The effect of loss aversion on congestion games

Guy Incognito

April 24, 2008

Abstract

Expected utility maximization is not a good predictive model of how people value outcomes. One of the ways in which people systematically deviate from expected utility maximization is loss aversion, where losses (relative to some reference point) are more highly weighted than gains. I investigate whether and how the theoretical guarantees of congestion games change when the agents are assumed to be loss averse. I also introduce and investigate the effects of assuming loss averse agents on a stochastic version of congestion games.

1 Introduction

Game theoretic models generally assume that all agents are perfectly rational expected utility maximizers who care only about their utilities for final outcomes. Experimental results from behavioural game theory do not bear this assumption out [3, 2, 10, 5]. Empirically, human agents exhibit a number of systematic deviations from rational utility maximization. Two of the most important such deviations are non-linear probability weighting and loss aversion.

Congestion games are a widely-studied class of game that can be used for modelling traffic flows (both vehicular and network), factory usage, and other situations in which the cost of a resource is primarily depends upon the number of agents who use it [7, 1]. Congestion games have a number of properties that make them attractive models. They always possess at least one pure strategy Nash equilibrium, which is often considered to be a more convincing solution concept than a mixed strategy equilibrium. In addition, a simple procedure in which agents iteratively best-respond to the strategy profile of the other agents always converges to a pure-strategy Nash equilibrium in a finite number of steps.

This paper investigates whether adding the assumption that agents maximize loss averse valuations rather than expected utility of outcomes significantly changes any of the desirable properties of congestion games. I consider whether a pure strategy equilibrium is still guaranteed to exist, whether the iterative procedure still terminates, and whether the social costs of an equilibrium of loss averse agents differs significantly from a "rational" equilibrium.

Standard congestion games do not involve any risk or uncertainty about payoffs given action profiles. Since one of the main places that loss aversion manifests itself is in situations involving risk, I also consider a stochastic version of congestion games where the costs associated with a resource may be a distribution rather than a deterministic value.

In section 2, I provide some background, describing the model of loss aversion used in this paper, standard congestion games, and a stochastic extension of congestion games. I describe the properties that are preserved or lost by assuming loss averse agents in standard and stochastic congestion games in section 3. I describe some possible future work and sum up in section 4.

2 Background

2.1 Loss aversion

Expected utility maximizing agents care only about their utilities for final outcomes, weighted by the probability of each outcome. Although Von Neumann and Morgenstern [11] showed that this is in some sense normative, behavioural studies show that it is not a good model of how people actually value outcomes [5, 10, 2]. Prospect theory [5, 10] models some of the systematic observed deviations from expected utility maximization in two ways: non-linear decision weights and reference-dependent valuations. This paper does not consider non-linear decision weights.

Reference-dependent valuations account for the fact that people do not value outcomes in their own right, but rather value them relative to a *reference point*. The reference point is usually interpreted as the agent's valuation of the status quo, so any outcome that is better than the reference point is experienced as a gain, and any outcome that is worse than the reference point is experienced as a loss. Experimental studies on the "endowment effect" show that this reference point can be updated extremely quickly for a new status quo [4].

It has been consistently shown that people tend to be *loss averse* [5, 4, 10]; that is, they are more sensitive to losses than to gains. So for example, if an agent has a reference point of 10, then an outcome of 5 and an outcome of 11 will have a greater difference in value than an outcome of 12 and an outcome of 18, because the former is a comparison of a loss of 5 to a gain of 1, while the latter is a comparison of a gain of 2 to a gain of 8.

2.1.1 Model

I primarily follow the model of Shalev [9] for loss aversion and equilibrium in the presence of loss aversion. Given a game G, an *extended game* is a pair $(G, (\lambda_i)_{i \in N})$, where each agent i has an associated *loss aversion coefficient* $\lambda_i \in \mathbb{R}_+$. The loss aversion coefficient indicates the degree to which an agent is loss averse. A value of $\lambda_i = 0$ specifies that i is an expected utility maximizer (i.e., his utility does not depend on a reference point).

The *basic value* of an outcome is the utility of the outcome when the reference point is equal to the outcome (i.e., how much an agent values the outcome when that outcome is the status quo). The final utility (taking loss aversion into account) for an agent i with a reference point of κ_i of an outcome with a basic value x is defined as

$$v_i(x, \kappa_i) = \begin{cases} x & \text{if } x \geq \kappa_i \\ x - \lambda_i(\kappa_i - x) & \text{if } x < \kappa_i. \end{cases} \tag{1}$$

The final utility for an agent i for a lottery over outcomes with base utilities x_j and probabilities p_j, relative to a reference point κ_i is

$$w_i([p_1 : x_1, \ldots, p_k : x_k], \kappa_i) = \sum_{j=1}^{k} p_j v(x_j, \kappa_i), \tag{2}$$

or more generally for a distribution p over base utilities,

$$w_i(p, \kappa_i) = \int_{-\infty}^{\infty} p(x) v_i(x, \kappa_i) dx. \tag{3}$$

For distributions over base utilities, a reference point is *consistent* with the distribution if the value of the distribution relative to that reference point equals the reference point. So a reference point κ_i is consistent with distribution p if

$$w_i(p, \kappa_i) = \kappa_i.$$

$\kappa_i(p)$ denotes the reference point that is consistent with distribution p. This value exists and is unique for all p [9]. Define by $\kappa(\sigma)$ the reference point that is consistent with the distribution over outcomes induced by strategy profile σ.

The analogous concept to the Nash equilibrium for an extended game is the *loss aversion equilibrium.* I consider only non-myopic loss aversion equilibria, since they are equivalent to myopic loss aversion equilibria for simultaneous games [9]. A strategy profile σ is a non-myopic loss aversion equilibrium if and only if for all possible strategies σ_i' and all $i \in N$,

$$\kappa_i(\sigma) \geq \kappa_i(\sigma_i', \sigma_{-i}).$$

In a loss aversion equilibrium, as in a Nash equilibrium, no agent can improve his utility by switching to a different strategy. Strategy profiles are evaluated by comparing their consistent reference points, because the non-myopic agent is aware that his reference point will shift to the new strategy profile's consistent reference point if the strategy profile changes, and takes that into account.

2.2 Standard congestion games

A standard congestion game is a non-cooperative simultaneous-move game in which n agents choose sets of resources to use [7]. Each resource has a cost associated with it that depends solely on how many agents chose it. Each agent has the same utility function, which simply sums the costs of the resources that the agent has chosen.

Congestion games can be used to model a variety of situations where resource costs are load-dependent: traffic flows, where the resources are roads or portions of routes; communication networks, where the resources are individual network links; and production problems, where the resources are facilities [7, 1].

Formally, a congestion game is a tuple (N, R, A, c), where N is a set of n agents, R is a set of k resources, $A = A_1 \times \cdots \times A_n$ is a set of action profiles, and $c = (c_1, \ldots, c_k)$ is a tuple of cost functions for each resource. For each agent i, $R_i \subseteq \mathcal{P}(R)\backslash\emptyset$ consists of the set of feasible resource choices.[1] Each cost function c_j depends only upon $\#(a, j)$, the number of agents using resource $j \in R$ in action profile a. For each agent i, the utility for an action profile $a = (a_1, \ldots, a_n) \in A$ is

$$u_i(a) = \sum_{j \in a_i} c_j(\#(a, j)). \tag{4}$$

Note that each cost function is assumed to return a non-positive number.

Part of what makes congestion games so attractive for modelling is that they are guaranteed to always possess a pure strategy Nash equilibrium [7]. Furthermore, since they are potential games, the procedure in figure 1 is guaranteed to both terminate and find a pure strategy Nash equilibrium [6]:

input : Game G, action profile a
output: Pure strategy Nash equilibrium a
while $\exists i \in N$ *and* $a'_i \in A_i$ *such that* $u_i(a'_i, a_{-i}) > u_i(a)$ **do**
  $a \leftarrow (a'_i, a_{-i})$
end
return a

Figure 1: The MyopicBestResponse procedure

Clearly, when MyopicBestResponse terminates, a describes a pure strategy Nash equilibrium. And since each step of the procedure strictly increases the potential function, and there are only a finite number of possible action profiles, the procedure will terminate.

2.3 Stochastic congestion games

A *stochastic congestion game* is identical to a standard congestion game, except that each cost function is now a mapping $c_j : \mathbb{N} \mapsto \Pi(\mathbb{R})$ from the number of agents using resource j to a probability distribution over costs, rather than to a fixed cost. The intuition behind this extension is that the cost of a resource may have more or less variance when it is under high load than when it is under low load. For example, a busy freeway may have best-case and worst-case travel times that significantly differ (e.g. depending upon whether or not an accident occurs), whereas an empty freeway may have a very low variance in travel times.

[1] For a traffic or networking model, a feasible resource choice might consist of a path through the routing graph. For a production problem, a feasible resource choice might consist of a process that requires certain inputs, or a factory that will be chosen for ordering parts from.

The utility function for each agent is changed to take a distribution over costs rather than a single cost. The utility function's argument is the distribution over total costs induced by the cost distributions of each resource. Formally,

$$u_i(a) = w_i(p_{a,i}, \kappa_i(p_{a,i})), \tag{5}$$

where $p_{a,i}$ is a distribution over costs defined as follows. For some ordering $r_1, r_2, \ldots, r_\ell$ of resources in a_i, let

$$\begin{aligned} g^j &= c_{r_j}(\#(a, r_j)) && \text{for } 1 \leq j \leq \ell \\ p_{a,i}^{(1)}(x) &= g^1(x) \\ p_{a,i}^{(j)}(x) &= \int_{-\infty}^{\infty} p_{a,i}^{(j-1)}(y) g^j(x-y) dy && \text{for } 1 < j \leq \ell \\ p_{a,i}(x) &= p_{a,i}^{\ell}(x). \end{aligned}$$

Implicit in this definition is the assumption that the cost distributions of the resources are mutually independent.

Note that when the agents are expected utility maximizers, a stochastic congestion game is equivalent to a standard congestion game where each cost distribution is replaced by its expected value. However, this is not true when agents are loss averse, because the consistent reference point for a distribution is not in general the expected value of the distribution, and the consistent reference point for a sum of distributions is not in general the sum of the distributions' consistent reference points.

3 Results

3.1 Standard congestion games with heterogeneously loss averse agents

I first show that any pure-strategy Nash equilibrium of a standard congestion game is also a pure-strategy non-myopic loss aversion equilibrium of the extended game with heterogeneously loss averse agents.

Proposition 1. *Let $a^* \in A$ be a pure-strategy Nash equilibrium of a standard congestion game G. Then a^* is also a non-myopic loss averse equilibrium of the extended game $(G, (\lambda_i)_{i \in N})$.*

Proof. Since the costs to each agent i are deterministic in a standard congestion game, $\kappa_i(a)$ in $(G, (\lambda_i)_{i \in N})$ equals $u_i(a)$ in G for all agents i and action profiles a. So under a^* each agent i has a reference point $\kappa_i(a^*)$ that is equal to $u_i(a^*)$ in G. Now assume that a^* is not a non-myopic loss aversion equilibrium. That means that there exists some action a_i' such that $\kappa_i(a_i', a_{-i}^*) > \kappa_i(a_i^*, a_{-i}^*)$. But that implies that $u_i(a_i', a_{-i}^*) > u_i(a_i^*, a_{-i}^*)$ in G, contradicting a^*'s being a Nash equilibrium. Hence there is no such a_i', and a^* is a non-myopic loss aversion equilibrium. $\square$

Since all standard congestion games G have at least one pure-strategy Nash equilibrium, this means that all extended standard congestion games $(G, (\lambda_i)_{i\in N})$ have at least one pure-strategy loss aversion equilibrium.

Next I consider two versions of MyopicBestResponse for operating on extended standard congestion games. In the first version, myopic LA-MyopicBestResponse, each potential new action profile is considered relative to the consistent reference point for the current action profile.

input : Extended game $(G, (\lambda_i)_{i\in N})$, action profile a
output: Loss aversion equilibrium a
while $\exists i \in N$ *and* $a'_i \in A_i$ *such that* $w_i((a'_i, a_{-i}), \kappa_i(a)) > w_i(a, \kappa_i(a))$ **do**
  $a \leftarrow (a'_i, a_{-i})$
end
return a

Figure 2: The myopic LA-MyopicBestResponse procedure

In the second version, non-myopic LA-MyopicBestResponse, each potential new action profile is considered relative to its own consistent reference point. This reflects more sophisticated agents who are aware that their reference points will change if the action profile changes.

input : Extended game $(G, (\lambda_i)_{i\in N})$, action profile a
output: Loss aversion equilibrium a
while $\exists i \in N$ *and* $a'_i \in A_i$ *such that*
$w_i((a'_i, a_{-i}), \kappa_i(a'_i, a_{-i})) > w_i(a, \kappa_i(a))$ **do**
  $a \leftarrow (a'_i, a_{-i})$
end
return a

Figure 3: The non-myopic LA-MyopicBestResponse procedure

It turns out that both of these procedures are guaranteed to terminate and return a loss aversion equilibrium when applied to an extended standard congestion game.

Lemma 2. *If $v_i(x, \kappa) > \kappa$ for some $x, \kappa \in \mathbb{R}$, then $x > \kappa$.*

Proof. Assume the converse, that $v_i(x, \kappa) > \kappa$ but $x \le \kappa$. Then from equation (1),

$$\begin{aligned} & v_i(x,\kappa) > \kappa \\ \iff & x - \lambda_i(\kappa - x) > \kappa \\ \iff & x > \kappa, \end{aligned}$$

a contradiction. □

Proposition 3. *Let G be a standard congestion game, and $(G, (\lambda_i)_{i \in N})$ be an extended game. Then both myopic* LA-MyopicBestResponse *and non-myopic* LA-MyopicBestResponse *terminate and return a loss aversion equilibrium.*

Proof. The payoffs are deterministic for each action profile, so $\kappa_i(a'_i, a_{-i}) > \kappa_i(a)$ is equivalent to $u_i(a'_i, a_{-i}) > u_i(a)$ in G, so the loop proceeds only so long as there exists a'_i such that $u_i(a'_i, a_{-i}) > u_i(a)$ in G. But these are just the same conditions under which MyopicBestResponse proceeds in G. Since we know that MyopicBestResponse terminates in G, that means that non-myopic LA-MyopicBestResponse also terminates in the extended version of G. Clearly a is a non-myopic loss aversion equilibrium at termination.

Similarly, if $w_i((a'_i, a_{-i}), \kappa_i(a)) > w_i(a, \kappa_i(a))$, then by lemma 2 and the fact that each action profile has a deterministic value, $\kappa((a'_i, a_{-i})) > \kappa(a)$, which again means $u_i(a'_i, a_{-i}) > u_i(a)$ in G. So the loop in myopic LA-MyopicBestResponse also proceeds in the extended version of G only when the MyopicBestResponse loop would proceed in G, and hence myopic LA-MyopicBestResponse always terminates in the extended version of G. Since $\kappa(a) \geq \kappa((a'_i, a_{-i}))$ for all $a'_i \in A_i$ at termination, a is a non-myopic loss aversion equilibrium at termination. □

3.2 Stochastic congestion games with heterogeneously loss averse agents

Unfortunately, extended stochastic congestion games do not in general have the desirable properties of extended standard congestion games.

Example 1. Consider the following extended stochastic congestion game:

$$\begin{aligned} N &= \{1,2\} & c_1(1) = c_2(1) &= [-110 : 0.5, -90 : 0.5] & \lambda_1 &= 0 \\ R &= \{1,2\} & c_1(2) = c_2(1) &= [-130 : 0.5, -50 : 0.5] & \lambda_2 &= 2 \\ A_1 = A_2 &= \{\{1\},\{2\}\}. \end{aligned}$$

Agent 1 is an expected utility maximizer, agent 2 is loss averse, and they each choose one of two resources. The two resources have the same cost function. When one agent uses a resource, its cost has a relatively small standard deviation (10) and an expected value of 100. When two agents use a resource, it has a smaller expected cost (90) but a larger standard deviation (40).[2]

Now note that for agent 2, the cost distribution $[-110 : 0.5, -90 : 0.5]$ has a consistent reference point of -110 (i.e., $w_2([-110 : 0.5, -90 : 0.5], -110) = -110$), and the cost distribution $[-130 : 0.5, -50 : 0.5]$ has a consistent reference point of -105 (i.e., $w_2([-130 : 0.5, -50 : 0.5], -105) = -105$). Agent 1's consistent reference points are just the expected values of the cost distributions. So

$$\kappa_1(\{1\},\{1\}) = \kappa_1(\{2\},\{2\}) > \kappa_1(\{1\},\{2\}) = \kappa_1(\{2\},\{1\}),$$

[2]Note that cost functions are not required to be increasing in the number of agents using a resource. This lower expected cost under higher load could arise through economies of scale.

but

$$\kappa_2(\{1\},\{1\}) = \kappa_2(\{2\},\{2\}) < \kappa_2(\{1\},\{2\}) = \kappa_2(\{2\},\{1\}),$$

and therefore there is no pure-strategy non-myopic loss aversion equilibrium for this game.

We can see this in another way by considering the normal-form game in which each agent's payoff is his consistent reference point for the cost distribution induced by the action profile. Table 1 contains this induced normal-form game.

	{1}	{2}
{1}	(-90,-110)	(-100,-105)
{2}	(-100,-105)	(-90,-110)

Table 1: Normal form of $((N, R, A, c), (\lambda_1, \lambda_2))$

This game has no pure-strategy Nash equilibrium, and therefore there is no non-myopic loss aversion equilibrium in the original game.

Since there is no pure strategy loss aversion equilibrium in this example, it is also clear that neither myopic nor non-myopic LA-MyopicBestResponse will terminate in general for extended stochastic congestion games.

3.2.1 Social cost

Define the social cost of an action profile as the expected cost of the induced distribution over total costs.[3] Then the game of example 1, when both players are rational expected utility maximizers, would have equilibrium $(\{1\},\{1\})$ or $(\{2\},\{2\})$, with a social cost of 180. On the other hand, if example 1 were played by loss averse agents with $\lambda_1 = \lambda_2 = 2$, then the equilibrium would be either $(\{1\},\{2\})$ or $(\{2\},\{1\})$, with a social cost of 200. Hence it is possible for an equilibrium of loss averse agents to have a worse social cost than the equilibrium that expected utility maximizers would reach in the same game.

The following example shows that it is also possible for an equilibrium of loss averse agents to have a better social cost than an equilibrium of expected utility maximizers.

Example 2. Consider the following extended stochastic congestion game:

$$\begin{aligned} N &= \{1,2\} & c_1(1) &= -75 & c_2(1) &= [-45 : 0.5, -5 : 0.5] \\ R &= \{1,2\} & c_1(2) &= -30 & c_2(2) &= -50 \\ A_1 = A_2 &= \{\{1\},\{2\}\}. \end{aligned}$$

When $\lambda_1 = \lambda_2 = 0$, the induced normal form for this game is shown in table 2, with a unique equilibrium of $(\{2\},\{2\})$, with a social cost of 100.

However, when $\lambda_1 = \lambda_2 = 2$, the induced normal form is as shown in table 3, with an additional equilibrium of $(\{1\},\{1\})$, for a better social cost of 60.

[3] This is fair even when the agents themselves are loss averse, because it allows us to average over the outcomes of the various agents.

	{1}	{2}
{1}	(-30,-30)	(-75,-25)
{2}	(-25,-75)	(-50,-50)

Table 2: Normal form when agents are both expected utility maximizers

	{1}	{2}
{1}	(-30,-30)	(-75,-35)
{2}	(-35,-75)	(-50,-50)

Table 3: Normal form when agents are both loss averse with $\lambda = 2$

So it is possible for an equilibrium of loss averse agents to be have either better (as in example 1) or worse (as in example 2) social cost than an equilibrium that would result from expected utility maximizers playing the same game.

4 Conclusions and future work

Given the experimental and empirical evidence that people are loss averse rather than being expected utility maximizers, it is natural to ask whether and how incorporating loss aversion into standard models changes their theoretical guarantees. This paper examined the impact of incorporating loss aversion into standard congestion games. It also examined the impact of loss aversion on a generalization of congestion games that assigns each resource a distribution over costs rather than a fixed cost based on its load.

The extended version of every standard congestion game that assigns a loss aversion coefficient to each agent always has a pure strategy loss aversion equilibrium. Furthermore, both myopic and non-myopic LA-MyopicBestResponse are guaranteed to terminate and return a pure strategy loss aversion equilibrium. In fact, this pure strategy loss aversion equilibrium will be exactly the same as the pure strategy Nash equilibrium of the underlying game, and will therefore have the same social cost.

Unfortunately, the same is not true of the extended version of stochastic congestion games. These games are not guaranteed to have any pure strategy loss aversion equilibrium. Of course that also means that the simple LA-MyopicBestResponse procedure is not guaranteed to terminate. The social cost of a loss aversion equilibrium in an extended stochastic congestion game may be either higher, lower, or identical to the social cost of a Nash equilibrium of the same underlying game played by expected utility maximizing agents.

The fact that adding stochastic costs and loss averse agents can cause the loss of pure strategy equilibria suggests that modelling situations where costs do vary stochastically by congestion games may produce misleading results.

There are many potential avenues of future research in this area. I considered the effects of loss aversion given linear decision weights. One possible extension would be to investigate the effects of nonlinear decision weights [5, 10] on the properties of congestion games (and other games), both with and without the presence of loss aversion. It could also be valuable to investigate the effects of loss aversion on the nonatomic congestion games of Schmeidler [8].

Finally, although we know that stochastic congestion games always have pure-strategy equilibria when agents are homogeneously loss neutral, and although I showed that assuming heterogeneous loss aversion in stochastic congestion games causes this guarantee of a pure strategy equilibrium to be lost, I did not consider the case where agents are homogeneously loss averse (i.e., where every agent has the same nonzero loss aversion coefficient). Models where every agent is loss averse, but to the same degree, could potentially provide a better balance between expressive power and tractability.

References

[1] E. Altman and L. Wynter. Equilibrium, Games, and Pricing in Transportation and Telecommunication Networks. *Networks and Spatial Economics*, 4(1):7–21, 2004.

[2] C. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press Princeton:, 2003.

[3] C.F. Camerer. Prospect theory in the wild: Evidence from the field. *Choices, values, and frames*, pages 288–300, 2000.

[4] D. Kahneman, J.L. Knetsch, and R.H. Thaler. Experimental Tests of the Endowment Effect and the Coase Theorem. *Journal of Political Economy*, 98(6):1325, 1990.

[5] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

[6] D. Monderer and L.S. Shapley. Potential Games. *Games and Economic Behavior*, 14(1):124–143, 1996.

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

[8] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, 1973.

[9] J. Shalev. Loss aversion equilibrium. *International Journal of Game Theory*, 29(2):269–287, 2000.

[10] A. Tversky and D. Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992.

[11] J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Wiley, 1944.