

The Eleventh Annual Game Design Think Tank

Project Horseshoe 2016



Group Report: Game Theory Applied to Actual Games: The Nash Equilibrium Payoff Matrix as a Model for Multiplayer Behavior

Participants: A.K.A. *"The Sleeping Matrices"*

Michael Austin, Hidden Path Entertainment

Kyle Brink, Mercenary Technology

Crystin Cox, ArenaNet

Tim Fowers, Fowers Games

Facilitator: Zack Hiwiler, Full Sail University

Jason Vandenberghe, Ubisoft

Link Hughes, ArenaNet

Mario Izquierdo, Twitch

Problem Statement

Players often play very differently in multiplayer games than they do in single-player games. A model for exactly how these changes manifest, and which provides some predictive power when applied to a known multiplayer game system, would be of great value in creating sustainably entertaining multiplayer games.

Game theory provides a possible answer in the reward matrix used to determine a Nash equilibrium. By exploring the applicability of this model to multiplayer gameplay, with known models of player motivation on hand, we determine what visibility it might provide into player behavior as well as the implications of implementing such a model in practice.

Payoff Matrices and the Nash Equilibrium

A payoff matrix depicts the possible benefits to each player in a game. Each position in the matrix is a possible game state consisting of a specific set of player choices. The payoffs are numerical weights given to the perceived benefit of making that choice.

A Nash Equilibrium is one such state where there is no incentive for any player to change strategies. In such a state, any deviation from the current strategy by any player provides lower expected payoffs.

In the classic Prisoner's Dilemma, for example, there is a Nash equilibrium at the state of both prisoners betraying one another, because for either player to choose cooperation worsens that player's outcome.

[This short video from Khan Academy](#) explains these concepts well.

[A Cornell student blog post from 2011](#) views the set of player choices in a team-based shooter as a form of Prisoner's Dilemma where the choices are "play to win" or "play for fun." The payoff matrix in this post presumes a) the greatest payoff lies in winning, achievable only through team play; and b) playing for fun precludes both team play and winning.

This is clearly not always the case; for example, playing to win and playing for fun are not mutually exclusive choices for players whose primary fun comes from winning. But the idea that the choices of players affects our own payoff options, that this can include whether or not we have fun, and that can all be modeled using payoff matrices, is salient.

Heterogeneous Payoff Matrices

In games, a significant portion of the payoff is in the form of enjoyment. Therefore, the game's rules do provide a full picture of the player payoff matrix until they are informed by player motivations and preferences.

For example, a player who enjoys mastering the game and a player who enjoys working together may have common ground and enjoyment in a team challenge, but feel very differently about spending time in a game hub with teammates. Their payoff matrices would have different values for choosing the "hub" activity in the game, but might have nearly identical values for choosing "team challenge."

The Bayesian Nash Equilibrium

In situations where player payoff matrices differ, those payoff matrices can be viewed as player "types" which predict player choices in a game, and players do not have perfect knowledge of what their fellow players' types are (though they may have some idea of the likelihood or prevalence of various player types). In game theory, this is known as a [Bayesian Game](#).

Conceptually, this is quite similar to how gamers approach multiplayer games. They have some experience with the gaming community as a whole, and with their friends and frequent compatriots (guildmates) in particular, which provides an informal mental model of player types and probabilities which gamers (consciously or unconsciously) use to modify their own play styles.

Player Type Models

Applying payoff matrices to multiplayer games, then, demands that we have a means of modeling player types in our payoff matrices. We considered several strong candidates.

Big 5 / Domains of Play

[Jason Vandenberghe's Domains of Play](#) translates the Big 5 (OCEAN) motivation theory from psychology into a form that applies to gamer motivations. It measures the core motivations of players across 30 factors.

Self-Determination Theory and PENS

Human beings are motivated by feelings of competence, autonomy, and relatedness. These three motivations are the core of [Self-Determination Theory](#). SDT is the basis for [the PENS \(Player Experience Need Satisfaction\) model](#) of player motivation.

Quantic Foundry's Gamer Motivation Model

Quantic Foundry analyzed their large data set from the Gamer Motivation Profile survey, and created [the Gamer Motivation Model](#) to describe the clusters of motivations they found. The model groups player motivations into Action, Social, Mastery, Achievement, Immersion, and Creativity (with subcategories in each).

Matrices Within Matrices

Putting together comprehensive payoff matrices for multiple players, across all the possible game action choices, is far beyond the scope of this conference.

In fact, to model it most accurately, each game action choice node in a game's overall payoff matrix should itself be populated with a payoff matrix of all the possible motivations a given player could be feeling for that choice. Two players could have vastly different reasons for liking a particular choice, and end up liking it the same -- but you could not know that without having the full model in place.

In looking at a prototype of such a model, it becomes clear that a practical approach requires an extensive body of game analytics describing the possible trackable game choices. This can then be the framework of player payoff matrices, populated with payoff values weighted according to known player motivation types (and the prevalence of each, from existing player analytics data) to predict probable outcomes and make design decisions.

In effect, this gives the game designer a level of specific insight necessary to achieve Bayesian Game levels of knowledge about player types and probabilities.

The Game is the Variable

In putting such a model together, it becomes clear that -- from the game designer's perspective -- **player types and motivations are the constant, and the game is the variable**. The players bring their motivations to us, and we bring our games to them.

In the following simplified matrix (using three Gamer Motivations and three broad strategies), we see that as each player type encounters different player choices by teammates, they receive different levels of satisfaction (payoff).

	Other Player Choice	Rush	Explore	Hub
My Gamer Motivation	My Chosen Action			
Action	Rush	3	1	
	Explore	1	2	
	Hub	0	0	
Social	Rush	2	1	
	Explore	1	2	
	Hub	1	1	
Mastery	Rush	2	2	
	Explore	2	2	
	Hub	1	1	

The game cannot change the player type, only the support for different choices and extrinsic rewards for those choices. The intrinsic reward is largely provided by the player's personality.

Thus, when player behavior begins to settle into equilibriums that are detrimental to the game (for example leading to mass abandonment), we as designers should not look to change players, but instead to change the game's payoffs as seen through the player type lens.

Multiplayer Games are Iterative Bayesian Games

Each individual choice in a multiplayer game is itself a Bayesian game. The player chooses based on her personal payoff matrix, informed by knowledge of what other players are likely to do. Other players do the same. Then, as those choices are observed by all, each player makes another choice based on new understandings of what other players will do.

Importantly, players will act against their "type" if they believe the actions of other players will reward them enough for doing so.

Testing the Predictive Power of a Player-Typed Payoff Matrix in a Multiplayer Game

As a back-of-the-envelope test, we built a payoff matrix for player choices among three known players and two basic game choices in Diablo III multiplayer.

The players, J, K, and C, can choose to "rush," heading directly for objectives and bosses and overcoming them, or they can "explore," clearing out each map thoroughly.

Rapport represents the additional payoff for each player in doing what someone else is also doing. This turned out to be instrumental in capturing the real cooperative game dynamics among friends. Note that among strangers, these same players might have different Rapport payoffs as they do when playing with friends.

(The payoff values in this chart were filled in by the actual players themselves, or by people who knew them very well.)

Player	Explore	Rush	Rapport
J	+5	+2	+2
K	+1	+4	+2
C	+2	+4	+5

This creates the following payoff matrix when all players are mapped against each other.

	C Exp. / K Exp.	C Exp. / K Rush	C Rush / K Rush	C Rush / K Exp.
J Explore	J+7 / C+7 / K+3	J+7 / C+7 / K+4	J+5 / C+9 / K+6	J+7 / C+4 / K+3
J Rush	J+2 / C+7 / K+3	J+4 / C+2 / K+6	J+4 / C+9 / K+6	J+4 / C+9 / K+1

Predictions from this model:

1. A migration towards the optimal state of K and C rushing while J explores.
2. When one of the three players goes back to town (effectively leaving the equation), the equilibrium point would shift due to that player no longer contributing to the Rapport payoff.
3. If, once reaching the C Rush / K Rush / J Explore state, K returns to town, suddenly C is not having as much fun rushing as she would if she were Exploring with J and getting back into Rapport.

All of the above predictions match actual observed play among these three players. This one anecdotal exploration suggests that, as long as the model is accurate and captures the intrinsic as well as extrinsic payoffs, it could be predictive.

Closing Thoughts

The lens of Nash equilibriums and reward structures is really useful to game designers as a means of discussing the behavior of players in multiplayer games.

While creating a comprehensive payoff matrix model for a multiplayer game seems on the surface to be a merit, a truer test of this model would require significant effort and access to a body of player data from an existing multiplayer game.

Using a player dataset from an existing multiplayer game, and a Gamer Motivation distribution across the player base, one could look back over past player behavior arcs and see how much players' actual behavior diverges from that predicted by the Gamer Motivation Model alone.

Then one could examine that delta to see how it has developed over time, and map that change to changes in aggregate player behavior. One could seek correlations between rises in one set of player choices and declines in others to spot where player interaction is changing the payoff for other players driving them to make different choices (taking into account the timing of game updates to see where game design changes may be driving these behavior deltas).

Additional Points to Consider

Keep these in mind when approaching game design from the perspective of multiplayer payoff matrices and gamer motivations.

- We stop playing when we feel our goals can't be achieved using the behavior we are falling into, we don't have goals, or when we are satiated.
- One approach could be to create a reward matrix of players to players and use that to structure a game
- Asymmetric evaluation is what drives trading between players in board games. Perhaps there's a way to model heterogeneous payoff matrices via asymmetric evaluation.
- We feel best when the optimal personal rewards come from operating in a manner consistent with our personalities, but we're willing to change our behavior to get the best personal rewards.
- The payoff matrix for a multiplayer game is about how strategies **combine**.
- Be sure to include points for intrinsic rewards ("Having people like me", "Maximizing my overall revenue") when setting up the payoff matrix.
- There appear to be a limited number of multiplayer games; perhaps there are some commonalities in the payoff matrices among games within each type of multiplayer game:
 - Free-for-all
 - Team deathmatch
 - Coop small
 - Coop big
 - Parallel small
 - Parallel big
 - Many vs one
 - Team objective
- In every game, one game choice made by each player is deciding to play or not play.
- The more robust the in-game communication, the greater impact players can have on each other's payoffs.
- People give up Autonomy by being in a group so it now comes to Mastery vs. Relatedness (in the SDT model).
 - Freedom isn't autonomy. Autonomy satisfaction comes from self-expression.
 - Relatedness satisfaction can come from the ability to tell a story - to share things with a group that confirm hierarchy and groupings
 - Mastery - demonstrating control over your environment. I have a model of the world; when I take an action, the thing I predict happens. The greater the risk with that prediction, the higher the payoff.
- Understand other people's payoff matrices to maximize your chances of making choices that bring joy.
- Communication is how people learn others aren't like them. The human base assumption is that people are like each other.