THE PENNSYLVANIA STATE UNIVERSITY
SCHREYER HONORS COLLEGE

DEPARTMENT OF CHEMICAL ENGINEERING

CHEMICAL GAME THEORY AND COOPERATION IN EXPERIMENTAL PRISONER'S DILEMMA GAMES

LAURA COOK
SPRING 2018

A thesis
submitted in partial fulfillment
of the requirements
for a baccalaureate degree
in Chemical Engineering
with honors in Chemical Engineering

Reviewed and approved* by the following:

Darrell Velegol
Distinguished Professor of Chemical Engineering
Thesis Supervisor

Ali Borhan
Professor of Chemical Engineering
Honors Adviser

** Signatures are on file in the Schreyer Honors College.

ABSTRACT

The objective of this thesis is to describe a new framework for solving game theory problems, and to analyze experimental data from the literature using this framework. The analysis will evaluate how well the framework, which we call "Chemical Game Theory" (CGT), addresses three key challenges with traditional game theory: (1) the consideration of a player's payoffs or pains only as relative values, (2) the discrepancy between classical predictions of how "rational" players should behave and how actual players behave in experimental situations, and (3) the inability to account for a player's pre-bias. Chemical game theory will be applied to each of these challenges, using examples such as the "epsilon problem" and selections from the literature, and the results from each will be compared to results from classical game theory. The hypothesis is that chemical game theory will address these challenges more accurately than classical game theory, potentially providing a useful new framework for games and a method for determining what players *will* do, instead of what they *should* do.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

This thesis has been adapted and expanded from the manuscript “Chemical Game Theory” by D. Velegol et al, which has been submitted for publication to *Industrial & Engineering Chemistry Research* as of the time of writing.

I would like to thank Dr. Darrell Velegol for his help, guidance, and encouragement over the years. Working in the CGT research group has expanded my horizons and changed the way I think about problems, and that would not have been possible without Dr. Velegol. I would also like to thank the other members of the research group, especially Natalie, Sneha, and Frank, for creating a supportive and innovative environment that I will be sad to leave. I look forward to seeing how the newer members of the group will propel CGT to new levels. Additional thanks to Dr. Ali Borhan for his help as my honors advisor and for reviewing this thesis.

Finally, I am very grateful to my family and friends for their love and support.

Chapter 1

Introduction

1.1 Introduction

The objective of this thesis is to describe a new framework for solving game theory problems, and to analyze experimental data from the literature using this framework. The analysis will evaluate how well the framework, which we call "Chemical Game Theory" (CGT), addresses three key challenges with traditional game theory: (1) the consideration of a player's payoffs or pains only as relative values, (2) the discrepancy between classical predictions of how "rational" players should behave and how actual players behave in experimental situations, and (3) the inability to account for a player's pre-bias.

In this thesis, chemical game theory will be applied to each of these challenges, and the results from each will be compared to results from classical game theory. The goal is to determine if chemical game theory can address the challenges better than classical game theory can. For the first challenge, I will use the "epsilon game," in which a player's pains are all roughly equivalent, to examine the difference between chemical and classical results. Classical game theory considers only the relative value of the pains; chemical game theory, however, takes absolute values of the pains into account, and so predicts a different outcome. These results will be compared to a common sense analysis.

To address the discrepancy between classical game theory and experimental results, I will also take two examples of experimental Prisoner's Dilemma games from the literature and solve them with CGT methods. These data will be used to test the hypothesis that CGT methods can explain the cooperation rates of the experimental data, which classical game theory methods generally do not.

I will also show how chemical game theory can account for pre-bias. CGT considers pre-bias explicitly when solving games, which may explain differences in cooperation for which classical game theory has no answer. I will analyze an experimental Prisoner's Dilemma game to test the effect of pre-bias on chemical game theory results.

Even if the hypothesis holds for these data, many further experimental tests would be required to assess this new framework fully. Nevertheless, CGT has the potential to transform how games are framed and analyzed, giving a new tool to those who wish to describe not how people *should* act, but how they *do* act.

Chapter 2

Classical Game Theory

2.1 A Brief Introduction to (Classical) Game Theory

What I will refer to in this thesis as “classical game theory,” to distinguish it from this new chemically based model, has its roots in a 1944 book by mathematician and physicist John von Neumann and economist Oskar Morgenstern, *Theory of Games and Economic Behavior*.[1] Since then, classical game theory has been expanded to mathematically describe human and animal behavior in many decision-making situations, including (but not limited to) climate change,[2] overfishing,[3] and politics,[4] as well as economics.

Classical game theory deals with strategic decision-making, which is a unique kind of optimization problem. Whereas in mathematics you can optimize a function by simply taking the derivative and setting it equal to zero, strategic decision-making is much more complicated. There is no one definitively correct answer—A’s optimal situation is for A to win, and B’s is for B to win. Additionally, the eventual outcome depends on both players’ choices. In creating a playing strategy, therefore, the players will consider their own preferences and what they think their opponent will do.

One of the most famous and well-studied games is the Prisoner’s Dilemma (PD). Picture this: two accomplices who have just committed a crime are taken to the police station and put in separate rooms. Each is approached with a choice: they can either stay quiet, or tell on their partner. Depending on the choices both players make, they will each receive a certain amount of prison time. The choices and respective punishments can be represented in “normal form” in Table 1, below. Each number in the table represents the years in prison time (a type of “pain”) that the player will receive for that certain outcome. Player A is denoted in bold italics, and Player B in normal font.

Table 1. Normal form representation of a classic Prisoner's Dilemma.

The pain values (here, years of prison time) for Player A are in bold italics, and those for Player B are in normal font.

	b1 = quiet	b2 = tell
a1 = quiet	***+1***, +1	***+3***, 0
a2 = tell	***0***, +3	***+2***, +2

The table is read by finding the choice that A makes on the left and the choice B makes on the top, and following these across and down to the corresponding decision block (in blue). Assuming that neither of the players can refuse to play, or change their decision once the outcome has been determined, the game has four possible outcomes:

1. Both stay quiet, and each receives 1 year of jail time
2. Both tell on the other, and each receives 2 years of jail time
3. A tells and B stays quiet, resulting in 3 years of jail time for B and 0 for A
4. B tells and A stays quiet, resulting in 3 years of jail time for A and 0 for B

It seems clear that it is better for both players to stay quiet than for both to tell, as the {quiet, quiet} decision results in fewer years of prison for both players. However, if you stay quiet, you open yourself up to the possibility of taking the fall for the other person—you get all of the prison time, and they get none. Therein lies the dilemma of the game’s title.

There are many such games under the heading of classical game theory, including Battle of the Sexes (Bach or Stravinsky), Tragedy of the Commons, Chicken, and even one with which most children will be familiar: Rock, Paper, Scissors. For the purposes of this thesis, focus will be restricted to the Prisoner's Dilemma; however, it is important to note that both classical game theory and chemical game theory can be applied to a plethora of other situations.

2.2 Classically Solving the Prisoner's Dilemma

The classical game theory solution to the Prisoner's Dilemma is determined via a Nash equilibrium,[5] a revolutionary concept developed by mathematician John Nash in the early 1950s. A Nash equilibrium is reached if neither player can get a better outcome by making a different choice. This can be applied to the PD game in Table 1 to determine what players "should" play.

For example, assume Player B chooses "quiet." Then, Player A will receive fewer years of jail time if they pick "tell" (0 years) than if they also choose "quiet" (1 year). Alternatively, if Player B chooses "tell," then Player A will do better to pick "tell" (2 years) over "quiet" (3 years). Thus, in both cases, Player A will receive less jail time if they pick "tell." Classical game theory therefore asserts that Player A will always pick "tell." Player B, in this game, has exactly the same choices and "pains" (years of jail time), and so will play similarly. According to this analysis "by inspection," the players will always end up in the {tell, tell} block, 100% of the time.

2.3 Challenges of Classical Game Theory

This thesis will address three important questions regarding the Prisoner's Dilemma game, specifically, for which classical game theory does not appear to have a definite answer. First, consider the game matrix in Table 2, which follows the classic Prisoner's Dilemma layout but has pains that are very similar to each other.

Table 2. The "ε problem": a PD game with pains that are similar to each other in value.

A common sense analysis would predict that players would not perceive much difference between 1.99 years of prison and 2.00 years of prison.

	b1 = quiet	b2 = tell
a1 = quiet	***+1.99***, +1.99	***+2.01***, +1.98
a2 = tell	***+1.98***, +2.01	***+2.00***, +2.00

The original game described earlier in Table 1 has differences of an entire point between the choices: 0, 1, 2, or 3 years of prison, in that case. Here, in what we call the "ε problem," the differences are ε = 0.01 years—or about 4 days.

If you are already going to be in prison for 2 years, do 4 extra days matter? Is the possibility of removing those extra days from your sentence a big enough incentive to tell on your partner? Classical game theory says that players (if, as is always assumed, they are "rational," or self-interested, actors) will behave here in exactly the same way as they would in the original 0-1-2-3 game. However, doesn't common sense tell us that the difference between 2.00 and 2.01 years of prison would be perceived differently than the difference between 2 and 3 years? It seems that the classical game theory result is counterintuitive for this game. Chemical game theory will give another approach for this problem, in which this dilemma is taken into account.

The second important question regards experimental results. When the Prisoner's Dilemma is played in an experimental setting with actual human beings, it turns out that—contrary to what classical game theory says they *should* do—players choose to keep quiet about half of the time, on average.[6]

For the purposes of comparison, we will analyze two selections from the literature. In one study, the relation between choice of major (economics or non-economics) and cooperation in several types of games was examined.[7] Each college student in the study played a one-shot (not repeated) PD game against two other students. As is common for experimental PD situations, the students met each other and were then taken to separate rooms to choose their strategies—cooperate (stay quiet) or defect (tell). In this case, the outcome of the game was not a punishment but a monetary reward, summarized in Table 3. As before, both players get a better outcome (more money, or less prison) if they choose the {quiet, quiet} block.

Table 3. Monetary rewards for an experimental Prisoner's Dilemma.

The matrix is now given in terms of payoffs instead of pains, but the structure of the PD game remains the same; {tell, tell} is the classically predicted outcome.

	b1 = quiet	b2 = tell
a1 = quiet	***$2***, $2	***$0***, $3
a2 = tell	***$3***, $0	***$1***, $1

When players were not allowed to make promises to cooperate—which is a closer approximation of a real life decision-making process—the cooperation rates (or how often "quiet" was picked) were 28.2% for economics majors and 52.7% for non-economics majors.

This particular study is not the only one to find such rates of cooperation where classical game theory would expect 0%. Another experiment, set at the U.S. Naval Academy, had students play a one-shot Prisoner's Dilemma game with monetary payoffs of $2, $5, $10, and $20.[8] The students cooperated 54% of the time—in line with the previous example.

Much research has been done to account for this kind of result. One possible explanation is the theory of bounded rationality,[9] which says that players have limited time and cognitive ability compared to a perfectly rational being, so they will often choose a "good enough" decision over the most optimal one. Other studies have proposed models in which players experience a "warm glow" from staying quiet[10] (that is, "cooperating" —in the sense that the player is cooperating with the other player by not telling on them), or build a reputation for cooperation over repeated iterations of a game.[11] These models might help explain the discrepancy between classical game theory and experimental results, but CGT offers another perspective and allows quantitative analysis of decision-making in a way most other methods do not.

The third question concerns pre-bias. Classical game theory predicts a certain outcome for a "rational" player, but does "rational" mean "exactly the same as other players"? It seems possible that a rational player can come into a game with a different initial mindset than another rational player. This pre-bias might take the form of prior experiences, beliefs, prejudices, or thoughts informing the way the player feels about the game. If the player has played a similar game before, that could have a strong effect on how they view it when playing again. CGT explicitly allows for incorporation of pre-bias in the form of initial concentrations, allowing this framework to be used to solve types of problems that classical game theory alone appears ill equipped to address.

CGT is a powerful tool for making predictions based on inputted parameters, and using it allows a different perspective on decision-making situations—one in which it is natural, even expected, for players to stay quiet.

Chapter 3

Chemical Game Theory

3.1 Background on Chemical Game Theory

In order to adapt existing chemical theory to complex strategic decision-making problems, the concept of "knowlecules" was developed.[12,13] Knowlecules are metaphorical chemical molecules which can be used to represent players and their choices in the decision-making process. Just as we would represent an atom of carbon with the letter C, here we can represent Player A—their personality, beliefs, experiences, and biases—with the letter A. It may seem strange that something as complex as a person's mind can be boiled down to a single character, but this representation is actually drawn from chemistry. An atom of carbon is incredibly complex. It has electrons, protons, and neutrons; it is described with shells and orbitals; it can form various bonds with many other elements. Even though this method of chemical representation is simplistic, it is useful for analyzing the behavior of molecules in chemical reactions, as is done in CGT.

Essentially, CGT takes both players and their possible choices and treats these concepts as knowlecules. The players and the choices react in "decision reactions" to form decisions as products. This reaction is shown in Figure 1 at the molecular level, where the possible choices (a1, a2, b1, and b2) can collide with a catalyst (in black, with pockets for each player's choice) on the surface of the molecule representing Player A (yellow). A reaction occurs, forming the decision molecule A21.

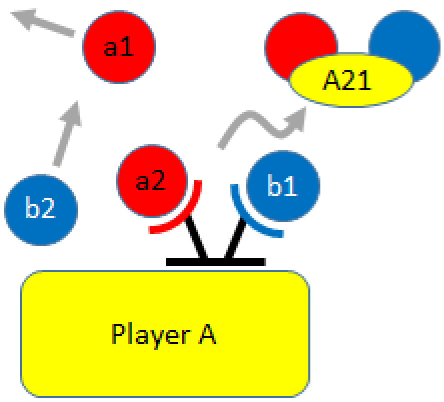

Figure 1. Player A picks "tell" (a2), and Player B picks "quiet" (b1).

The two choices react with a catalyst (in black) on the surface of the molecule representing Player A (yellow) to give the decision A21. Figure from Velegol et al (Ref 12).

To model the overall process, we use a block flow diagram for each player. Here, Player A's information will be identical to Player B's, so each player's block flow diagram will be the same. In Player A's diagram (Figure 2), Reactor A represents how Player A will play. However, Player A must also predict how Player B will play, represented by Reactor B. Finally, Player A must also consider a new participant in the decision-making process—the Decider. In the case of the Prisoner's Dilemma, the Decider may be the police officer attempting to extract the confessions from the players, or the district attorney. The Decider could be neutral or even biased—perhaps the Decider really wants B to take the fall for A, for example. This is an important distinction from classical game theory, where the Decider is not often considered an active participant in the decision.

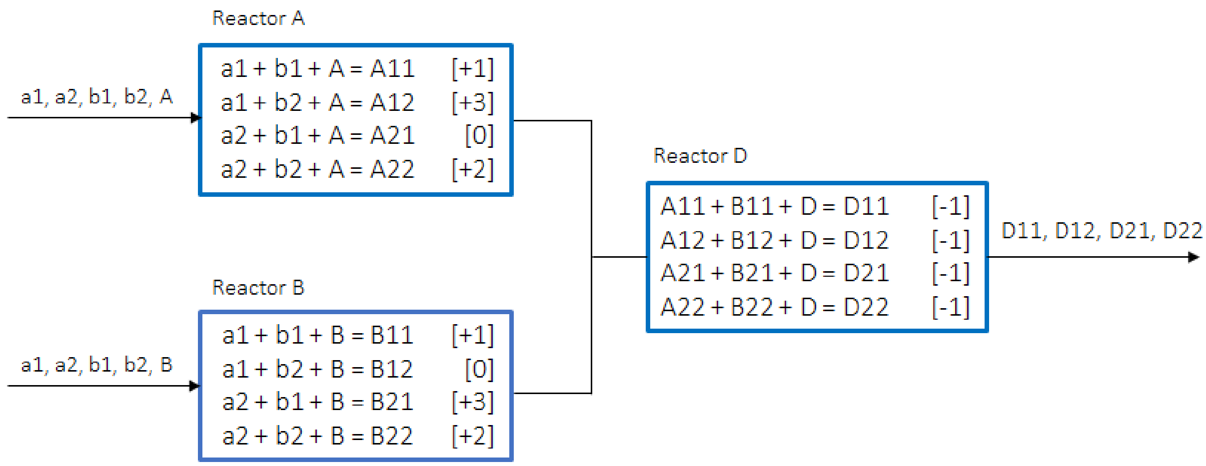

Figure 2. Block flow diagram representing the decision-making process for Player A.

Player A must consider how they themselves will play (Reactor A), how B will play (Reactor B), and how the Decider will affect the decision (Reactor D). At the end of the process, the final products are four decision outcomes with different probabilities of occurring. Figure from Velegol et al (Ref 12).

In CGT notation, the players' two possible choices are 1 (quiet) and 2 (tell). For example, in Reactor A, when Player A picks "tell" and Player B picks "quiet," the knowlecules a2 and b1 react with Player A to create an intermediate decision product A21. This product proceeds to Reactor D, where it combines with the equivalent decision product from Reactor B (B21) to produce a final decision product, D21. After each reactor, the remaining reactants are separated from the products; this makes sense, as in these games you are prohibited from going back on your choice once you have made it.

3.2 Solving PD Games in CGT

As with many chemical engineering problems, an important step in solving a game in CGT is to set up a mass balance. This balance can be displayed in the form of a stoichiometric ("SPICEY") table, like the one in Table 4, which represents the decision reactions occurring in

Reactor A (reactions 1 - 4). Stoichiometric tables are used to tabulate information about each species in the reaction, including its initial concentration, the change in concentration over the course of the reaction, the final concentration, and the final mole fraction.

Table 4. SPICEY table for Reactor A.

Each player's choices react on the surface of the solid representing Player A to produce the intermediate decision products A11, A12, A21, and A22. Table from Velegol et al (Ref 12).

species	initial	change	end	y (mole fraction)
a1 a2 b1 b2	0.50 0.50 0.50 0.50	$-(\varepsilon_1+\varepsilon_2)$ $-(\varepsilon_3+\varepsilon_4)$ $-(\varepsilon_1+\varepsilon_3)$ $-(\varepsilon_2+\varepsilon_4)$	$0.50-(\varepsilon_1+\varepsilon_2)$ $0.50-(\varepsilon_3+\varepsilon_4)$ $0.50-(\varepsilon_1+\varepsilon_3)$ $0.50-(\varepsilon_2+\varepsilon_4)$	$[0.50-(\varepsilon_1+\varepsilon_2)]/\sum$ $[0.50-(\varepsilon_3+\varepsilon_4)]/\sum$ $[0.50-(\varepsilon_1+\varepsilon_3)]/\sum$ $[0.50-(\varepsilon_2+\varepsilon_4)]/\sum$
A11 A12 A21 A22	0 0 0 0	$+\varepsilon_1$ $+\varepsilon_2$ $+\varepsilon_3$ $+\varepsilon_4$	ε_1 ε_2 ε_3 ε_4	$\varepsilon_1/\sum$ $\varepsilon_2/\sum$ $\varepsilon_3/\sum$ $\varepsilon_4/\sum$
inert	0	0	0	0
total	$\sum_0 = 2.00$	$-(\varepsilon_1+\varepsilon_2+\varepsilon_3+\varepsilon_4)$	$\sum = 2.00-(\varepsilon_1+\varepsilon_2+\varepsilon_3+\varepsilon_4)$	1.00

In Reactor A, the reactants are the possible choices, which have certain initial concentrations. Often, players are approximated as "unbiased," with the concentration of a1 (quiet) and a2 (tell) both equal to 0.5. However, many people might have preconceptions when going into a game—whether they are about the game itself, the other player, the Decider, or any other aspect of the situation. The player might even be in a particularly giving mood that day that biases them towards cooperating with the other player. Accordingly, the initial concentrations, or pre-biases, can be changed to reflect this variety of attitudes toward the game.

The other species in Reactor A are the intermediate decision products that form over the course of the reaction, whose initial concentrations are therefore zero. We can also add inert species—distractions that are not a part of the decision, but that interfere with the reaction simply by being present—if desired.

In this type of problem, it is unknown how far the reactions will proceed when setting up the problem. None will go to completion, as entropy dictates that having the greater number of species present is more favorable energetically. In order to solve the system, then, each reaction is assigned an "extent of reaction," ε.

Reactor B (reactions 5 - 8) and Reactor D (reactions 9 - 12) can be represented in the same way. The stoichiometric table for Reactor D can be seen in Table 5, where the intermediate decision products from the eight reactions in the two previous reactors combine to form final decision products.

Table 5. SPICEY table for Reactor D.

The intermediate decision products from Reactors A and B combine to form final decision products D11, D12, D21, and D22. Table from Velegol et al (Ref 12).

species	initial	change	end	y (mole fraction)
A11	ε_1	$-\varepsilon_9$	$\varepsilon_1 - \varepsilon_9$	$(\varepsilon_1 - \varepsilon_9) / \sum$
A12	ε_2	$-\varepsilon_{10}$	$\varepsilon_2 - \varepsilon_{10}$	$(\varepsilon_2 - \varepsilon_{10}) / \sum$
A21	ε_3	$-\varepsilon_{11}$	$\varepsilon_3 - \varepsilon_{11}$	$(\varepsilon_3 - \varepsilon_{11}) / \sum$
A22	ε_4	$-\varepsilon_{12}$	$\varepsilon_4 - \varepsilon_{12}$	$(\varepsilon_4 - \varepsilon_{12}) / \sum$
B11	ε_5	$-\varepsilon_9$	$\varepsilon_5 - \varepsilon_9$	$(\varepsilon_5 - \varepsilon_9) / \sum$
B12	ε_6	$-\varepsilon_{10}$	$\varepsilon_6 - \varepsilon_{10}$	$(\varepsilon_6 - \varepsilon_{10}) / \sum$
B21	ε_7	$-\varepsilon_{11}$	$\varepsilon_7 - \varepsilon_{11}$	$(\varepsilon_7 - \varepsilon_{11}) / \sum$
B22	ε_8	$-\varepsilon_{12}$	$\varepsilon_8 - \varepsilon_{12}$	$(\varepsilon_8 - \varepsilon_{12}) / \sum$
D11	0	ε_9	ε_9	$\varepsilon_9 / \sum$
D12	0	ε_{10}	ε_{10}	$\varepsilon_{10} / \sum$
D21	0	ε_{11}	ε_{11}	$\varepsilon_{11} / \sum$
D22	0	ε_{12}	ε_{12}	$\varepsilon_{12} / \sum$
inert	0	0	0	0
total	$\sum_0$	Δ	$\sum = \sum_0 + \Delta$	1.00

The final concentrations of each species can be calculated using either thermodynamics or kinetics. CGT currently uses the thermodynamic approach—specifically, the following equation, which allows us to calculate the standard state change in molar Gibbs free energy of the reaction:

(1) $$\Delta g^0 = -RT \ln\left(p / p^0\right) \sum \nu_i - RT \sum \nu_i \ln\, y_i$$

The full derivation[14] can be found in Appendix A, but for practical purposes, the most salient point is the second term on the right hand side, which represents the entropy of mixing. This concept is missing from classical game theory, whereas CGT takes into account this entropy of mixing term.

A system of one or two reactions could be solved on paper, but with twelve reactions, it is easier to use a computer program like Microsoft Excel's Solver add-in or the software GAMS. The system is modeled using the SPICEY table as a foundation, and a pain is assigned to each reaction according to the game's pain matrix. Each pain is a non-dimensionalized value found by dividing Δg^0 (from Eqn 1) by RT. We can also refer to these pains as g_{Aij}. Pains are related to the equilibrium constant K for the reaction by the following equation, using a reaction from Reactor A as an example:

(2) $$K_{A21} = \frac{y_{A21}}{y_{a2} y_{b1}} = \exp\left(-g_{A21}\right)$$

This relation also connects the Gibbs free energy and the mole fractions of the chemical species in that reaction. In this way, the entropy of mixing is accounted for in the CGT solution, and we can calculate extents of reaction and the final chemical concentrations that result.

Chapter 4

Results

4.1 The ε Problem Game

The first challenge with classical game theory is the ε problem, laid out previously in Table 2. As stated in Chapter 2, because all of the decision outcomes have similar pains, a common sense analysis would expect that unbiased players would not perceive much difference between their choices, and therefore each outcome should be roughly equally likely. We can test this hypothesis by solving the game in Excel.

Using the equations from the SPICEY tables for the three reactors and Eqn 2 from Chapter 3, the system can be modeled in an Excel spreadsheet, with a section for each reactor. The spreadsheet is solved with the Excel add-in Solver, which can run iterative loops to arrive at a solution based on inputted parameters. Essentially, Solver will vary the extents of reaction in each reactor to give the desired non-dimensionalized Gibbs free energy values as determined by the pain matrix. Detailed instructions on this process can be found in Appendix B. For now, we are interested in the result, shown in Table 6, assuming unbiased players and a Decider with pain values of -1 for each possible outcome.

Table 6. Results for the CGT solution of the ε problem.

Each represents about 25% of the overall decision, as predicted by the common sense analysis.

	D11	D12	D21	D22
Concentration	0.003554	0.003527	0.003527	0.0035
Probability	25.2%	25.0%	25.0%	24.8%

The final concentrations of the four decision outcomes are all roughly 0.0035, indicating that unbiased players do not have much of a preference between the possible outcomes. The {quiet, quiet} decision is slightly favored, with 25.2% of the total decision concentration; {quiet, tell} and {tell, quiet} follow with 25.0%; and lastly {tell, tell} with 24.8%. This result confirms our initial hypothesis: the final outcomes are almost equally likely.

CGT games can also be solved in other software that is able to minimize the energetics of the reactions. One example is GAMS,[i] for which the solution to the ε problem game for unbiased players can be seen in Figure 3. The GAMS solution, when rounded to the same number of decimal places, is the same solution as was found in Excel.

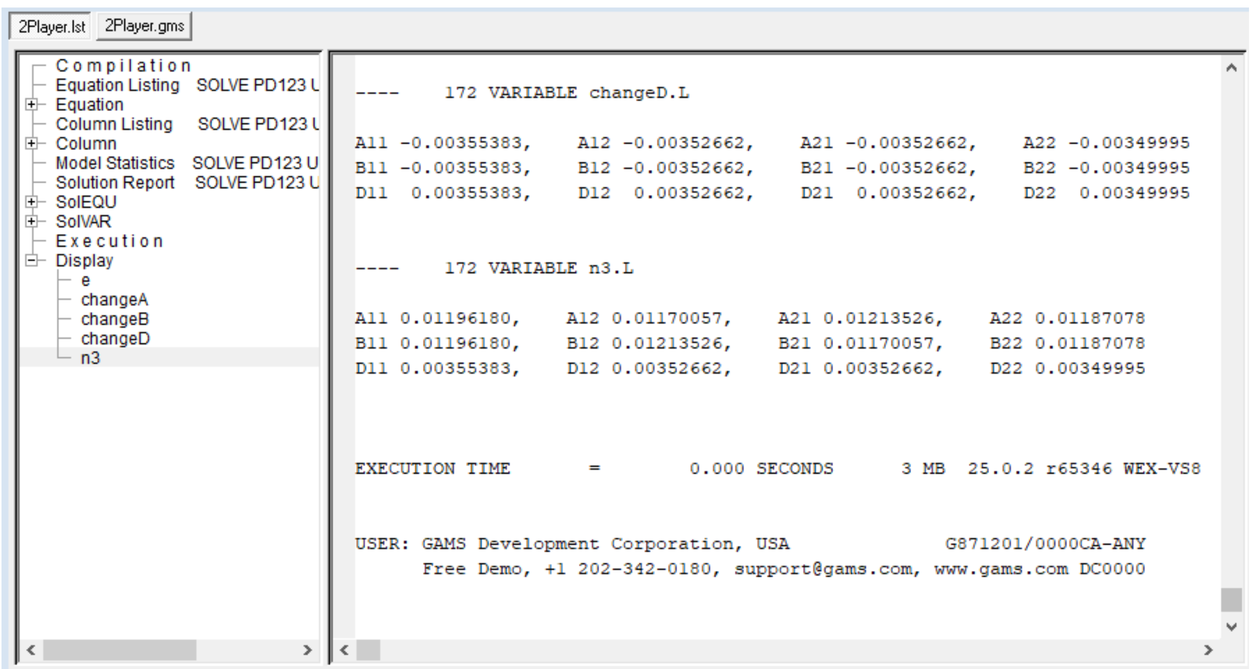

Figure 3. Final decision outcomes from GAMS for the ε problem game (unbiased players).

The final concentrations (D11, D12, D21, and D22) are the same, when rounded to the same number of decimal places, as the concentrations found in Excel.

[i] The GAMS representation of the CGT framework was developed by Frank Gentile and Miras Katenov.

Of course, not every player will be unbiased. In the CGT model, the initial concentrations can be changed to reflect a predisposition to tell or stay quiet, and this will affect the probability of each decision outcome. These results for the ε problem game (panel a) and, for comparison, the traditional 0-1-2-3 game from Table 1 (panel b) are plotted in Figure 4.

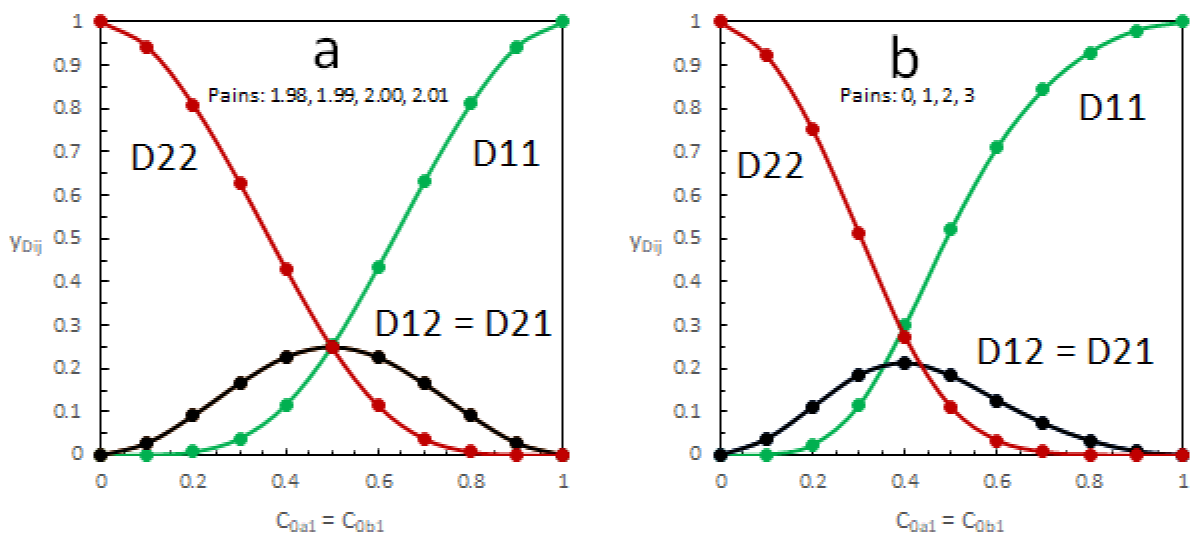

Figure 4. Normalized decision outcomes over a range of initial concentrations.

When c0a1 = c0b1 = 0.1, the players are biased 90% towards "tell," and when c0a1 = c0b1 = 0.9, the players are biased 90% towards "quiet." There is a point in the ε problem game (panel a) where all decisions are roughly equally likely, but there is no such point in the 0-1-2-3 game (panel b).

As shown previously, the ε problem game has a point, when both players are unbiased, where all four decision outcomes are just about equally likely. This does not occur in the 0-1-2-3 game, which makes sense—most players would perceive a difference between 0, 1, 2, and 3 years of prison. The other important difference is that unbiased players playing the 0-1-2-3 game are predicted by this model to favor the {quiet, quiet} block; the decision D11 dominates when c_{0a1} = c_{0b1} = 0.5. This result directly contradicts classical game theory, but it agrees with experimental results.[6]

4.2 Comparison to Experimental Results

Having shown how CGT addresses one of the key challenges with classical game theory, we now turn to the second: the discrepancy between what classical game theory predicts and experimental results. Two experimental studies and their rates of cooperation, which were higher than classical game theory would dictate, were described previously in this thesis; both can be modeled with CGT to calculate a predicted cooperation rate.

Both papers framed their games in terms of monetary payoffs. This type of game matrix could be analyzed by arbitrarily assuming, as with the calculations for the ε problem, that the number in dimensional units (such as dollars) equals the number of pain units; however, pains can also be determined from data by constructing a Weber-Fechner perception function.[15] Perception functions are important because the same dimensional amounts can mean different things to different players. For example, a fine of $200 would be huge to a broke college student, but to a millionaire it would be a drop in the bucket. The perception function constructed here is logarithmic in form because that is often how humans perceive stimuli—see sound loudness (decibel scale), earthquake strength (Richter scale), and acidity of solutions (pH scale). Other models could be used, but this one is common.

This perception function is based on a survey of the members of the CGT research group, which said: "A professor of sociology issues a request for participants in a study that will take one hour of time. What is the absolute minimum amount of money you would consider to do it (level 0), a small amount (1), a medium amount (2), a large amount (3), and a huge amount such that you would say, 'Wow, really?' (4)?" The data from this survey was used to construct the following perception function:

$$p = -(0.98 \pm 0.16)\ln\left(\frac{m}{\$3.06 \pm 1.76}\right) \quad (3)$$

where p is the pain caused to the player by each amount of money m. The statistics are from a 90% confidence interval. This equation follows the notation that pain is represented by a positive number and negative pain (pleasure) is negative. This is the reverse of the classical game theory convention, but it is done here to correspond with chemical systems, where negative Gibbs free energy values are favorable.

For each CGT solution, the average values in Eqn 3 were used to calculate pains (Gibbs free energy values) for Reactors A and B, which were then solved, along with Reactor D, with the method detailed previously. The cooperation rate for the game was calculated by dividing the total amount of "quiet" played (the sum of the change in a1 and b1 in Reactors A and B) by the total amount of any choice knowlecule played (the sum of the changes in a1, b1, a2, and b2 in Reactors A and B).

When the game from the study comparing economics majors and non-economics majors[7] was modeled in CGT with unbiased players, the cooperation rate was calculated to be 57.4%—much closer than classical game theory to the actual result of 52.7% for non-economics majors.

Similarly, the study with the game played at the U.S. Naval Academy[8] was modeled in CGT with unbiased players. This analysis resulted in a calculated cooperation rate of 54.5%, which is similar to the experimental cooperation rate of 54%. CGT appears to predict that for many games, approximating players as unbiased will result in a cooperation rate of roughly 50-60%, which falls in line with experimental data.

These "one-number" solutions were obtained by using only the average values from Eqn 3, but the statistics added with a 90% confidence interval can be used to calculate a range of solutions according to the Monte Carlo method. In a Monte Carlo simulation, random numbers for a variable of interest are generated from a probability distribution. The desired computation is then performed, giving a range of results.

Here, twenty sets of pains, representing twenty players, were generated for the economics majors game from a normal distribution. The twenty players were paired up to play ten games, each of which resulted in a cooperation rate, assuming $c_{0a1} = c_{0b1} = 0.5$ (unbiased). The mean of these cooperation rates was found to be (57.0 ± 0.5)%—again, similar to the experimental rate of 52.7% and the one-number CGT prediction of 57.4%. The one-number and Monte Carlo CGT predictions are not exactly the same, likely due to the effect of the "Flaw of Averages," where the average of a function is not equal to the function of the average.[16] However, the Monte Carlo results hold when checked for consistency; the average of the first five games is within a 90% confidence interval of the average of the second five games.

Monte Carlo analysis can also be applied to the USNA game, resulting in a cooperation rate of (54.0 ± 0.4)%. Even with pains randomly generated from the distribution in Eqn 3, the average cooperation rate in the games was close to the experimental and one-number CGT values. These results are in contrast to the classical game theory prediction, which is 0% cooperation. This is because classical game theory dictates what players "should" do, if they are rational; on the other hand, CGT only attempts to describe what players actually do.

4.3 Incorporation of Pre-Bias

It makes sense that not all players will come into a game with the same view of the situation. Prior experiences might make them wary of cooperation, or their system of beliefs might encourage helping others. They might simply have had a bad morning. All of these factors could affect the player's pre-bias. Classical game theory assumes a "rational" player, which does not offer much of a way to account for pre-bias. Chemical game theory, on the other hand, includes this concept explicitly, in the form of initial concentrations. Initial concentrations reflect a player's initial preference for "quiet" or "tell," and these preferences will affect the outcome of the game.

For an example, we return to the study of the economics majors and non-economics majors. While CGT's unbiased player approximation seems to predict the behavior of the non-economics majors well, the economics major cooperation rate of 28.2% is very different. Why is this the case? It may be that economics majors learn in their classes how the game "should" be played, and are therefore biased towards "tell." It is also possible that choosing to major in economics is correlated with having a personality type that prefers "tell" to "quiet." While there are many potential explanations for this phenomenon, it seems that one commonality is some type of pre-bias, which CGT can be used to analyze.

The pre-bias (or initial concentration) for both players was varied between $c_{0a1} = c_{0b1} = 0.1$ (10/90 quiet/tell) and $c_{0a1} = c_{0b1} = 0.9$ (90/10 quiet/tell). The cooperation rates this analysis produced are summarized in Figure 5. The experimental result for economics majors, 28.2%, falls close to the value for a pre-bias of 20% quiet and 80% tell.

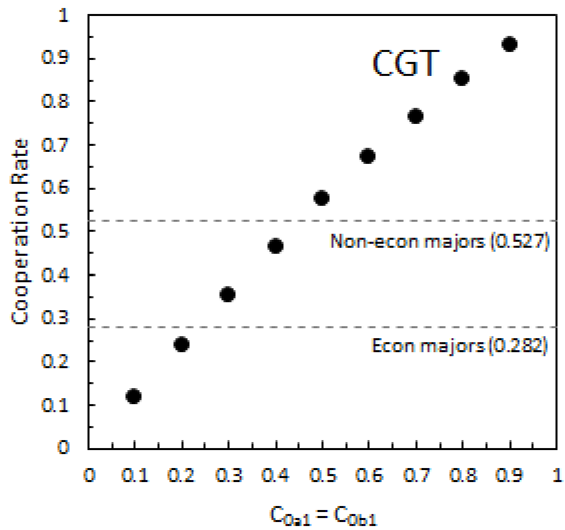

Figure 5. Cooperation rates for varying pre-bias values in the game played with economics majors and non-economics majors.

Both experimental values fall within the range of cooperation rates predicted by CGT. Figure from Velegol et al (Ref 12).

It is not possible to definitively call these economics majors 80% biased towards "tell," but it is worthwhile to note that both experimental values fall within the range of cooperation rates predicted by the CGT model; that is, CGT could be used to explain both, based on a difference in pre-bias.

Pre-bias is present in many experimental studies of Prisoner's Dilemma games, and it is often used to manipulate cooperation rate, which corresponds to CGT results. For example, when players were shown a smiling picture of the person they were supposedly playing against, cooperation increased.[17] While classical game theory does not have much of a method for analyzing that situation, chemical game theory can be readily applied to the problem by accounting for the players' pre-bias.

Chapter 5

Conclusions and Future Work

5.1 Conclusions

Classical game theory analysis of the Prisoner's Dilemma game predicts what players should do, if they are rational actors: they should always choose "tell," every time they play the game. However, there are three potential challenges with this result. One is that it does not consider the value of the pains; classical solutions see no difference between comparing \$5 to \$10 and \$2 to \$2.01. The second issue is that this analysis does not accurately describe experimental results: players do not usually play "tell" 100% of the time. Finally, there is no mechanism to account for pre-bias, the beliefs and preconceptions that players bring to the game.

There have been numerous attempts to develop new strategies to address these challenges, but none seem to include the effects of entropy, as Chemical Game Theory does. Entropy appears to be an important component in solving these games to predict not what players should do, but what they will do—just as chemistry does not prescribe how a reaction should proceed, but simply describes what happens in actuality.

With this inclusion of entropy, CGT can address the three key challenges to classical game theory. The CGT solution to the ε problem predicts that if all decision outcomes give very similar pains, a player who comes into the game unbiased will not perceive much difference between the possible choices. The result is that each outcome is predicted to occur about 25% of the time. This solution contradicts classical game theory, but it agrees with the common sense analysis of the ε problem.

CGT can also be used to model games that have been played in experimental situations. Two games from the literature were analyzed with CGT to produce both a "one-number" average cooperation rate and a normal distribution of cooperation rates from the Monte Carlo method. The CGT results were much closer to the experimental cooperation rate than classical game theory would predict, in both cases.

Finally, CGT includes an easy method for accounting for the pre-bias of players: initial concentrations. Some games may result in cooperation rates that are much higher or lower than the "unbiased player" approximation would expect. These games can be explained by a difference in pre-bias. A game from the literature was analyzed over a range of initial concentrations to determine if the experimental cooperation rates could be predicted by CGT. The cooperation rate of the non-economics majors was predicted fairly well by the unbiased player approximation, and the cooperation rate of the economics majors appeared to correlate with a pre-bias of 80% "tell." Classical game theory, which assumes a "rational player," does not have much of a mechanism for addressing this discrepancy, but chemical game theory does so readily.

More analysis is required to definitively test the validity of this new game-solving framework, but these initial results are promising. They show that CGT has the capacity to model and analyze a wide variety of decision-making situations, as well as addressing some of the challenges inherent in classical game theory. CGT, with its essential inclusion of entropy, has the potential to transform the process of framing games, solving them, and modeling the behavior of real human beings in contested decisions.

5.2 Future Work

The CGT model is currently being expanded into the areas of Tragedy of the Commons problems and agenda-setting. In addition, more research will need to be done in the future in order to define other chemical concepts, like temperature and pressure, in a game theory context. Temperature may be related to the words used in describing the game to the players; for example, framing the game as a cold "business transaction" has been shown to inhibit cooperation.[18] The research group also hopes to explore a kinetic model in addition to the thermodynamic method currently being used. These developments will enable CGT to be applied to even more diverse and complex decision-making situations.

Appendix A

Derivation of Free Energy Equation

We begin with the definition of the Gibbs free energy:

$$dG = -SdT + Vdp + \Sigma\mu_i dn_i \quad (1)$$

where G is the Gibbs free energy, S is entropy, T is temperature, p is pressure, V is volume, μ_i is the chemical potential of a given species i, and n_i is the moles of the species i. At constant temperature and pressure, dT = dp = 0, so that:

$$dG = \Sigma\mu_i dn_i \quad (2)$$

With stoichiometric coefficients ν_i for each chemical species in the reaction, the change in moles dn_i can be substituted with a change in extent of reaction ε:

$$dn_i = \nu_i d\varepsilon \quad (3)$$

The chemical potential can also be substituted with a function of the molecule's activity a_i:

$$\mu_i = \mu_i{}^0 + RT \ln a_i \quad (4)$$

The activity of a gas can be described by the following equation:

$$a_i = y_i \phi_i \frac{p}{p^0} \quad (5)$$

where y_i is the mole fraction of species i, ϕ_i is the fugacity coefficient of species i, p is the pressure, and p^0 is a standard pressure of reference. For ideal gases, the fugacity coefficient is equal to 1, which is a good approximation at low pressures. Combining equations (2)-(5) gives the following:

$$dG = \Sigma[(\mu_i{}^0 + RT \ln y_i \frac{p}{p^0})(\nu_i d\varepsilon)] \quad (6)$$

To minimize the Gibbs free energy, we find the point at which its derivative equals zero. After some simplification, this produces:

$$\frac{\partial G}{\partial \varepsilon} = 0 = \Sigma\mu_i{}^0\nu_i + RT \ln \frac{p}{p^0}\, \Sigma\nu_i + RT\, \Sigma\nu_i \ln y_i \quad (7)$$

Because the sum of the products of ${\mu_i}^0$ and ν_i is defined as the standard state Gibbs free energy change of reaction Δg^0, one final substitution can be made:

$$\Delta g^0 = -RT \ln \frac{p}{p^0}\, \Sigma\nu_i - RT\, \Sigma\nu_i \ln y_i \quad (8)$$

For the purposes of this research, game theory "pains" are equivalent to non-dimensionalized molar Gibbs free energy values (i.e. Eqn 8 divided by RT to produce values without units attached).

Appendix B

Details on Using Solver

This tutorial for the Solver add-in uses the ε problem game as an example.

The SPICEY table for Reactor A is shown in Figure 6. The players are both approximated as unbiased (with initial concentrations of the choice knowlecules all set to 0.5), and placeholder values have been entered for the extents of reaction ε_1 through ε_4. The spreadsheet has been set up with the equations from Table 4, as well as Eqn 2 from Chapter 3, which governs the relation of the mole fractions to the Gibbs free energy value.

	A	B	C	D	E	F	G	H	I	J
10										
11										
12										
13	Reactor A									
14	SP	I	C	E	Y				K	Δg0/RT
15	a1	0.5	-0.2	0.3	0.1875			Eqn 1	1.777778	-0.57536
16	b1	0.5	-0.2	0.3	0.1875			Eqn 2	1.777778	-0.57536
17	a2	0.5	-0.2	0.3	0.1875			Eqn 3	1.777778	-0.57536
18	b2	0.5	-0.2	0.3	0.1875			Eqn 4	1.777778	-0.57536
19	A11	0	0.1	0.1	0.0625					
20	A12	0	0.1	0.1	0.0625					
21	A21	0	0.1	0.1	0.0625					
22	A22	0	0.1	0.1	0.0625					
23				1.6	1					
24										
25	ε1	0.1								
26	ε2	0.1								
27	ε3	0.1								
28	ε4	0.1								

Figure 6. Excel spreadsheet showing the SPICEY table representing Reactor A in the ε problem game.

The equations from Table 4 have been used to set up this SPICEY table, with placeholder values for the extents of reaction.

Excel's Solver add-in has the ability to solve systems using linear, nonlinear, or evolutionary methods. It works by treating some cells as variables, which it can change, to achieve a desired outcome in another cell—a maximum, minimum, or certain explicitly defined value. Here, the desired outcome can be defined in several ways. One method is to subtract the calculated Gibbs free energy from the desired Gibbs free energy and set the difference to a minimum. Another method, which will be used in this example, is to set the energy cells directly to their desired pain values.

The parameters used to solve the spreadsheet above are shown in Figure 7. Solver will vary the variable cells (the extents of reaction, B25 - B28) to satisfy the objectives and constraints given. The goal is for the cells representing Player A's pains (J15 - J18) to equal the values given by the game's pain matrix: +1.99, +2.01, +1.98, and +2.00. (For simplicity's sake, I will assume that 2 years of prison is equivalent to 2 pain units.) To satisfy the laws of thermodynamics and physics, we have also constrained the extents of reaction to positive values, and the mole fractions of each species to be between 0 and 1.

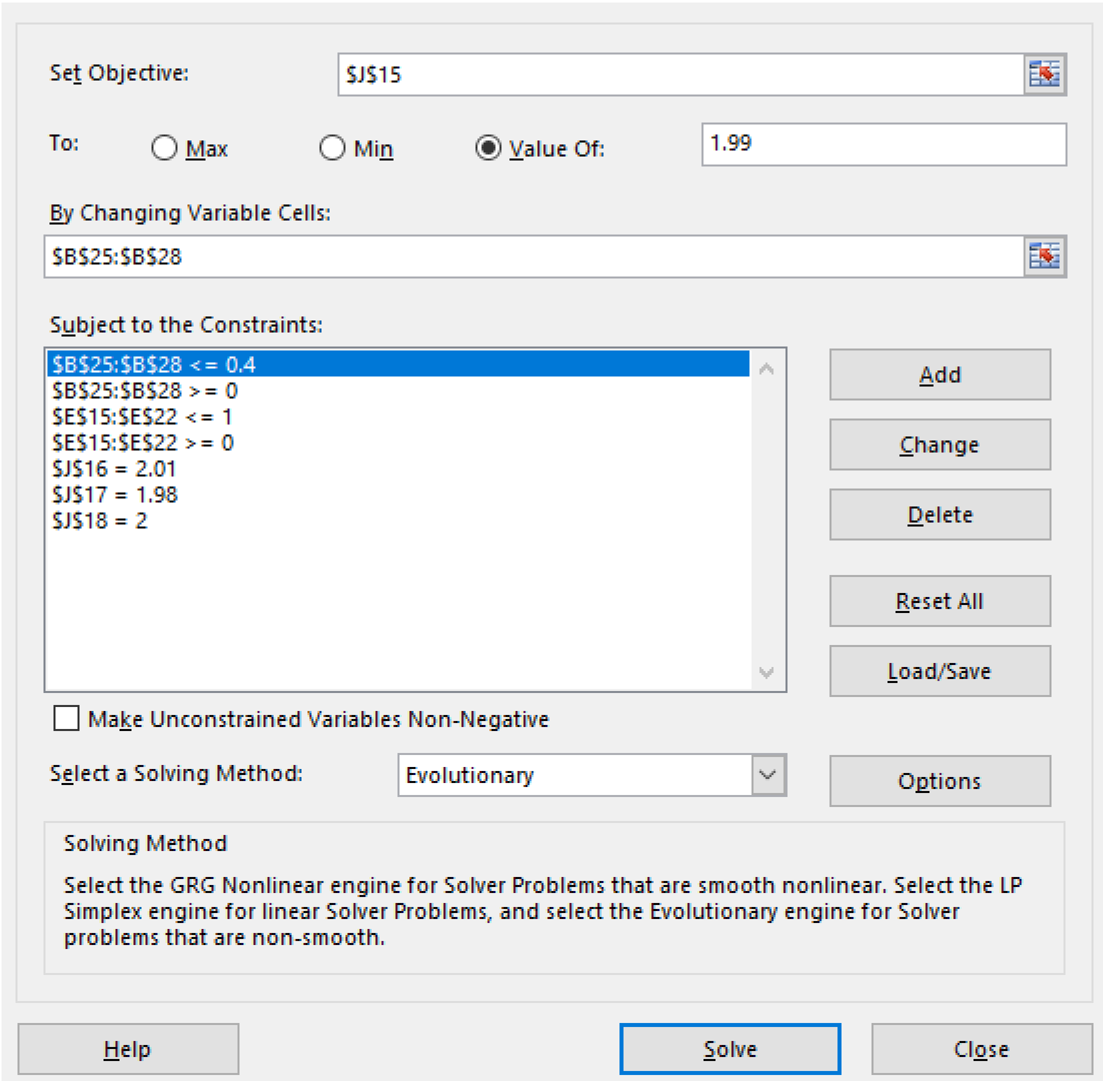

Figure 7. The Solver parameters for solving Reactor A of this game.

Solver will vary the extents of reaction to try to satisfy the constraints given: the pains should equal +1.99, +2.01, +1.98, and +2.00; the extents of reaction should not be negative; and the mole fractions should be between 0 and 1.

The solving method is “Evolutionary,” which means that Solver will perform many iterations as it attempts to get the system to converge. When you click “Solve,” Solver will run, and one of several outcomes can occur.

Solver may have solved the system, which is the ideal situation. If Solver cannot find a feasible solution, however, you can choose to keep its attempt and then adjust your placeholder values for the extents of reaction to make your “first guess” a bit closer to the actual solution. As with many iterative programs, it is easier to converge on a correct answer when the first guess is close to it.

Eventually, Solver will find a solution it cannot improve upon. This may look something like the outcome in Figure 8. The pain values are close to what we wanted them to be, but they are not exactly the same.

	A	B	C	D	E	F	G	H	I	J
10										
11										
12										
13	Reactor A									
14	SP	I	C	E	Y				K	Δg0/RT
15	a1	0.5	-0.03073	0.469268	0.242113			Eqn 1	0.136842	1.988929
16	b1	0.5	-0.03118	0.468816	0.24188			Eqn 2	0.133742	2.011839
17	a2	0.5	-0.03105	0.468952	0.24195			Eqn 3	0.137981	1.980637
18	b2	0.5	-0.0306	0.469404	0.242183			Eqn 4	0.135567	1.998291
19	A11	0	0.015532	0.015532	0.008014					
20	A12	0	0.0152	0.0152	0.007842					
21	A21	0	0.015651	0.015651	0.008075					
22	A22	0	0.015397	0.015397	0.007944					
23				1.93822	1					
24										
25	ε1	0.015532								
26	ε2	0.0152								
27	ε3	0.015651								
28	ε4	0.015397								

Figure 8. Results after running an evolutionary solution method that converged.
The pain values are close to our desired numbers, but they are not exact.

This can be remedied by running Solver again with the exact same parameters, changing only the solving method: from “Evolutionary” to “GRG Nonlinear.” Now that the pains are close enough to the desired answer, the nonlinear solving method can handle this system and give the solution for Reactor A, shown in Figure 9.

	A	B	C	D	E	F	G	H	I	J
10										
11										
12										
13	Reactor A									
14	SP	I	C	E	Y				K	Δg0/RT
15	a1	0.5	-0.03074	0.469257	0.242107			Eqn 1	0.136695	1.99
16	b1	0.5	-0.03118	0.468822	0.241882			Eqn 2	0.133989	2.01
17	a2	0.5	-0.03103	0.468967	0.241957			Eqn 3	0.138069	1.98
18	b2	0.5	-0.0306	0.469402	0.242181			Eqn 4	0.135335	2
19	A11	0	0.015516	0.015516	0.008005					
20	A12	0	0.015227	0.015227	0.007856					
21	A21	0	0.015662	0.015662	0.008081					
22	A22	0	0.015371	0.015371	0.00793					
23				1.938225	1					
24										
25	ε1	0.015516								
26	ε2	0.015227								
27	ε3	0.015662								
28	ε4	0.015371								

Figure 9. Results after running Solver a second time, with the nonlinear solving method.

The Gibbs free energy values are now exactly as determined by the game's pain matrix, meaning these extents of reaction are the desired answer.

Reactor B can be solved in a similar way, though now the pains are +1.99, +1.98, +2.01, and +2.00, in that order. Because Players A and B have the same pain values (but with the middle two flipped), the extents of reaction will be the same as in Reactor A, with the appropriate change in order.

Finally, Reactor D can also be solved by running Solver yet another time. However, for the Decider, the pains will be different than for Players A and B. We usually set them to -1.00 for each reaction in Reactor D, approximating a Decider who is unbiased with respect to the players, but who wants a decision to be made. As this reactor depends on the results from the previous two, it must be solved last. The final solution for the ε problem game, assuming unbiased players, is shown in Figure 10.

	A	B	C	D	E	F	G	H	I
48									
49	Reactor D								
50	SP	I	C	E	Y			K	Δg0/RT
51	A11	0.015516	-0.00355	0.011962	0.109296		Eqn 9	2.718282	-1
52	B11	0.015516	-0.00355	0.011962	0.109296		Eqn 10	2.718283	-1
53	A12	0.015227	-0.00353	0.011701	0.106909		Eqn 11	2.718283	-1
54	B12	0.015662	-0.00353	0.012135	0.110881		Eqn 12	2.718282	-1
55	A21	0.015662	-0.00353	0.012135	0.110881				
56	B21	0.015227	-0.00353	0.011701	0.106909				
57	A22	0.015371	-0.0035	0.011871	0.108465				
58	B22	0.015371	-0.0035	0.011871	0.108465				
59	D11	0	0.003554	0.003554	0.032472				
60	D12	0	0.003527	0.003527	0.032223				
61	D21	0	0.003527	0.003527	0.032223				
62	D22	0	0.0035	0.0035	0.031979				
63				0.109444	1				
64									
65	ε9	0.003554							
66	ε10	0.003527							
67	ε11	0.003527							
68	ε12	0.0035							

Figure 10. Final decision outcomes for the ε problem game with unbiased players.

The final concentrations (cells D59 - D62) for each decision are roughly equal, as expected.

The final concentrations of the four decision outcomes are all roughly 0.0035, indicating that unbiased players do not have much of a preference between the possible outcomes. The {quiet, quiet} decision is slightly favored, with 25.2% of the total decision concentration; {quiet, tell} and {tell, quiet} follow with 25.0%; and lastly {tell, tell} with 24.8%.

References

1. Von Neumann, J.; Morgenstern, O. *Theory of Games and Economic Behavior.* Princeton University Press, 2007.
2. Lange, A.; Vogt, C. “Cooperation in international environmental negotiations due to a preference for equity.” *J. Public Econ.* **2003**, *87*, 2049-2067.
3. Ostrom, E. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.
4. Ordeshook, P. *Game Theory and Political Theory: An Introduction*. Cambridge University Press, 1986.
5. Nash, J. “Non-cooperative Games.” *Annals of Mathematics* **1951**, *54* (2), 286-295.
6. Sally, D. “Conversation and Cooperation in Social Dilemmas: A Meta-Analysis of Experiments from 1958 to 1992.” *Ration. and Soc.* **1995**, *7*, 58-92.
7. Frank, R.; Gilovich, T.; Regan, D. “Does Studying Economics Inhibit Cooperation?” *J. Econ. Perspect.* **1993**, *7*, 159-171.
8. Swope, K.; Cadigan, J.; Schmitt, P.; Shupp, R. “Personality preferences in laboratory economics experiments.” *Journal of Socio-Economics* **2008**, *37*, 998-1009.
9. Simon, H. *Administrative Behavior*, 4th ed. The Free Press, a division of Simon & Schuster, 1997. Originally published 1945.
10. Andreoni, J. “Warm-Glow versus Cold-Prickle: The Effects of Positive and Negative Framing on Cooperation in Experiments.” *Q. J. Econ.* **1995**, *110*(1), 1-21.

11. Cooper, R.; DeJong, D.; Forsythe, R.; Ross, T. "Cooperation without Reputation: Experimental Evidence from Prisoner's Dilemma Games." *Games and Economic Behavior* **1996**, *12* (2), 187-218.

12. Velegol, D.; Suhey, P.; Connolly, J.; Morrissey, N.; Cook, L. "Chemical Game Theory." *Industrial & Engineering Chemical* Research. Submitted 2017.

13. Velegol, D. *Physics of Community Course Notes for Fall 2015*. Wild Scholars Media, 2015.

14. Denbigh, K. *The Principles of Chemical Equilibrium*, 4th ed. Cambridge University Press, 1981.

15. Gintis, Herbert. *The Bounds of Reason*. Revised ed. Princeton University Press, 2009.

16. Savage, S. *The Flaw of Averages: Why We Underestimate Risk in the Face of Uncertainty*. Wiley, 2009.

17. Scharlemann, J.; Eckel, C.; Kacelnik, A.; Wilson, R. "The value of a smile: Game theory with a human face." *J. Econ. Psychol.* **2001**, *22*(5), 617-640.

18. Batson, C.; Moran, T. "Empathy-induced altruism in a prisoner's dilemma." *Eur. J. Soc. Psychol.* **1999**, *29*, 909-924.

Academic Vita

Laura Cook

lcookchem@gmail.com

EDUCATION

08.2014 – 05.2018 **THE PENNSYLVANIA STATE UNIVERSITY**
BS Chemical Engineering, Schreyer Honors College

- Schreyer Scholar
- Dean's List, President Sparks Award, President's Freshman Award

EXPERIENCE

05.2017 – 08.2017 **MONDELĒZ INTERNATIONAL, INC.**
Manufacturing Intern – Continuous Improvement

- Used tools like Pareto charts, capability graphs, measurement system analyses, and root cause analyses to assess process
- Assisted in completion of 12-Step Kobetsu Kaizen and prepared core team to present it for IL6S Audit
- Investigated possible causes of unexplained speed loss using methodologies from the Focused Improvement Pillar

06.2014 – 08.2016 **WALMART STORES, INC.**
Cashier (Seasonal)

- Conducted monetary transactions such as purchases and utility payments
- Addressed customer complaints, questions, and requests

RESEARCH

09.2016 – PRESENT **VELEGOL LABORATORY, CHEMICAL GAME THEORY**
Undergraduate Research Assistant

- Performed research in Chemical Game Theory, a new field combining economics and chemical physics
- Compared CGT solutions of Prisoner's Dilemma games to experimental results to test new framework
- Assisted in writing of manuscript for submission to journals

SPRING 2017 **EXPERIMENTAL PHYSICAL CHEMISTRY**
Group Project on Photophysics of Pyranine in PVA Films

- Conducted experiments on pyranine in PVA films at varied pHs and temperatures to test emission, absorption, and kinetics
- Found evidence of solid-state proton transfer to OH^-