Game Theory with Sparsity-Based Bounded Rationality*

Xavier Gabaix

NYU Stern, CEPR and NBER

June 16, 2011

Extremely preliminary and incomplete.

Abstract

This paper proposes a way to enrich traditional game theory via an element of bounded rationality. It builds on previous work by the author in the context of one-agent decision problems, but extends it to several players and a generalization of the concept of Nash equilibrium. Each decision maker builds a simplified representation of the world, which includes only a partial understanding of the equilibrium play of the other agents. The agents' desire for simplificity is modeled as "sparsity" in a tractable manner. The paper formulates the model in the concept of one-shot games, and extends it to sequential games and Bayesian games. It applies the model to a variety of exemplary games: $p-$beauty contests, matching pennies, centipede, traveler's dilemma, and dollar auction games. The model's predictions are congruent with salient facts of the experimental evidence. Compared to previous succesful proposals (e.g., models with different levels of thinking or noisy decision making), the model is particularly tractable and yields simple closed forms where none were available. A sparsity-based approach gives a portable and parsimonious way to inject bounded rationality into models with strategic agents.

Keywords: inattention, boundedly rational game theory, behavioral modeling, sparse representations, ℓ_1 minimization.

*xgabaix@stern.nyu.edu. For very good research assistance I am grateful to Tingting Ding and Farzad Saidi. For helpful comments I thank Colin Camerer, Vincent Crawford and seminar participants at NYU and Caltech. I also thank the NSF for support under grant DMS-0938185. I apologize for the typos that this preliminary version contains.

1 Introduction

A key open question in economics is the formulation of a workable alternative to the rational agent. In recent work (Gabaix 2011), working with a solo-agent model, I proposed a "sparsity-based" model of bounded rationality which is quite tractable and can account for a variety of realistic departures from the traditional model. In the present paper, I continue this first step by proposing a formulation with several agents that allows to consider games. I also propose an extension of the concept of Nash equilibrium with sparsity-based bounded rationality. The model accounts reasonably well and sensibly for a variety of important games that were found to be at odds with the rational model. In addition, the model is quite tractable.

It builds in spirit on a series of important contributions: level-k thinking (e.g., Nagel 1995, Stalh and Wilson 19966, Crawford and Iriberri 2007; see also the excellent survey by Crawford, Costa-Gomes, and Iriberri 2010), cognitive hierarchy (Camerer, Ho, and Chong 2004), and – more loosely – quantal response equilibrium (McKelvey and Palfrey 1995), cursed equilibrium (Eyster and Rabin 2005), and analogy-based equilibrium (Jéhiel 2005). These models have proven extremely useful for making sense of systematic deviations from Nash equilibrium such that they are both psychologically reasonable and quantitatively successful. Despite this success of existing models, the Sparse Bounded Rationality (Sparse BR) in games model proposed here may still be useful for two reasons.

First, conceptually, most of the existing models (level-k thinking, cognitive hierarchy, cursed equilibrium) assume that the agent is fully rational (and maximizes perfectly) but that the other players do not fully maximize. This is clearly an important limitation. We do need a model in which, even when playing alone (e.g., solving a complicated solo problem), the agent is less than fully rational. Such a model has been proposed in the sparsity-based approach of Gabaix (2011). Hence, the present model tries to glue a one-person boundedly rational model to other models with several agents. Second, the above-mentioned models are typically easy to use with a computer (which is fine to analyze particular experimental games) but harder to use for paper-and-pencil analysis. Here, I strive to always obtain tractability and closed-form solutions. This is important for applications, and useful for a host of economic questions (e.g., in the basic models we teach) alike.

In the model, agents form a sense of the other players' actions (like in level-k models).

However, they form only an imperfect sense: they still anchor on a default model of the other agents (e.g., that the other agents will play randomly). How much they simulate the other agent is given by an approximate cost-benefit analysis of the gain from such a simulation.

It turns out that, under some ancillary assumptions, the resulting model can generate cursed equilibrium as a particular case (in which it is very much indebted to cursed equilibrium for the progress made by the latter in modeling Bayesian games). However, unlike cursed equilibrium (which assumes perfect rationality), it does not assume perfect rationality in solo situations. Thus, it offers some unification of departures from rationality, including those visible in simple solo games, to multi-player games of full information such as p-beauty and to other games with asymmetric information such as auctions.

Technically, a difficulty was to have a framework that handles "smooth" decision spaces such as macro decisions in $\mathbb{R}$ with convex decision problems, and discrete decision spaces such as the discrete actions in game theory. To this end, I relied for inspiration on the basic idea in applied mathematics of distances between variables, in particular the Wasserstein distance (Villani 2003).

The paper proceeds as follows. In Section 2, I lay out the framework and illustrate it with some simple games, particularly differentiable games. Section 3 illustrates it in contexts with less structure, such as games with a discrete action space. Section 4 discusses further extensions to extensive-form games and learning.

2 Basic Ideas

2.1 Model Statement with Pure Strategies

There are K types of players, indexed $k = 1...K$. Each player k plays $a_k \in A_k$. His utility is:

$$u_k\left(a_k, a_{-k}, X\left(a\right)\right).$$

We denote $a = \left(a_i\right)_{i=1...K}$. The value $X\left(a\right)$ depends on all the a's of the players. For instance, it could be a vector of prices, as in a market game.

The decision maker (DM) observes $X\left(a\right)$ (e.g., a vector of prices), and forms a belief about the other players' actions a.

We start with a vector of default values a_k^d, X_k^d, which corresponds to player k's default model of the other players' actions, vector X's value. Depending on the context, default policies could be very random, e.g., "pick each action with equal probability."

Player k, at the cost of some cognitive cost modeled below, forms the state vector $x = (a_{-k}, X(a))$, and interpolates it with his default vector $x^d = \left(a^d, X^d\right)$:

$$x_i(m_i) = (1 - m_i)\, x_i^d + m_i x_i^*. \tag{1}$$

Hence, if $m_i = 0$, $x_i(m_i) = x_i^d$ is simply the default, while if $m_i = 1$, $x_i(m_i) = x_i^*$ is the true (equilibrium) state of the variable.

His utility function under the model-in-model (MIM) m is:

$$V_k(a_k, m) = u_k(a_k, x(m)).$$

The DM applies the Sparse BR model to $V_k(a_k, m)$.

There is a distance $D(a, a')$ between strategies. When strategies are pure strategies on the real line, $D(a, a') = |a - a'|$. When strategies are pure strategies in an unstructured set (e.g., {Up, Down}), then $D(a, a') = 1_{a \neq a'}$. The more general case of mixed strategies (with a distance between random strategies) is a bit more technical, and will be discussed later in Section 7.1.

The Sparse BR algorithm for player k is a transposition of the algorithm previously proposed in Gabaix (2011). There is a distribution $q_k(\cdot)$ on player A's strategy (intuitively, centered around a_k^d). For a function $f : A \to \mathbb{R}$, we denote $\sigma[f(a)]$ the standard deviation of $f(a)$ using probability q.

Algorithm 1 *(Sparse BR Algorithm for One Player) A player (here, k) decides according to the following steps.*

Step 1: *The DM chooses a parsimonious representation of the world $m \in R^{n_m}$ by solving:*

$$\min_{m_i} \sum_i (m_i - \mu_i)^2 \frac{\sigma\left[V_m\left(a_k, m^d\right)\right]^2}{\sigma\left[V\left(a_k, m^d\right)\right]} + \kappa^m \sum_i |m_i - \mu_i|\, \sigma\left[V_m\left(a_k, m^d\right)\right]. \tag{2}$$

Step 2: *The DM chooses the best action:*

$$\max_a V(a, m) - \kappa^a D\left(a, a^d\right) \left\|V_m\left(a_k, m^d\right) \eta_m\right\|. \tag{3}$$

We define

$$BR \max_{a,m} V(a, m) \tag{4}$$

as the set of solutions to (3).

This algorithm is rather streamlined compared to the one in Gabaix (2011). This is because we cannot use convenient objects like second derivatives. In this paper, we wish to avoid the smoothness assumption.[1]

Step 1 leads to the following Lemma.

Lemma 1 *In Step 1 of Algorithm 1, the solution is:*

$$m_i = m_i^d + \tau\left(\mu_i - m_i^d, \kappa^m \frac{\sigma\left[V\left(a_k, m^d\right)\right]}{\sigma\left[V_m\left(a_k, m^d\right)\right]}\right) \tag{6}$$

using the "anchoring and adjustment" function: $\tau(\mu, \kappa) = (|\mu| - |\kappa|)_+ \, sign(\mu)$, *i.e.,*

$$\tau(\mu, \kappa) = \begin{cases} \mu + \kappa & \text{if } \mu \le -\kappa \\ 0 & \text{if } |\mu| < \kappa \\ \mu - \kappa & \text{if } \mu \ge \kappa \end{cases} \tag{7}$$

for $\kappa > 0$.

Figure 1 plots the anchoring and adjustment function $\tau(\mu, \kappa)$.

We are at an (BR) equilibrium when a_k^* is a BR best response to $V_k(a_k, m)$, with $x^* = (a^*, X(a^*))$ the equilibrium action. The key feature is that the agent may not think through all the other players' actions, and instead anchors them on the default.

I define a BR Nash equilibrium in the following way:

Definition 1 *(Sparse BR Nash Equilibrium) With* K *types, an equilibrium action is given by* $a^* = (a_k^*)_{k=1\ldots K}$ *(where* a_k^* *is possibly a mixed strategy rather than a pure one) such that: forming* $x = (a, X)$*, the vector of actions and states of the outside world, given the defaults*

[1] However, Step 1 modifies it somewhat: because a continuum support is not generally available, expressions such as V_{aa} are also not available. However, in some sense the new Step 1 and the Step 1 in the basic Sparse BR model are very related. First, they share the same structure, and minimimize a quadratic $(m_i - \mu_i)^2$ loss and a first-order deviation from lack of sparsity $|m_i - \mu_i|$. Next, Step 1 is de facto equivalent to:

$$m_i = m_i^d + \tau\left(\mu_i - m_i^d, \kappa^m \frac{stdev_k\left(V\left(a_k, m^d\right)\right)}{stdev_k\left(V_m\left(a_k, m^d\right)\right)}\right) \tag{5}$$

In the basic Sparse BR model with one-dimensional action, we have

$$m_i = m_i^d + \tau\left(\mu_i - m_i^d, \kappa^m \frac{|\eta_a V_{aa} \eta_a|}{|V_{ma}\eta_a|}\right),$$

but in a local sense $stdev_k\left(V\left(a_k, m^d\right)\right) \simeq |\eta_a V_{aa}\eta_a|$ and $stdev_k\left(V_m\left(a_k, m^d\right)\right) \simeq |V_{ma}\eta_a|$.

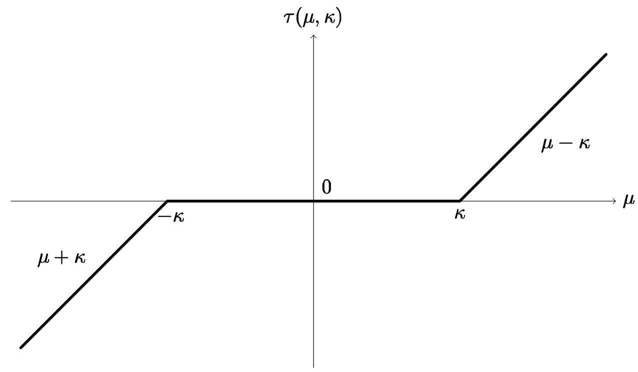

Figure 1: The anchoring and adjustment function τ

$x^d = \left(a^d, X^d\right)$, and defining $x_i\left(m_i\right) = x_i^d + m\left(x_i - x_i^d\right)$ (the interpolation between the default and the actual states of the world) and $\eta_x = \left(\eta_a, \eta_X\right)$, we have

$$a_k^* \in BR\max_{a_k, m} V_k\left(a_k, m\right) \tag{8}$$

for all players k.

I next define a "natural" notion, based on level-k models.

Definition 2 *(SparseBR Cognitive Steps Model) Define the function $SBR^k\left(a_{-k}^e\right)$ as the best response to a sense that the other players might play a_{-k}^e rather than the default a^d. Define $a^e\left(t=0\right) = a^d$ and, for $t \geq 0$, define iteratively:*

$$a_k\left(t+1\right) = SBR^k\left(a_{-k}^e\left(t\right)\right).$$

Agent iterate (mentally) and stop at the first time t when:

$$\forall k, \left|u_k\left(a\left(t-1\right)\right) - u_k\left(a\left(t\right)\right)\right| < \xi\sigma_k \tag{9}$$

for some given $\xi > 0$.

The intuition is that the agent progressively updates his model of the other players, and also his model of what other players think. His reasoning stops when all players' increment in utility is small.

I recommend using the Definition 1 of the BR Nash equilibrium. Definition 2 with level-k is useful conceptually, but more cumbersome to use in practice.

I next comment on the equilibrium. First, we have the following simple lemma.

Lemma 2 *(Three Limits of the BR Nash equilibrium). We have the following limits for the BR Nash equilibrium:*

(i) When $\kappa^m = \kappa^a = 0$, the definition of a Sparse BR equilibrium is simply that of a standard Nash equilibrium.

(ii) When $\kappa^a = \infty$ (i.e., no response of actions), the model is a level-0 model

(iii) When $\kappa^a = 0$ and $\kappa^m \gg 1$, and no Nature, then the model is a level-1 model.

Proof. Indeed, in the SparseBR algorithm yields $m = \mu$ when $\kappa^m = 0$, so that agents have the correct model of the world, and when $\kappa^a = 0$, $BR\max$ is simply the full maximization. Similar reasoning yields the other points. ■

With this lemma, the paper is in good company, as arguably level-k models have been quite successful. They can be yielded through a continuous parameter.

This definition uses a fixed point approach. Intuitively, the agents have some sense of the equilibrium behavior, but do not wish to take that sense fully into account. Instead, they anchor on a simpler default model of the others' behavior, e.g., random play. This is of course an idealization (in some cases a "round of thinking" model à la level-k might seem more intuitive).

This model is a extension of the basic Sparse BR model with just one player. One player could be Nature. Take the limit of one player against Nature: the player would like to $\max_a V^1(a, x(m))$ with $m = \mu$, as in the basic Sparse BR model. Here, it is represented as some MIM for Nature's model, $a^{(2)} = m$.

Comparision with iterated depth of reasoning models We can also note that when $m = 0$, $\kappa^a = 0$, and there is no move by Nature, then the model is just the best response to the default model. The Sparse BR agent is identical to the level-1 thinking agent. On the other hand, if κ^a is large enough, the agent does not deviate from the default policy (which in a typical game is to randomize uniformly over the possible actions), so that the agent is a level-0 thinker.

Hence, the model espouses important features of level-k models. In addition, it predicts an endogenous level of k: agents stop thinking when the perceived marginal benefits are small. It would be interesting to investigate directly how well the model "predicts the level k," but this will not be pursued here.

More methodologically, the Sparse BR model works with continuous parameters (κ^m, κ^a) rather than discrete parameters (the level k of thinking) or a distribution of k's. The model maintains homogenous actions across agents. This is useful for pure theory. For experimental economics, it is easy to destroy homogeneity if one wishes to, e.g., by assuming that there is a dispersion of κ's in the population.

Before studying concrete example, let us record that SparseBR Nash equilibria exist.

Proposition 1 *Suppose that the support of actions a is compact, and for each $u^k(a)$ is continuous. Then, a Sparse BR Nash equilibrium exists.*

Proof. Step 1 gives (6), hence, with $\kappa_k := \kappa^m \frac{\sigma\left[V^k\left(m^d\right)\right]}{\sigma\left[V_m^k\left(m^d\right)\right]}$,

$$\begin{aligned} V_k\left(a_k, a_{-k}\right) &= u_k\left(a_k, x\left(m\right)\right) \\ &= u_k\left(a_k, \left(a_i^d + \tau(a_i - a_i^d, \kappa_k)\right)_{i \neq k}\right) \end{aligned}$$

Step 2 implies that player k maximizes over a_k:

$$W_k\left(a_k, a_{-k}\right) := V_k\left(a_k, a_{-k}\right) - \kappa^a D\left(a, a^d\right) \left\|V_m^k\left(a_k, m^d\right)\eta_m\right\| \tag{10}$$

Hence, the SparseBR Nash equilibrium is simple the regular Nash equilibrium of a game with payoffs $W_k(a)$. The value function $W_k(a)$ are continuous, by the standard theorem on the existence of a Nash equilibrium, there exists a Nash equilibrium a^* for the game with value functions $W_k(a)$. This a^* is the SparseBR Nash equilibrium of the original game. ■

2.2 Application: Differentiable Games

Consider as an application a symmetric game with pure strategies. Players maximize $\max_a u(a, b)$ where b is the average action of other players. We have the parametrization: Player 1 uses the model

$$b(m) = a^d + m,$$

a smooth interpolation between the default and the actual equilibrium behavior.

Then, the optimality condition on action a for a player is:

$$\partial_1 u\left(a^*, a^d + m\right) = 0. \tag{11}$$

By (6), we can use:

$$m = \tau\left(1, \kappa^m \sigma_a \frac{u_{11}}{u_{12}}\right). \tag{12}$$

Hence, we have the following.

Proposition 2 *For symmetric, differentiable games with utility $u(a,b)$, the equilibrium a^* solves the following equation:*

$$\partial_1 u\left(a^*, a^d + \tau\left(a^* - a^d, \kappa^m \sigma_a \frac{u_{11}}{u_{12}}\right)\right) = 0 \tag{13}$$

where $\frac{u_{11}}{u_{12}}$ is evaluated at $\left(a^d, a^d\right)$.

Let us apply this equation.

2.2.1 Beauty Contest

With support $[0,1]$, $a^d = 1/2$ and $\sigma_a = stdev\,(U\,[0,1]) = 1/\sqrt{12} \simeq 0.3$, and we use the payoff:

$$u(a,b) = -\frac{1}{2}(a - pb)^2.$$

Proposition 3 *Play in the p-beauty contest game (with $\kappa^m > 0$, $\kappa^a = 0$) with goal "p times the mean" is:*

$$a^* = \min\left(\frac{p}{2}, \frac{\kappa^m \sigma_a}{1-p}\right) \text{ when } p < 1 \tag{14}$$

where $\sigma_a = 1/\sqrt{12}$, while for $p \geq 1$, we have

$$\begin{aligned} a^* &= \frac{p}{2} \text{ when } 1 \leq p \leq p^* \\ &= 1 \text{ when } p > p^* = \frac{1 + \sqrt{1 + 8\kappa^m \sigma_a}}{2} \end{aligned}$$

In the more general game where the goal is "p× (mean plus μ)" with $p < 1$, the action is:

$$a^* = \min\left(p\left(\frac{1}{2} + \mu\right), \frac{p\mu + \kappa^m \sigma_a}{1-p}\right).$$

The above numbers for a^* holds for an action space $A = [0,1]$. If $A = [0,100]$, the all a^* are multiplied by 100.

Proof. We have:

$$m = \tau\left(1, \kappa^m \frac{u_{11}}{u_{12}} \frac{\sigma_a}{a^* - a^d}\right) = \tau\left(1, \kappa^m \frac{\sigma_a}{p\,(a^* - 1/2)}\right).$$

If $m \neq 0$, we guess $a^* \leq 1/2$, and

$$m = 1 - \frac{\kappa^m \sigma_a}{p\left(\frac{1}{2} - a^*\right)},$$

which leads to:

$$\begin{aligned} a^* &= p\left(\frac{1}{2} + m\left(a^* - \frac{1}{2}\right) + \mu\right) \\ &= \frac{p}{2} + p\left(a^* - \frac{1}{2} + \mu\right) + \kappa^m \sigma_a \\ &= p\left(a^* + \mu\right) + \kappa^m \sigma_a, \end{aligned}$$

i.e.,

$$a^* = \frac{p\mu + \kappa^m \sigma_a}{1-p}. \tag{15}$$

If $m = 0$, $a^* = p\left(\frac{1}{2} + m\left(a^* - \frac{1}{2}\right) + \mu\right) = p\left(\frac{1}{2} + \mu\right)$. So,

$$a^* = \min\left(p\left(\frac{1}{2} + \mu\right), \frac{p\mu + \kappa^m \sigma_a}{1-p}\right). \tag{16}$$

The case $p > 1$ is in the appendix. ■

We see that we have the intuitive comparative static that the equilibrium action increases with p.

It is worth comparing this solution to others, particularly those with k levels of thinking, that espouse very well the psychology of the p-beauty contest. Here, I have sought a concept that preserves the representative agent throughout: as agents are identical ex ante, if at all possible, in the model agents play identical strategies ex post (this is in contrast to level-k, including cognitive hierarchy, models where there are different strategies). The reason is that I wish the model to be portable to applied economics such as macroeconomics or finance where the representative agent model is too useful to be discarded.[2]

That given, one can formulate a level-k thinking analogue of the above concept. However, it is more involved, and may be useful only when the core model proposed above fails to sufficiently capture the economics of the situation.

[2]Note that some models dispense with it with a small number of categories. But within a category of agents (say entrepreneurs vs consumers), those models keep homogeneity as much as possible. This is preserved by the present model. On the other hand, for some purposes (e.g. predicting the amount of heterogeneity in a game), have a model that does necessarily predict heterogeneity is a good thing.

Game	Parameter	BR Theory	Empirical	Nash	Source for Empirical
p-beauty contest	$1/2\times$Mean	23.1	24.4	0	Nagel (2008)
	$2/3\times$Mean	33.3	36.7	0	idem
	$4/3\times$Mean	66.7	58.9	$100^{(a)}$	idem
	$0.7\times(\text{Median}+18)$	47.6	49.3	42	idem
	$0.8\times(\text{Median}+18)$	54.4	52.4	72	idem
Traveler's dilemma	$R=5$	187	180	80	Capra et al. (1999)
	$R=10$	178	177	80	idem
	$R=50$	145	155	80	idem
	$R=80$	80	120	80	idem
Matching Pennies	$x=320$	.91	.94	$.5^{(b)}$	Goeree and Holt (2001)
	$x=80$	.5	.48	.5	idem
	$x=40$	.06	.08	.5	idem

Table 1: Collection of experimental and theoretical results.

Notes: In all models, the cost $\kappa^m = 0.4$ was selected. (a) I select the Nash equilibrium that is robust to trembles. (b) The numbers are the probability that the row player plays Top.

Comparison to Other Models It is instructive to compare the predictions to those of other models. Level-k predicts $a(k) = \frac{1}{2}p^k$, hence if the distribution of levels is given by $f(k)$, the average response will be:

$$a^* = \frac{1}{2}\sum_{k=0}^{\infty} f(k)\, p^k. \tag{17}$$

This is not a particularly tractable quantity, except when $f(k)$ is Poisson with expected value τ (such that $f(k) = e^{-\tau}\tau^k/k!$), so $a^* = \frac{1}{2}e^{(p-1)\tau}$. In general though, it is quite arduous to derive $a(k)$ with paper and pencil.

The cognitive hierarchy model predicts:

$$a(0) = \frac{1}{2}, a(k+1) = p\frac{\sum_{i=0}^{k} f(i)\, a(i)}{\sum_{i=0}^{k} f(i)}.$$

Because there is no known closed form for $\sum_{i=0}^{k} f(i)$ for finite k, there is no known closed form for $a(k)$, nor for the average behavior $a^* = \sum_{k=0}^{\infty} a(k)\, f(k)$. I conclude that the cognitive hierarchy model is very useful for economics relying on computer assistance, but it is hard to yield closed-form solutions with it.

2.2.2 The Traveler's Dilemma

This game was invented by Basu (1994). The claims are a, b in the interval $[C, D] = [80, 200]$ (Capra et al. 1999). Both players submit a claim a, b. Call $h = \min(a, b)$ the lowest claim. If the players submit the same claim, they get h. Otherwise, the lowest claimant gets $h + R$ ($R > 0$ for a reward), and the highest claimant gets $h - R$ (a penalty for claiming too much). Formally,

$$u(a, b) = (a + R) 1_{a<b} + (b - R) 1_{a>b} + a 1_{a=b}. \tag{18}$$

There is just one Nash equilibrium: $a = b = C$, the minimum possible claim. This is because of unraveling: if other players play b, it is best to play $b - \varepsilon$ for a small $\varepsilon > 0$.

However, the experimental evidence runs counter to this when R is small: players submit claims close to D, perhaps in the hope that the other player does the same. This is what the BR algorithm gives.

Proposition 4 *In the traveler's dilemma, the BR bid is (using $\kappa^a = 0$): $a^* = C$ when $R > (D - C)/2$, and when $R < (C - D)/2$,*

$$a^* = \max\left(D - \frac{2R\sigma_a}{\kappa^m \left(\frac{D-C}{2} + 2R\right)}, C\right). \tag{19}$$

where $\sigma_a = (D - C)/\sqrt{12}$.

As expected, when $\kappa^m \to 0$, $a^* = C$, the Nash result. The model predicts that for a low reward R the play will indeed be close to the maximum, while for high reward it is at the minimum. The evidence, reported in Table 1, essentially supports that conjecture. The main misfit is that the model predicts that subjects will play the minimum value $a^* = C$ for R large enough, while in empirical data some subjects play above the minimum. This could be remedied by replacing $\kappa^a = 0$ (which was chosen for parsimony) by some $\kappa^a > 0$, which will lead subjects to react less extremely to incentives.

Capra et al. (1999) propose a model with logit noisy best response to explain the facts. At some broad level, the psychology is not too different. Again, the present model is deterministic, and has the advantage of analytical simplicity and good behavior in one-person games.

3 Some Games with a Discrete Action Space

Let us start with two actions. For an action space with two actions x and y ($A = \{x, y\}$), call p the probability of playing x while $1-p$ is the probability of play y. Step 1 becomes:

$$m_i = m_i^d + \tau \left(\mu_i - m_i^d, \kappa^m \frac{\left|V\left(x, m^d\right) - V\left(y, m^d\right)\right|}{\left|V_m\left(x, m^d\right) - V_m\left(y, m^d\right)\right|} \right). \tag{20}$$

Step 2 is simply a maximization on the probability of playing the first action $p(x)$:

$$\max_{p \in [0,1]} p\left[V(x, m) - V(y, m)\right] - \kappa^a D\left(p, p^d\right) \left|V_m\left(x, m^d\right) - V_m\left(y, m^d\right)\right|. \tag{21}$$

When $D\left(p, p^d\right) = \left|p - p^d\right|$, the problem is piecewise linear, so the solution must be equal to $0, p^d$, or 1. This generates in some games a too brutal response of p. A "soft max" is better than a full maximum. One parametrization is particularly tractable, the quadratic loss:

$$D^{Quad}(p, q) = (p - q)^2. \tag{22}$$

The responses are summarized in the following lemma:

Lemma 3 *(Maximization with Quadratic Penalty) Consider the maximization problem*

$$\max_{p \in [0,1]} pv - \frac{1}{2\beta}\left(p - p^d\right)^2$$

with $p^d \in [0, 1]$. *The solution is*

$$p = T_{[0,1]}\left(p_d + \beta v\right)$$

where $T_{[0,1]}(x) = \max(0, \min(1, x))$ *is the projection in* $[0, 1]$.

Proof. Define $p_0 = p^d + \beta v$. If the constraint $p \in [0, 1]$ does not bind, then the optimum is naturally $p = p_0$. Otherwise, because the objective function is quadratic, the optimum is the closest point to p_0 in $[0, 1]$, i.e., $T_{[0,1]}(p_0)$. ■

Another distance is the Kullback–Leibler divergence:

$$D^{Logit}(p, q) = p \ln \frac{p}{q}. \tag{23}$$

It yields a logit response function.

Lemma 4 *(Maximization with Logit Penalty) Consider the maximization problem*

$$\max_{p_1,\dots,p_n} \sum_i p_i v_i - \frac{1}{\beta} \sum_i p_i \ln \frac{p_i}{q_i}$$

subject to $\sum p_i = 1$ *and* $p_i \geq 0$. *The solution is*

$$p_i = \frac{q_i e^{\beta v_i}}{\sum_j q_j e^{\beta v_j}}.$$

Proof. The Lagrangian is:

$$L = \sum_i p_i v_i - \frac{1}{\beta} \sum_i p_i \ln \frac{p_i}{q_i} + \lambda \sum p_i.$$

The f.o.c. is:

$$0 = \frac{\partial L}{\partial p_i} = v_i - \frac{1}{\beta} (\ln p_i - \ln q_i + 1) - \lambda,$$

so

$$p_i = q_i e^{\beta v_i} / Q$$

for a constant $Q = e^{\beta \lambda + 1}$. The constraint $\sum p_i = 1$ yields the value of $Q = \sum_j q_j e^{\beta v_j}$, and allows to conclude. ■

For empirical work, the logit penalty is very useful (as in the quantal response equilibrium of McKelvey and Palfrey 1985). However, it is generally impossible to solve in closed form. The quadratic penalty (22) however brings a tractable way to model soft best responses.

3.1 Matching Pennies

Consider the game of matching pennies, with $x > 1$:

	L	R
U	$x, 0$	$0, 1$
D	$0, 1$	$1, 0$

(24)

Given $x > 1$, the "pull" of U is strong. Indeed, experimental subjects play U with a probability greater than 1/2. However, Nash equilibrium requires that the row player choose U and D with equal probability: this is to make the column player indifferent between L and R.

Let us apply the Sparse BR technology to this game. Call p the probability that the row player plays U, q the probability that the column player plays L, and $\widehat{p} = p - 1/2$, $\widehat{q} = p - 1/2$.

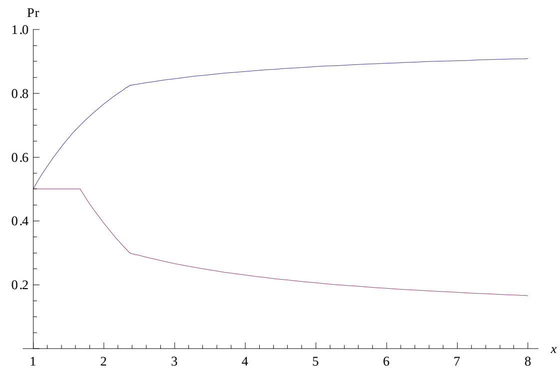

Figure 2: The top curve is the probability that the row player plays U, and the bottom curve is the probability that the column player plays L. Both are given as a function of x, the top left payment in payoff matrix (24).

The moduli are calculated to be $M_1 = (1+x)/2$ and $M_2 = 1$. Using the algorithm, we obtain:

$$\widehat{p}/\beta = \frac{v_1\left(q^d\right)}{M_1} + \tau\left(\frac{(1+x)\,\widehat{q}}{M_1}, \kappa\right) = \frac{x-1}{x+1} + \tau\left(2\widehat{q}, \kappa\right)$$
$$\widehat{q}/\beta = \tau\left(-2\widehat{p}, \kappa\right).$$

Proposition 5 *The solution of the matching pennies game is as follows. Define* $x_1 = \frac{2\beta+\kappa^m}{2\beta-\kappa^m}$ *and* $x_2 = \frac{4\beta^2+\kappa^m+2\beta\kappa^m}{4\beta^2-\kappa^m-2\beta\kappa^m} \geq x_1$. *The probability that the row player plays* U *is*

$$p(x) = \begin{cases} \frac{1}{2} + \beta\frac{x-1}{x+1} & \text{if } x \leq x_1 \\ \frac{1}{2} + \frac{\beta}{1+4\beta^2}\left(\frac{x-1}{x+1} + \kappa^m(2\beta+1)\right) & \text{if } x \geq x_1 \end{cases}$$

while the probability that the column player plays L *is*

$$q(x) = \begin{cases} \frac{1}{2} + \beta\frac{x-1}{x+1} & \text{if } x \leq x_2 \\ \frac{1}{2} - 2\beta p(x) & \text{if } x \geq x_2 \end{cases}.$$

It may be interesting to compare the predictions to those of the cognitive hierarchy model, as shown in Figure 3. We see that the predictions are somewhat similar. The Sparse BR model is, however, smoother (it is piecewise smooth) whereas the cognitive hierarchy model is more jagged. Technically, this is because the cognitive hierarchy model has a non-smooth best response. Also, the Sparse BR model is arguably more tractable as it obtains solutions in closed form.

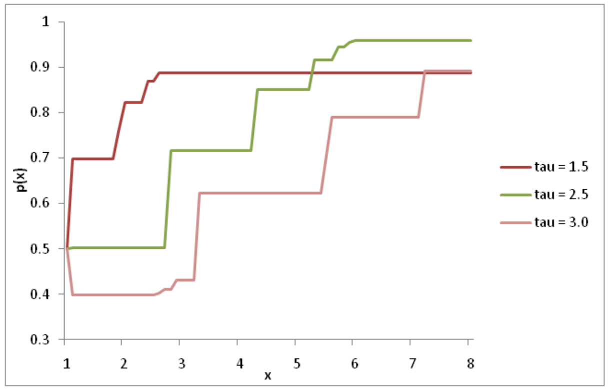

Figure 3: Predictions of the cognitive hierarchy model for the probability that the row player plays U as a function of x, the top left payment in payoff matrix (24). The predictions are shown for three values of τ, which is the average degree of thinking in the cognitive hierarchy model.

4 Other Game-theoretic Concepts: Bayesian Games and Sequential Games

This section is particularly preliminary.

4.1 Bayesian Games

Eyster and Rabin's (2005) notion of cursed equilibrium has proven to be a simple yet attractively versatile proposal for the BR modeling of Bayesian games. It turns out that, applied to Bayesian games, the Sparse BR model can generate cursed equilibrium. To avoid the somewhat heavy notations required here, I first sketch the argument verbally. Suppose the default MIM for the other players' actions is their average action given the payoff structure. Then, the Sparse BR model is just the cursed equilibrium with a parameter $\chi = 1 - m$. The Sparse BR model predicts some comparative statics on the cursing parameter on χ: basically, when the incentives to get things right are higher (subject to the average stakes in the particular games), then agents will be closer to the truth.

4.2 Sequential Games

Sequential games are done the same way, with a default MIM of the future players' actions. I develop two applications.

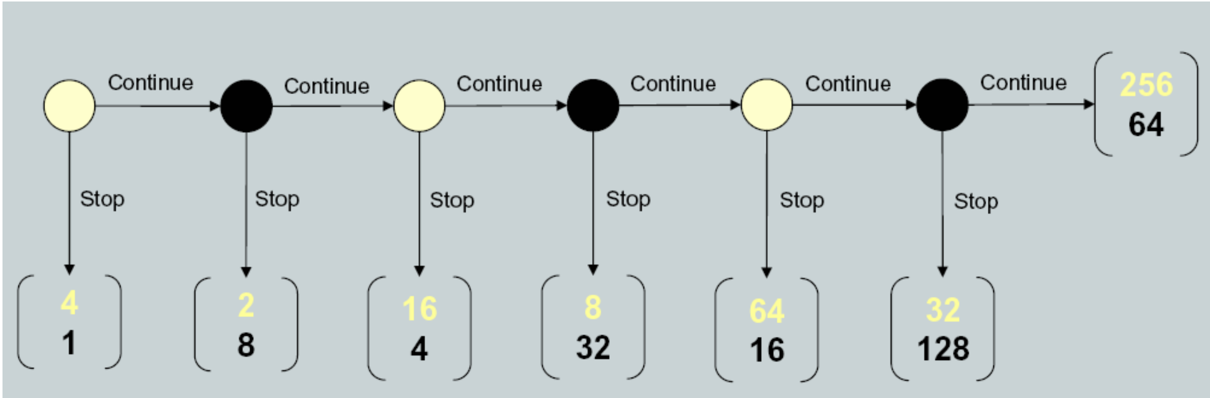

Figure 4: A centipede game. Reproduced from Palacios-Huerta and Volij (2009).

4.2.1 Centipede Game

Figure 4 shows a centipede game (Rosenthal 1981), reproduced from Palacios-Huerta and Volij (2009). The game-theoretic prediction is that Player 1 stops and gets 4, but in practice very few non-professional players do that.

We assume that the default action is "randomize at each node."[3] Under that scenario, Player 1 considers his payoff. A simple calculation shows that the expected payoff from continuing is 19 while the payoff from stopping is 4. Hence, Player 1 continues. Indeed, at all nodes, players continue.

This explanation of the centipede game shares similarities with that by McKelvey and Palfrey (1992). In their interpretation, the opponent can be selfish with some probability M, or altruist (and always continue) with some probability $1 - M$. Using this structure, their model still requires backward induction. However, in the present model, only forward induction is needed. Hence, the model is simpler to use in the centipede situation. Also, depending on the costs and benefits of thinking, the models will make different predictions (a future study, beyond the scope of this paper, would be required to do justice to this issue).

4.2.2 Dollar Auction Game

The dollar auction game (Shubik 1971) is an amusing and enlightening game. There is an amount of money to auction, say a $D = \$20$ bill. There are two players (for simplicity in this paper) who can participate in an ascending auction with increments of \$1. The person who drops out and bid b pays b. The person who stayed and bid $b' = b + 1$ gets the bill and pays b',

[3] We could have a more robust rule, e.g., "randomize at each node, but with a weight 0 on the nodes that can lead to a payoff less than M for some M."

for a total payoff of $D - b'$. The initial bid, if any, is at an arbitrary level.

In practice, one gets an initial bid $b < D$ (say $b = 1$), and then the game escalates without limit (or, until a limit imposed by the instructor is reached). The paradox is that it always seems better to go "one more step" even though the expected payoff is infinitely negative. That is not the equilibrium predicted by orthodox game theory, which is the following: someone bids $b = D$, and the other player does not bid.

The model delivers the classroom result. Suppose that the MIM is that people will continue with probability π and stop with probability $1 - \pi$ (as a baseline, $\pi = 1/2$). In the model, the right thing to do is to best-respond given those future events. To analyze this formally, consider $V(b)$, the expected payoff of continuing with a bid b. It is easy to show that it is equal to:[4]

$$V(b) = D - b - \frac{2\pi}{1-\pi}. \tag{25}$$

Thus, it is better to continue and bid $b + 2$ rather than give up iff $V(b+2) \geq -b$, i.e., iff $D \geq \frac{2\pi}{1-\pi}$. This is the case, for instance, if $D = \$20$ and $\pi = 1/2$. Indeed, for a wide range of parameters, agents will continue forever. They will stop only if they see that the probability of continuing is $\pi > D/(2+D)$, which is 0.91 when $D = \$20$.

4.3 Learning

The model is a model of players confronted with relatively novel situations. If the game is repeated, players might learn. The SparseBR framework provides a tempting way to model learning via the default action.

5 Conclusion

This paper presented a simple way to model boundedly rational behavior in games. It uses the tractable machinery developed in Gabaix (2011), and builds in spirit on earlier successful approaches (e.g., level-k thinking, cognitive hierarchy, quantal response equilibrium). It has

[4] We have the Bellman equation

$$V(b) = (1-\pi)(D-b) + \pi \max(V(b+2), -b-2)$$

Indeed, with probability $1-\pi$, the other player drops out, and the payoff is $D - b$. With probability π, he continues, and then the player chooses the best option between dropping out (which yields $-b-2$) or continuing (which yields $V(b+2)$). To solve the Bellman equation, we seek a solution of the type $V(b) = Ab + B$, and solve for A and B by plugging that solution in the Bellman equation. This yields (25).

several virtues that can make it appealing: (i) It is very tractable, and can be solved with paper and pencil, yielding closed-form solutions of various games. (ii) It is the n-player extension of a model of bounded rationality. Hence, it brings us closer to having a rather unified tractable model with boundedly rational features.

There are natural next steps. One is to model learning with the technology presented here. Another is using this tractable model in macroeconomic "games," so we can see when aggregate economic outcomes are robust to injecting (some form of) bounded rationality and when they are not.

References

Basu, Kaushik. "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory." *American Economic Review, Papers and Proceedings*, May 1994, 84(2), pp. 391-95

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong, "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics*, 119 (2004), 861-898.

Candès, Emmanuel, and Terence Tao, "Near-optimal signal recovery from random projections: universal encoding strategies?" *IEEE Transactions on Information Theory*, 52 (2006), 5406-5425.

Capra, Monica, Jacob K. Goeree, Rosario Gomez, Charles A. Holt, "Anomalous Behavior in a Traveler's Dilemma?" *American Economic Review*, Vol. 89, No. 3 (Jun., 1999), pp. 678-690

Carroll, Gabriel D., James Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick, "Optimal Defaults and Active Decisions," *Quarterly Journal of Economics*, 124 (2009), 1639-1674.

Charness, Gary, and Dan Levin, "The Origin of the Winner's Curse: A Laboratory Study," *American Economic Journal: Microeconomics*, 1 (2009), 207–236

Chetty, Raj, "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply, "NBER Working Paper # 15616, 2009.

Conlisk, John, "Why Bounded Rationality?" *Journal of Economic Literature*, 34 (1996), 669-700.

Crawford, Vincent, Miguel Costa-Gomes and Nagore Iriberri, "Strategic Thinking," Working Paper, Oxford, 2010.

Crawford, Vincent, and Nagore Iriberri, "Level-k Auctions: Can Boundedly Rational Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica,* 75 (2007), 1721-1770.

Dellavigna, Stefano, and Ulrike Malmendier, "Paying not to go to the gym, " *American Economic Review*, 96 (2006), 694-719.

Donoho, David, "Compressed Sensing," *IEEE Transactions on Information Theory*, 52 (2006), 1289-1306.

Esponda, Ignacio, "Behavioral equilibrium in economies with adverse selection," *American Economic Review*, 98 (2008), 1269-91.

Eyster, Erik, and Matthew Rabin, "Cursed Equilibrium," *Econometrica*, 73 (2005), 1623-1672.

Eyster, Erik, and Georg von Weizecker, "Correlation Neglect in Financial Decision-Making," Working Paper, London School of Economics, 2010.

Gabaix, Xavier, 2011. “A Sparsity-Based Model of Bounded Rationality,” Working Paper, NYU.

Goeree, Jacob K. and Charles A. Holt (2001), "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions," *American Economic Review*, 91(5), 1402-1422.

Jéhiel, Philippe, “Analogy-Based Expectation Equilibrium,” *Journal of Economic Theory,* 123 (2005), 81–10.

Kahneman, Daniel, “Maps of Bounded Rationality: Psychology for Behavioral Economics,” *American Economic Review*, 93 (2003), 1449-1475.

Kahneman, Daniel, and Amos Tversky, “Prospect Theory: An Analysis of Decision under Risk” *Econometrica*, 47 (1979), 263-292.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler, “Experimental Tests of the Endowment Effect and the Coase Theorem,” *Journal of Political Economy*, 97 (1990), 1325-1348.

Kawagoe, Toshiji, and Hirokazu Takizawa. 2008. “Level-K Analysis of Experimental Centipede Games.” http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1289514

Koszegi, Botond, and Matthew Rabin, “A Model of Reference-Dependent Preferences,” *Quarterly Journal of Economics*, 121 (2006), 1133-1166.

Laibson, David, “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics* 62 (1997), 443-77.

McKelvey, Richard; Palfrey, Thomas (1995), "Quantal Response Equilibria for Normal Form Games", Games and Economic Behavior 10: 6–38.

Mullainathan, Sendhil, “A Memory-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 117 (2002), 735-774.

Nagel, Rosemarie. 1995. “Unraveling in Guessing Games: An Experimental Study.” *American Economic Review*, 85(5): 1313-1326.

Nagel, Rosemarie, 2008 “Experimental Beauty Contest Games: Levels of reasoning and convergence to equilibrium”, Chapter 45 in Plott and Smith (2008).

Palacios-Huerta, Ignacio, and Oscar Volij, “Field Centipedes,” *American Economic Review*, 99 (2009), 1619-1635.

Plott, Charles and Vernon Smith, *Handbook of Experimental Economics Results*, Vol. 1, 2008, Elsevier.

Rabin, Matthew, “Inference by Believers in the Law of Small Numbers,” *Quarterly Journal* *of Economics*, 117 (2002), 775–816.

Rabin, Matthew, and Joel L. Schrag, "First Impressions Matter: A Model of Confirmatory Bias," *Quarterly Journal of Economics*, 114 (1999), 37–82.

Radner, Roy, and Timothy Van Zandt, Real-Time Decentralized Information Processing and Returns to Scale," *Economic Theory,* 17 (2001), 545-575.

Rosenthal, Robert W., "Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox," *Journal of Economic Theory*, 25 (1981), 92-100.

Rubinstein, Ariel, *Modeling Bounded Rationality* (Cambridge: MIT Press, 1998).

Samuelson, William F., and Max H. Bazerman, "The Winner's Curse in Bilateral Negotiations," in *Research in Experimental Economics: A Research Annual vol 3*, Vernon Smith, ed. (Greenwich: JAI Press, 1985).

Samuelson, William F., and Richard Zeckhauser, "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, 1 (1988), 7-59.

Sargent, Thomas, *Bounded rationality in macroeconomics* (Oxford: Oxford University Press, 1993).

Shubik, Martin, "The Dollar Auction Game: A Paradox in Noncooperative Behavior and Escalation," *The Journal of Conflict Resolution*, 15 (1971), 109-111.

Sims, Christopher, "Implications of Rational Inattention," *Journal of Monetary Economics*, 50 (2003), 665–690.

Stahl, Dale O., and Paul R. Wilson, "On Players' Models of Other Players: Theory and Experimental Evidence." *Games and Economic Behavior*, 1995 (10), 218-254

Thaler, Richard H., "Mental Accounting and Consumer Choice," *Marketing Science*, 4 (1985), 199-214.

Tibshirani, Robert, "Regression shrinkage and selection via the lasso." *Journal of the Royal Statistical Society B*, 58 (1996), 267–288.

Tversky, Amos, and Daniel Kahneman, "Judgment under uncertainty: Heuristics and biases," *Science*, 185 (1974), 1124–1130.

Tversky, Amos, and Daniel Kahneman. "The Framing of Decisions and the Psychology of Choice," *Science*, 211 (1981), 453-458.

6 Appendix: Additional derivations

Second part of the proof of Proposition 3 When $p > 1$, as in some experimental games, then, using the notation $x \wedge y = \min(x, y)$,

$$a^* = \phi(a^*) \tag{26}$$

$$\phi(a) := 1 \wedge p(a(m)) = 1 \wedge p\left(\frac{1}{2} + \left(a - \frac{1}{2}\right)\tau\left(1, \frac{\kappa^m \sigma_a}{p\left(\frac{1}{2} - a\right)}\right)\right)$$

$$\phi(a) = 1 \wedge \left(\frac{p}{2} + \tau\left(p\left(a - \frac{1}{2}\right), \kappa^m \sigma_a\right)\right) \tag{27}$$

If $\tau\left(p\left(a - \frac{1}{2}\right), \kappa^m \sigma_a\right) = 0$, then $a = p/2$. This is possible if:

$$\kappa^m \sigma_a \geq p\left(\frac{p}{2} - \frac{1}{2}\right)$$

i.e.

$$G(p) \equiv p^2 - p - 2\kappa^m \sigma_a \leq 0$$

i.e $p \leq p^*$, where p^* is the positive root the $G(p) = 0$:

$$p^* = \frac{1 + \sqrt{1 + 8\kappa^m \sigma_a}}{2}$$

If $p > p^*$, there is an unstable solution a_1 of $a^* = \phi(a^*)$ (unstable because $\phi'(a_1^*) = p > 1$), and a stable one: $a^* = 1$. Indeed,

$$\begin{aligned}\phi(1) &= 1 \wedge \left(\frac{p}{2} + \tau\left(\frac{p}{2}, \kappa^m \sigma_a\right)\right) \\ &= 1 \wedge (p - \kappa^m \sigma_a)\end{aligned}$$

so $\phi(1) = 1$ if $p \geq p_2 \equiv 1 + \kappa^m \sigma_a$. Given that

$$G(p_2) = \kappa^m \sigma_a (1 - \kappa^m \sigma_a)$$

If $\kappa^m \sigma_a \leq 1$. Then, $G(p_2) \geq 0$, and we have $p_2 \leq p^*$. Hence, if $p > p^*$, we have $p \geq p_2$ hence $\phi(1) = 1$.

If $\kappa^m \sigma_a > 1$, then $p^* > \frac{1+\sqrt{1+8}}{2} = 2$, and then for $p > p^*$, $p/2 > 1$ and $\phi(1) = 1$.

In both cases, $\phi(1) = 1$, so that the stable solution of $a^* = \phi(a^*)$ is $a^* = 1$.

7 Appendix: Technical Complements

7.1 Wasserstein Distance

To talk about distances between probabilities on the action space A, the structure on A matters a great deal. There are two polar cases. One case is the traditional game-theoretic one where the action space is "amorphous," there is no particular order or metric between actions. Another case is that of a Euclidean action space $A \subset \mathbb{R}$. Then, for two actions x and y, we wish to use the distance $|x-y|$.

Now we generalize that to probabilities p and q with a distance between p and q, $D(p,q)$. We want that when $p = \delta_x$ (i.e., a probability 1 of playing x) and $q = \delta_y$ for some points x, y, the distance $D(p,q)$ is just the ordinary distance $|x-y|$:

$$D(\delta_x, \delta_y) = |x-y|. \tag{28}$$

A convenient notion that has this property is the Wasserstein distance (see Villani 2003), defined as

$$W_\alpha(p,q) = \inf_\pi \left(\int d(x,y)^\alpha \pi(x,y) \right)^{1/\alpha} \text{ subject to}$$
$$\int \pi(x,y)\, dy = p(x) \text{ and } \int \pi(x,y)\, dx = q(y).$$

The typical values are $\alpha = 1$ (then it is also called the Kantorovich-Rubinstein distance) and $\alpha = 2$ (quadratic distance). This distance is the "minimum transport cost" to move from distribution p to distribution q. It has many nice properties (e.g., it metrizes weak convergence), and has been the topic of great recent interest (Villani 2003). It can be explicitly solved in some cases of interest.

I propose to use that distance even in the "amorphous" game-theoretic case. Then, formally $d(x,y) = 1_{x \neq y}$, and one can show that the $W_\alpha(p,q)$ is just the total variation between distributions p and q:

$$TV(p,q) = \frac{1}{2}\sum_x |p(x) - q(x)|. \tag{29}$$

In what follows, I will take the distance between distributions $D = W_\alpha$ with $\alpha = 1$.

One might think of always using $TV(p,q)$ as a distance. However, that would violate the desideratum that the distance be $|x-y|$ when the probabilities are masses at x and y (equation 28). The Wasserstein distance, however, does satisfy that desideratum.