma}_\nu^*$, rather than the team-optimal policy $\bar{\gamma}_\mu^*$, and the second inequality holds by the definition of total variation distance. The result follows. □

Theorem 4.8. *If DM action spaces are compact and Assumptions A2, A3, A4 and A6 hold, the value function is continuous under weak convergence of the priors.*

Proof. Let ζ_m be a sequence of priors on $\mathbb{X}$ converging weakly to ζ. Without loss of generality we assume that $J^*(g, Q^1 \ldots Q^n\zeta) - J^*(g, Q^1 \ldots Q^n\zeta_m) \geq 0$ for some game g. Following the same procedure as in Theorem 4.7, we get

$$
\begin{aligned}
&J^*(g, Q^1 \ldots Q^n\zeta) - J^*(g, Q^1 \ldots Q^n\zeta_m) \\
&\quad \leq \int (\zeta - \zeta_m)(dx) \int c(x, \gamma_m^{1,*}(y^1), \ldots, \gamma_m^{n,*}(y^n))Q^1(dy^1|x) \ldots Q^n(dy^n|x).
\end{aligned}
$$

Continuing as in the proof of Theorem 4.6, using equicontinuity, the result follows. □

Theorem 4.9. *Assume Assumptions A1, A3, A4, A5, and A7 hold. Let μ_m be a sequence of information structures converging weakly to an information structure μ. If the prior is not fixed, then we also impose continuity in x (under A2). Under these conditions the value function is upper semicontinuous in μ under weak convergence.*

Proof. Let $\bar{\Gamma} = \Gamma^1 \times \cdots \times \Gamma^n$. Let $\mu_{\mathbb{Y}^j}$ be the marginal of μ on its $(j+1)$th component. Let γ^j be an arbitrary policy for DM j. The action spaces for DMs are convex by

Assumption A4. Let $M = \|c\|_\infty$. Then, by Lusin's theorem, for any $\epsilon > 0$, there exists a continuous function $f^j : \mathbb{Y}^j \to \mathbb{U}^j$ such that $\mu_{\mathbb{Y}^j}(\{y^j : f^j(y^j) \neq \gamma^j(y^j)\}) < \epsilon$.

Letting B^j denote the set $B^j = \{y^j : f^j(y^j) \neq \gamma^j(y^j)\}$, and proceeding with this for $j = 1, \ldots, n$, we have:

$$\int c(x, f^1(y^1), \ldots, f^n(y^n))\mu(dx, dy^1, \ldots, dy^n)$$
$$< \int c(x, \gamma^1(y^1), \ldots, \gamma^n(y^n))\mu(dx, dy^1, \ldots, dy^n) + nM\epsilon.$$

Denote by Γ_C^j the policy space for DM j in which the DM's policy is continuous from $\mathbb{Y}^j$ to $\mathbb{U}^j$. By the above, it follows that the value for a game will be the same if each DM uses Γ or Γ_C. Thus, applying the fact that μ_m converges weakly to μ:

$$\begin{aligned}
&\limsup_{m\to\infty} \inf_{\bar{\gamma}\in\bar{\Gamma}} \int c(x, \gamma^1(y^1), \ldots, \gamma^n(y^n))\mu_m(dx, dy^1, \ldots, dy^n) \\
&= \limsup_{m\to\infty} \inf_{\bar{\gamma}_C\in\bar{\Gamma}_C} \int c(x, \gamma_C^1(y^1), \ldots, \gamma_C^n(y^n))\mu_m(dx, dy^1, \ldots, dy^n) \\
&\leq \inf_{\bar{\gamma}_C\in\bar{\Gamma}_C} \limsup_{m\to\infty} \int c(x, \gamma_C^1(y^1), \ldots, \gamma_C^n(y^n))\mu_m(dx, dy^1, \ldots, dy^n) \\
&= \inf_{\bar{\gamma}_C\in\bar{\Gamma}_C} \int c(x, \gamma_C^1(y^1), \ldots, \gamma_C^n(y^n))\mu(dx, dy^1, \ldots, dy^n) \\
&= \inf_{\bar{\gamma}\in\bar{\Gamma}} \int c(x, \gamma^1(y^1), \ldots, \gamma^n(y^n))\mu(dx, dy^1, \ldots, dy^n).
\end{aligned}$$

We note that if the prior is fixed, in the inequality above we do not need continuity of c in x since weak convergence of a product measure with a fixed marginal is equivalent to w-s convergence of the joint measure. If the prior is not fixed, then we impose continuity in x also under A2 and the inequality holds. □

Following a similar argument, without Lusin's theorem, we have the following.

Theorem 4.10. *Assume Assumptions A1, A3, A4 and A5 hold. Let μ_m be a sequence of information structures converging setwise to an information structure μ. Then the value function is upper semicontinuous in μ under setwise convergence.*

Proof. By interchanging the limit superior and the infimum, and then applying the fact that $\{\mu_m\} \to \mu$ setwise, we get:

$$\begin{aligned}
&\limsup_{m\to\infty} \inf_{\bar{\gamma}\in\bar{\Gamma}} \int c(x,\gamma^1(y^1),\ldots,\gamma^n(y^n))\mu_m(dx,dy^1,\ldots,dy^n) \\
&\leq \inf_{\bar{\gamma}\in\bar{\Gamma}} \limsup_{m\to\infty} \int c(x,\gamma^1(y^1),\ldots,\gamma^n(y^n))\mu_m(dx,dy^1,\ldots,dy^n) \\
&= \inf_{\bar{\gamma}\in\bar{\Gamma}} \int c(x,\gamma^1(y^1),\ldots,\gamma^n(y^n))\mu(dx,dy^1,\ldots,dy^n).
\end{aligned}$$

□

4.5 Continuity Results for General Games

While we have shown that zero-sum games and team problems both exhibit continuity under total variation convergence of information structures for games with measurable and bounded cost functions, here we present a counterexample that reveals this is not true for general non-zero-sum games.

Example 4.2. *Consider a two-DM game with state space $\mathbb{X} = [-1,1]$ endowed with the continuous uniform distribution as ζ. Let $\mathbb{Y}^1 = \{1,2,3\}$ and $\mathbb{Y}^2 = \{0\}$, and $\mathbb{U}^1 = \mathbb{U}^2 = [-1,1]$. For $m \in \mathbb{Z}_{\geq 1}$, define information structure μ_m by channels Q_m^1 and Q_m^2 for DMs 1 and 2 respectively, where Q_m^1 is a quantizer with three bins:*

$$y^1_m = \begin{cases} 1, & x \in [-1, -\frac{1}{2} - \frac{1}{8m}) \\ 2, & x \in [-\frac{1}{2} - \frac{1}{8m}, \frac{1}{2} + \frac{1}{4m}] \\ 3, & x \in (\frac{1}{2} + \frac{1}{4m}, 1] \end{cases}$$

and Q^2_m returns $y^2 = 0$ for all $x \in \mathbb{X}$.

Following [66, Theorem 5.7], since our information structure is defined by quantizers that converge setwise at input ζ, $\{\mu_m\} \to \mu$ in total variation, where μ is defined by the quantizer for DM 1 which sorts x into the bins $[-1, -1/2)$, $[-1/2, 1/2]$, and $(1/2, 1]$, and DM 2 has the same channel that always returns $y^2 = 0$.

Now we define the following cost functions for the DMs:

$$c^1(x, u^1, u^2) = (x - u^1)^2 - (u^2)^2, \qquad c^2(x, u^1, u^2) = \begin{cases} (u^2)^2, & u^1 = 0 \\ (u^2 - 1)^2, & u^1 \neq 0 \end{cases}$$

For each m, DM 1's optimal strategy is to minimize $(x - u^1)^2$, and thus his optimal policy is to play $u^1 = E_{\mu_n}[x|y^1]$. Due to the asymmetry of DM 1's measurement channel, $u^1 \neq 0\ \forall m$. Thus, DM 2's optimal strategy is to play $u^2 = 1$, and her expected cost is 0.

However, under μ, DM 1 will play 0 with probability $1/2$. Thus, DM 2's optimal strategy is to play $u^2 = 1/2$, and her expected cost is $1/4$. Thus, the value function for DM 2 is not continuous under total variation convergence of the information structure.

Going one step further, we can introduce a third DM with cost function c^3 whose cost is A when DM 2 plays $u^2 = 0$, and $-A$ otherwise, for some $A \in \mathbb{R}_{>0}$. For this

DM, the difference in expected value between μ_n *and* μ *is* $2A$ *for every* n*, which is the maximum possible performance difference given the bounds of the cost function. Thus, for any* $\epsilon > 0$*, there exists an* $\tilde{m}$ *such that* $\|\mu_{\tilde{m}} - \mu\|_{TV} < \epsilon$*, while the value distance for DM 3 between* $\mu_{\tilde{m}}$ *and* μ *is* $2A = 2\|c^3\|_\infty$*. It follows that total variation distance cannot be used to usefully bound performance change for general games.*

We conclude that general non-zero-sum games do not necessarily exhibit continuity under total variation, weak, or setwise convergence of information structures for games with measurable and bounded cost functions (as total variation is a stronger notion of convergence than both weak and setwise convergence).

Chapter 5

Conclusions

5.1 Summary

In this thesis we investigated properties of information structures in stochastic games. In particular, we focused on the comparison of information structures and continuity of the equilibrium value when perturbing the information structure.

In Chapter 2, we presented conditions for the existence of equilibria in zero-sum games with incomplete information in standard Borel measurement and action spaces. These results were closely related to existing results in the literature, although derived independently.

In Chapter 3, we presented an ordering of information structures for a broad class of zero-sum Bayesian games with incomplete information in standard Borel spaces, demonstrating necessary and sufficient conditions for an information structure to be "better" than another. We also provided a supporting result on a partial converse to Blackwell's ordering of information structures in this general setting.

In Chapter 4, we presented continuity properties of value functions and equilibrium solutions in zero-sum, team, and general game problems with respect to information structures. It was shown that the value function for both zero-sum games and team problems is continuous under total variation convergence of information structures. In both cases, the change in expected value when switching between two information structures is bounded above by the product of the total variation distance and the L^∞ norm of the cost function. For zero-sum games, the value function is upper semicontinuous for minimizer-garbling sequences of information structures, and lower semicontinuous for sequences of maximizer-garbling information structures under weak (if the cost function is continuous and bounded and the action spaces are convex) or setwise (if the cost function is measurable and bounded) convergence. For team problems, the value function is upper semicontinuous under setwise convergence for measurable and bounded cost functions, and upper semicontinuous under weak convergence for bounded and continuous cost functions when the players' action spaces are convex. A counterexample revealed players in general non-zero-sum games may not have value functions that are continuous under total variation convergence of information structures, even when the cost functions are bounded.

5.2 Future Work

One possible direction for future research is to investigate the topics presented here in greater depth for dynamic games. For regularity problems, while we studied static games and teams in this thesis, since it is known that under absolute continuity conditions there is an isomorphism relationship between equilibrium solutions to dynamic teams/games and their static reductions (which turns out to be policy independent)

[53], the results also apply to such dynamic games with absolutely continuous information structures. Dynamic problems which do not admit static reductions may be worth investigating in more detail in this context.

Another problem is characterizing properties of games in which additional information cannot hurt decision-makers. Sufficient conditions were provided in [9], however the authors demonstrate the results do not provide necessary conditions. In a similar vein, it may also be interesting to investigate sufficient conditions for a broad class of general games to demonstrate continuity in information structures under various notions of convergence.

Finally, a promising problem is to generalize the equilibrium and comparison results of Chapters 2 and 3 to team-against-team zero-sum games in standard Borel spaces, which would strictly generalize existing results for both zero-sum games and team problems.

Bibliography

[1] E. F. Abaya and G. L. Wise. Convergence of vector quantizers with applications to optimal quantization. *SIAM Journal on Applied Mathematics*, 44:183–189, 1984.

[2] R. J. Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, 55:1–18, 1987.

[3] T. Başar. Stochastic differential games and intricacy of information structures. In V. M. Haunschmied, J.and Veliov and S. Wrzaczek, editors, *Dynamic Games in Economics*, volume 16 of *Dynamic Modeling and Econometrics in Economics and Finance*, pages 23–49. Springer Berlin Heidelberg, 2014.

[4] T. Başar and Y.-C. Ho. Informational properties of the Nash solutions of two stochastic nonzero-sum games. *Journal of Economic Theory*, 7(4):370–387, 1974.

[5] T. Başar and G. J. Olsder. *Dynamic Noncooperative Game Theory*. SIAM Classics in Applied Mathematics, Philadelphia, PA, 1999.

[6] E. J. Balder. Generalized equilibrium results for games with incomplete information. *Mathematics of Operations Research*, 13(2):265–276, 1988.

[7] E. J. Balder. On ws-convergence of product measures. *Mathematics of Operations Research*, 26(3):494–518, 2001.

[8] T. Başar. Time consistency and robustness of equilibria in non-cooperative dynamic games. In *Contributions to economic analysis*, volume 181, pages 9–54. Elsevier, 1989.

[9] B. Bassan, O. Gossner, M. Scarsini, and S. Zamir. Positive value of information in games. *International Journal of Game Theory*, 32:17–31, December 2003.

[10] D. Blackwell. Equivalent comparison of experiments. *Annals of Mathematical Statistics*, 24:265–272, 1953.

[11] C. Boll. *Comparison of experiments in the infinite case.* PhD Dissertation, Stanford University, 1955.

[12] V. S. Borkar. White-noise representations in stochastic realization theory. *SIAM Journal on Control and Optimization*, 31:1093–1102, 1993.

[13] E. Boylan. Equiconvergence of martingales. *The Annals of Mathematical Statistics*, 42(2):552–559, 1971.

[14] L. Le Cam. Sufficiency and approximate sufficiency. *The Annals of Mathematical Statistics*, 35:1419–1455, 1964.

[15] L. Le Cam. Comparison of experiments - A short review. In *Ferguson, T., Shapley, L. (Eds.), Statistics, Probability and Game Theory Papers in Honor of David Blackwell*. IMS Lecture Notes Monograph Ser., 1996.

[16] P. Cartier, J. M. G. Fell, and P.-A. Meyer. Comparaison, des mesures portées par un ensemble convexe compact. *Bulletin de la Société Mathématique de France*, 92:435–445, 1964.

[17] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, New York, 1991.

[18] I. Csiszár. Information-type measures of difference of probability distributions and indirect observation. *studia scientiarum Mathematicarum Hungarica*, 2:229–318, 1967.

[19] B. De Meyer, , E. Lehrer, and D. Rosenberg. Evaluating information in zero-sum games with incomplete information on both sides. *Mathematics of Operations Research*, 34:851–863, November 2010.

[20] R. M. Dudley. *Real Analysis and Probability*. Cambridge University Press, Cambridge, 2nd edition, 2002.

[21] E. Shmaya E. Lehrer, D. Rosenberg. Signaling and mediation in games with common interests. *Games and Economic Behaviour*, 68:670–682, 2006.

[22] E. Einy, O. Haimanko, D. Moreno, and B. Shitovitz. Uniform continuity of the value of zero-sum games with differential information. *Mathematics of Operations Research*, 33(3):552–560, 2008.

[23] E. Einy, O. Haimanko, and B. Tumendemberel. Continuity of the value and optimal strategies when common priors change. *International Journal of Game Theory*, 41:829–849, 2012.

[24] K. Fan. Minimax theorems. *Proceedings of the National Academy of Sciences of the United States of America*, 39:42–47, 1953.

[25] F. Forges. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35:277–310, 1993.

[26] F. Gensbittel, M. Pęski, and J. Renault. Value-based distance between information structures. *arXiv preprint arXiv:2109.06656*, 2021.

[27] I. I. Gihman and A. V. Skorohod. *Controlled stochastic processes*. Springer Science & Business Media, 2012.

[28] O. Gossner and JF. Mertens. The value of information in zero-sum games. *preprint*, 2001.

[29] R. M. Gray and L. D. Davisson. Quantizer mismatch. *IEEE Transactions Communications*, 23:439–443, 1975.

[30] J. Hirshleifer. The private and social value of information and the reward to inventive activity. *The American Economic Review*, pages 561–574, 1971.

[31] I. Hogeboom-Burr and S. Yüksel. Comparison of information structures for zero-sum games and a partial converse to blackwell ordering in standard borel spaces. *SIAM Journal on Control and Optimization*, 59, 2021.

[32] A. Kajii and S. Morris. Payoff continuity in incomplete information games. *Journal of Economic Theory*, 82(1):267–276, 1998.

[33] M. I. Kamien, Y. Tauman, and S. Zamir. On the value of information in a strategic conflict. *Games and Economic Behavior*, 2(2):129–153, 1990.

[34] M. Kandori. The use of information in repeated games with imperfect monitoring. *The Review of Economic Studies*, 59(3):581–593, 1992.

[35] M. Kandori and I. Obara. Less is more: an observability paradox in repeated games. *International Journal of Game Theory*, 34(4):475–493, 2006.

[36] A. D. Kara and S. Yüksel. Robustness to incorrect priors in partially observed stochastic control. *SIAM Journal on Control and Optimization*, 57(3):1929–1964, 2019.

[37] A. Kloosterman. Public information in Markov games. *Journal of Economic Theory*, 157(C):28–48, 2015.

[38] H.J. Langen. Convergence of dynamic programming models. *Mathematics of Operations Research*, 6(4):493–512, Nov. 1981.

[39] E. Lehrer, D. Rosenberg, and E. Shmaya. Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68:670–682, 2010.

[40] J. W. Mamer and K. E. Schilling. A zero-sum game with incomplete information and compact action spaces. *Mathematics of Operations Research*, 11(4):627–631, 1986.

[41] J.-F. Mertens, S. Sorin, and S. Zamir. *Repeated games*, volume 55. Cambridge University Press, 2015.

[42] B. De Meyer, E. Lehrer, and D. Rosenberg. Two remarks on blackwell's theorem. *Journal of Applied Probability*, 45:580–586, June 2008.

[43] P. R. Milgrom and R. J. Weber. Distributional strategies for games with incomplete information. *Mathematics of operations research*, 10(4):619–632, 1985.

[44] M. Pęski. Comparison of information structures in zero-sum games. *Games and Economic Behavior*, 62(2):732–735, 2008.

[45] I. Petersen, M. R. James, and P. Dupuis. Minimax optimal control of stochastic uncertain systems with relative entropy constraints. *IEEE Transactions on Automatic Control*, 45(3):398–412, 2000.

[46] M. Pęski. Comparison of information structures in zero-sum games. *Games and Economic Behaviour*, 62:732–735, 2008.

[47] M. Raginsky. Shannon meets Blackwell and Le Cam: Channels, codes, and statistical experiments. In *2011 IEEE International Symposium on Information Theory Proceedings*, pages 1220–1224. IEEE, 2011.

[48] R. T. Rockafellar. *Convex analysis*. Number 28. Princeton university press, 1970.

[49] C. G. Rothschild. Payoff continuity in incomplete information games: A comment. *Journal of Economic Theory*, 120(2):270–274, 2005.

[50] A. Rubinstein. The electronic mail game: Strategic behavior under almost common knowledge. *The American Economic Review*, 79(3):385–391, 1989.

[51] W. Rudin. *Functional Analysis*. McGraw-Hill, New York, 1991.

[52] N. Saldi and S. Yüksel. Geometry of information structures, strategic measures and associated control topologies. *arXiv preprint arXiv:2010.07377*, 2020.

[53] S. Sanjari, T. Başar, and S. Yüksel. Isomorphism properties of optimality and equilibrium solutions under equivalent information structure transformations ii: Stochastic dynamic games. *arXiv preprint arXiv:2104.05787*, 2021.

[54] M. Schäl. On dynamic programming: compactness of the space of policies. *Stochastic Processes and their Applications*, 3(4):345–364, 1975.

[55] R. Serfozo. Convergence of lebesgue integrals with varying measures. *Sankhyā: The Indian Journal of Statistics, Series A*, 44:380–402, 1982.

[56] C. Shannon. A note on a partial ordering for communication channels. *Information and Control*, 1:390–397, 1958.

[57] V. Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, 1965.

[58] E. Torgersen. *Comparison of statistical experiments*, volume 36. Cambridge University Press, 1991.

[59] V. A. Ugrinovskii and I. R. Petersen. Minimax LQG control of stochastic partially observed uncertain systems. *SIAM Journal of Control and Optimization*, 40:1189–1226, 2001.

[60] J. von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.

[61] H. S. Witsenhausen. On the relations between the values of a game and its information structure. *Information and Control*, 19(3):204–215, 1971.

[62] H. S. Witsenhausen. Equivalent stochastic control problems. *Mathematics of Control, Signals, and Systems*, 1:3–11, 1988.

[63] Y. Wu and S. Verdú. Functional properties of minimum mean-square error and mutual information. *IEEE Transactions on Information Theory*, 58(3):1289–1301, 2012.

[64] S. Yüksel. A universal dynamic program and refined existence results for decentralized stochastic control. *SIAM Journal on Control and Optimization*, 58(5):2711–2739, 2020.

[65] S. Yüksel and T. Başar. *Stochastic Networked Control Systems: Stabilization and Optimization under Information Constraints*. Springer, New York, 2013.

[66] S. Yüksel and T. Linder. Optimization and convergence of observation channels in stochastic control. *SIAM Journal on Control and Optimization*, 50:864–887, 2012.

[67] S. Yüksel and N. Saldi. Convex analysis in decentralized stochastic control, strategic measures and optimal solutions. *SIAM Journal on Control and Optimization*, 55:1–28, 2017.