The social payoff in differentiation games

Gilad Bavly*† and Amnon Schreiber*‡

May 21, 2018

Abstract

We consider a class of games in which players have a common set of choices, asymmetric information about the values of these choices, and an incentive to differ from each other. We ask to what extent collisions can be avoided in equilibrium. We give exact answers for some special information structures, and a lower bound for the general case.

Keywords: incomplete information, partitions, collisions, social payoff, differentiation.
JEL classification: C72, D82.

1 Introduction

In many situations agents prefer to make choices that are different from the choices made by others. Often, they also have some private information about the values of different choices. For example, consider a group of scientists who are each about to start working on a new project, and not communicating their plans to each other. Each one of them may prefer to choose a project not chosen by the others, and may have her own assessment of the "worthiness" of candidate projects. Will the scientists all choose the same project? Will they all choose differently? The answer of course depends on the kind of information they have.

Our paper studies a class of games we call *differentiation games*, in which players have an incentive to differ from each other, and have asymmetric information about which choices could be good and which could not. We ask to what extent collisions (i.e., players making the same choice) can be avoided in equilibrium.

In our model, there is a finite set Ω of "states." Nature randomly chooses one of them, "the true state of the world," with uniform probability over Ω. The private information is described by knowledge partitions, as in Aumann (1976). Each player is associated with a partition of Ω (we call the elements

*Department of Economics, Bar-Ilan University, 5290002 Ramat-Gan, Israel.
†Email: gilad.bavly@gmail.com
‡Email: amnonschr@gmail.com

of a partition “cells”). A player’s private information consists of the cell of her partition that contains the true state of the world. Thus, each player knows that the true state is within a subset of Ω, yet that subset may differ from one player to the next.

Each player chooses one state within her subset. A dollar is equally shared between the players who choose the true state, if there are such players; other states yield nothing.

Example 1.1. Consider the game depicted in Figure 1. There are two players

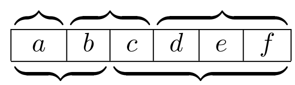

Figure 1: Two-player Game

and six states, denoted a, b, c, d, e, f. Player 1’s partition, represented by the curly brackets above the states, consists of three cells, $\{a\}$, $\{b, c\}$, and $\{d, e, f\}$. Player 2’s partition, represented by the curly brackets below the states, consists of two cells, $\{a, b\}$ and $\{c, d, e, f\}$.

A player’s strategy chooses one state within every cell in her partition. For example, the strategy of Player 1 depicted in Figure 2 chooses a within the cell

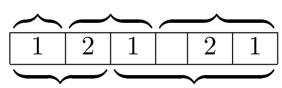

Figure 2: A Pair of Strategies

$\{a\}$, c within $\{b, c\}$, and f within $\{d, e, f\}$. Player 2 chooses b within $\{a, b\}$, and e within $\{c, d, e, f\}$.

If, for example, the true state of the world is c, then Player 1 knows that the true state is either b or c, and Player 2 knows that it is either c, d, e, or f. Player 1 chooses c and gets a dollar. Player 2 chooses e and gets nothing.

Overall, Player 1 gets a dollar in three out of the six states, and Player 2 gets a dollar in two states. Therefore, Player 1’s expected payoff is $1/2$ and Player 2’s expected payoff is $1/3$. The expectation of the social payoff (namely, the sum of the players’ payoffs) is $5/6$. The players never collide; hence they never share the dollar. This implies that this pair of strategies (1) is an equilibrium, and (2) yields the socially optimal payoff, i.e., no pair of strategies yields a higher expected social payoff.

In general, unlike in Example 1.1, collisions may be unavoidable. The socially optimal payoff is attained by strategy profiles that minimize the waste of collisions. We show (Theorem 3.1) that the socially optimal payoff is attained by some (pure) equilibrium in any differentiation game, not just in games in which collisions can be avoided.

Our main result gives a lower bound for the socially optimal payoff. The bound is in terms of a property of states we call "visibility," which is an aggregate measure of the accuracy of information in a state. In particular, this result implies that if the visibility of all states is at least one then the socially optimal payoff equals one, and if the visibility of all states is at most one then collisions can be avoided.

We then show how in cases where there is some regularity in the size of the cells, our bound allows us to pin down the socially optimal payoff. Another special case is where players' information is independent: in any "reasonable" game of this kind, collisions can be avoided.

Related Literature. There are many papers highlighting the role of differentiation in industrial organization settings; see, e.g., Eaton and Lipsey (1975), Gilbert and Matutes (1993), Katz (1984), and Shaked and Sutton (1982). Chen et al. (2015) study a scenario of privately informed scientists who repeatedly choose projects, and look for a fair mechanism that allows the players to share information. Kleinberg and Oren (2011) study mechanisms of allocating credit for performed research to scientists without asymmetric information. In Matros and Smirnov (2016), two players without asymmetric information compete in a repeated costly search, with a strong incentive to avoid searching the same spot simultaneously.

Differentiation games bear some similarity to congestion games (Rosenthal, 1973), which are complete information games in which each player chooses a set of items to use, and players may prefer less loaded items.[1]

The rest of the paper is organized as follows. Section 2 contains the model, and Section 3 the results. Section 4 discusses possible extensions of our model, and also explains how some popular puzzles about information (sometimes called "hat puzzles") can be formulated as differentiation games. Section 5 contains the proofs, except for very short ones which are given on the spot.

2 Model

Let N be the set of players, $n = |N|$, and let Ω be a finite set of *states*. Nature randomly chooses one state, *the true state of the world*, with uniform probability over Ω (in other words, the common prior of the players is the uniform probability over Ω).

The private information of players is represented by partitions (see, e.g., Aumann, 1976): for each player i there is a partition Π_i of Ω, namely, a list of disjoint subsets of Ω whose union is the whole Ω. We refer to the elements of player i's partition (i.e., the subsets) as player i's *cells*. When the true state of the world is $\omega \in \Omega$, player i knows that the true state is one of the states in

[1] As noted by Igal Milchtaich, for any differentiation game G there exists a congestion game whose payoff matrix is isomorphic to the expected payoff matrix of G.

$P_i(\omega)$, which is the cell of Π_i that includes[2] ω.

Each player chooses a state; i.e., her set of actions is Ω. Therefore, a pure strategy of player i is a function $\sigma_i : \Pi_i \to \Omega$ from i's cells to states. W.l.o.g. we assume that for every cell $\pi \in \Pi_i$, $\sigma_i(\pi) \in \pi$ (player i knows that the true state is within the cell π, and would be foolish to choose anything outside π). Similarly, a mixed strategy of player i is a function $\sigma_i : \Pi_i \to \Delta(\Omega)$ from i's cells to probability distributions over states, where for every $\pi \in \Pi_i$, $\sigma_i(\pi) \in \Delta(\pi)$.

The payoff of a player is $1/m$ if the state she chose is the true state of the world, where m is the overall number of players who chose that state; and 0 if she chose a different state.

A strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is an *equilibrium* if for any player i and any cell $\pi \in \Pi_i$, player i's expected payoff does not strictly increase by altering $\sigma_i(\pi)$, given that every other player $j \neq i$ sticks to σ_j.

2.1 Collisions and social payoff

The *social payoff* is the sum of players' payoffs. It equals 1 if the dollar is retrieved (i.e., at least one player chose the true state), and 0 otherwise. Therefore, the expected social payoff equals the probability that the dollar is retrieved. The *socially optimal payoff* is the maximum of the expected social payoff over all strategy profiles.

Player i's strategy σ_i chooses one state within every cell. Let $\pi \in \Pi_i$ be a cell and $\omega = \sigma_i(\pi)$ be the state she chooses within that cell. Note that when the true state happens to be this ω then indeed she chooses ω (because when the true state is ω she knows that the true state is within the cell that contains ω, which is the cell we called π, and within π her strategy chooses ω). Thus, in every cell $\pi \in \Pi_i$ there is one state in which player i is going to be right. Since the prior is uniform, her overall probability of choosing the true state is $|\Pi_i| \,/\, |\Omega|$, where $|\Pi_i|$ is the number of cells in her partition.

Suppose, for a moment, that player i is the only player. Then her overall probability of getting a dollar is $|\Pi_i| \,/\, |\Omega|$. This, then, is her expected payoff, no matter what her strategy σ_i is.

Let us turn back to our game, where player i is not the only player. She can only lose from the presence of other players, because she may have to share the dollar.

Definition 2.1. A profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ of pure strategies is *collision-free* if for any two players i, j and any $\pi_i \in \Pi_i, \pi_j \in \Pi_j$, we have $\sigma_i(\pi_i) \neq \sigma_j(\pi_j)$. A game is *collision-free* if it admits a collision-free profile of strategies.

In words, σ is collision-free if no two players ever choose the same state.

Suppose that our game admits a collision-free profile σ. Under σ, the expected payoff of player i equals $|\Pi_i| \,/\, |\Omega|$, the same as if she were alone. This is as high as her expected payoff under any strategy profile. Hence, in particular,

[2] As is usual in such models, the partitions themselves are common knowledge among all players.

σ is an equilibrium. The expected social payoff under σ equals $\sum_{i=1}^{n} |\Pi_i| / |\Omega|$, which is definitely the socially optimal payoff of the game.

In general, the game need not be collision-free. In a non-collision-free game, the socially optimal payoff is strictly less[3] than $\sum_{i=1}^{n} |\Pi_i| / |\Omega|$. It is attained in strategy profiles that minimize collisions.

3 Results

3.1 Socially optimal equilibria

There exists an equilibrium that yields the socially optimal payoff in any game, not just in collision-free games.

Theorem 3.1. *In any game, the socially optimal payoff is supported by a pure equilibrium.*

Still, not all equilibria yield the socially optimal payoff. Consider the game of Example 1.1. In Figure 2 we saw an equilibrium yielding an expected social payoff of 5/6, which is the socially optimal payoff. Figure 3 depicts an equilib-

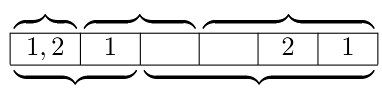

Figure 3: A Non-Optimal Equilibrium

rium yielding only 4/6. Under this equilibrium, the two players share the dollar when the state of the world is a.

Player 2 chooses a within the cell $\{a, b\}$, and she cannot improve upon this choice, since switching to b results in the two players colliding in b. Within the cell $\{a\}$ Player 1 cannot switch: a is her only choice. There are no collisions within other cells of either player; hence switching in those cells cannot improve a player's payoff. Therefore, this strategy profile is indeed an equilibrium.

3.2 Main result

Definition 3.2. The *visibility* of a state $\omega \in \Omega$, denoted $v(\omega)$, is defined by

$$v(\omega) = \sum_{i \in N} 1/|P_i(\omega)|.$$

When the state of the world is ω, the posterior belief of player i, after she has learned her information, is a uniform probability over $P_i(\omega)$. Therefore,

[3]This does not mean that, for example, refining a partition is bad for the social payoff. Although a refinement may increase the chances of a collision, it can never decrease the socially optimal payoff.

$v(\omega)$ is the sum of the posterior probabilities that players assign to ω being the state of the world, when the true state is[4] indeed ω. Thus, $v(\omega)$ is an aggregate measure of the accuracy of the information the players have when ω is the true state. Visibility is a local property of ω, in the sense that it depends only on cells that contain ω.

Theorem 3.3. *The socially optimal payoff is at least*

$$\sum_{\omega \in \Omega} \min \{v(\omega), 1\} / |\Omega| .$$

The visibility values in the game of Example 1.1 are presented in Figure 4. For instance, the state a is contained in $P_1(a) = \{a\}$ and $P_2(a) = \{a, b\}$; hence

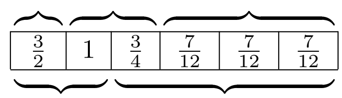

Figure 4: Visibilities

$v(a) = 1/1 + 1/2 = 3/2$. Theorem 3.3 tells us that the socially optimal payoff is at least $(1 + 1 + 3/4 + 7/12 + 7/12 + 7/12)/6 = (9/2)/6 = 3/4$.

By the way, we can deduce the exact socially optimal payoff in this example: there are six states; hence it must be an integer multiple of $1/6$. Therefore, being at least $3/4$ implies that it is at least $5/6$. On the other hand, $5/6$ is what you get if collisions are completely avoided. Hence, the socially optimal payoff equals $5/6$ and the game is collision-free (as we have already seen in the Introduction).

Corollary 3.4.

(i) *If $v(\omega) \geq 1$ for every $\omega \in \Omega$ then the socially optimal payoff equals one.*

(ii) *If $v(\omega) \leq 1$ for every $\omega \in \Omega$ then the game is collision-free.*

Recall from Section 2.1 that $\sum_{i=1}^{n} |\Pi_i| / |\Omega|$ is an upper bound of the socially optimal payoff, achieved if and only if the game is collision-free. We note that the sum of visibilities equals the number of cells:

Observation 3.5. $\sum_{\omega \in \Omega} v(\omega) = \sum_{i=1}^{n} |\Pi_i|$.

Proof. By definition, $\sum_{\omega \in \Omega} v(\omega) = \sum_{\omega \in \Omega} \sum_{i=1}^{n} 1/|P_i(\omega)|$. Every cell $\pi \in \Pi_i$ of every player i appears $|\pi|$ times in this sum, one time for each state within π (i.e., each state ω such that $P_i(\omega) = \pi$). Each of these times the cell contributes $1/|\pi|$ to the sum. Therefore, π's total contribution is one; hence this sum equals the number of cells $\sum_{i=1}^{n} |\Pi_i|$. □

[4]Equivalently, $v(\omega)$ is the expected number of times that ω would have been chosen, had the players chosen uniformly at random within every cell and ω been the true state.

Proof of Corollary 3.4. (i) By Theorem 3.3, the socially optimal payoff is at least $\sum_{\omega\in\Omega} 1/|\Omega| = 1$.
(ii) By Theorem 3.3, the socially optimal payoff is at least $\sum_{\omega\in\Omega} v(\omega)/|\Omega|$. Hence, by Observation 3.5, the socially optimal payoff is at least $\sum_{i=1}^{n} |\Pi_i|/|\Omega|$. Since this is an upper bound, this must be exactly the socially optimal payoff, and the game is collision-free. □

3.3 Special cases

We now consider a few special cases where there is some regularity in the size of cells that allows us to pin down the socially optimal payoff as a corollary of the main result.

Proposition 3.6. *For any k, if there are at least k players whose cells are at most of size k, then the socially optimal payoff equals one.*

A simple case is $k = 1$; i.e., there is a player whose every cell consists of a single state. This player always chooses the true state; hence the socially optimal payoff is one.

Proof. Let $K \subseteq N$ be a set of k players with cells of size $\leq k$. For any $\omega \in \Omega$, $v(\omega) \geq \sum_{i\in K} 1/P_i(\omega) \geq \sum_{i\in K} 1/k = 1$. By Corollary 3.4(i), the socially optimal payoff is one. □

Proposition 3.7. *If all cells of all players are at least of size n, then the game is collision-free.*

Proof. For any $\omega \in \Omega$, $v(\omega) \leq \sum_{i=1}^{n} 1/n = 1$. By Corollary 3.4(ii), the game is collision-free. □

Proposition 3.8. *If every player i has cells of uniform size d_i, then the socially optimal payoff equals* $\min\{1, \sum_{i=1}^{n} 1/d_i\}$.

Proof. The uniformity means that for every $\omega \in \Omega$, $v(\omega) = \sum_{i=1}^{n} 1/d_i$. Denote this number by α. If $\alpha > 1$, then by Corollary 3.4(i) the socially optimal payoff is one.

If $\alpha \leq 1$, then by Corollary 3.4(ii) the game is collision-free; hence the socially optimal payoff is $\sum_{i=1}^{n} |\Pi_i|/|\Omega|$, which, by Observation 3.5, equals $\sum_{\omega\in\Omega} v(\omega)/|\Omega| = (|\Omega| \cdot \alpha)/|\Omega| = \alpha$. □

The following claim is an instance of all three propositions above.

Corollary 3.9. *If all cells of all players are of size n, then the game is collision-free and the socially optimal payoff is one.*

	L	R
T	1	2
B	2	1

Figure 5: Independent Signals

3.3.1 Independent signals

We now consider another kind of special case, that of games in which the signals of all players are mutually independent. We will see that such games are"normally" collision-free.

Example 3.10. In the game depicted in Figure 5 there are two players, and Ω is a 2×2 matrix. Player 1's partition consists of the rows, i.e., one cell of 1 is the row $T = \{(T, L), (T, R)\}$ and the other cell is the row B. Player 2's partition consists of the columns.

The figure also depicts a collision-free pair of strategies, where Player 1 chooses the state (T, L) in row T, Player 2 chooses (B, L) in column L, etc.

Incidentally, collision-freeness is guaranteed by Corollary 3.9, but the point we are making here is that the game exhibits independence of signals: the probability of row T is 1/2 given that the column is L, and likewise it is 1/2 given that the column is R. Similarly, the probability of any column is the same within any row. That is, the rows and columns are independent as random variables over the space Ω.

In other words, the actual signals that Players 1 and 2 receive, namely, the cells $P_1(\omega)$ and $P_2(\omega)$ that contain the true state of the world ω, are independent random variables.

We say that the signals of n players are mutually independent if the probability of a profile of cells is the product of the marginal probabilities, i.e., for any profile of cells $(\pi_1, \dots, \pi_n) \in \Pi_1 \times \dots \times \Pi_n$, $\mathbb{P}(\pi_1, \dots, \pi_n) = \prod_{i=1}^n \mathbb{P}(\pi_i)$. Since the prior is uniform, this condition on probabilities is equivalent to

$$\left|\cap_{i=1}^n \pi_i\right| / |\Omega| = |\pi_1| \dots |\pi_n| / |\Omega|^n . \tag{1}$$

If Ω is an n-dimensional matrix and player i knows the i-th coordinate, as in Example 3.10, then the players' signals are mutually independent. Mutually independent signals need not be of this form, however, as the following game shows. There are two players, and $\Omega = \{1, \dots, 10\}$. The first player has two cells: one contains the odd numbers and the other contains the evens. The second player has two cells: one is $\{1, \dots, 4\}$ and the other is $\{5, \dots, 10\}$. Equation (1) is easily verified.

Proposition 3.11. *If signals are mutually independent then the socially optimal payoff equals* $\min\{1 , \sum_{i=1}^n |\Pi_i| / |\Omega|\}$.

Note that the RHS of Equation (1) cannot be zero. Therefore, mutual independence of signals implies that every profile of cells $(\pi_1, \dots, \pi_n) \in \Pi_1 \times$

$\ldots \times \Pi_n$ has a nonempty intersection. Proposition 3.11, then, follows from the following more general claim.

Proposition 3.12. *If every profile of cells has a nonempty intersection, then the socially optimal payoff equals* $\min\{1\,,\sum_{i=1}^{n}|\Pi_i|\,/\,|\Omega|\}$.

In particular, such games are collision-free if and only if $\sum_{i=1}^{n}|\Pi_i| \leq |\Omega|$. Yet, the nonempty intersection property also implies that $|\Omega|$ is at least $\prod_{i=1}^{n}|\Pi_i|$. Therefore, if collision-freeness does not obtain then there must be relatively many players with trivial information, i.e., players whose partition consists of a single cell.

4 Remarks

A *common knowledge component* of Ω is a minimal subset $D \subseteq \Omega$ such that for any player i and any $\omega \in D$, $P_i(\omega) \subseteq D$. Thus, the event D is common knowledge among the players at any state $\omega \in D$ (Aumann, 1976). Clearly, our results hold for any common knowledge component of Ω by itself, since one can think of each component as a separate game, unrelated to the rest of Ω.

4.1 Extensions

Our model can be extended in several natural ways. One is not to assume that the common prior is a uniform probability. Then any result that guaranteed a socially optimal payoff of one still holds. Likewise, results about collision-freeness still hold, but note that different collision-free strategy profiles may yield different social payoffs. Also, the socially optimal payoff may not be supported by an equilibrium; i.e., Theorem 3.1 may fail.

Another natural extension is to let the play go on for a few stages, instead of just one, so that the game ends either after a fixed number of stages or when the true state of the world is chosen by some player.

We may also think of games where not only one action is good (pays a dollar), but several actions are (or maybe some actions even pay something between zero and one). This is not a straightforward extension of the model, though. To model this, we need to enrich the information structure, while keeping the set of actions fixed. Apropos, even with just a single good action, one may consider lots of ways to enrich the information structure.

4.2 Puzzles

Some popular puzzles about information, sometimes called "hat puzzles," can be formulated as differentiation games. For example, two agents sit in separate rooms, where they each flip a coin and observe the outcome. Then both guess what the outcome of the coin flip in the *other* room was. They work as a team, and have agreed on their strategies beforehand. They succeed if at least one of them guesses correctly. What is the highest probability of success they can achieve?[5]

We can present this puzzle as a differentiation game: every possible pair of outcomes of the two coin flips is a state, i.e., $\Omega = \{H, T\} \times \{H, T\}$. A player's partition consists of two cells: one contains the two states where her own flip outcome is H (Heads) and the other contains the two states where her own flip outcome is T (Tails). She saw her own coin; hence she knows the cell that contains the true state of the world. She chooses a state within that cell, which amounts to guessing the outcome of the other coin. She gets one if she was the only player to guess correctly, one half if both guessed correctly, and zero otherwise.[6]

This game is actually the game of Example 3.10, though we may relabel the rows and columns as Heads and Tails. As we saw there, the socially optimal payoff is one; i.e., success can be guaranteed in the puzzle. Figure 5 depicts a pair of strategies that achieve that. In terms of the puzzle, they do the following. Player 1 guesses that the other coin landed the same as her own coin, and Player 2 guesses that the other coin landed differently than her own coin. Since (exactly) one of them is right, success is guaranteed.

The following popular puzzle is an extension of the two-coin puzzle. There are n agents, each wearing a hat whose color is one of n given colors (different hats may be the same color). An agent observes the colors of all hats but her own, and has to guess what her own color is. Again, they work as a team, and have agreed on their strategies beforehand. They succeed if at least one of them guesses correctly. Can they guarantee success?[7]

Let us present this as a differentiation game. A state is a list of n colors; hence there are n^n states. A player's cell contains n states that differ only in her own color (i.e., there is one cell for every combination of $n-1$ colors of the others). By Corollary 3.9, the game is collision-free and the socially optimal payoff is one. Thus, the team in the puzzle can guarantee success.[8]

The two-coin puzzle can also be extended to n agents in a different way: each agent observes only the coin she flipped, and has to guess the outcome of all the other coins. What is the highest probability of success, namely, that at least one of them guesses correctly?

In the corresponding game, $\Omega = \{H, T\}^n$. It is an n-dimensional matrix, and each player knows one coordinate; therefore, signals are mutually independent.

---

[5]We thank Sergiu Hart, who brought this puzzle to our attention.

[6]This individual payoff was not part of the story, which related only to the team's payoff. Still, we remain true to the story as long as the social payoff is the same, especially in view of Theorem 3.1.

[7]The relation to the two-coin puzzle may be obscured by the wording, as there an agent observes only her own item, while here she observes everything but her own item. What really matters, though, is that in both puzzles there is exactly one item that an agent does not observe. We could as well have said that an agent does not observe her own coin but does observe the coin in the other room.

[8]A nice and simple solution of the puzzle goes as follows. Index the set of colors by numbers between zero and $n-1$. Agent i chooses a color (an index) such that the sum of colors of all agents is $i \pmod{n}$. Thus, whatever the sum in the true state is, exactly one agent guesses correctly.

By Proposition 3.11, the game is collision-free. With two cells for each player, the socially optimal payoff is $2n/2^n$ (e.g., for $n = 3$ the probability is 3/4).

5 Proofs

5.1 Proof of Theorem 3.1

For a pure strategy profile $\sigma = (\sigma_1, \dots .\sigma_n)$ and a state ω, let $c_\sigma(\omega)$ denote the number of players that choose ω under σ, i.e., players i such that $\sigma_i(P_i(\omega)) = \omega$. The "imbalance" of σ is defined by

$$I(\sigma) = \sum_{\{\omega_1,\omega_2\}} |c_\sigma(\omega_1) - c_\sigma(\omega_2)|$$

summing over all (unordered) pairs of states in Ω.

Let σ be a pure strategy profile. If σ is not an equilibrium, then there is a profitable deviation; i.e., there is a player i and a cell $\pi \in \Pi_i$ such that switching from choosing $\omega = \sigma_i(\pi)$ to choosing some other state θ within π is profitable for i. Denote the new strategy profile, after this switch, by τ. Since the deviation is profitable, it must consist of switching from a crowded state to a less crowded one, i.e., $c_\tau(\theta) < c_\sigma(\omega)$. Hence, $c_\tau(\theta) \leq c_\tau(\omega) = c_\sigma(\omega) - 1$. Overall, we get

$$c_\sigma(\theta) + 1 = c_\tau(\theta) \leq c_\tau(\omega) = c_\sigma(\omega) - 1.$$

We claim that $I(\tau) < I(\sigma)$. First, switching from σ to τ strictly decreases $|c(\omega) - c(\theta)|$. Second, for any state $\rho \neq \omega, \theta$, switching does not affect $c(\rho)$, and:

(i) If $c(\rho) \leq c_\sigma(\theta)$ then switching increases $|c(\theta) - c(\rho)|$ by 1, but decreases $|c(\omega) - c(\rho)|$ by 1.

(ii) Similarly, if $c(\rho) \geq c_\sigma(\omega)$, the first term is decreased by 1, and the second term is increased by 1.

(iii) Otherwise $c_\sigma(\theta) < c(\rho) < c_\sigma(\omega)$, and switching decreases both terms.

Of course, ω is still chosen by some players under τ (namely, $c_\tau(\omega) \geq 1$). Therefore, switching from σ to τ does not decrease the social payoff. To sum up: we start with a profile σ, and if σ is not an equilibrium, we get a new profile τ while strictly decreasing the imbalance I, and without decreasing the social payoff.

Now, let σ_0 be some pure strategy profile that yields the socially optimal payoff. Starting with σ_0, we apply the above operation iteratively. Since I is integer-valued and nonnegative, this procedure has to stop after a finite number of steps, yielding a strategy profile σ^*. Then, σ^* must be an equilibrium, and as the social payoff did not decrease during the procedure, σ^* yields the socially optimal payoff.

5.2 Proof of Theorem 3.3

For a subset of states $\Theta \subseteq \Omega$, denote all cells[9] that intersect with Θ by $\varphi(\Theta) = \bigcup_{i=1}^{n} \{\pi \in \Pi_i : \pi \cap \Theta \neq \emptyset\}$.

Lemma 5.1. *Let*

$$d = \max_{\Theta \subseteq \Omega} \left(|\Theta| - |\varphi(\Theta)| \right).$$

The socially optimal payoff equals $1 - d/|\Omega|$.

We note that while Lemma 5.1 is an exact characterization of the socially optimal payoff, it is not in itself the kind of result we are looking for, as it requires going over many subsets.

Proof. Note that d is nonnegative since $\varphi(\emptyset) = \emptyset$. First, we see that the social payoff cannot exceed $1 - d/|\Omega|$. Let $\Theta \subseteq \Omega$ be a subset that achieves the maximum, i.e., $|\Theta| - |\varphi(\Theta)| = d$. For any strategy profile there are at least d states of Θ not chosen within any cell of any player, since there are only $|\varphi(\Theta)| = |\Theta| - d$ cells that intersect with Θ. Therefore, the social payoff is at most $1 - d/|\Omega|$.

To see that a social payoff of $1 - d/|\Omega|$ can be achieved, we define an auxiliary game g_* by adding d new players whose information is trivial (each has a partition that consists of a single cell) to the original game g_0. For the game g_*, define a bipartite graph $(\Omega, \varphi_{g_*}(\Omega); E)$ as follows. One side of the graph is the states Ω, and the other side is $\varphi_{g_*}(\Omega)$, which is simply all cells of all players $\bigcup_{i=1}^{n+d} \Pi_i$. The edges between the two sides are $E = \{(\omega, \pi) : \omega \in \pi\}$; i.e., a state $\omega \in \Omega$ is connected to a cell iff the cell contains it.

For a bipartite graph $(V_1, V_2; E)$ and a subset $S \subseteq V_1$, let $\Gamma(S)$ denote the set of all neighbors of the nodes in S, namely, every $v_2 \in V_2$ for which there exists $v_1 \in S$ such that $(v_1, v_2) \in E$. Note that in our graph, for any $\Theta \subseteq \Omega$, $\Gamma(\Theta) = \varphi_{g_*}(\Theta)$.

The classic Hall's theorem (see, e.g., Bollobás, 2012, p. 54) says that a bipartite graph $(V_1, V_2; E)$ contains a *complete matching* from V_1 to V_2 (a one-to-one matching between any $v_1 \in V_1$ and some neighbor of v_1) iff for any $S \subseteq V_1$, $|\Gamma(S)| \geq |S|$.

In our graph, for any subset $\Theta \subseteq \Omega$, $\varphi_{g_*}(\Theta)$ is the union of $\varphi_{g_0}(\Theta)$ with the cells of players $n+1, \ldots, n+d$, since any of these players has a single cell containing all states. Hence, $|\varphi_{g_*}(\Theta)| = |\varphi_{g_0}(\Theta)| + d \geq (|\Theta| - d) + d = |\Theta|$. Therefore, there is a complete matching, which amounts to a set Ψ of $|\Omega|$ many cells and a choice of a distinct ω within any cell in Ψ.

We construct a profile of strategies σ_* (in the game g_*), namely, a choice within every cell, by this complete matching combined with an arbitrary choice within any cell that is not in Ψ, if there is any. Under σ_*, any state in Ω is

[9] If the exact same cell appears in the partitions of two different players, these are considered as two distinct cells. Strictly speaking, within a collection of cells of more than one player, we can think of a cell as a pair (i, π), where i is the player to whom the cell belongs and π is the actual cell.

chosen at least once. Denote the projection of σ_* on the original game g_0 (i.e., simply omitting the extra players $n+1,\ldots,n+d$) by σ_0. Since every extra player made only one choice, the number of states not chosen under σ_0 is at most d. Therefore, σ_0 yields at least a social payoff of $1-d/|\Omega|$. □

Proceeding to the proof of Theorem 3.3, we claim that for any $\Theta \subseteq \Omega$,

$$|\varphi(\Theta)| \geq \sum_{\omega\in\Theta} v(\omega).$$

By definition, $\sum_{\omega\in\Theta} v(\omega) = \sum_{\omega\in\Theta}\sum_{i=1}^{n} 1/|P_i(\omega)|$. Similarly to the proof of Observation 3.5, any cell $\pi \in \varphi(\Theta)$ of any player i appears $|\pi\cap\Theta|$ times in this sum, one time for each state within $\pi\cap\Theta$ (i.e., each state $\omega\in\Theta$ such that $P_i(\omega)=\pi$). Each of these times the cell contributes $1/|\pi|$ to the sum. Therefore, π's total contribution is at most one; hence this sum is at most the number of cells $|\varphi(\Theta)|$.

Therefore,

$$\begin{aligned}
&|\Theta| - |\varphi(\Theta)| \;\leq\; |\Theta| - \sum_{\omega\in\Theta} v(\omega)\\
&\leq |\Theta| - \sum_{\omega\in\Theta} \min\{v(\omega),1\} \;=\; \sum_{\omega\in\Theta}(1-\min\{v(\omega),1\})\\
&\leq \sum_{\omega\in\Omega}(1-\min\{v(\omega),1\}) \qquad \text{(since all summands are nonnegative)}\\
&= |\Omega| - \sum_{\omega\in\Omega}\min\{v(\omega),1\}.
\end{aligned}$$

Thus, the number d of Lemma 5.1 is at most $|\Omega| - \sum_{\omega\in\Omega}\min\{v(\omega),1\}$. By the lemma, the socially optimal payoff is at least

$$1-\left(|\Omega| - \sum_{\omega\in\Omega}\min\{v(\omega),1\}\right)/|\Omega| = \sum_{\omega\in\Omega}\min\{v(\omega),1\}/|\Omega|.$$

5.3 Proof of Proposition 3.12

Let $M=\{i\in N : |\Pi_i|\geq 2\}$ be the set of players whose partition is not trivial, and denote $m=|M|$. For any player $i\in N$, $\omega\in\Omega$, and any tuple of cells $(\pi_j)_{j\neq i}$, the profile of n cells $(P_i(\omega),(\pi_j)_{j\neq i})$ has a nonempty intersection. Therefore, $P_i(\omega)$ contains at least $\prod_{j\in N\setminus\{i\}}|\Pi_j|$ such intersection points, and

$$\prod_{j\in N\setminus\{i\}}|\Pi_j| = \prod_{j\in M\setminus\{i\}}|\Pi_j| \geq 2^{m-1}.$$

Hence, $|P_i(\omega)|\geq 2^{m-1}$.

Now consider a smaller, auxiliary game created by omitting all players in $N\setminus M$, leaving only members of M to play. We claim that the smaller game is collision-free.

If $m = 0$ then this claim is empty. If $m \geq 1$ then for any $\omega \in \Omega$, the visibility of ω *in the smaller game* is

$$\sum_{i \in M} 1/\left|P_i(\omega)\right| \leq m \cdot \frac{1}{2^{m-1}} \leq 1.$$

Therefore, by Corollary 3.4(ii), the game is collision-free.

Let σ' be a collision-free strategy profile in the smaller game. Under σ', $\sum_{i \in M} |\Pi_i|$ distinct states are chosen. Going back to the original game, we define a strategy profile σ by complementing σ' with strategies of the members of $N \setminus M$ as follows. We let them choose, one by one (within their single cell, namely, the whole Ω), some arbitrary state that has not already been chosen, as long as there is one.

If $\sum_{i=1}^{n} |\Pi_i| \leq |\Omega|$ then this procedure continues until all members of $N \setminus M$ have chosen. We end up with a strategy profile σ that is collision-free; hence the socially optimal payoff is $\sum_{i=1}^{n} |\Pi_i| / |\Omega|$. If $\sum_{i=1}^{n} |\Pi_i| > |\Omega|$ then the procedure cannot be completed, as at some stage all the states will already have been chosen before all the players have been able to choose. Let the remaining players choose arbitrarily. We end up with a strategy profile σ under which all states are chosen; hence the socially optimal payoff is one.

Acknowledgments

We thank Yuval Heller, Ron Peretz, Dov Samet, and Eilon Solan for helpful comments. G. Bavly acknowledges support from the Department of Economics at Bar-Ilan University, the Sir Isaac Wolfson Chair in Economics and Business Administration, ISF grant 1188/14, and ERC starting grant #677057.

References

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

Aumann, R. J. and S. Hart (2006). Agreeing on decisions. Technical report.

Bollobás, B. (2012). *Graph theory: an introductory course*, Volume 63. Springer Science & Business Media.

Chen, Y., K. Nissim, and B. Waggoner (2015). Fair information sharing for treasure hunting. In *AAAI 2015*, pp. 851–857.

Eaton, B. C. and R. G. Lipsey (1975). The principle of minimum differentiation reconsidered: Some new developments in the theory of spatial competition. *The Review of Economic Studies 42*(1), 27–49.

Gilbert, R. J. and C. Matutes (1993). Product line rivalry with brand differentiation. *The Journal of Industrial Economics 41*(3), 223–240.

Katz, M. L. (1984). Firm-specific differentiation and competition among multi-product firms. *The Journal of Business 57*(1), S149–S166.

Kleinberg, J. and S. Oren (2011). Mechanisms for (mis) allocating scientific credit. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pp. 529–538. ACM.

Matros, A. and V. Smirnov (2016). Duplicative search. *Games and Economic Behavior 99*, 1–22.

Rosenthal, R. W. (1973). A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory 2*(1), 65–67.

Shaked, A. and J. Sutton (1982). Relaxing price competition through product differentiation. *The Review of Economic Studies 49*(1), 3–13.