Introduction: The Rules of the Game

Objective: Understand the distinguishing characteristics of a game.

Games are situations where an individual's welfare depends on the actions of others and an individual's actions affect the welfare of others. Furthermore, individuals know this is the case. What this all means is that individuals must consider what others are likely to do before they can decide what is best. Individuals must also think about what others think they are likely to do. This is contrary to the model of perfect competition and many other economic models where the decisions of others are simply treated as exogenous.

Objective: Understand the common elements of a game.

For all games we will need to answer some basic questions about the rules of the game:

a) Who are the players?
b) Who can do what when?
c) Who knows what when?
d) How are players rewarded based on what they do?

The answers to these questions constitute the rules of the game and are **common knowledge** among all players.

To illustrate, let us begin with the simple game of matching pennies. The game has two players, *Mason* and *Spencer*. Players simultaneously choose either *Heads* or *Tails*. *Mason* wins when both make the same choice. *Spencer* wins when both make different choices. The winner pays the loser $1, so this is an example of a special class of games called zero sum games.

Who are the players?

Mason and *Spencer*.

Who can do what when?

Mason can choose *Heads* or *Tails*.
Spencer can choose *Heads* or *Tails*.
Mason and *Spencer* make their choices at the same time or simultaneously.

Who knows what when?

Spencer does not know *Mason*'s choice when he makes his.
Mason does not know *Spencer*'s choice when he makes his.

How are players rewarded based on what they do?

Mason earns $1 and *Spencer* loses $1 if both choose *Heads* or both choose *Tails*.
Mason loses $1 and *Spencer* earns $1 if one chooses *Heads* and the other chooses *Tails*.

Mason likes a slightly different variation of this game that takes advantage of *Spencer*'s youth and inexperience. Instead of choosing simultaneously, he likes *Spencer* to go first so he can see what *Spencer* has done before making his own choice. With this variation some of the answers to our questions and the rules of the game change.

Who are the players?
    Mason and *Spencer*.

Who can do what and when?
    Spencer can choose *Heads* or *Tails*.
    After *Spencer* chooses, *Mason* can choose *Heads* or *Tails*.

Who knows what and when?
    Spencer does not know *Mason*'s choice when he makes his.
    Mason knows *Spencer*'s choice when he makes his.

How are players rewarded based on what they do?
    Mason earns $1 and *Spencer* loses $1 if both choose *Heads* or both choose *Tails*.
    Mason loses $1 and *Spencer* earns $1 if one chooses *Heads* and the other chooses *Tails*.

Objective: Understand the distinction between noncooperative and cooperative game theory.

Both cooperative and noncooperative game theory start by defining the rules of the game. The primary distinction between the two is how the solution to a game is reached. For cooperative games, the solution is reached by following the implications of some set of axioms that frequently appeal to such notions as Pareto optimality, fairness, and equity. For noncooperative games, the solution is determined assuming players act individualistically to maximize their utility within the constraints set forth by the rules of the game.

The approach we will take in this class is that of noncooperative games.

Objective: Understand the description of a game in the extensive form.

The most informative way to describe a game is using a game tree, which is referred to as the extensive form.

To describe a game in the extensive form, we must describe the tree. A tree consists of three essential elements: a set of nodes (V), a set of branches (A), and a root (r). A branch connects two distinct nodes. A path is a sequence of nodes connected by branches. Any two nodes are connected by exactly one path. The root is where the tree begins. The root also provides a sense of direction and dictates the ends to the tree. A branch to a node that is on the path to the root is incoming. If it is not on the path to the root it is outgoing. Nodes without outgoing branch are called terminal or leaves and represent an end to the tree and game. Nodes with outgoing branches are called non-terminal or decision nodes. A branch can then be thought of as a particular choice or action.

Figure 1 shows an example of a tree. The nodes are v_1, v_2, v_3, v_4, v_5, v_6, and v_7. Branches can be identified by the nodes they connect: (v_1, v_2), (v_1, v_3), (v_2, v_4), (v_2, v_5), (v_3, v_6), and (v_3, v_7). The sequence v_3, v_1, v_2, v_4 is an example of a path, while v_3, v_1, v_2, v_6 is not because there is no branch connecting v_2 and v_6. Figure 2 shows an example that is not a tree because node v_1 and v_5 are connected by more than one path: v_1, v_2, v_5 and v_1, v_3, v_5.

Once the tree is described, we must also

(i) describe the players: $N = \{0,1,2,\ldots,n\}$ where player 0 is a special player referred to as chance or Nature.

(ii) partition the non-terminal nodes in the tree to the players: $P^0, P^1, \ldots, P^n$.

(iii) assign actions or choices to each branch of the tree.

(iv) assign probability distributions over outgoing branches for each node in P^0.

(v) partition the nodes into $k(i)$ information sets for each player i: $U_1^i, U_2^i, \ldots, U_{k(i)}^i$.

(vi) Describe an n-dimensional vector, $g(t) = (g^1(t), g^2(t), \ldots, g^n(t))$, of payoffs for each terminal node t.

Note: Information sets are used to describe what players know when. The idea is that a player cannot distinguish between nodes in the same information set.

Let us consider *Mason*'s preferred version of matching pennies in Figure 3.

Spencer moves first at v_1. If *Spencer* chooses *Heads*, (v_1,v_2), *Mason* knows it and gets to choose between *Heads*, (v_2,v_4), or *Tails*, (v_2,v_5). If *Spencer* chooses *Tails*, (v_1,v_3), *Mason* knows it and gets to choose between *Heads*, (v_3,v_6), or *Tails*, (v_3,v_7). The payoffs are summarized in parentheses, with the first number corresponding to *Spencer* and the second to *Mason*.

Formally,
$N = \{M, S\}$ where M is *Mason* and S is *Spencer*.

$P^0 = \varnothing$, $P^S = \{v_1\}$, and $P^M = \{v_2, v_3\}$

$(v_1, v_2) = (v_2, v_4) = (v_3, v_6) =$ *Heads*
$(v_1, v_3) = (v_2, v_5) = (v_3, v_7) =$ *Tails*

$U_1^M = \{v_2\}$, $U_2^M = \{v_3\}$, and $U_1^S = \{v_1\}$

$g(t) = (g^S(t), g^M(t))$
$g(v_4) = g(v_7) = (-\$1, \$1)$
$g(v_5) = g(v_6) = (\$1, -\$1)$

Figure 4 is the more traditional version of matching pennies where *Mason* does not get to see *Spencer*'s choice. What this means is that v_2 and v_3 will be in the same information set. When *Mason* makes his choice, he does not know what *Spencer* does. This is represented graphically in Figure 4 by connecting v_2 and v_3 with a dashed line.

Formally,

$N = \{M, S\}$

$P^0 = \emptyset$, $P^S = \{v_1\}$, and $P^M = \{v_2, v_3\}$

$(v_1, v_2) = (v_2, v_4) = (v_3, v_6) = Heads$
$(v_1, v_3) = (v_2, v_5) = (v_3, v_7) = Tails$

$U_1^M = \{v_2, v_3\}$, and $U_1^S = \{v_1\}$

$g(t) = (g^S(t), g^M(t))$
$g(v_4)= g(v_7) = (-\$1, \$1)$
$g(v_5) = g(v_6) = (\$1, -\$1)$

We should now come back to information sets for a moment and ask ourselves what restrictions make sense for these information sets. Figure 5 presents two examples of information sets that don't make much sense. On the left hand side *Mason* is choosing between *Heads* and *Tails* if *Spencer* chooses *Heads* and *Heads* and *Feet* if *Spencer* chooses *Tails*. Obviously, the differences in choices should tip *Mason* off to where he is at in the game tree. On the right hand side, *Mason* is choosing between *Heads* and *Tails* if *Spencer* chooses *Heads* and *Heads*, *Tails*, and *Feet* if *Spencer* chooses *Tails*. Again, a pretty big tip to *Mason* about where he is at in the game tree. Specifically, for all nodes in the same information set, the outgoing branches must correspond to the same set of actions.

Figure 6 presents another example. What this information set says is that *Mason* doesn't know if he is at v_2, v_3, or v_7, but to get to v_7 he had to make a choice at v_3, which should tip him off to the fact that he is at v_7 unless he is forgetful. This case can be eliminated by restricting paths from passing through an information set more than once.

Another example of forgetfulness is illustrated in Figure 7. In Figure 7, *Mason* starts the game by choosing *R* or *L*, *Spencer* then chooses *R* or *L* knowing what *Mason* did, finally *Mason* must end the game by choosing *R* or *L* not knowing what *Spencer* did or he did originally. Again, this makes sense only if *Mason* is forgetful. While recognizing this type of forgetfulness in a game is pretty straightforward, describing the restrictions required to eliminate it is not so simple. However, if you are interested, the formalities can be found in Fundenberg and Tirole pp. 77-82. While people are sometimes forgetful, the games we will study do not allow for forgetfulness.

Objective: Understand what a strategy in an extensive form game represents.

A strategy is a player's description of how to play the game. There are two types of strategies: pure and mixed. For now, we will focus on pure strategies.

A pure strategy is a complete description of play for all contingencies. That is, a pure strategy specifies a course of action for a player at every information set. Let S_j^i be the set of available actions for information set U_j^i, where $s_j^i \in S_j^i$ is a particular action. A pure strategy is $s^i = \{s_1^i, s_2^i, \ldots, s_{k(i)}^i\}$. The set of all possible pure strategies is $S^i = S_1^i \times S_2^i \times \ldots \times S_{k(i)}^i$.

For the traditional version of the matching pennies game, $S^M = S^S = S_1^M = S_1^S = \{Heads, Tails\}$.

For *Mason*'s preferred version, $S_1^M = S_2^M = S_1^S = \{Heads, Tails\}$, $S^S = \{Heads, Tails\}$, and $S^M = \{(Heads, Heads), (Heads, Tails), (Tails, Heads), (Tails, Tails)\}$. A representative strategy for *Mason* is (s_1^M, s_2^M) where the first element describes what *Mason* will do under the contingency that *Spencer* chooses *Heads* and the second describes what *Mason* will do under the contingency that *Spencer* chooses *Tails*.

Objective: Understand the normal form representation of a game.

An alternative way of describing a game is the normal form, which is also referred to as the strategic or matrix form.

A strategic form game is described by:

(i) A set $N = \{1,2,\ldots,n\}$ of players.

(ii) A finite set of pure strategies S^i for each $i \in N$ where $S = S^1 \times S^2 \times \ldots \times S^n$ is the set of all possible pure strategy outcomes.

(iii) A payoff function $g^i: S \rightarrow \Re$ for each $i \in N$.

For our traditional matching pennies game, we can represent the normal form game in a matrix.

		s^S	
		Heads	*Tails*
s^M	*Heads*	-$1 $1	$1 -$1
	Tails	$1 -$1	-$1 $1

Note: *Mason*'s payoff is in the lower left hand corner, while *Spencer*'s is in the upper right hand corner.

For *Mason*'s preferred variation, the normal form game looks like

		s^S	
		Heads	*Tails*
	(*Heads*,*Heads*)	-$1 $1	$1 -$1
s^M	(*Heads*, *Tails*)	-$1 $1	-$1 $1
	(*Tails*,*Heads*)	$1 -$1	$1 -$1
	(*Tails*, *Tails*)	$1 -$1	-$1 $1

Finite extensive form games can always be written in the strategic form.

Extensive form games are more informative because many different extensive form games can produce essentially the same strategic form with a simple relabeling of strategies.

Objective: Understand common classifications of games.

Simultaneous Move Games: Games where players must make choices knowing the incentives of other players, but not their choices.

Sequential Move Games: Games where players make choices knowing the incentives of other players and at least one player knows at least one choice of another before making one of his own.

Simultaneous move games are static in nature and easily represented in the normal form. Sequential games are dynamic in nature and often cumbersome to represent in the normal form.

Perfect Information Games: Games where all information sets contain a single node.

Imperfect Information Games: Games where at least one information set contains two or more nodes.

Everyone always knows precisely where they are at in the game tree in perfect information games. In imperfect information games, some or all players do not always know for certain where they are at in the game tree.

Complete Information Games: Games where all players know each other's payoffs.

Incomplete Information Games: Games where some players' payoffs are unknown.

To deal with incomplete information games, we will take advantage of our chance player and let it assign probabilities to alternative sets of possible payoff outcomes. This transforms the game into an imperfect information game.

Objective: Understand how experimental methods are being used to improve game theory.

For a long time, positive economics was besieged by the perspective that it was a purely observational and theoretical science, much like astronomy and paleontology. During this time, economists satisfied their curiosity by observing the world around them and theorizing about why things were the way they were. Economists built an industrial sized tool kit of statistical models to torture their observations into confessing which theories were likely responsible. Of course, the challenge with all observational science is that observations are confounded by a multitude of factors that must be carefully untangled.

Vernon Smith has spent his lifetime changing this perspective by showing that economists can use controlled experimentation in a laboratory environment to gain important insights into market and human behavior. Experimental methods now flourish in economics due to his pioneering work for which he shared the 2002 Nobel prize in economic science with Daniel Kahneman.

Before experimental methods were widely accepted in economics, the study of game theory was primarily theoretical and full of introspection and guessing. Noisy observations of real world data are often not well suited for testing competing theories or validating the subtle assumptions upon which they were constructed. When a theory fails to explain our observations, it is natural to ask why. Did the game we hypothesize accurately characterize the strategic environment that produced the data? Did the theory fail due to the violation of a fundamental assumption upon which it was built?

Experimental methods are good for testing game theoretic principles because confounding factors can be controlled in order to reduce observational noise. In a sufficiently controlled experiment, we do not have to hypothesize about the game people are playing because we tell then specifically. By carefully constructing and playing games in the lab, we can link specific deviations from the theory to specific assumptions regarding behavior.

The successful application of experimental methods to game theory has led to a proliferation of research that has identified both strengths and weaknesses in traditional theory. Empirical regularities identified with experimental methods that we will explore include reciprocal, fairness, and altruistic behavior; limited reasoning; and the effect of learning on equilibrium outcomes.

Figure 1: A Basic Tree

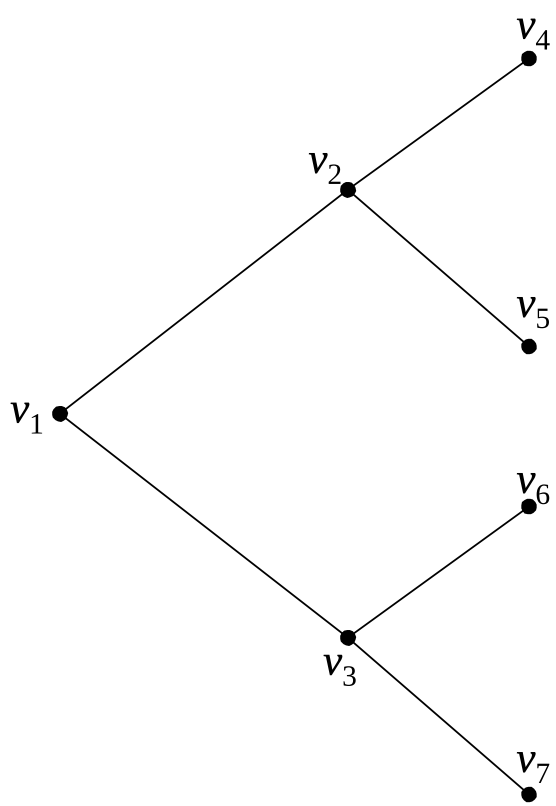

Figure 2: Not a Tree

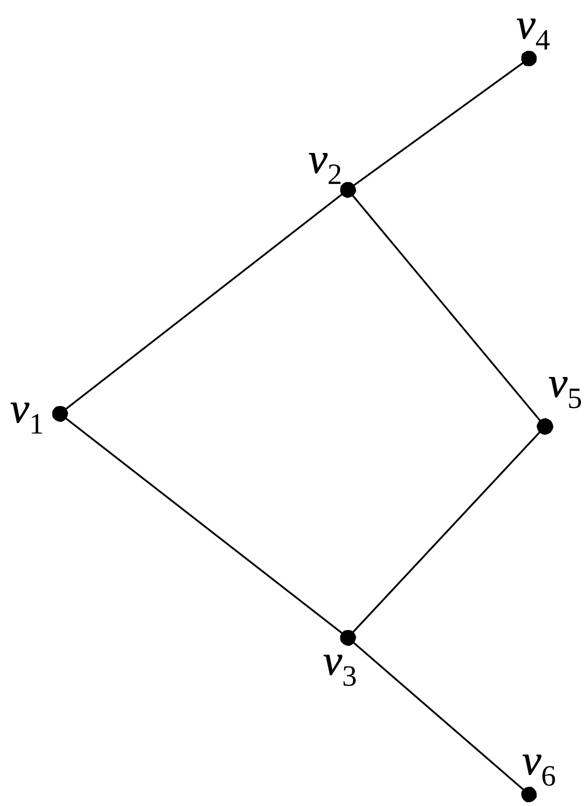

Figure 3: Mason's Preferred Matching Pennies Game

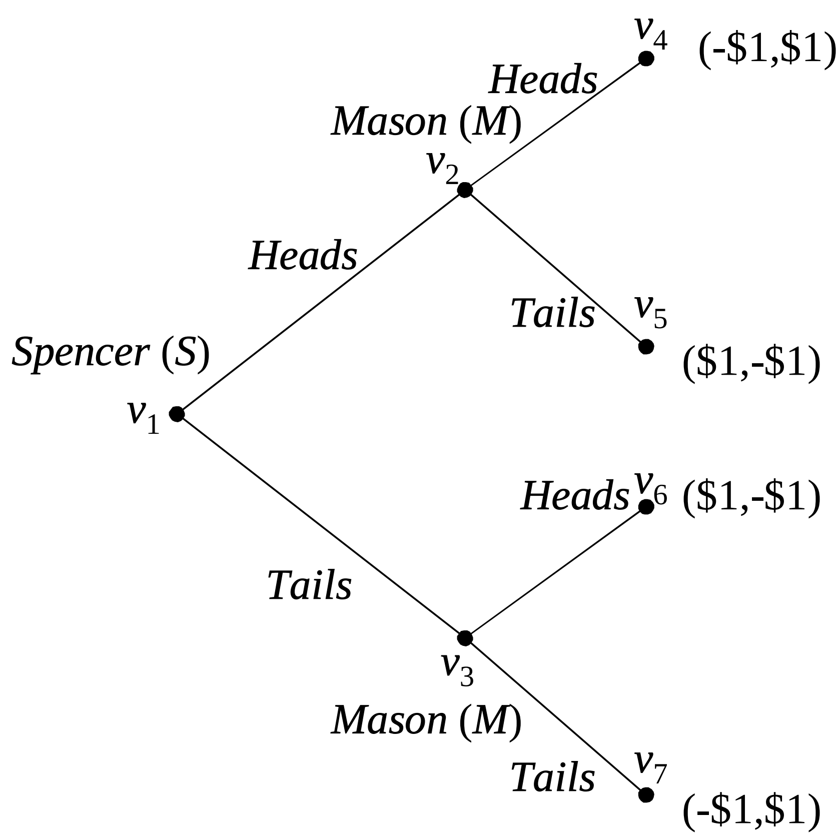

Figure 4: Traditional Matching Pennies Game

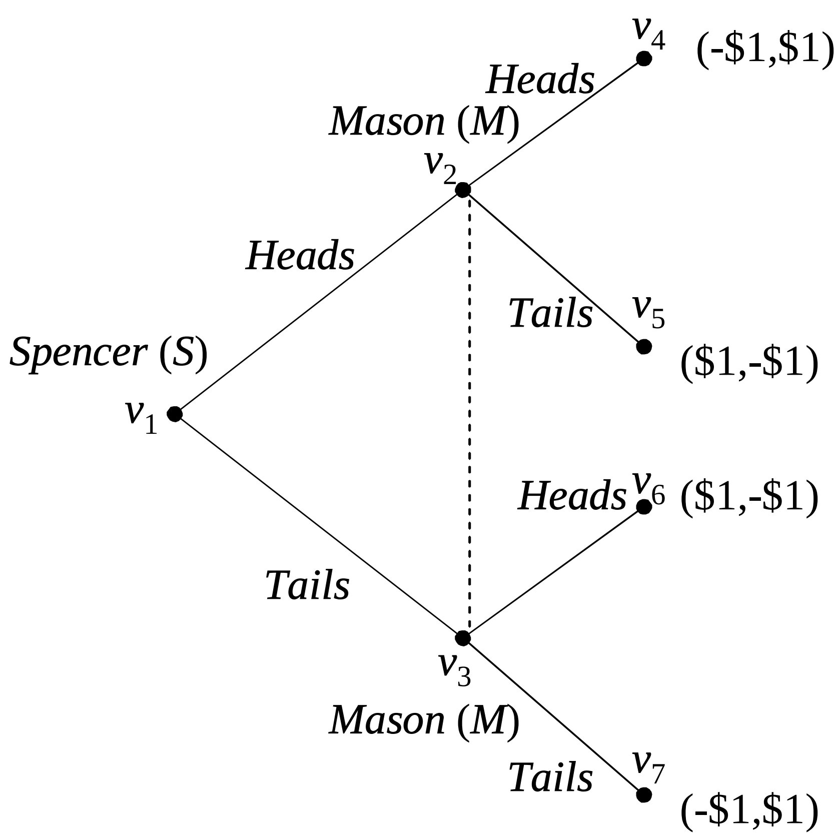

Figure 5: Contradictory Information Sets

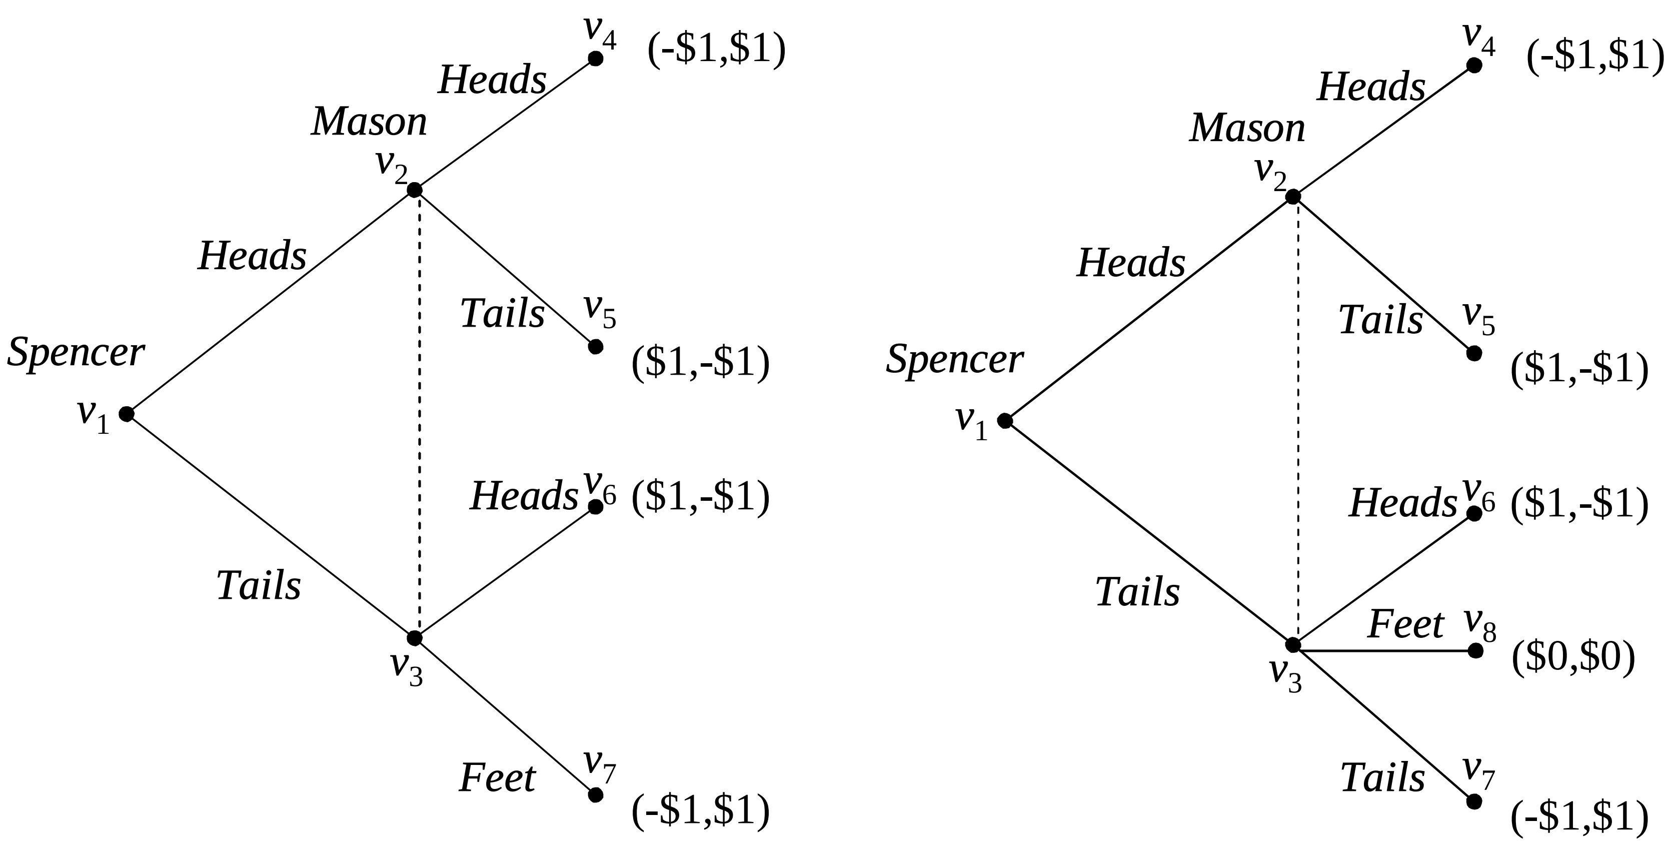

Figure 6: Forgetful Information Set

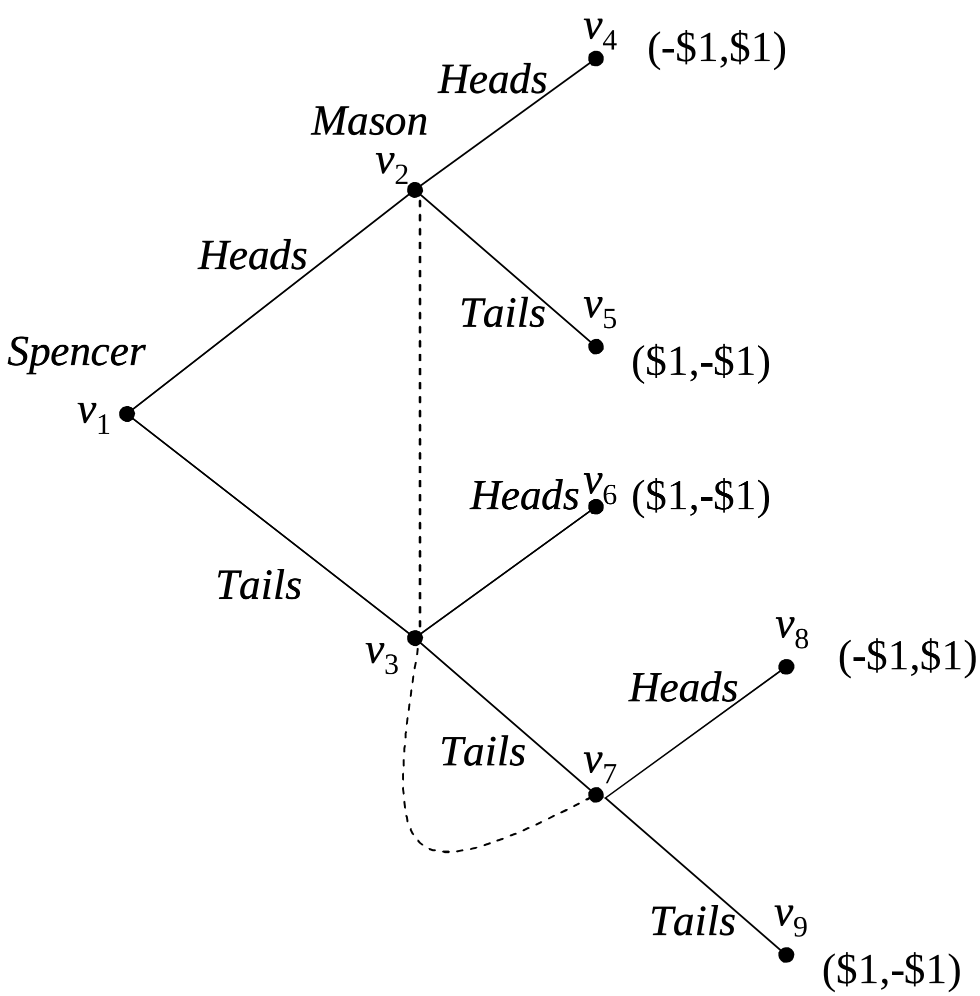

Figure 7: Another Forgetful Information Set